



Quartus II Handbook Version 10.1 Volume 1: Design and Synthesis



101 Innovation Drive
San Jose, CA 95134
www.altera.com

QI15V1-10.1.0

© 2010 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, HARDCOPY, MAX, MEGACORE, NIOS, QUARTUS and STRATIX are Reg. U.S. Pat. & Tm. Off. and/or trademarks of Altera Corporation in the U.S. and other countries. All other trademarks and service marks are the property of their respective holders as described at www.altera.com/common/legal.html. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.



Chapter Revision Dates	xxiii
-------------------------------------	-------

Section I. Design Flows

Chapter 1. Design Planning with the Quartus II Software

Creating Design Specifications	1-2
Intellectual Property Selection	1-2
System Design	1-2
Device Selection	1-3
Device Migration Planning	1-4
Planning for Device Programming or Configuration	1-4
Early Power Estimation	1-5
Early Pin Planning and I/O Analysis	1-6
Simultaneous Switching Noise Analysis	1-8
Selecting Third-Party EDA Tool Flows	1-9
Synthesis Tools	1-9
Simulation Tools	1-9
Formal Verification Tools	1-10
Planning for On-Chip Debugging Options	1-10
Design Practices and HDL Coding Styles	1-11
Design Recommendations	1-12
Recommended HDL Coding Styles	1-12
Managing Metastability	1-13
Planning for Hierarchical and Team-Based Design	1-13
Flat Compilation Flow with No Design Partitions	1-14
Incremental Compilation with Design Partitions	1-14
Planning Design Partitions and Floorplan Location Assignments	1-15
Fast Synthesis and Early Timing Estimation	1-16
Conclusion	1-16
Document Revision History	1-17

Chapter 2. Quartus II Incremental Compilation for Hierarchical and Team-Based Design

Deciding Whether to Use an Incremental Compilation Flow	2-1
Flat Compilation Flow with No Design Partitions	2-2
Incremental Capabilities Available When A Design Has No Partitions	2-2
Incremental Compilation Flow With Design Partitions	2-3
Team-Based Design Flows and IP Delivery	2-6
Incremental Compilation Summary	2-7
Steps for Incremental Compilation	2-8
Preparing a Design for Incremental Compilation	2-8
Compiling a Design Using Incremental Compilation	2-9
Creating Design Partitions	2-9
Creating Design Partitions in the Project Navigator	2-9
Creating Design Partitions in the Design Partitions Window	2-9
Creating Design Partitions With the Design Partition Planner	2-10
Creating Design Partitions With Tcl Scripting	2-10
Automatically-Generated Partitions	2-10
Common Design Scenarios Using Incremental Compilation	2-10

Reducing Compilation Time When Changing Source Files for One Partition	2-11
Optimizing a Timing-Critical Partition	2-11
Adding Design Logic Incrementally or Working With an Incomplete Design	2-12
Debugging Incrementally With the SignalTap II Logic Analyzer	2-13
Deciding Which Design Blocks Should Be Design Partitions	2-14
Impact of Design Partitions on Design Optimization	2-16
Design Partition Assignments Compared to Physical Placement Assignments	2-17
Using Partitions With Third-Party Synthesis Tools	2-17
Synopsys Synplify Pro/Premier and Mentor Graphics Precision RTL Plus	2-17
Other Synthesis Tools	2-18
Assessing Partition Quality	2-18
Partition Statistics Reports	2-18
Partition Timing Reports	2-19
Incremental Compilation Advisor	2-19
Specifying the Level of Results Preservation for	
Subsequent Compilations	2-21
Netlist Type for Design Partitions	2-21
Fitter Preservation Level for Design Partitions	2-22
Where Are the Netlist Databases Saved?	2-23
Deleting Netlists	2-23
What Changes Initiate a Partition's Automatic Resynthesis?	2-24
Resynthesis Due to Source Code Changes	2-25
Forcing Use of the Compilation Netlist When a Partition has Changed	2-26
Exporting Design Partitions from Separate Quartus II Projects	2-26
Preparing the Top-Level Design	2-27
Empty Partitions	2-28
Project Management—Making the Top-Level Design Available to Other Designers	2-28
Distributing the Top-Level Quartus II Project	2-28
Generating Design Partition Scripts	2-30
Exporting Partitions	2-31
Viewing the Contents of a Quartus II Exported Partition File (.qxp)	2-31
Integrating Partitions into the Top-Level Design	2-32
Integrating Assignments from the .qxp	2-32
Integrating Encrypted IP Cores from .qxp Files	2-33
Advanced Importing Options	2-33
Team-Based Design Optimization and Third-Party IP Delivery Scenarios	2-35
Using an Exported Partition to Send to a Design Without	
Including Source Files	2-35
Creating Precompiled Design Blocks (or Hard-Wired Macros) for Reuse	2-36
Designing in a Team-Based Environment	2-38
Enabling Designers on a Team to Optimize Independently	2-39
Resolving Assignment Conflicts During Integration	2-41
Importing a Partition to be Instantiated Multiple Times	2-42
Performing Design Iterations With Lower-Level Partitions	2-42
Creating a Design Floorplan With LogicLock Regions	2-44
Creating and Manipulating LogicLock Regions	2-45
Changing Partition Placement with LogicLock Changes	2-46
Taking Advantage of the Early Timing Estimator	2-46
Incremental Compilation Restrictions	2-47
When Timing Performance May Not Be Preserved Exactly	2-47
When Placement and Routing May Not Be Preserved Exactly	2-47
Using Incremental Compilation With Quartus II Archive Files	2-48
Limitations for HardCopy Compilation and Migration Flows	2-48
Formal Verification Support	2-49

SignalProbe Pins and Engineering Change Orders	2-49
SignalTap II Logic Analyzer in Exported Partitions	2-49
External Logic Analyzer Interface in Exported Partitions	2-50
Assignments Made in HDL Source Code in Exported Partitions	2-50
Design Partition Script Limitations	2-50
Warnings About Extra Clocks Due to Design Partition Scripts	2-50
Synopsys Design Constraint Files for the TimeQuest Timing Analyzer in Design Partition Scripts	2-51
Wildcard Support in Design Partition Scripts	2-51
Derived Clocks and PLLs in Design Partition Scripts	2-51
Pin Assignments for GXB and LVDS Blocks in Design Partition Scripts	2-52
Virtual Pin Timing Assignments in Design Partition Scripts	2-52
Top-Level Ports that Feed Multiple Lower-Level Pins in Design Partition Scripts	2-52
Restrictions on Megafunction Partitions	2-52
Register Packing and Partition Boundaries	2-53
I/O Register Packing	2-53
Scripting Support	2-54
Creating Design Partitions	2-54
Enabling or Disabling Design Partition Assignments During Compilation	2-55
Setting the Netlist Type	2-55
Setting the Fitter Preservation Level for a Post-fit or Imported Netlist	2-56
Preserving High-Speed Optimization	2-56
Specifying the Software Should Use the Specified Netlist and Ignore Source File Changes	2-56
Generating Design Partition Scripts	2-56
Exporting a Partition	2-57
Importing a Partition into the Top-Level Design	2-57
Makefiles	2-58
Scripting and Command-Line Application Examples	2-58
Reducing Opening a Project, Creating Design Partitions, and Performing an Initial Compilation	2-59
Reducing Compilation Time When Changing a Source File for One Partition—Command-Line Example	2-59
Optimizing the Placement for a Timing-Critical Partition	2-59
Conclusion	2-60
Document Revision History	2-61

Chapter 3. Quartus II Support for HardCopy Series Devices

HardCopy Series Design Benefits	3-1
Quartus II Features for HardCopy Planning	3-2
HardCopy Development Flow	3-2
Designing the FPGA First	3-3
Designing the HardCopy Device First	3-5
HardCopy Companion Device Selection	3-6
HardCopy Utilities	3-7
Companion Revisions	3-8
Compiling the HardCopy Companion Revision	3-9
Comparing HardCopy and FPGA Companion Revisions	3-9
Generating a HardCopy Handoff Report	3-9
Archiving HardCopy Handoff Files	3-10
HardCopy Advisor	3-10
HardCopy Device Resource Guide	3-11
HardCopy Recommended Settings in the Quartus II Software	3-12
Limit DSP and RAM to HardCopy Device Resources	3-12

Enabling Design Assistant to Run During Compile	3–12
Timing Settings	3–13
Constraints for Clock Effect Characteristics	3–13
Quartus II Software Features Supported for HardCopy Designs	3–14
Physical Synthesis Optimization	3–14
LogicLock Regions	3–15
PowerPlay Power Analyzer	3–15
Incremental Compilation	3–15
HardCopy Design Readiness Check	3–15
Turning the HardCopy Design Readiness Check On and Off	3–16
Setting Check	3–16
Summary	3–16
Global Setting	3–16
Instance Setting	3–16
Operating Setting	3–16
I/O Check	3–17
Input Pin Placement for Global and Regional Clock	3–17
PLL Usage Check	3–17
PLL Real-Time Reconfigurable Check	3–18
PLL Clock Outputs Driving Multiple Clock Network Types Check	3–18
PLL with No Compensation Mode Check	3–18
PLL with Normal or Source Synchronous Mode Feeding Output Pin Check	3–18
RAM Usage Check	3–18
Initialized Memory Dependency Testing	3–19
ALTGX Usage Check	3–20
Performing ECOs with Quartus II Engineering Change Management with the Chip Planner	3–20
Migrating One-to-One Changes	3–20
Migrating Changes that Must Be Implemented Differently	3–21
Changes that Cannot be Migrated	3–22
Overall Migration Flow	3–22
Preparing the Revisions	3–22
Applying ECO Changes	3–22
Formal Verification of FPGA and HardCopy Revisions	3–23
HardCopy Floorplan View	3–24
Document Revision History	3–25

Chapter 4. Quartus II Design Separation Flow

Design Flow Overview	4–2
Creating Design Partitions for the Design Separation Flow	4–4
Merging PLL Resources	4–5
Avoiding Multiple Design Partitions With a Secured Region	4–6
Creating a Design Floorplan with Secured Regions	4–6
Using Security Attributes	4–7
Using Secured Regions	4–9
Adding I/O Pins as Members of Secured Regions	4–9
Using Security Routing Interfaces	4–9
Making Design Separation Flow Location Assignments in the Chip Planner	4–10
Understanding Fencing Regions	4–11
Creating Non-Rectangular Regions	4–13
Guidelines for the Relative Placement of Secured LogicLock Regions	4–14
Creating a Complete Floorplan	4–14
Ensuring Routability Between Regions	4–16
Ensuring Planarity	4–17
Placing Physical Resources	4–19

Making Signal Security Assignments	4-19
Understanding Signal Names	4-20
Working with Global Signals	4-21
Assigning I/O Pins	4-25
Making Post Compilation Edits	4-26
Routing Restrictions	4-26
Number of Signals in Routing Interfaces	4-28
Application Example: Modifying a Fitter-Generated Floorplan for the Design Separation Flow ...	4-31
Report Panels	4-33
Secured LogicLock Region Summary	4-33
Security Routing Interfaces	4-34
Secured LogicLock Region Inputs and Outputs	4-34
Security I/O Bank Usage	4-35
Quartus Settings File Syntax	4-35
LL_SECURITY_ROUTING_INTERFACE	4-35
LL_REGION_SECURITY_LEVEL	4-35
LL_MEMBER_OF_SECURITY_ROUTING_INTERFACE	4-35
LL_SIGNAL_SECURITY_LEVEL	4-36
Document Revision History	4-36

Section II. System Design with Qsys

Chapter 5. Creating a System with Qsys

Qsys GUI	5-2
Qsys Component Library	5-3
Integrating Custom Components	5-3
Integrating Third-Party Components	5-4
Adding System Contents	5-4
Adding Components	5-4
Connecting Components	5-4
Filtering Components	5-5
Using the System Inspector	5-5
Defining the Address Map	5-6
Specifying Clock Settings	5-7
Specifying Project Settings	5-7
System Generation	5-8
Viewing the HDL Example	5-8
Qsys Design Flow	5-8
Generating Output Files	5-10
Simulating a Qsys System	5-11
Example Hierarchical System	5-12
Using Pipeline Bridges	5-16
Creating Hierarchical Components	5-16
Document Revision History	5-17

Chapter 6. Creating Qsys Components

Qsys Components	6-1
Component Providers	6-2
Component Interfaces	6-2
Component Types	6-2
Component Structure	6-3
Component Description File (_hw.tcl)	6-3
Component File Organization	6-4

Component Versioning	6-4
Component Search Path	6-4
Adding Components to the Library	6-5
Copy to the IP Root Directory	6-5
Reference Components in an .ipx File	6-6
Understanding IPX File Syntax	6-8
Component Editor	6-9
Component Hardware Structure	6-9
Starting the Component Editor	6-10
HDL Files Tab	6-10
Bottom-Up Component Design	6-10
Top-Down Component Design	6-11
Signals Tab	6-11
Naming Signals for Automatic Type and Interface Recognition	6-11
Templates for Interfaces to External Logic	6-12
Interfaces Tab	6-13
HDL Parameters Tab	6-13
Library Info	6-14
Saving a Component	6-15
Editing a Component	6-15
Registering Software Assignments	6-15
Component Parameterization	6-15
Document Revision History	6-16

Chapter 7. Qsys Interconnect

Avalon-MM Interface Components	7-2
Component Interconnect Domains	7-5
Using Two Separate Domains	7-6
Using One Domain with Width Adaptation	7-6
Qsys Transformations	7-7
Master Command and Slave Response Networks	7-7
Merlin Master Translator	7-8
Merlin Master Agent	7-8
Merlin Router	7-9
Merlin Traffic Limiter	7-9
Merlin Slave Translator	7-9
Merlin Slave Agent	7-10
Arbitration	7-10
Arbitration Examples	7-11
Merlin Arbiter	7-11
Interconnect Pipelining	7-13
Additional Qsys Interconnect Components	7-14
Clock Bridge	7-15
Avalon-MM Clock Crossing Bridge (Qsys)	7-15
Avalon-MM Pipeline Bridge (Qsys)	7-15
Merlin Width Adapter	7-16
Burst Transfers	7-17
Merlin Burst Adapter	7-17
Burst Types	7-18
Avalon-ST Interfaces	7-18
Avalon-ST Examples	7-18
Avalon-ST Components	7-19
Avalon-ST Handshake Clock Crosser	7-19
Avalon-ST Pipeline Stage	7-19

Merlin Multiplexer	7-20
Merlin Demultiplexer	7-20
Avalon-ST and Avalon-MM Interfaces	7-20
Tristate Conduit Components	7-21
Generic Tristate Controller	7-24
Tristate Conduit Pin Sharer	7-26
Tristate Conduit Bridge	7-26
Timing	7-27
Interrupt Interfaces	7-27
Assigning IRQs in Qsys	7-27
IRQ Bridge	7-28
Merlin IRQ Mapper	7-28
Merlin IRQ Clock Crosser	7-29
Clock Interfaces	7-29
Reset Interfaces	7-29
Single Global Reset Signal Implemented by Qsys	7-29
Multiple Reset Signals	7-29
Merlin Reset Controller	7-29
Reset Bridge	7-30
Conduits	7-30
Summary: Qsys Interconnect Components	7-30
Document Revision History	7-32

Chapter 8. Component Interface Tcl Reference

Information in a Hardware Component Description File	8-1
Component Phases	8-2
Writing a Hardware Component Description File	8-3
Providing Basic Information	8-3
Declaring Parameters	8-4
User Parameters	8-4
Derived Parameters	8-4
SYSTEM_INFO Parameters	8-4
Declaring Interfaces	8-5
Adding Files and Guiding Generation	8-5
Default Behaviors	8-6
Validation Phase Behavior	8-6
Elaboration Phase Behavior	8-6
Automatic Port Widths	8-6
Parameterized Parameter Widths	8-7
Generation Phase Behavior	8-7
Edit Phase Behavior	8-7
Overriding Default Behaviors	8-8
Validation Callback	8-9
Elaboration Callback	8-9
Generation Callback	8-10
Compose Callback	8-11
Editor Callback	8-13
Hardware Tcl Command Reference	8-14
Module Definition	8-17
package	8-17
get_module_properties	8-17
get_module_property	8-19
set_module_property	8-19
get_module_ports	8-20

get_module_assignments	8–20
get_module_assignment	8–20
set_module_assignment	8–21
get_files	8–21
add_file	8–21
add_documentation_link	8–22
get_file_properties	8–22
get_file_property	8–22
set_file_property	8–23
send_message	8–23
Parameters	8–24
add_parameter	8–24
get_parameters	8–25
get_parameter_properties	8–25
get_parameter_property	8–30
set_parameter_property	8–30
get_parameter_value	8–31
set_parameter_value	8–31
decode_address_map	8–32
Display Items	8–33
add_display_item	8–33
get_display_items	8–34
get_display_item_properties	8–35
get_display_item_property	8–35
set_display_item_property	8–35
Interfaces and Ports	8–36
add_interface	8–36
get_interfaces	8–37
get_interface_properties	8–38
get_interface_property	8–38
set_interface_property	8–39
add_interface_port	8–39
get_interface_ports	8–40
get_port_properties	8–40
get_port_property	8–41
set_port_property	8–42
get_interface_assignments	8–42
get_interface_assignment	8–42
set_interface_assignment	8–43
Compose	8–43
add_instance	8–43
get_instances	8–44
get_instance_parameters	8–44
set_instance_parameter_value	8–44
get_instance_parameter_value	8–44
get_instance_parameter_properties	8–45
get_instance_parameter_property	8–45
get_instance_interfaces	8–46
get_instance_interface_properties	8–46
get_instance_interface_property	8–46
get_instance_interface_ports	8–47
get_instance_port_property	8–47
add_connection	8–47
get_connections	8–48

get_connection_parameters	8-48
get_connection_parameter_value	8-49
set_connection_parameter_value	8-49
Generation	8-49
get_generation_properties	8-49
get_generation_property	8-50
Document Revision History	8-50

Section III. Design Guidelines

Chapter 9. Recommended Design Practices

Synchronous FPGA Design Practices	9-2
Fundamentals of Synchronous Design	9-2
Hazards of Asynchronous Design	9-3
Design Guidelines	9-4
Combinational Logic Structures	9-4
Combinational Loops	9-4
Latches	9-5
Delay Chains	9-5
Pulse Generators and Multivibrators	9-6
Clocking Schemes	9-7
Internally Generated Clocks	9-8
Divided Clocks	9-8
Ripple Counters	9-8
Multiplexed Clocks	9-9
Gated Clocks	9-10
Synchronous Clock Enables	9-11
Recommended Clock-Gating Methods	9-11
Power Optimization	9-12
Metastability	9-13
Incremental Compilation	9-13
Checking Design Violations With the Design Assistant	9-13
Quartus II Design Flow with the Design Assistant	9-14
Enabling and Disabling Design Assistant Rules	9-15
Viewing Design Assistant Results	9-15
Custom Rules	9-15
Custom Rules Coding Examples	9-16
Targeting Clock and Register-Control Architectural Features	9-19
Clock Network Resources	9-20
Reset Resources	9-21
Synchronous Reset	9-21
Asynchronous Reset	9-21
Synchronized Asynchronous Reset	9-22
Register Control Signals	9-24
Targeting Embedded RAM Architectural Features	9-24
Conclusion	9-25
Document Revision History	9-26

Chapter 10. Recommended HDL Coding Styles

Quartus II Language Templates	10-1
Using Altera Megafunctions	10-2
Instantiating Altera Megafunctions in HDL Code	10-3
Instantiating Megafunctions Using the MegaWizard Plug-In Manager	10-3

Creating a Netlist File for Other Synthesis Tools	10-4
Instantiating Megafunctions Using the Port and Parameter Definition	10-4
Inferring Multiplier and DSP Functions from HDL Code	10-5
Inferring Multipliers from HDL Code	10-5
Inferring Multiply-Accumulators and Multiply-Adders from HDL Code	10-8
Inferring Memory Functions from HDL Code	10-13
Inferring RAM functions from HDL Code	10-13
Use Synchronous Memory Blocks	10-14
Avoid Unsupported Reset and Control Conditions	10-14
Check Read-During-Write Behavior	10-16
Controlling Inference and Implementation in Device RAM Blocks	10-18
Single-Clock Synchronous RAM with Old Data Read-During-Write Behavior	10-18
Single-Clock Synchronous RAM with New Data Read-During-Write Behavior	10-20
Simple Dual-Port, Dual-Clock Synchronous RAM	10-22
True Dual-Port Synchronous RAM	10-24
Mixed-Width Dual-Port RAM	10-28
RAM with Byte-Enable Signals	10-31
Specifying Initial Memory Contents at Power-Up	10-33
Inferring ROM Functions from HDL Code	10-36
Shift Registers—Inferring the ALTSHIFT_TAPS Megafunction from HDL Code	10-40
Simple Shift Register	10-41
Shift Register with Evenly Spaced Taps	10-42
Coding Guidelines for Registers and Latches	10-43
Register Power-Up Values in Altera Devices	10-43
Secondary Register Control Signals Such as Clear and Clock Enable	10-45
Latches	10-49
Unintentional Latch Generation	10-49
Inferring Latches Correctly	10-50
General Coding Guidelines	10-53
Tri-State Signals	10-54
Clock Multiplexing	10-54
Adder Trees	10-58
Architectures with 4-Input LUTs in Logic Elements	10-58
Architectures with 6-Input LUTs in Adaptive Logic Modules	10-59
State Machines	10-60
Verilog HDL State Machines	10-61
VHDL State Machines	10-65
Multiplexers	10-67
Quartus II Software Option for Multiplexer Restructuring	10-67
Multiplexer Types	10-67
Implicit Defaults in If Statements	10-69
Default or Others Case Assignment	10-69
Cyclic Redundancy Check Functions	10-70
If Performance is Important, Optimize for Speed	10-70
Use Separate CRC Blocks Instead of Cascaded Stages	10-70
Use Separate CRC Blocks Instead of Allowing Blocks to Merge	10-71
Take Advantage of Latency if Available	10-71
Save Power by Disabling CRC Blocks When Not in Use	10-71
Use the Device Synchronous Load (sload) Signal to Initialize	10-72
Comparators	10-72
Counters	10-73
Designing with Low-Level Primitives	10-73
Conclusion	10-74
Document Revision History	10-74

Chapter 11. Managing Metastability with the Quartus II Software

Introduction	11-1
Metastability Analysis in the Quartus II Software	11-2
Synchronization Register Chains	11-2
Identifying Synchronizers for Metastability Analysis	11-4
How Timing Constraints Affect Synchronizer Identification and Metastability Analysis	11-4
Metastability and MTBF Reporting	11-5
Metastability Reports	11-5
MTBF Summary Report	11-5
Synchronizer Summary Report	11-6
Synchronizer Chain Statistics Report in the Timing Analyzer	11-7
Synchronizer Data Toggle Rate in MTBF Calculation	11-7
MTBF Optimization	11-8
Synchronization Register Chain Length	11-8
Reducing Metastability Effects	11-9
Apply Complete System-Centric Timing Constraints for the Timing Analyzer	11-9
Force the Identification of Synchronization Registers	11-9
Set the Synchronizer Data Toggle Rate	11-10
Optimize Metastability During Fitting	11-10
Increase the Length of Synchronizers to Protect and Optimize	11-10
Set Fitter Effort to Standard Fit instead of Auto Fit	11-10
Increase the Number of Stages Used in Synchronizers, If Possible	11-10
Select a Faster Speed Grade Device, if Possible	11-11
Scripting Support	11-11
Identifying Synchronizers for Metastability Analysis	11-11
Synchronizer Data Toggle Rate in MTBF Calculation	11-12
report_metastability and Tcl Command	11-12
MTBF Optimization	11-12
Synchronization Register Chain Length	11-12
Conclusion	11-13
Document Revision History	11-13

Chapter 12. Best Practices for Incremental Compilation Partitions and Floorplan Assignments

Overview: Incremental Compilation	12-2
Recommendations for the Netlist Type	12-2
Design Flows Using Incremental Compilation	12-3
Project Management in Team-Based Design Flows	12-4
Why to Plan Partitions and Floorplan Assignments	12-5
Partition Boundaries and Optimization	12-6
General Partitioning Guidelines	12-7
Plan Design Hierarchy and Source Design Files	12-8
Using Partitions with Third-Party Synthesis Tools	12-8
Partition Design by Functionality and Block Size	12-9
Partition Design by Clock Domain and Timing Criticality	12-9
Consider What Is Changing	12-9
Design Partition Guidelines	12-10
Register Partition Inputs and Outputs	12-10
Minimize Cross-Partition-Boundary I/O	12-11
Avoid the Need for Logic Optimization Across Partitions	12-12
Keep Logic in the Same Partition for Optimization and Merging	12-13
Keep Constants in the Same Partition as Logic	12-14
Avoid Unconnected Partition I/O	12-15
Avoid Signals That Drive Multiple Partition I/O or Connect I/O Together	12-16

Invert Clocks in Destination Partitions	12-17
Connect I/O Pin Directly to I/O Register for Packing Across Partition Boundaries	12-18
Do Not Use Internal Tri-States	12-22
Include All Tri-State and Enable Logic in the Same Partition	12-23
Include Bidirectional I/O Registers in the Same Partition (For Older Device Families)	12-24
Summary of Guidelines Related to Logic Optimization Across Partitions	12-24
Consider a Cascaded Reset Structure	12-25
Design Partition Guidelines for Third-Party IP Delivery	12-26
Allocate Logic Resources	12-27
Allocate Global Routing Signals and Clock Networks if Required	12-28
Assign Virtual Pins	12-29
Perform Timing Budgeting if Required	12-30
Drive Clocks Directly	12-30
Recreate PLLs for Lower-Level Partitions if Required	12-31
Checking Partition Quality	12-31
Incremental Compilation Advisor	12-32
Design Partition Planner	12-32
Viewing Design Partition Planner and Floorplan Side-by-Side	12-34
Partition Statistics Report	12-35
Report Partition Timing in the TimeQuest Timing Analyzer	12-36
Check if Partition Assignments Impact the Quality of Results	12-36
Including SDC Constraints from Lower-Level Partitions for Third-Party IP Delivery	12-37
Creating an .sdc File With Project-Wide Constraints	12-38
Creating an .sdc with Partition-Specific Constraints	12-39
Consolidating the .sdc in the Top-Level Design	12-40
Introduction to Design Floorplans	12-41
The Difference between Logical Partitions and Physical Regions	12-41
Why Create a Floorplan?	12-42
When to Create a Floorplan	12-44
Early Floorplan	12-44
Late Floorplan	12-44
Design Floorplan Placement Guidelines	12-44
Assigning Partitions to LogicLock Regions	12-45
How to Size and Place Regions	12-46
Modifying Region Size and Origin	12-46
I/O Connections	12-47
LogicLock Resource Exclusions	12-48
Creating Non-Rectangular Regions	12-50
Checking Floorplan Quality	12-50
Incremental Compilation Advisor	12-50
LogicLock Region Resource Estimates	12-50
LogicLock Region Properties Statistics Report	12-50
Locate the Quartus II TimeQuest Timing Analyzer Path in the Chip Planner	12-51
Inter-Region Connection Bundles	12-51
Routing Utilization	12-51
Ensure Floorplan Assignments Do Not Significantly Impact Quality of Results	12-51
Recommended Design Flows and Application Examples	12-52
Create a Floorplan for Major Design Blocks	12-52
Create a Floorplan Assignment for One Design Block with Difficult Timing	12-53
Create a Floorplan as the Project Lead in a Team-Based Flow	12-53
Conclusion	12-54
Document Revision History	12-55

Section IV. Synthesis

Chapter 13. Quartus II Integrated Synthesis

Design Flow	13-1
Language Support	13-3
Verilog HDL Support	13-4
SystemVerilog Support	13-5
Initial Constructs and Memory System Tasks	13-7
Verilog HDL Macros	13-8
VHDL Support	13-8
VHDL-2008 Support	13-10
VHDL Standard Libraries and Packages	13-10
VHDL wait Constructs	13-10
AHDL Support	13-11
Schematic Design Entry Support	13-11
State Machine Editor	13-11
Design Libraries	13-12
Specifying a Destination Library Name in the Settings Dialog Box	13-13
Specifying a Destination Library Name in the Quartus II Settings File or Using Tcl	13-13
Specifying a Destination Library Name in a VHDL File	13-13
Mapping a VHDL Instance to an Entity in a Specific Library	13-14
Using Parameters/Generics	13-16
Setting Default Parameter Values and BDF Instance Parameter Values	13-16
Passing Parameters Between Two Design Languages	13-18
Incremental Compilation	13-20
Partitions for Preserving Hierarchical Boundaries	13-20
Parallel Synthesis	13-21
Quartus II Exported Partition File as Source	13-22
Quartus II Synthesis Options	13-22
Setting Synthesis Options	13-24
Analysis & Synthesis Settings Page of the Settings Dialog Box	13-24
Quartus II Logic Options	13-24
Synthesis Attributes	13-25
Synthesis Directives	13-27
Optimization Technique	13-28
Auto Gated Clock Conversion	13-28
Timing-Driven Synthesis	13-30
SDC Constraint Protection	13-31
PowerPlay Power Optimization	13-31
Limiting Resource Usage in Partitions	13-32
Creating LogicLock Regions	13-32
Using Assignments to Limit the Number of RAM and DSP Blocks	13-33
Restructure Multiplexers	13-33
Synthesis Effort	13-35
Synthesis Seed	13-35
State Machine Processing	13-35
Manually Specifying State Assignments Using the <code>syn_encoding</code> Attribute	13-37
Manually Specifying Enumerated Types Using the <code>enum_encoding</code> Attribute	13-38
Safe State Machines	13-39
Power-Up Level	13-41
Inferred Power-Up Levels	13-41
Power-Up Don't Care	13-42
Remove Duplicate Registers	13-42

Preserve Registers	13–43
Disable Register Merging/Don't Merge Register	13–43
Noprune Synthesis Attribute/Preserve Fan-out Free Register Node	13–44
Keep Combinational Node/Implement as Output of Logic Cell	13–45
Disabling Synthesis Netlist Optimizations with dont_retime Attribute	13–46
Disabling Synthesis Netlist Optimizations with dont_replicate Attribute	13–47
Maximum Fan-Out	13–48
Controlling Clock Enable Signals with Auto Clock Enable Replacement and direct_enable	13–49
Inferring Multiplier, DSP, and Memory Functions from HDL Code	13–50
Multiply-Accumulators and Multiply-Adders	13–51
Shift Registers	13–51
RAM and ROM	13–51
Resource Aware RAM, ROM, and Shift-Register Inference	13–52
Auto RAM to Logic Cell Conversion	13–53
RAM Style and ROM Style—for Inferred Memory	13–53
Turning Off the Add Pass-Through Logic to Inferred RAMs no_rw_check Attribute	13–55
RAM Initialization File—for Inferred Memory	13–59
Multiplier Style—for Inferred Multipliers	13–59
Full Case Attribute	13–61
Parallel Case	13–62
Translate Off and On / Synthesis Off and On	13–63
Ignore translate_off and synthesis_off Directives	13–64
Read Comments as HDL	13–65
Use I/O Flipflops	13–66
Specifying Pin Locations with chip_pin	13–67
Using altera_attribute to Set Quartus II Logic Options	13–69
Analyzing Synthesis Results	13–72
Analysis & Synthesis Section of the Compilation Report	13–72
Project Navigator	13–72
Analyzing and Controlling Synthesis Messages	13–72
Quartus II Messages	13–73
VHDL and Verilog HDL Messages	13–73
Setting the HDL Message Level	13–75
Enabling or Disabling Specific HDL Messages by Module/Entity	13–76
Node-Naming Conventions in Quartus II Integrated Synthesis	13–77
Hierarchical Node-Naming Conventions	13–77
Node-Naming Conventions for Registers (DFF or D Flipflop Atoms)	13–78
Register Changes During Synthesis	13–79
Synthesis and Fitting Optimizations	13–79
State Machines	13–80
Inferred Adder-Subtractors, Shift Registers, Memory, and DSP Functions	13–80
Packed Input and Output Registers of RAM and DSP Blocks	13–81
Preserving Register Names	13–81
Node-Naming Conventions for Combinational Logic Cells	13–81
Preserving Combinational Logic Names	13–82
Scripting Support	13–83
Adding an HDL File to a Project and Setting the HDL Version	13–84
Assigning a Pin	13–86
Creating Design Partitions for Incremental Compilation	13–86
Quartus II Synthesis Options	13–87
Conclusion	13–89
Document Revision History	13–90

Chapter 14. Synopsys Synplify Support

Altera Device Family Support	14-1
Design Flow	14-2
Specifying the Output Netlist File Name and Result Format	14-5
Specifying the Quartus II Software Version	14-5
Synplify Optimization Strategies	14-6
Using Synplify Premier to Optimize Your Design	14-6
Using Implementations in Synplify Pro or Premier	14-7
Timing-Driven Synthesis Settings	14-7
Clock Frequencies	14-7
Multiple Clock Domains	14-8
Input and Output Delays	14-8
Multicycle Paths	14-8
False Paths	14-8
FSM Compiler	14-9
FSM Explorer in Synplify Pro and Premier	14-9
Optimization Attributes and Options	14-10
Retiming in Synplify Pro and Premier	14-10
Maximum Fan-Out	14-10
Preserving Nets	14-10
Register Packing	14-10
Resource Sharing	14-10
Preserving Hierarchy	14-11
Register Input and Output Delays	14-11
syn_direct_enable	14-12
I/O Standard	14-12
Altera-Specific Attributes	14-12
altera_chip_pin_lc	14-12
altera_io_powerup	14-13
altera_io_opendrain	14-13
Exporting Designs to the Quartus II Software Using NativeLink Integration	14-13
Running the Quartus II Software from within the Synplify Software	14-14
Using the Quartus II Software to Run the Synplify Software	14-15
Running the Quartus II Software Manually With the Synplify-Generated Tcl Script	14-15
Passing TimeQuest SDC Timing Constraints to the Quartus II Software	14-15
Individual Clocks and Frequencies	14-16
Input and Output Delay	14-16
Multicycle Path	14-16
False Path	14-16
Guidelines for Altera Megafunctions and Architecture-Specific Features	14-16
Instantiating Altera Megafunctions With the MegaWizard Plug-In Manager	14-17
Instantiating Megafunctions With MegaWizard Plug-In Manager-Generated Verilog HDL Files	14-18
Instantiating Megafunctions with MegaWizard Plug-In Manager-Generated VHDL Files ..	14-18
Changing Synplify's Default Behavior for Instantiated Altera Megafunctions	14-18
Instantiating Intellectual Property With the MegaWizard Plug-In Manager and IP Toolbench	14-19
Instantiating Black Box IP Functions With Generated Verilog HDL Files	14-20
Instantiating Black Box IP Functions With Generated VHDL Files	14-20
Other Synplify Software Attributes for Creating Black Boxes	14-21
Including Files for Quartus II Placement and Routing Only	14-22
Inferring Altera Megafunctions from HDL Code	14-22
Inferring Multipliers	14-23
Inferring RAM	14-25

RAM Initialization	14-27
Inferring ROM	14-28
Inferring Shift Registers	14-28
Incremental Compilation and Block-Based Design	14-29
Creating a Design with Separate Netlist Files for Incremental Compilation	14-30
Using MultiPoint Synthesis with Incremental Compilation	14-31
Set Compile Points and Create Constraint Files	14-31
Additional Considerations for Compile Points	14-33
Creating a Quartus II Project for Compile Points and Multiple .vqm Files	14-33
Creating Multiple .vqm Files for a Incremental Compilation Flow With Separate Synplify Projects	14-35
Manually Creating Multiple .vqm Files With Black Boxes	14-35
Creating a Quartus II Project for Multiple .vqm Files	14-39
Performing Incremental Compilation in the Quartus II Software	14-40
Conclusion	14-41
Document Revision History	14-41

Chapter 15. Mentor Graphics Precision Synthesis Support

Altera Device Family Support	15-1
Design Flow	15-2
Creating and Compiling a Project in the Precision Synthesis Software	15-4
Mapping the Precision Synthesis Design	15-5
Setting Timing Constraints	15-6
Setting Mapping Constraints	15-6
Assigning Pin Numbers and I/O Settings	15-6
Assigning I/O Registers	15-8
Disabling I/O Pad Insertion	15-8
Preventing the Precision Synthesis Software from Adding I/O Pads	15-8
Preventing the Precision Synthesis Software from Adding an I/O Pad on an Individual Pin	15-9
Controlling Fan-Out on Data Nets	15-9
Synthesizing the Design and Evaluating the Results	15-9
Obtaining Accurate Logic Utilization and Timing Analysis Reports	15-10
Exporting Designs to the Quartus II Software Using NativeLink Integration	15-10
Running the Quartus II Software from within the Precision Synthesis Software	15-10
Running the Quartus II Software Manually Using the Precision Synthesis-Generated Tcl Script	15-11
Using the Quartus II Software to Run the Precision Synthesis Software	15-12
Passing Constraints to the Quartus II Software	15-12
create_clock	15-12
set_input_delay	15-13
set_output_delay	15-13
set_max_delay and set_min_delay	15-14
set_false_path	15-14
set_multicycle_path	15-15
Guidelines for Altera Megafunctions and Architecture-Specific Features	15-15
Instantiating Altera Megafunctions Using the MegaWizard Plug-In Manager	15-16
Instantiating Megafunctions With MegaWizard Plug-In Manager-Generated Verilog HDL Files	15-16
Instantiating Megafunctions With MegaWizard Plug-In Manager-Generated VHDL Files ..	15-17
Instantiating Intellectual Property With the MegaWizard Plug-In Manager and IP Toolbench	15-17
Instantiating Black Box IP Functions With Generated Verilog HDL Files	15-18
Instantiating Black Box IP Functions With Generated VHDL Files	15-18

Inferring Altera Megafunctions from HDL Code	15–19
Multipliers	15–19
Setting the Use Dedicated Multiplier Option	15–20
Setting the dedicated_mult Attribute	15–20
Multiplier-Accumulators and Multiplier-Adders	15–21
Controlling DSP Block Inference	15–22
RAM and ROM	15–24
Incremental Compilation and Block-Based Design	15–24
Creating a Design with Precision RTL Plus Incremental Synthesis	15–24
Creating Partitions with the incr_partition Attribute	15–25
Creating Multiple Mapped Netlist Files With Separate Precision Projects or Implementations	15–26
Creating Black Boxes to Create EDIF Netlists	15–28
Creating Black Boxes in Verilog HDL	15–28
Creating Black Boxes in VHDL	15–29
Creating Quartus II Projects for Multiple EDIF Files	15–30
Creating a Single Quartus II Project for a Standard Incremental Compilation Flow	15–31
Creating Multiple Quartus II Projects for a Bottom-Up Flow	15–32
Hierarchy and Design Considerations	15–32
Conclusion	15–33
Document Revision History	15–33

Chapter 16. Mentor Graphics LeonardoSpectrum Support

Altera Device Family Support	16–1
Design Flow	16–2
LeonardoSpectrum Optimization Strategies	16–4
Timing-Driven Synthesis	16–4
Global PowerTab	16–4
Clock PowerTab	16–5
Input and Output PowerTabs	16–5
Other Constraints	16–5
Encoding Style	16–5
Resource Sharing	16–6
Mapping I/O Registers	16–6
Timing Analysis with the LeonardoSpectrum Software	16–6
Exporting Designs Using NativeLink Integration	16–7
Generating Netlist Files	16–7
Including Design Files for Black Boxed Modules	16–7
Passing Constraints with Scripts	16–8
Integration with the Quartus II Software	16–8
Guidelines for Altera Megafunctions and LPM Functions	16–8
Instantiating Altera Megafunctions	16–9
Inferring Altera Memory Elements	16–9
Inferring Multipliers and DSP Functions	16–10
Simple Multipliers	16–10
Multiplier Accumulators	16–10
Multiplier Adders	16–11
Controlling DSP Block Inference	16–11
Global Attribute	16–11
Module Level Attributes	16–12
Signal Level Attributes	16–13
Guidelines for Using DSP Blocks	16–15
Block-Based Design with the Quartus II Software	16–16
Hierarchy and Design Considerations	16–16
Creating a Design with Multiple .edf Files	16–17

Generating Multiple .edf Files Using the LogicLock Option	16–17
Creating a Quartus II Project for Multiple .edf Files Including LogicLock Regions	16–19
Generating Multiple .edf Files Using Black Boxes	16–20
Black Box Methodology in Verilog HDL	16–21
Black Boxing in VHDL	16–22
Creating a Quartus II Project for Multiple .edf Files	16–24
Incremental Synthesis Flow	16–25
Modifications Required for the LogicLock_Incremental.tcl Script File	16–25
Running the Tcl Script File in LeonardoSpectrum	16–26
Conclusion	16–27
Document Revision History	16–27

Chapter 17. Analyzing Designs with Quartus II Netlist Viewers

When to Use the Netlist Viewers: Analyzing Design Problems	17–1
Quartus II Design Flow with the Netlist Viewers	17–2
RTL Viewer Overview	17–4
State Machine Viewer Overview	17–5
Technology Map Viewer Overview	17–5
Introduction to the User Interface	17–6
Schematic View	17–7
Schematic Symbols	17–7
Selecting an Item in the Schematic View	17–14
Moving and Panning in the Schematic View	17–15
Netlist Navigator Pane	17–15
State Machine Viewer	17–16
State Diagram View	17–17
State Transition Table	17–18
State Encoding Table	17–18
Selecting an Item in the State Machine Viewer	17–18
Switching Between State Machines	17–18
Global Options	17–18
Display Settings	17–19
Tracing	17–20
Customize View	17–21
Shortcut Commands	17–22
Navigating the Schematic View	17–22
Traversing and Viewing the Design Hierarchy	17–22
Flattening the Design Hierarchy	17–22
Viewing the Contents of a Design Hierarchy in the Current Schematic	17–22
Viewing Contents of Atom Primitives	17–23
Viewing the Properties of Instances and Primitives	17–24
Viewing LUT Representations in the Technology Map Viewer	17–24
Grouping Combinational Logic into Logic Clouds	17–26
Logic Clouds in the RTL Viewer	17–26
Logic Clouds in the Technology Map Viewer	17–27
Grouping and Ungrouping Logic Clouds	17–28
Changing the Constant Signal Value Formatting	17–28
Zooming and Magnification	17–28
Schematic Debugging and Tracing Using the Bird’s Eye View	17–29
Partitioning the Schematic into Pages	17–30
Moving Between Schematic Pages	17–31
Moving Back and Forward Through Schematic Pages	17–31
Following Nets Across Schematic Pages	17–31
Go to Net Driver	17–32

Filtering in the Schematic View	17-33
Filter Sources Command	17-33
Filter Destinations Command	17-34
Filter Sources and Destinations Command	17-34
Filter Between Selected Nodes Command	17-35
Filter Selected Nodes and Nets Command	17-36
Filter Bus Index Command	17-37
Filter Command Processing	17-37
Filtering Across Hierarchies	17-37
Expanding a Filtered Netlist	17-38
Reducing a Filtered Netlist	17-39
Probing to a Source Design File and Other Quartus II Windows	17-39
Moving Selected Nodes to Other Quartus II Windows	17-40
Probing to the Netlist Viewers from Other Quartus II Windows	17-40
Viewing a Timing Path	17-41
Other Features in the Schematic Viewer	17-42
Tooltips	17-42
Finding Design Elements in the Netlist Viewers	17-44
Exporting and Copying a Schematic Image	17-45
Printing	17-45
Conclusion	17-46
Document Revision History	17-46

Additional Information

How to Contact Altera	Info-1
Typographic Conventions	Info-2

The chapters in this document, Quartus II Handbook Version 10.1 Volume 1: Design and Synthesis, were revised on the following dates. Where chapters or groups of chapters are available separately, part numbers are listed.

- Chapter 1. Design Planning with the Quartus II Software
Revised: *December 2010*
Part Number: *QII51016-10.1.0*

- Chapter 2. Quartus II Incremental Compilation for Hierarchical and Team-Based Design
Revised: *December 2010*
Part Number: *QII51015-10.1.0*

- Chapter 3. Quartus II Support for HardCopy Series Devices
Revised: *December 2010*
Part Number: *QII51004-10.1.0*

- Chapter 4. Quartus II Design Separation Flow
Revised: *December 2010*
Part Number: *QII51019-10.1.0*

- Chapter 5. Creating a System with Qsys
Revised: *December 2010*
Part Number: *QII51020-10.1.0*

- Chapter 6. Creating Qsys Components
Revised: *December 2010*
Part Number: *QII51022-10.1.0*

- Chapter 7. Qsys Interconnect
Revised: *December 2010*
Part Number: *QII51021-10.1.0*

- Chapter 8. Component Interface Tcl Reference
Revised: *December 2010*
Part Number: *QII51023-10.1.0*

- Chapter 9. Recommended Design Practices
Revised: *December 2010*
Part Number: *QII51006-10.1.0*

- Chapter 10. Recommended HDL Coding Styles
Revised: *December 2010*
Part Number: *QII51007-10.1.0*

- Chapter 11. Managing Metastability with the Quartus II Software
Revised: *December 2010*
Part Number: *QII51018-10.0.1*

- Chapter 12. Best Practices for Incremental Compilation Partitions and Floorplan Assignments
Revised: *December 2010*
Part Number: *QII51017-10.1.0*
- Chapter 13. Quartus II Integrated Synthesis
Revised: *December 2010*
Part Number: *QII51008-10.1.0*
- Chapter 14. Synopsys Synplify Support
Revised: *December 2010*
Part Number: *QII51009-10.1.0*
- Chapter 15. Mentor Graphics Precision Synthesis Support
Revised: *December 2010*
Part Number: *QII51011-10.1.0*
- Chapter 16. Mentor Graphics LeonardoSpectrum Support
Revised: *December 2010*
Part Number: *QII51010-10.1.0*
- Chapter 17. Analyzing Designs with Quartus II Netlist Viewers
Revised: *December 2010*
Part Number: *QII51013-10.0.1*

The Altera® Quartus® II design software provides a complete design environment that easily adapts to your specific design requirements. This handbook is arranged in chapters, sections, and volumes that correspond to the major stages in the overall design flow. For a general introduction to features and the standard design flow in the software, refer to the *Introduction to the Quartus II Software* manual.

This section is an introduction to design planning. It documents various specialized design flows in the following chapters:

- **Chapter 1, Design Planning with the Quartus II Software**

This chapter is an overview of various design planning considerations: device selection, early power estimation, I/O pin planning, and design planning. To help you improve design productivity, it provides recommendations and describes various tools available for Altera FPGAs.

- **Chapter 2, Quartus II Incremental Compilation for Hierarchical and Team-Based Design**

This chapter documents Altera's incremental design and compilation flow, which allows you to preserve the results and performance for unchanged logic in your design as you make changes elsewhere, reduces design iteration time by up to 70% so you achieve timing closure efficiently, and facilitates modular hierarchical and team-based design flows using top-down or bottom-up methodologies.

- **Chapter 3, Quartus II Support for HardCopy Series Devices**

With the Quartus II software, you can use an FPGA device as a prototype and seamlessly migrate your design to a HardCopy ASIC to reduce cost for volume production. This chapter describes the Quartus II support for HardCopy design flows.

- **Chapter 4, Quartus II Design Separation Flow**

This chapter describes rules and guidelines for creating a floorplan with the Design Separation flow. The Quartus II Design Separation flow provides the ability to design physically independent structures on a single device. This allows system designers to achieve a higher level of integration on a single FPGA, and alleviates increasingly strict Size Weight and Power (SWaP) requirements.

This chapter discusses key FPGA design planning considerations, provides recommendations, and describes various tools available for you to improve your design productivity with Altera® FPGAs.

Because of the significant increase in FPGA device densities, designs are complex and can sometimes involve multiple designers. System architects must resolve design issues when integrating design blocks, often leading to problems that affect the overall time to market and thereby increasing cost. You can solve potential problems early in the design cycle by following the design planning considerations provided in this chapter.

This chapter contains the following sections:

- “Creating Design Specifications” on page 1–2
- “Intellectual Property Selection” on page 1–2
- “System Design” on page 1–2
- “Device Selection” on page 1–3
- “Planning for Device Programming or Configuration” on page 1–4
- “Early Power Estimation” on page 1–5
- “Early Pin Planning and I/O Analysis” on page 1–6
- “Selecting Third-Party EDA Tool Flows” on page 1–9
- “Planning for On-Chip Debugging Options” on page 1–10
- “Design Practices and HDL Coding Styles” on page 1–11
- “Planning for Hierarchical and Team-Based Design” on page 1–13
- “Fast Synthesis and Early Timing Estimation” on page 1–16

 This chapter provides only an introduction to various design planning features in the Quartus® II software. For a general overview of the Quartus II design flow and features, refer to the *Introduction to the Quartus II Software* manual. For more information about specific Quartus II features and methodologies, this chapter provides references to other appropriate chapters in the *Quartus II Handbook*.

Before reading the design planning guidelines discussed in this chapter, consider your design priorities. More device features, density, or performance requirements can increase system cost. Signal integrity and board issues can impact I/O pin locations. Power, timing performance, and area utilization all affect each other, and compilation time is affected when optimizing these priorities.

The Quartus II software optimizes designs for the best results, but you can change the settings to intensify optimization of one aspect of your design. Certain tools or debugging options can lead to restrictions in your design flow. Your design priorities help you choose the tools, features, and methodologies to use for the design.

- After you select a device family, to check if additional guidelines are available, refer to the design guidelines section of the appropriate device handbook.

Creating Design Specifications

Before you create your logic design or complete your system design, create detailed design specifications that define the system, specify the I/O interfaces for the FPGA, identify the different clock domains, and include a block diagram of basic design functions.

Creating a test plan also helps you to design for verification and manufacturability. For example, you might need to validate interfaces incorporated in the design. To perform any built-in self-test functions to drive interfaces, you can use a UART interface with a Nios® II processor inside the FPGA device. For guidelines related to analyzing and debugging the device after it is in the system, refer to “[Planning for On-Chip Debugging Options](#)” on page 1–10.

If more than one designer works on your design, you should consider a common design directory structure or source control system to make design integration easier. For more suggestions on team-based designs, refer to “[Planning for Hierarchical and Team-Based Design](#)” on page 1–13.

Intellectual Property Selection

Altera and its third-party intellectual property (IP) partners offer a large selection of off-the-shelf IP cores optimized for Altera devices. The IP you select often affects system design, especially if the FPGA interfaces with other devices in the system. Consider which I/O interfaces or other blocks in your system design are implemented using IP cores, and plan to incorporate these cores in your FPGA design.

The OpenCore Plus feature, which is available for many IP cores, allows you to program the FPGA to verify your design in the hardware before you purchase the IP license. The evaluation supports the following modes:

- Untethered—the design runs for a limited time.
- Tethered—the design requires an Altera serial JTAG cable connected between the JTAG port on your board and a host computer running the Quartus II Programmer for the duration of the hardware evaluation period.

- For descriptions of available IP cores, refer to the [Intellectual Property](#) page of the Altera website.

System Design

You can use the Quartus II SOPC Builder or Qsys system integration tools to create your design. With SOPC Builder and Qsys, you can specify system components in a GUI and generate the required interconnect logic automatically, along with adapters for clock crossing and width adaptation. Because system design tools change the design entry methodology, you should plan to start developing your design within the tool and ensure all design blocks use appropriate standard interfaces from the beginning of the design cycle so that you do not need to make changes later.

SOPC Builder and Qsys components use Avalon® standard interfaces for the physical connection of components, and you can connect any logical device (either on-chip or off-chip) that has an Avalon interface. The Avalon Memory-Mapped interface allows a component to use an address mapped read or write protocol that enables flexible topologies for connecting master components to any slave components. The Avalon Streaming interface enables point-to-point connections between streaming components that send and receive data using a high-speed, unidirectional system interconnect between source and sink ports.

-  For more information about SOPC Builder, refer to the *SOPC Builder User Guide*.
-  For information about using Qsys to improve your productivity, refer to the *System Design with Qsys* section in volume 1 of the *Quartus II Handbook*.

Device Selection

The device you choose affects board specification and layout. This section provides guidelines in the device selection process.

Choose the device family that best suits your design requirements. Families differ in cost, performance, logic and memory density, I/O density, power utilization, and packaging. You should also consider feature requirements, such as I/O standards support, high-speed transceivers, global or regional clock networks, and the number of phase-locked loops (PLLs) available in the device.

-  You can use the [Altera Product Selector](#) available on the Altera website to help you choose your device. You can also review important features of each device family in the [Selector Guides](#) page of the Altera website. Each device family also has a device handbook or set of data sheets that documents the device features in detail. You can also see a summary of the resources for each device in the **Device** dialog box in the Quartus II software.

Carefully study the device density requirements for your design. Devices with more logic resources and higher I/O counts can implement larger and potentially more complex designs, but might cost more. Smaller devices use lower static power. Select a device that has some extra capacity than what meets your design requirements, in case you want to add more logic later in the design cycle to upgrade or expand your design, and reserve logic and memory for on-chip debugging (refer to “[Planning for On-Chip Debugging Options](#)” on page 1–10). Consider requirements for specific types of dedicated logic blocks, such as memory blocks of different sizes, or digital signal processing (DSP) blocks to implement certain arithmetic functions.

If you have older designs that target an Altera device, you can use their resource utilization as an estimate for your design. Compile existing designs in the Quartus II software with the **Auto device selected by the Fitter** option in the **Settings** dialog box. Review the resource utilization to learn which device density fits the design. Consider coding style, device architecture, and the optimization options used in the Quartus II software, which can significantly affect the resource utilization and timing performance of your design.

-  To obtain resource utilization estimates for certain configurations of Altera's IP designs, refer to the user guides for Altera megafunctions and IP MegaCores on the [IP and Megafunctions](#) literature page of the Altera website.

Device Migration Planning

Determine whether you want the option to migrate your design to another device density to allow flexibility when your design nears completion, or whether you want to migrate to a HardCopy[®] ASIC when your design reaches volume production. In some cases, designers may target a smaller (and less expensive) device and then move to a larger device if necessary to meet their design requirements. Other designers may prototype their design in a larger device to reduce optimization time and achieve timing closure more quickly, and then migrate to a smaller device after prototyping. Similarly, many designers compile and optimize their design for an FPGA device and then migrate to a HardCopy ASIC when the design is complete and ready for higher-volume production. If you want the flexibility to migrate your design, you should specify these migration options in the Quartus II software at the beginning of your design cycle.

-  For more information about specifying the target migration devices, refer to [Specifying Devices for Device Migration](#) in Quartus II Help.

Selecting a migration device impacts pin placement because some pins may serve different functions in different device densities or package sizes. If you make pin assignments in the Quartus II software, the Pin Migration View in the Pin Planner highlights pins that change function between your migration devices. (For more information, refer to [“Early Pin Planning and I/O Analysis”](#) on page 1–6.) Selecting a companion device might restrict logic utilization to ensure that your design is compatible with a selected HardCopy device. Adding migration or companion devices later in the design cycle is possible, but requires extra effort to check pin assignments, and might require design changes to fit into the new target device. Consider these issues early in the design cycle rather than at the end, when the design is near completion and ready for migration.

Additionally, if you plan to migrate your design to a HardCopy ASIC, review HardCopy guidelines early in the design cycle for any Quartus II settings that you should use or other restrictions you should consider. You must use complete timing constraints if you want to migrate to a HardCopy ASIC because of the rigorous verification requirements for ASIC devices.

-  For more information about timing requirements and analysis for HardCopy designs, refer to the [HardCopy Series Handbook](#), and the [Quartus II Support for HardCopy Series Devices](#) chapter in volume 1 of the [Quartus II Handbook](#).

Planning for Device Programming or Configuration

Planning how you want to program or configure the device in your system allows system and board designers to determine what companion devices, if any, your system requires. Your board layout also depends on the type of programming or configuration method you plan to use for programmable devices. Many programming options require a JTAG interface to connect to the devices, so you might have to set up a JTAG chain on the board. Additionally, the Quartus II software uses

the settings for the configuration scheme, configuration device, and configuration device voltage to enable the appropriate dual purpose pins as regular I/O pins after you complete configuration. The Quartus II software performs voltage compatibility checks of those pins during I/O assignment analysis and compilation of your design. You can use the **Configuration** tab of the **Device and Pin Options** dialog box to select your configuration scheme.

-  The device family handbooks describe the configuration options available for a given device family. For more details about configuration options, refer to the *Configuration Handbook*. For information about programming CPLD devices, refer to your device data sheet or handbook.

Early Power Estimation

You can use the Quartus II power estimation and analysis tools to provide information to PCB board and system designers. Power consumption in FPGA devices depend on the logic design, which can make planning difficult. You can perform early power estimation before you create any source code, or when you have a preliminary version of the design source code, and then perform the most accurate analysis with the PowerPlay Power Analyzer when you complete the design.

You must accurately estimate device power consumption to develop an appropriate power budget and to design the power supplies, voltage regulators, heat sink, and cooling system. Power estimation and analysis helps you satisfy two important planning requirements:

- Thermal—ensure that the cooling solution is sufficient to dissipate the heat generated by the device. The computed junction temperature must fall within normal device specifications.
- Power supply—ensure that the power supplies provide adequate current to support device operation.

The PowerPlay Early Power Estimator (EPE) spreadsheet allows you to estimate power utilization for your design.

You can enter data about the design manually, or you can use the tools in the Quartus II software to assist you in generating the device resources usage information for your design.

To manually enter data into the EPE spreadsheet, enter the device resources, operating frequency, toggle rates, and other parameters for your design. If you do not have an existing design, estimate the number of device resources used in your design, and then enter them manually.

If you have an existing design or a partially completed design, you can use the Quartus II software to generate the PowerPlay EPE File to assist you in completing the PowerPlay EPE spreadsheet.

-  For more information about generating the PowerPlay EPE File, refer to *Performing an Early Power Estimate Using the PowerPlay Early Power Estimator* in Quartus II Help.

The PowerPlay EPE spreadsheet includes the Import Data macro that parses the information in the PowerPlay EPE File and transfers the information into the spreadsheet. If you do not want to use the macro, you can manually transfer the data into the EPE spreadsheet. For example, after importing the PowerPlay EPE File information into the PowerPlay EPE spreadsheet, you can add additional devices resource information. If the existing Quartus II project represents only a portion of your full design, manually enter the additional device resources you use in the final design.

Estimating power consumption early in the design cycle allows planning of power budgets and avoids unexpected results for designers developing the PCB.

-  The PowerPlay EPE spreadsheets for each supported device family are available on the [PowerPlay Early Power Estimator and Power Analyzer](#) page of the Altera website.

When you complete the design, perform a complete power analysis to check the power consumption more accurately. The PowerPlay Power Analyzer tool in the Quartus II software provides an accurate estimation of power, ensuring that thermal and supply limitations are not violated. For the most accurate power estimation, use gate-level simulation results from a Verilog Value Change Dump File (.vcd) with the PowerPlay Power Analyzer.

-  For more information about power estimation and analysis, refer to the [PowerPlay Power Analysis](#) chapter in volume 3 of the *Quartus II Handbook*.

Early Pin Planning and I/O Analysis

In many design environments, FPGA designers want to plan the top-level FPGA I/O pins early to help board designers begin the PCB design and layout. The I/O capabilities and board layout guidelines of the FPGA device influence pin locations and other types of assignments. If the board design team specifies an FPGA pin-out, it is crucial that the pin locations are verified in the FPGA placement and routing software to avoid board design changes.

You can create a preliminary pin-out for an Altera FPGA with the Quartus II Pin Planner before you develop the source code, based on standard I/O interfaces (such as memory and bus interfaces) and any other I/O-related assignments defined by system requirements. The Quartus II I/O Assignment Analysis checks that the pin locations and assignments are supported in the target FPGA architecture. You can then use I/O Assignment Analysis to validate I/O-related assignments that you create or modify throughout the design process. When you compile your design in the Quartus II software, I/O Assignment Analysis is run automatically in the Fitter to validate that the assignments meet all the device requirements and generates messages if there are problems.

This section describes pin planning and I/O analysis features for different stages of the design flow.

Early in the design process, before the source code is created, the system architect has information about the standard I/O interfaces (such as memory and bus interfaces), the IP cores that are used in the design, and any other I/O-related assignments defined by system requirements. You can use this information with the **Create/Import Megafunction** feature in the Pin Planner to specify details about the design I/O interfaces. Specifying these details allows you to create a top-level design file that includes all your I/O information, so that you can analyze the I/O assignments in the Quartus II software.

The Pin Planner interfaces with the MegaWizard™ Plug-In Manager, and allows you to create or import custom megafunctions and IP cores that use I/O interfaces. You can configure how the functions and cores are connected to each other by specifying matching node names for selected ports in the **Set Up Top-Level Design File** dialog box. Create any other I/O-related assignments for these interfaces or other design I/O pins in the Pin Planner, as described in this section. When you have entered as much I/O-related information as possible, generate a top-level design file using the **Create Top-Level Design File** command. The Pin Planner creates virtual pin assignments for internal nodes, so internal nodes are not assigned to device pins during compilation. After analysis and synthesis of the newly generated top-level wrapper file, use the generated netlist to perform I/O Analysis with the **Start I/O Assignment Analysis** command.

- ② For more information about setting up the nodes in your design, refer to [Set Up Top-Level Design File Window \(Edit Menu\)](#) in Quartus II Help.

You can use the I/O analysis results to change pin assignments or IP parameters even before the design is created, and repeat the checking process until the I/O interface meets your design requirements and passes the pin checks in the Quartus II software. When you complete initial pin planning, you can create a revision based on the Quartus II-generated netlist. You can then use the generated netlist to develop the top-level design file for the actual design, or disregard the generated netlist and use the generated Quartus II Settings File (.qsf) with the actual design.

During this initial pin planning, after you have generated a top-level design file, or when you have developed your design source code, you can assign pin locations and assignments using the Pin Planner.

The Pin Planner enables easy I/O pin assignment planning, assignment, and validation. You can use the View menu in the Pin Planner to create pin location and other assignments using a device package view instead of pin numbers.

With the Pin Planner, you can identify I/O banks, voltage reference (VREF) groups, and differential pin pairings to help you through the I/O planning process. If migration devices are selected (including HardCopy devices) as described in [“Device Migration Planning” on page 1–4](#), the **Pin Migration View** highlights the pins that have changed functions in the migration device when compared to the currently selected device. Selecting the pins in the Device Migration view cross-probes to the rest of the Pin Planner, so that you can use device migration information when planning your pin assignments. You can also configure board trace models of selected pins for use in “board-aware” signal integrity reports generated with the **Enable**

Advanced I/O Timing option. This option ensures that you get very accurate I/O timing analysis. You can use a Microsoft Excel spreadsheet to start the I/O planning process if you normally use a spreadsheet in your design flow, and you can export a Comma-Separated Value File (.csv) containing your I/O assignments for spreadsheet use when you assign all pins.

When you complete your pin planning, you can pass pin location information to PCB designers. The Pin Planner is tightly integrated with certain PCB design EDA tools, and can read pin location changes from these tools to check suggested changes. Your pin assignments must match between the Quartus II software and your schematic and board layout tools to ensure the FPGA works correctly on the board, especially if you must make changes to the pin-out. The system architect uses the Quartus II software to pass pin information to team members designing individual logic blocks, allowing them to achieve better timing closure when they compile their design.

Start FPGA planning before you complete the HDL design to improve the confidence in early board layouts, reduce the chance of error, and improve the overall time to market of the design. When you complete the design, use the Fitter reports for the final sign-off of pin assignments. After compilation, the Quartus II software generates the Pin-Out File (.pin), and you can use this file to verify that each pin is correctly connected in board schematics.

 For more information about I/O assignment and analysis, refer to the *I/O Management* chapter in volume 2 of the *Quartus II Handbook*. For more information about passing I/O information between the Quartus II software and third-party EDA tools, refer to the *Mentor Graphics PCB Design Tools Support* and *Cadence PCB Design Tools Support* chapters in the *I/O and PCB Tools* section in volume 2 of the *Quartus II Handbook*.

Simultaneous Switching Noise Analysis

Simultaneous switching noise (SSN) is a noise voltage inducted onto a victim I/O pin of a device due to the switching behavior of other aggressor I/O pins in the device. SSN often leads to the degradation of signal integrity by causing signal distortion, thereby reducing the noise margin of a system. The best approach to resolving this issue is to address SSN with estimation early in your system design, to reduce the chance of any later board design changes. When the design is complete, perform a complete SSN analysis of your FPGA in the Quartus II software to verify the board design.

Altera provides tools for SSN analysis and estimation, including SSN characterization reports, an Early SSN Estimator (ESE) tool, and the SSN Analyzer in the Quartus II software.

You can use the ESE tool to estimate SSN in your FPGA design, which is available for various device families.

 For more information and device support for the ESE spreadsheet tool, refer to [Altera's Signal Integrity Center](#) on the Altera website. For more information about the SSN Analyzer, refer to the *Simultaneous Switching Noise (SSN) Analysis and Optimizations* chapter in volume 2 of the *Quartus II Handbook*.

Selecting Third-Party EDA Tool Flows

Your complete FPGA design flow may include third-party EDA tools in addition to the Quartus II software. Determine which tools you want to use with the Quartus II software to ensure that they are supported and set up correctly, and that you are aware of any useful features or undesired limitations.

Synthesis Tools

The Quartus II software includes integrated synthesis that supports Verilog HDL, VHDL, Altera Hardware Description Language (AHDL), and schematic design entry. You can also use supported standard third-party EDA synthesis tools to synthesize your Verilog HDL or VHDL design, and then compile the resulting output netlist file in the Quartus II software. Different synthesis tools may give different results for each design. To assess the best-performing tool for your application, you can experiment by synthesizing typical designs for your specific application and coding style. Perform placement and routing in the Quartus II software to get accurate timing analysis and logic utilization results.

Because tool vendors frequently add new features, fix tool issues, and enhance performance for Altera devices, you should use the most recent version of third-party synthesis tools. The *Quartus II Software Release Notes* lists the version of each synthesis tool that is officially supported by that version of the Quartus II software.

The synthesis tool you choose may allow you to create a Quartus II project and pass constraints, such as the EDA tool setting, device selection, and timing requirements that you specified in your synthesis project. You can save time when setting up your Quartus II project for placement and routing.

If you want to take advantage of an incremental compilation methodology, you should partition your design for synthesis and generate multiple output netlist files. For more information, refer to [“Incremental Compilation with Design Partitions” on page 1-14](#).



For more information about synthesis tool flows, refer to the appropriate chapter in the *Synthesis* section in volume 1 of the *Quartus II Handbook*.

Simulation Tools

Altera provides the ModelSim®-Altera Starter Edition with the Quartus II software. You can also purchase the ModelSim-Altera Edition to support large designs and achieve faster simulation performance. The Quartus II software can generate both functional and timing netlist files for ModelSim and other third-party simulators.

Use the simulator version that is supported with your Quartus II software version for best results. You should also use the model libraries provided with your Quartus II software version. Libraries can change between versions, which might cause a mismatch with your simulation netlist. The *Quartus II Software Release Notes* list the version of each simulation tool that is supported with that particular version of the Quartus II software.

Specify your simulation tool in the **EDA Tools Settings** page of the **Settings** dialog box to generate the appropriate output simulation netlist.

-  For more information about simulation tool flows, refer to the appropriate chapter in the *Simulation* section in volume 3 of the *Quartus II Handbook*.

Formal Verification Tools

Consider whether the formal verification flow that you want to use is supported, and whether the flow impacts the design and compilation stages of your design.

-  For more information about formal verification flows and supported tools, refer to the appropriate chapter in the *Formal Verification* section in volume 3 of the *Quartus II Handbook*.

Using a formal verification flow can impact performance results because the flow requires turning off certain logic optimizations, such as register retiming, and forces you to preserve hierarchy blocks, which can restrict optimization. Formal verification treats memory blocks as black boxes. Therefore, you should keep memory in a separate hierarchy block so other logic does not get incorporated into the black box for verification. Other restrictions may limit your design, and you should consult the documentation for details. If formal verification is important to your design, plan for limitations and restrictions at the beginning of the design cycle rather than make changes later.

Specify your formal verification tool in the **EDA Tools Settings** page of the **Settings** dialog box to generate the appropriate output netlist.

Planning for On-Chip Debugging Options

In-system debugging tools offer different advantages and trade-offs. A particular debugging tool may work better for different systems and designers. You should evaluate on-chip debugging options early in your design process, to ensure that your system board, Quartus II project, and design are all set up to support the appropriate options. You can reduce debugging time and avoid later changes to accommodate your preferred debugging methodologies.

-  For more information about debugging tools, refer to *Section IV. In-System Debugging* in volume 3 of the *Quartus II Handbook*. For an overview of debugging options that can help you decide which option to use, refer to the *System Debugging Tools Overview* chapter in volume 3 of the *Quartus II Handbook*.

If you intend to use any of these features, you may have to plan for the features when developing your system board, Quartus II project, and design. Consider the following factors related to your debugging requirements when you are planning your design:

- JTAG connections—required to perform in-system debugging with JTAG tools. Plan your system and board with JTAG ports that are available for debugging.
- Additional logic resources—required to implement JTAG hub logic. If you set up the appropriate feature early in your design cycle, you can include these device resources in your early resource estimations to ensure that you do not overfill the device with logic.

- Reserve device memory—required if your tool uses device memory to capture data during system operation. To ensure that you have enough memory resources to take advantage of this debugging technique, consider reserving device memory to be used during debugging.
- Reserve I/O pins—required if you are using the logic analyzer interface (LAI) or SignalProbe feature, which require I/O pins for debugging. If you reserve I/O pins for debugging, you do not have to change the design or board later. Keep in mind that the LAI can multiplex signals with design I/O pins if required. Ensure that your board supports a debugging mode, where debugging signals do not affect system operation.
- Instantiate a megafunction in your HDL code—required if your debugging tool uses a Quartus II megafunction.

Table 1-1 lists which factors are important for each debugging tool.

Table 1-1. Factors to Consider When Using Debugging Tools During Design Planning Stages

Factor	SignalTap II Logic Analyzer	System Console	In-System Memory Content Editor	Logic Analyzer Interface (LAI)	SignalProbe	In-System Sources and Probes	Virtual JTAG Megafunction
JTAG connections	✓	✓	✓	✓	—	✓	✓
Additional logic resources	—	✓	—	—	—	—	✓
Reserve device memory	✓	✓	—	—	—	—	—
Reserve I/O pins	—	—	—	✓	✓	—	—
Instantiate a megafunction in your HDL code	— (1)	—	—	—	—	✓	✓

Notes to Table 1-1:

- (1) You can instantiate the SignalTap II Logic Analyzer as a megafunction, so that you can manually connect the SignalTap II Logic Analyzer to nodes in your design and ensure that the tapped node names do not change during synthesis. You can add the debugging block as a separate design partition for incremental compilation to minimize recompilation times.

Design Practices and HDL Coding Styles

When you develop complex FPGA designs, design practices and coding styles have an enormous impact on the timing performance, logic utilization, and system reliability of your device.

Design Recommendations

You can use synchronous design practices to consistently meet your design goals. Problems with asynchronous design techniques include reliance on propagation delays in a device, incomplete timing analysis, and possible glitches. In a synchronous design, a clock signal triggers all events. When you meet all register timing requirements, a synchronous design behaves in a predictable and reliable manner for all process, voltage, and temperature (PVT) conditions. You can easily target synchronous designs to different device families or speed grades.

Clock signals have a large effect on the timing accuracy, performance, and reliability of your design. Problems with clock signals can cause functional and timing problems in your design. Use dedicated clock pins and clock routing for best results, and if you have PLLs in your target device, use the PLLs for clock inversion, multiplication, and division. For clock multiplexing and gating, use the dedicated clock control block or PLL clock switchover feature instead of combinational logic if these features are available in your device. If you must use internally-generated clock signals, register the output of any combinational logic used as a clock signal to reduce glitches.

The Design Assistant in the Quartus II software is a design-rule checking tool that enables you to verify design issues. The Design Assistant checks your design for adherence to Altera-recommended design guidelines. You can also use third-party "lint" tools to check your coding style.

- ❓ For more information about running the Design Assistant, refer to *About the Design Assistant* in Quartus II Help.

You should also understand the target architecture of your device in order to take advantage of device-specific features. For example, the control signals should use the dedicated control signals in the device architecture. In some cases, you might need to limit the number of different control signals used in your design to achieve the best results.

- 🔗 For more information about design recommendations and using the Design Assistant, refer to the *Design Recommendations for Altera Devices and the Quartus II Design Assistant* chapter in volume 1 of the *Quartus II Handbook*. You can also refer to industry papers for more information about multiple clock design. For a good analysis, refer to *Synthesis and Scripting Techniques for Designing Multi-Asynchronous Clock Designs* under **Papers** (www.sunburst-design.com).

Recommended HDL Coding Styles

HDL coding styles can have a significant effect on the quality of results for programmable logic designs. If you design memory and DSP functions, you should understand the target architecture of your device so you can use the dedicated logic block sizes and configurations. Follow the coding guidelines for inferring megafunctions and targeting dedicated device hardware, such as memory and DSP blocks.

- 🔗 For specific HDL coding examples and recommendations, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*. For additional tool-specific guidelines, refer to the documentation of your synthesis tool.

Managing Metastability

Metastability problems can occur in digital design when a signal is transferred between circuitry in unrelated or asynchronous clock domains, because the designer cannot guarantee that the signal meets the setup and hold time requirements during the signal transfer. Designers commonly use a synchronization chain to minimize the occurrence of metastable events. Ensure that your design accounts for synchronization between any asynchronous clock domains. Consider using a synchronizer chain of more than two registers for high-frequency clocks and frequently-toggling data signals to reduce the chance of a metastability failure.

You can use the Quartus II software to analyze the average mean time between failures (MTBF) due to metastability when a design synchronizes asynchronous signals, and optimize the design to improve the metastability MTBF. The MTBF due to metastability is an estimate of the average time between instances when metastability could cause a design failure. A high MTBF (such as hundreds or thousands of years between metastability failures) indicates a more robust design. Determine an acceptable target MTBF given the context of your entire system and the fact that MTBF calculations are statistical estimates.

The Quartus II software can help you determine whether you have enough synchronization registers in your design to produce a high enough MTBF at your clock and data frequencies.

-  For information about the industry-leading metastability analysis, reporting, and optimization features in the Quartus II software, refer to the *Managing Metastability with the Quartus II Software* chapter in volume 1 of the *Quartus II Handbook*.

Planning for Hierarchical and Team-Based Design

If you want to create a hierarchical design so that you can use compilation-time savings and performance preservation with the Quartus II software incremental compilation feature, plan for an incremental compilation flow from the beginning of your design cycle. The following subsections describe the flat compilation flow, in which the design hierarchy is flattened without design partitions, and then the incremental compilation flow that uses design partitions. Incremental compilation flows offer several advantages but require more design planning to ensure good quality of results. The last subsections discuss factors to consider when planning an incremental compilation flow, planning design partitions, and optionally creating a design floorplan.

-  For information about using the incremental compilation flow methodology in the Quartus II software, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.

Flat Compilation Flow with No Design Partitions

In the flat compilation flow with no design partitions in the Quartus II software, the entire design is compiled together in a “flat” netlist. Your source code can have hierarchy, but the design is flattened during compilation and all the design source code is synthesized and fit in the target device whenever the design is recompiled after any change in the design. By processing the entire design, the software performs all available logic and placement optimizations on the entire design to improve area and performance. You can use debugging tools in an incremental design flow, such as the SignalTap II Logic Analyzer, but you do not specify any design partitions to preserve design hierarchy during compilation.

The flat compilation flow is easy to use; you do not have to plan any design partitions. However, because the entire design is recompiled whenever there are any changes to the design, compilation times can be relatively long for large devices. Additionally, you may find that the results for one part of the design change when you change a different part of your design. You can turn on the **Rapid Recompile** option to instruct the software to preserve compatible placement and routing results when the design changes in subsequent compilations. This option can reduce your compilation time in a flat or partitioned design when you make very small changes to the design.

Incremental Compilation with Design Partitions

In an incremental compilation flow, the system architect splits a large design into partitions. When hierarchical design partitions are well chosen and placed in the device floorplan, you can speed up your design compilation time while maintaining the quality of results.

Incremental compilation preserves the compilation results and performance of unchanged partitions in the design, greatly reducing design iteration time by focusing new compilations on changed design partitions only. New compilation results are then merged with the previous compilation results from unchanged design partitions. Additionally, you can target optimization techniques, such as physical synthesis, to specific design partitions while leaving other partitions untouched. You can also use empty partitions to indicate that parts of your design are incomplete or missing, while you compile the rest of the design.

Third-party IP designers can also export logic blocks to be integrated into the top-level design. Team members can work on partitions independently, which can simplify the design process and reduce compilation time. With exported partitions, the system architect must provide guidance to designers or IP providers to ensure that each partition uses the appropriate device resources. Because the designs may be developed independently, each designer has no information about the overall design or how their partition connects with other partitions. This lack of information can lead to problems during system integration. The top-level project information, including pin locations, physical constraints, and timing requirements, must be communicated to the designers of lower-level partitions before they start their design.

The system architect plans design partitions at the top level and allows third-party designs to access the top-level project framework. By designing within a copy of the top-level project (or by checking out the project files in a source control environment), the designers of the lower-level block have full information about the entire project, which helps to ensure optimal results.

When you are planning your design code and hierarchy, ensure that each design entity is created in a separate file so that the entities remain independent when you make source code changes in the file. If you use a third-party synthesis tool, create separate Verilog Quartus Mapping or EDIF netlists for each design partition in your synthesis tool. You may have to create separate projects within your synthesis tool, so that the tool synthesizes each partition separately and generates separate output netlist files. The netlists are then considered the source files for incremental compilation. For information about support for Quartus II incremental compilation, refer to your synthesis tool documentation.

Planning Design Partitions and Floorplan Location Assignments

Partitioning a design for an FPGA requires planning to ensure optimal results when the partitions are integrated, and ensure that each partition is placed well relative to other partitions in the device. Following Altera's recommendations for creating design partitions improves the overall quality of results. For example, registering partition I/O boundaries keeps critical timing paths inside one partition that can be optimized independently. When the design partitions are specified, you can use the Incremental Compilation Advisor to ensure that partitions meet Altera's recommendations.

If you have timing-critical partitions that are changing through the design flow, or partitions exported from another Quartus II project, you can create design floorplan assignments to constrain the placement of the affected partitions. Creating location assignments ensures that no location conflicts occur between partitions. Additionally, design floorplan assignments help to avoid a situation in which the Fitter is directed to place or replace a portion of the design in an area of the device in which most resources are claimed. You can use the Quartus II Chip Planner to create floorplan assignments using LogicLock™ region assignments for design partitions. With a basic design framework for the top-level design, you can view connections between regions, estimate physical timing delays on the chip, and move regions around the device floorplan. When you have compiled the full design, you can also view logic placement and locate areas of routing congestion to improve floorplan assignments. Good partition and floorplan design helps partitions meet top-level design requirements when integrated with the rest of the design, reducing time spent integrating and verifying the timing of the top-level design.

-  For detailed guidelines about creating design partitions and organizing your source code, as well as information about when and how to create floorplan assignments, refer to the *Best Practices for Incremental Compilation Partitions and Floorplan* chapter in volume 1 of the *Quartus II Handbook*.
-  For more information about creating floorplan assignments in the Chip Planner, refer to the *Analyzing and Optimizing the Design Floorplan* chapter in volume 2 of the *Quartus II Handbook*.

Fast Synthesis and Early Timing Estimation

It is more cost-effective to find design issues early in the design cycle than to find problems in the final timing closure stages. When the first version of the design source code is complete, you might want to perform a quick compilation to create a kind of silicon virtual prototype (SVP) that you can use to perform timing analysis.

If you synthesize with the Quartus II software, you can choose to perform a **Fast** synthesis, which reduces the compilation time but may give reduced quality of results.

- ② For more information about Fast synthesis, refer to *Synthesis Effort logic option* in Quartus II Help.

Regardless of your compilation flow, you can use the an Early Timing Estimate to perform a quick placement and routing, and a timing analysis of your design. The software chooses a device automatically if required, places any LogicLock regions used to create a floorplan, finds a quick initial placement for all the design logic, and provides a useful estimate of the final design performance. If you have entered timing constraints, timing analysis reports on these constraints.

- ② For more information about how to run an early timing estimate, refer to *Running a Timing Analysis* in Quartus II Help.

If you are designing individual design blocks or partitions separately, you can use these features as you develop the design. Any issues highlighted in the lower-level design blocks are communicated to the system architect. Resolving these issues might require allocating additional device resources to the individual partition, or changing its timing budget.

If you are a top-level designer, you can also use fast synthesis and early timing estimation to prototype the entire design. Incomplete partitions are marked as empty in an incremental compilation flow, while the rest of the design is compiled to get an early timing estimate and detect any problems with design integration.

Conclusion

Modern FPGAs support large, complex designs with fast timing performance. By planning several aspects of your design early, you can reduce time in later stages of the development cycle. Use features of the Quartus II software to quickly plan your design and achieve the best possible results. Following the guidelines presented in this chapter can improve productivity, which can reduce cost and development time.

Document Revision History

Table 1-2 shows the revision history for this chapter.

Table 1-2. Document Revision History (Part 1 of 2)

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Changed to new document template ■ Updated “System Design” on page 1-2 to include information about the Qsys system integration tool ■ Added link to the Altera Product Selector in “Device Selection” on page 1-3 ■ Converted information into new table (Table 1-1) in “Planning for On-Chip Debugging Options” on page 1-10 ■ Simplified description of incremental compilation usages in “Incremental Compilation with Design Partitions” on page 1-14 ■ Added information about the Rapid Recompile option in “Flat Compilation Flow with No Design Partitions” on page 1-14 ■ Removed details and linked to Quartus II Help in “Fast Synthesis and Early Timing Estimation” on page 1-16
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Added new section “System Design” on page 1-3 ■ Removed details about debugging tools from “Planning for On-Chip Debugging Options” on page 1-10 and referred to other handbook chapters for more information ■ Updated information on recommended design flows in “Incremental Compilation with Design Partitions” on page 1-14 and removed “Single-Project Versus Multiple-Project Incremental Flows” heading ■ Merged the “Planning Design Partitions” section with the “Creating a Design Floorplan” section. Changed heading title to “Planning Design Partitions and Floorplan Location Assignments” on page 1-15 ■ Removed “Creating a Design Floorplan” section ■ Removed “Referenced Documents” section ■ Minor updates throughout chapter

Table 1-2. Document Revision History (Part 2 of 2)

Date	Version	Changes
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Added details to “Creating Design Specifications” on page 1-2 ■ Added details to “Intellectual Property Selection” on page 1-2 ■ Updated information on “Device Selection” on page 1-3 ■ Added reference to “Device Migration Planning” on page 1-4 ■ Removed information from “Planning for Device Programming or Configuration” on page 1-4 ■ Added details to “Early Power Estimation” on page 1-5 ■ Updated information on “Early Pin Planning and I/O Analysis” on page 1-6 ■ Updated information on “Creating a Top-Level Design File for I/O Analysis” on page 1-8 ■ Added new “Simultaneous Switching Noise Analysis” section ■ Updated information on “Synthesis Tools” on page 1-9 ■ Updated information on “Simulation Tools” on page 1-9 ■ Updated information on “Planning for On-Chip Debugging Options” on page 1-10 ■ Added new “Managing Metastability” section ■ Changed heading title “Top-Down Versus Bottom-Up Incremental Flows” to “Single-Project Versus Multiple-Project Incremental Flows” ■ Updated information on “Creating a Design Floorplan” on page 1-18 ■ Removed information from “Fast Synthesis and Early Timing Estimation” on page 1-18
March 2009	9.0.0	<ul style="list-style-type: none"> ■ No change to content
November 2008	8.1.0	<ul style="list-style-type: none"> ■ Changed to 8-1/2 x 11 page size. No change to content.
May 2008	8.0.0	<ul style="list-style-type: none"> ■ Organization changes ■ Added “Creating Design Specifications” section ■ Added reference to new details in the In-System Design Debugging section of volume 3 ■ Added more details to the “Design Practices and HDL Coding Styles” section ■ Added references to the new Best Practices for Incremental Compilation and Floorplan Assignments chapter ■ Added reference to the Quartus II Language Templates

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter provides information and design scenarios to help you partition your design to take advantage of the Quartus® II incremental compilation feature.

The ability to iterate rapidly through FPGA design and debugging stages is critical. The Quartus II software introduced the FPGA industry's first true incremental design and compilation flow, with the following benefits:

- Preserves the results and performance for unchanged logic in your design as you make changes elsewhere.
- Reduces design iteration time by an average of 75% for small changes in large designs, so that you can perform more design iterations per day and achieve timing closure efficiently.
- Facilitates modular hierarchical and team-based design flows, as well as design reuse and intellectual property (IP) delivery.



Quartus II incremental compilation supports the Arria® GX, Stratix®, and Cyclone® series of devices, with limited support for HardCopy® ASICs (for details, refer to “[Limitations for HardCopy Compilation and Migration Flows](#)” on page 2–48).

This document contains the following sections:

- “[Deciding Whether to Use an Incremental Compilation Flow](#)”
- “[Incremental Compilation Summary](#)” on page 2–7
- “[Common Design Scenarios Using Incremental Compilation](#)” on page 2–10
- “[Deciding Which Design Blocks Should Be Design Partitions](#)” on page 2–14
- “[Specifying the Level of Results Preservation for Subsequent Compilations](#)” on page 2–21
- “[Exporting Design Partitions from Separate Quartus II Projects](#)” on page 2–26
- “[Team-Based Design Optimization and Third-Party IP Delivery Scenarios](#)” on page 2–35
- “[Creating a Design Floorplan With LogicLock Regions](#)” on page 2–44
- “[Incremental Compilation Restrictions](#)” on page 2–47
- “[Scripting Support](#)” on page 2–54

Deciding Whether to Use an Incremental Compilation Flow

The Quartus II incremental compilation feature enhances the standard Quartus II design flow by allowing you to preserve satisfactory compilation results and performance of unchanged blocks of your design.

This section outlines the flat compilation flow with no design partitions in “[Flat Compilation Flow with No Design Partitions](#)”, and the incremental flow when you divide the design into partitions in “[Incremental Compilation Flow With Design Partitions](#)” on page 2-3, and explains the differences. This section also explains when a flat compilation flow is satisfactory, and highlights some of the reasons you might want to create design partitions and use the incremental flow. A discussion about incremental and team design flows in “[Team-Based Design Flows and IP Delivery](#)” on page 2-6 describes how it is beneficial to keep your design within one project, as well as when it might be necessary for other team members or IP providers to develop particular design blocks or partitions separately, and then later integrate their partitions into the top-level design.

Flat Compilation Flow with No Design Partitions

In the flat compilation flow with no design partitions, all the source code is processed with the Analysis and Synthesis module, and all the logic is placed and routed by the Fitter module whenever the design is recompiled after a change in any part of the design. One reason for this behavior is to ensure optimal push-button quality of results. By processing the entire design, the Compiler can perform global optimizations to improve area and performance.

You can use a flat compilation flow for small designs, such as designs in CPLD devices or low-density FPGA devices, when the timing requirements are met easily with a single compilation. A flat design is satisfactory when compilation time and preserving results for timing closure are not concerns.

Refer to the next subsection for ways to reduce compilation time when you use a flat compilation for your design.

Incremental Capabilities Available When A Design Has No Partitions

The Quartus II software has incremental compilation capabilities available even when you do not partition your design, including Smart Compilation, incremental debugging, and Rapid Recompile. These features work with design partitions as well, if you do follow an incremental design flow.

In any Quartus II compilation flow, you can use Smart Compilation to allow the compiler to determine which compiler stages are required, based on the changes made to the design since the last smart compilation, and then skip any stages that are not required. For example, when Smart Compilation is turned on, the compiler skips the Analysis and Synthesis module if all the design source files are unchanged. When Smart Compilation is turned on, if you make any changes to the logic of a design, the Compiler uses all the modules during processing. You can turn on Smart Compilation in the **Settings** dialog box on the **Compilation Process Settings** page.

The Quartus II software also includes a Rapid Recompile feature that instructs the compiler to reuse the compatible compilation results if most of the design has not changed since the last compilation. This feature reduces compilation times for small and isolated design changes. You do not have control over which parts of the design are recompiled using this option; the compiler determines which parts of the design must be recompiled. The Rapid Recompile preserves performance and can save compile time by reducing the amount of changed logic that must be recompiled. You can turn on the **Rapid Recompile** option in the Quartus II software on the **Incremental Compilation** page in the **Settings** dialog box.

During the debugging stage of the design cycle, you can use incremental compilation to add the SignalTap® II Logic Analyzer incrementally to your design, even if the design does not have partitions. To preserve the compilation netlist for the entire design, instruct the software to reuse the compilation results for the automatically-created "Top" partition that contains the entire design. For more information, refer to [“Debugging Incrementally With the SignalTap II Logic Analyzer” on page 2-13](#).

Incremental Compilation Flow With Design Partitions

In the standard incremental compilation design flow, the top-level design is divided into design partitions, which can be compiled and optimized together in the top-level Quartus II project. You can preserve fitting results and performance for completed partitions while other parts of the design are changing, which reduces the compilation times for each design iteration.

Incremental compilation is recommended for large designs and high resource densities when preserving results is important to achieve timing closure. The incremental compilation feature also facilitates team-based design flows that allow designers to create and optimize design blocks independently, when necessary. Refer to the next section [“Team-Based Design Flows and IP Delivery” on page 2-6](#) for more information.

To take advantage of incremental compilation, start by splitting your design along any of its hierarchical boundaries into design blocks to be compiled incrementally, and assign each block as a design partition. The Quartus II software synthesizes each individual hierarchical design partition separately, and then merges the partitions into a complete netlist for subsequent stages of the compilation flow. When recompiling your design, you can use source code, post-synthesis results, or post-fitting results to preserve satisfactory results for each partition. Refer to [“Incremental Compilation Summary” on page 2-7](#) and subsequent sections for more details.

In a team-based environment, part of your design may be incomplete, or it may have been developed by another designer or IP provider. In this scenario, you can add the completed partitions to the design incrementally. Alternatively, other designers or IP providers can develop and optimize partitions independently and the project lead can later integrate the partitions into the top-level design. Refer to [“Team-Based Design Flows and IP Delivery” on page 2-6](#) for more details.

[Table 2-1](#) shows a summary of the impact the Quartus II incremental compilation feature has on compilation results.

Table 2-1. Impact Summary of Using Incremental Compilation (Part 1 of 2)

Characteristic	Impact of Incremental Compilation with Design Partitions
Compilation Time Savings	Typically saves an average of 75% compilation time for small design changes in large designs when post-fit netlists are preserved; there are savings in both Quartus II Integrated Synthesis and the Fitter. <i>(1)</i>
Performance Preservation	Excellent performance preservation when timing critical paths are contained within a partition, because you can preserve post-fitting information for unchanged partitions.
Node Name Preservation	Preserves post-fitting node names for unchanged partitions.

Table 2-1. Impact Summary of Using Incremental Compilation (Part 2 of 2)

Characteristic	Impact of Incremental Compilation with Design Partitions
Area Changes	The area (logic resource utilization) might increase because cross-boundary optimizations are no longer possible, and placement and register packing are restricted.
f_{MAX} Changes	The design's maximum frequency might be reduced because cross-boundary optimizations are no longer possible. If the design is partitioned and the floorplan location assignments are created appropriately, there might be no negative impact on f_{MAX} .

Note to Table 2-1:

- (1) Quartus II incremental compilation does not reduce processing time for the early "pre-fitter" operations, such as determining pin locations and clock routing, so the feature cannot reduce compilation time if runtime is dominated by those operations.

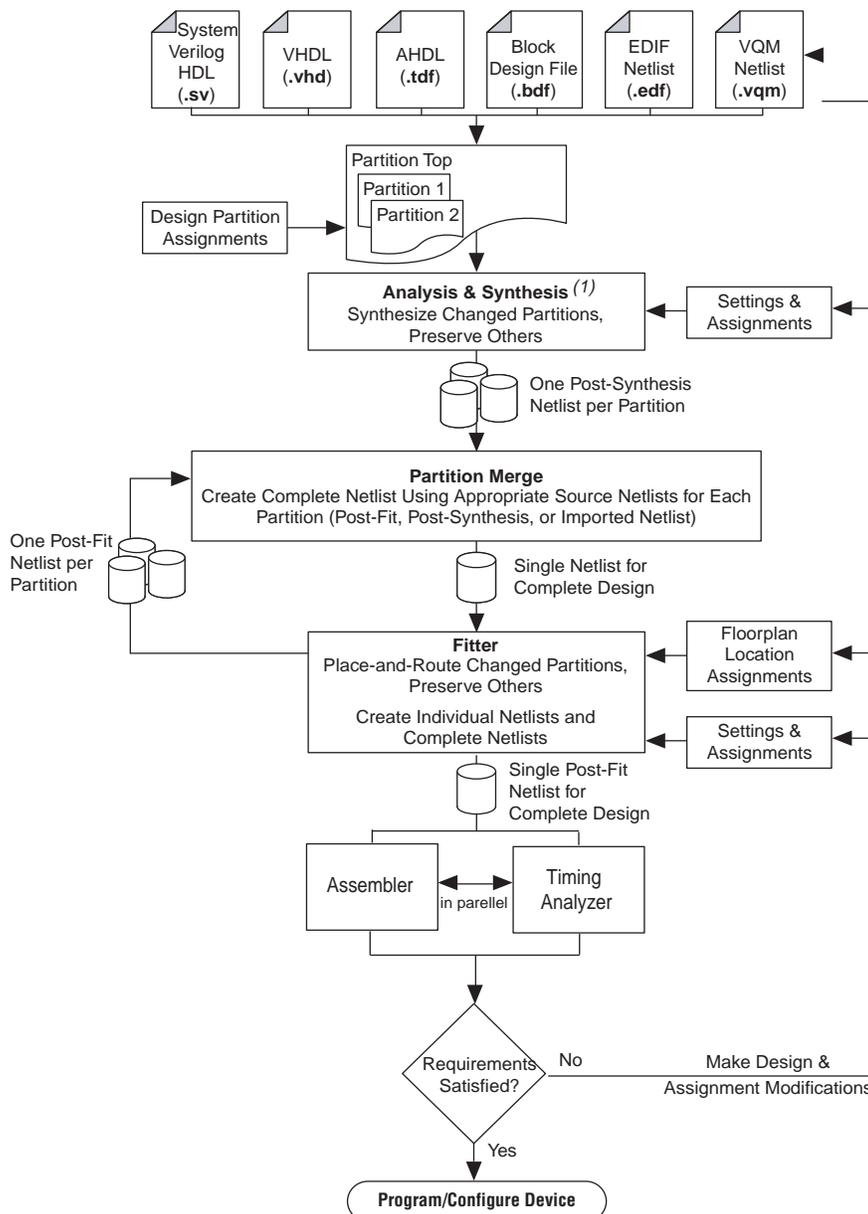
If you use the incremental compilation feature at any point in your design flow, it is easier to accommodate the guidelines for partitioning a design and creating a floorplan if you start planning for incremental compilation at the beginning of your design cycle.



For more information and recommendations on how to prepare your design to use the Quartus II incremental compilation feature, and how to avoid negative impact on your design results, refer to the *Best Practices for Incremental Compilation Partitions and Floorplan Assignments* chapter in volume 1 of the *Quartus II Handbook*.

Figure 2-1 shows a diagram of the Quartus II design flow using incremental compilation with design partitions.

Figure 2-1. Quartus II Design Flow Using Incremental Compilation



Note to Figure 2-1:

(1) When you use EDIF or VQM netlists created by third-party EDA synthesis tools, Analysis and Synthesis creates the design database, but logic synthesis and technology mapping are performed only for black boxes.

The diagram in Figure 2-1 shows a top-level partition and two lower-level partitions. If any part of the design changes, Analysis and Synthesis processes the changed partitions and keeps the existing netlists for the unchanged partitions. After completion of Analysis and Synthesis, there is one post-synthesis netlist for each partition.

The Partition Merge step creates a single, complete netlist that consists of post-synthesis netlists, post-fit netlists, and netlists exported from other Quartus II projects, depending on the netlist type that you specify for each partition.

The Fitter then processes the merged netlist, preserves the placement and routing of unchanged partitions, and refits only those partitions that have changed. The Fitter generates the complete netlist for use in future stages of the compilation flow, including timing analysis and programming file generation, which can take place in parallel if more than one processor is enabled for use in the Quartus II software. The Fitter also generates individual netlists for each partition so that the Partition Merge stage can use the post-fit netlist to preserve the placement and routing of a partition if you specify to do so in future compilations.

If you define partitions, but want to check your compilation results without partitions in a “what if” scenario, you can direct the compiler to ignore all partitions assignments in your project and compile the design as a “flat” netlist. When you turn on the **Ignore partitions assignments during compilation** option on the **Incremental Compilation** page, the Quartus II software disables all design partition assignments in your project and runs a full compilation ignoring all partition boundaries and netlists. Turning off the **Ignore partitions assignments during compilation** option restores all partition assignments and netlists for subsequent compilations.

- ❓ For more information on incremental compilation settings, refer to *Incremental Compilation Page* in Quartus II Help.

Team-Based Design Flows and IP Delivery

The Quartus II software supports various design flows to enable team-based design and third-party IP delivery. A top-level design can include one or more partitions that are designed or optimized by different designers or IP providers, as well as partitions that will be developed as part of a standard incremental methodology.

In a team-based environment, part of your design may be incomplete because it is being developed elsewhere. The project lead or system architect can create empty placeholders in the top-level design for partitions that are not yet complete. Designers or IP providers can create and verify HDL code separately, and then the project lead later integrates the code into the single top-level Quartus II project. In this scenario, you can add the completed partitions to the design incrementally, however, the design flow allows all design optimization to occur in the top-level design for easiest design integration. Altera recommends using a single Quartus II project whenever possible because using multiple projects can add significant up-front and debugging time to the development cycle.

Alternatively, partition designers can design their partition in a copy of the top-level design, or in a separate Quartus II project. Designers export their completed partition as either a post-synthesis netlist, or optimized placed and routed netlist, or both, along with assignments such as LogicLock™ regions, as appropriate. The project lead then integrates each design block as a design partition into the top-level design. Altera recommends that designers export and reuse post-synthesis netlists, unless optimized post-fit results are required in the top-level design, to simplify design optimization.

Teams with a bottom-up design approach often want to optimize placement and routing of design partitions independently and may want to create separate Quartus II projects for each partition. However, optimizing design partitions in separate Quartus II projects, and then later integrating the results into a top-level design, can have the following potential drawbacks that require careful planning:

- Achieving timing closure for the full design may be more difficult if you compile partitions independently without information about other partitions in the design. This problem may be avoided by careful timing budgeting and special design rules, such as always registering the ports at the module boundaries.
- Resource budgeting and allocation may be required to avoid resource conflicts and overuse. Creating a floorplan with LogicLock regions is recommended when design partitions are developed independently in separate Quartus II projects.
- Maintaining consistency of assignments and timing constraints can be more difficult if there are separate Quartus II projects. The project lead must ensure that the top-level design and the separate projects are consistent in their assignments.

A unique challenge of team-based design and IP delivery for FPGAs is the fact that the partitions being developed independently must share a common set of resources. To minimize issues that might arise from sharing a common set of resources, you can design partitions within a single Quartus II project, or a copy of the top-level design. A common project ensures that designers have a consistent view of the top-level project framework.

For timing-critical partitions being developed and optimized by another designer, it is important that each designer has complete information about the top-level design in order to maintain timing closure during integration, and to obtain the best results. When you want to integrate partitions from separate Quartus II projects, the project lead can perform most of the design planning, and then pass the top-level design constraints to the partition designers. Preferably, partition designers can obtain a copy of the top-level design by checking out the required files from a source control system. Alternatively, the project lead can provide a copy of the top-level project framework, or pass design information using Quartus II-generated design partition scripts. In the case that a third-party designer has no information about the top-level design, developers can export their partition from an independent project if required.

For more information about managing team-based design flows, refer to [“Exporting Design Partitions from Separate Quartus II Projects”](#) on page 2-26 and the subsection [“Project Management—Making the Top-Level Design Available to Other Designers”](#) on page 2-28.



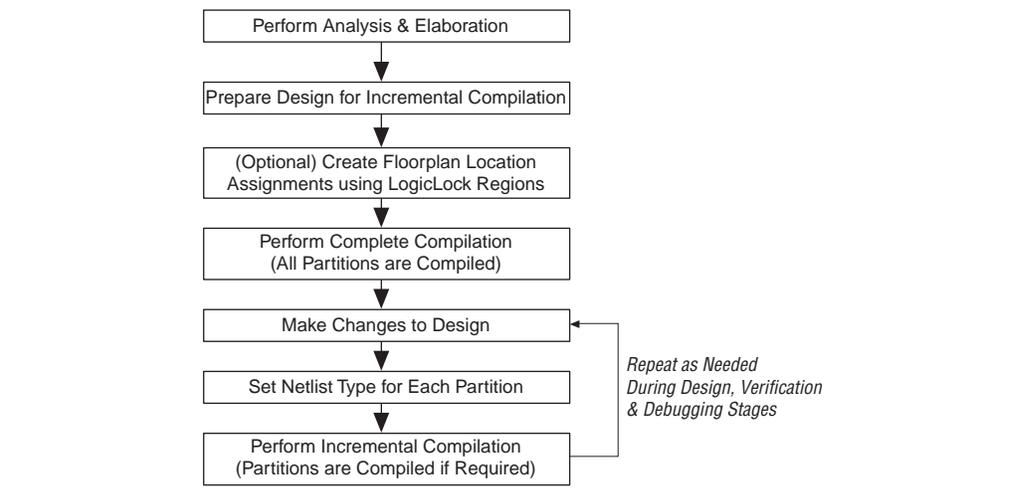
Exporting partitions is not supported in HardCopy or FPGA companion device compilations when there is a migration device setting. For details, refer to [“Limitations for HardCopy Compilation and Migration Flows”](#) on page 2-48.

Incremental Compilation Summary

This section provides a summary of the standard incremental compilation design flow and describes how to create design partitions.

Figure 2-2 illustrates the incremental compilation design flow when all partitions are contained in one top-level design.

Figure 2-2. Summary of Standard Incremental Compilation Design Flow



Steps for Incremental Compilation

This section summarizes the steps in an incremental compilation flow; preparing a design to use the incremental compilation feature, and then preserving satisfactory results and performance in subsequent incremental compilations.

- ④ For an interactive introduction to implementing an incremental compilation design flow, refer to the **Getting Started Tutorial** on the Help menu in the Quartus II software. For a step-by-step introduction on how to use incremental compilation, refer to *Using the Incremental Compilation Design Flow* in Quartus II Help.

Preparing a Design for Incremental Compilation

To begin, elaborate your design, or run any compilation flow (such as a full compilation) that includes the elaboration step. Elaboration is the part of the synthesis process that identifies your design's hierarchy.

Next, designate specific instances in the design hierarchy as design partitions, as described in *“Creating Design Partitions”* on page 2-9.

If required for your design flow, create a floorplan with LogicLock regions location assignments for timing-critical partitions that change with future compilations. Assigning a partition to a physical region on the device can help maintain quality of results and avoid conflicts in certain situations. Refer to *“Creating a Design Floorplan With LogicLock Regions”* on page 2-44 for details about LogicLock region assignments.

Compiling a Design Using Incremental Compilation

The first compilation after making partition assignments is a full compilation, and prepares the design for subsequent incremental compilations. In subsequent compilations of your design, you can preserve satisfactory compilation results and performance of unchanged partitions with the **Netlist Type** setting in the Design Partitions window. The **Netlist Type** setting determines which type of netlist or source file the Partition Merge stage uses in the next incremental compilation. You can choose to use the Source File, Post-Synthesis netlist, or Post-Fit netlist. For details about the **Netlist Type** setting, refer to “[Specifying the Level of Results Preservation for Subsequent Compilations](#)” on page 2-21.

Creating Design Partitions

There are several ways to designate a design instance as a design partition. This section provides an overview of tools you can use to create partitions in the Quartus II software. For information on selecting which design blocks to assign as partitions and how to analyze the quality of your partition assignments, refer to “[Deciding Which Design Blocks Should Be Design Partitions](#)” on page 2-14.

Creating Design Partitions in the Project Navigator

You can right-click an instance in the list under the **Hierarchy** tab in the Project Navigator and use the sub-menu to create and delete design partitions.

- For detailed information about how to create design partitions in the Project Navigator, refer to *Creating Design Partitions* in Quartus II Help.

Creating Design Partitions in the Design Partitions Window

The Design Partitions window, available from the Assignments menu, allows you to create, delete, and merge partitions, and is the main window for setting the netlist type to specify the level of results preservation for each partition on subsequent compilations. For information about how to set the netlist type and the available settings, refer to “[Netlist Type for Design Partitions](#)” on page 2-21.

The Design Partitions window also lists recommendations at the bottom of the window with links to the Incremental Compilation Advisor, where you can view additional recommendations about partitions. The **Color** column indicates the color of each partition as it appears in the Design Partition Planner and Chip Planner.

You can right-click a partition in the window to perform various common tasks, such as viewing property information about a partition, including the time and date of the compilation netlists and the partition statistics.

When you create a partition, the Quartus II software automatically generates a name based on the instance name and hierarchy path. You can edit the partition name in the Design Partitions Window so that you avoid referring to them by their hierarchy path, which can sometimes be long. This is especially useful when using command-line commands or assignments, or when you merge partitions to give the partition a meaningful name. Partition names can be from 1 to 1024 characters in length and must be unique. The name can only contain alphanumeric characters and the pipe (|), colon (:), and underscore (_) characters.

- ❓ For more information about how to create and manage design partitions in the Design Partitions window, refer to *Creating Design Partitions* in Quartus II Help.

Creating Design Partitions With the Design Partition Planner

The Design Partition Planner allows you to view design connectivity and hierarchy, and can assist you in creating effective design partitions that follow Altera's guidelines.

The Design Partition Planner displays a visual representation of design connectivity and hierarchy, as well as partitions and entity relationships. You can explore the connectivity between entities in the design, evaluate existing partitions with respect to connectivity between entities, and try new partitioning schemes in "what if" scenarios.

When you extract design blocks from the top-level design and drag them into the Design Partition Planner, connection bundles are drawn between entities, showing the number of connections existing between pairs of entities. In the Design Partition Planner, you can then set extracted design blocks as design partitions.

The Design Partition Planner also has an **Auto-Partition** feature that creates partitions based on the size and connectivity of the hierarchical design blocks.

-  For more details about how to use the Design Partition Planner, refer to *Using the Design Partition Planner* in Quartus II Help and the *Best Practices for Incremental Compilation Partitions and Floorplan Assignments* chapter in volume 1 of the *Quartus II Handbook*.

Creating Design Partitions With Tcl Scripting

You can also create partitions with Tcl scripting commands. For details about the command line and scripting flow, refer to "Scripting Support" on page 2-54.

Automatically-Generated Partitions

The compiler creates some partitions automatically as part of the compilation process, which appear in some post-compilation reports. For example, the `sld_hub` partition is created for tools that use JTAG hub connections, such as the SignalTap II Logic Analyzer. The `hard_block` partition is created to contain certain "hard" or dedicated logic blocks in the device that are implemented in a separate partition so that they can be shared throughout the design.

Common Design Scenarios Using Incremental Compilation

This section provides recommended applications of the incremental compilation flow after you have set up your design with partitions for incremental compilation as described in, "Steps for Incremental Compilation" on page 2-8.

This section contains the following design scenarios:

- "Reducing Compilation Time When Changing Source Files for One Partition" on page 2-11
- "Optimizing a Timing-Critical Partition" on page 2-11
- "Adding Design Logic Incrementally or Working With an Incomplete Design" on page 2-12

- “Debugging Incrementally With the SignalTap II Logic Analyzer” on page 2-13

Reducing Compilation Time When Changing Source Files for One Partition

Scenario background: You set up your design to include partitions for several of the major design blocks, and now you have just performed a lengthy compilation of the entire design. An error is found in the HDL source file for one partition and it is being fixed. Because the design is currently meeting timing requirements, and the fix is not expected to affect timing performance, it makes sense to compile only the affected partition and preserve the rest of the design.

Use the flow in this example to update the source file in one partition without having to recompile the other parts of the design. To reduce the compilation time, instruct the software to reuse the post-fit netlists for the unchanged partitions. This flow also preserves the performance for these blocks, which reduces additional timing closure efforts.

Perform the following steps to update a single source file:

1. Apply and save the fix to the HDL source file.
2. On the Assignments menu, open the **Design Partitions** window.
3. Change the netlist type of each partition, including the top-level entity, to **Post-Fit** to preserve as much as possible for the next compilation.



The Quartus II software recompiles partitions by default when changes are detected in a source file. You can refer to the Partition Dependent Files table in the Analysis and Synthesis report to see which partitions were recompiled. If you change an assignment but do not change the logic in a source file, you can set the netlist type to **Source File** for that partition to instruct the software to recompile the partition's source design files and its assignments.

4. Click **Start Compilation** to incrementally compile the fixed HDL code. This compilation should take much less time than the initial full compilation.
5. Simulate the design to ensure that the error is fixed, and use the Timing Analyzer report to ensure that timing results have not degraded.

Optimizing a Timing-Critical Partition

Scenario background: You have just performed a lengthy full compilation of a design that consists of multiple partitions. The Timing Analyzer reports that the clock timing requirement is not met, and you have to optimize one particular partition. You want to try optimization techniques such as raising the Placement Effort Multiplier, enabling Physical Synthesis, and running the Design Space Explorer. Because these techniques all involve significant compilation time, it makes sense to apply them to only the partition in question.

Use the flow in this example to optimize the results of one partition when the other partitions in the design have already meet their requirements. You can use this flow iteratively to lock down the performance of one partition, and then move on to optimization of another partition.

Perform the following steps to preserve the results for partitions that meet their timing requirements, and to recompile a timing-critical partition with new optimization settings:

1. Open the **Design Partitions** window.
2. For the partition in question, set the netlist type to **Source File**.



If you change a setting that affects only the Fitter, you can save additional compilation time by setting the netlist type to **Post-Synthesis** to reuse the synthesis results and refit the partition.

3. For the remaining partitions (including the top-level entity), set the netlist type to **Post-Fit**.



You can optionally set the **Fitter Preservation Level** on the **Advanced** tab in the **Design Partitions Properties** dialog box to **Placement** to allow for the most flexibility during routing.

4. Apply the desired optimization settings.
5. Click **Start Compilation** to perform incremental compilation on the design with the new settings. During this compilation, the Partition Merge stage automatically merges the critical partition's new synthesis netlist with the post-fit netlists of the remaining partitions. The Fitter then refits only the required partition. Because the effort is reduced as compared to the initial full compilation, the compilation time is also reduced.

To use the Design Space Explorer, perform the following steps:

1. Repeat steps 1–3 of the previous procedure.
2. Save the project and run the Design Space Explorer.

Adding Design Logic Incrementally or Working With an Incomplete Design

Scenario background: You have one or more partitions that are known to be timing-critical in your full design. You want to focus on developing and optimizing this subset of the design first, before adding the rest of the design logic.

Use this flow to compile a timing-critical partition or partitions in isolation, optionally with extra optimizations turned on. After timing closure is achieved for the critical logic, you can preserve its content and placement and compile the remaining partitions with normal or reduced optimization levels. For example, you may want to compile an IP block that comes with instructions to perform optimization before you incorporate the rest of your custom logic.

To implement this design flow, perform the following steps:

1. Partition the design and create floorplan location assignments. For best results, ensure that the top-level design includes the entire project framework, even if some parts of the design are incomplete and are represented by an empty wrapper file.
2. For the partitions to be compiled first, in the Design Partitions window, set the netlist type to **Source File**.

3. For the remaining partitions, set the netlist type to **Empty**.
4. To compile with the desired optimizations turned on, click **Start Compilation**.
5. Check Timing Analyzer reports to ensure that timing requirements are met. If so, proceed to step 6. Otherwise, repeat steps 4 and 5 until the requirements are met.
6. In the Design Partitions window, set the netlist type to **Post-Fit** for the first partitions. You can set the **Fitter Preservation Level** on the **Advanced** tab in the **Design Partitions Properties** dialog box to **Placement** to allow more flexibility during routing if exact placement and routing preservation is not required.
7. Change the netlist type from **Empty** to **Source File** for the remaining partitions , and ensure that the completed source files are added to the project.
8. Set the appropriate level of optimizations and compile the design. Changing the optimizations at this point does not affect any fitted partitions, because each partition has its netlist type set to **Post-Fit**.
9. Check Timing Analyzer reports to ensure that timing requirements are met. If not, make design or option changes and repeat step 8 and step 9 until the requirements are met.



The flow in this example is similar to design flows in which a module is implemented separately and is later merged into the top-level , such as in the team-based design flow described in “[Designing in a Team-Based Environment](#)” on page 2–38. Generally, optimization in this flow works only if each critical path is contained within a single partition due to the effects described in “[Deciding Which Design Blocks Should Be Design Partitions](#)” on page 2–14. Ensure that if there are any partitions representing a design file that is missing from the project, you create a placeholder wrapper file to define the port interface. Refer to “[Empty Partitions](#)” on page 2–28 for more information.

Debugging Incrementally With the SignalTap II Logic Analyzer

Scenario background: Your design is not functioning as expected, and you want to debug the design using the SignalTap II Logic Analyzer. To maintain low compilation times and to ensure that you do not negatively affect the current version of your design, you want to preserve the synthesis and fitting results and add the SignalTap II Logic Analyzer to your design without recompiling the source code.

Use this flow to reduce compilation times when you add the logic analyzer to debug your design, or when you want to modify the configuration of the SignalTap II File without modifying your design logic or its placement.

It is not necessary to create design partitions in order to use the SignalTap II incremental compilation feature. The SignalTap II Logic Analyzer acts as its own separate design partition.

Perform the following steps to use the SignalTap II Logic Analyzer in an incremental compilation flow:

1. Open the Design Partitions window.
2. Set the netlist type to **Post-fit** for all partitions to preserve their placement.



The netlist type for the top-level partition defaults to **Source File**, so be sure to change this “Top” partition in addition to any design partitions that you have created.

- If you have not already compiled the design with the current set of partitions, perform a full compilation. If the design has already been compiled with the current set of partitions, the design is ready to add the SignalTap II Logic Analyzer.
- Set up your SignalTap II File using the **SignalTap II: post-fitting** filter in the **Node Finder** to add signals for logic analysis. This allows the Fitter to add the SignalTap II logic to the post-fit netlist without modifying the design results.

To add signals from the pre-synthesis netlist, set the partition’s netlist type to **Source File** and use the **SignalTap II: pre-synthesis** filter in the **Node Finder**. This allows the software to resynthesize the partition and to tap directly to the pre-synthesis node names that you choose. In this case, the partition is resynthesized and refit, so the placement is typically different from previous fitting results.



For more information about setting up the SignalTap II Logic Analyzer, refer to the *Design Debugging Using the SignalTap II Embedded Logic Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

Deciding Which Design Blocks Should Be Design Partitions

The incremental compilation design flow requires more up-front planning than flat compilations. For example, you might have to structure your source code or design hierarchy to ensure that logic is grouped correctly for optimization.

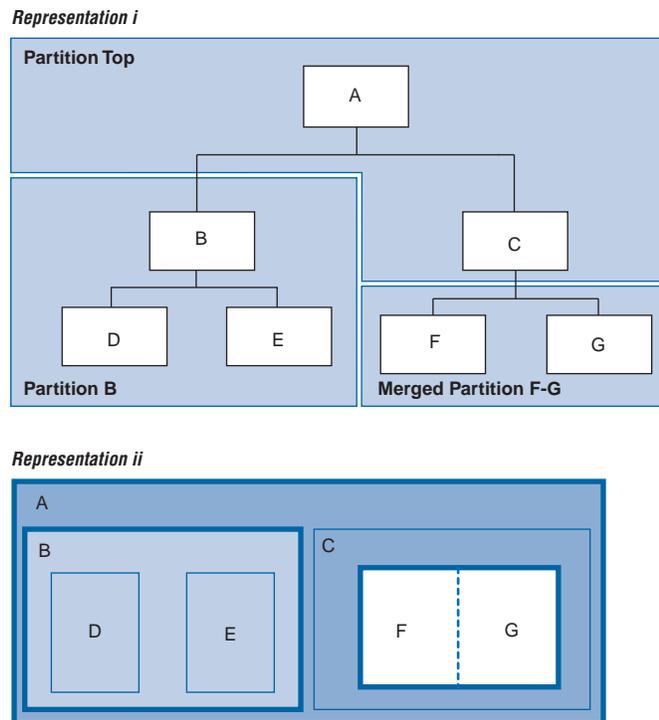
It is a common design practice to create modular or hierarchical designs in which you develop each design entity separately, and then instantiate them in a higher-level entity, forming a complete design. The Quartus II software does not automatically consider each design entity or instance to be a design partition for incremental compilation; instead, you must designate one or more design hierarchies below the top-level project as a design partition. Creating partitions prevents the compiler from performing optimizations across partition boundaries, as discussed in *“Impact of Design Partitions on Design Optimization”* on page 2-16. However, this allows for separate synthesis and placement for each partition, making incremental compilation possible.

Partitions must have the same boundaries as hierarchical blocks in the design because a partition cannot be a portion of the logic within a hierarchical entity. You can merge partitions that have the same immediate parent partition to create a single partition that includes more than one hierarchical entity in the design. When you declare a partition, every hierarchical instance within that partition becomes part of the same partition. You can create new partitions for hierarchical instances within an existing partition, in which case the instances within the new partition are no longer included in the higher-level partition, as described in the following example.

In [Figure 2-3](#), a complete design is made up of instances **A**, **B**, **C**, **D**, **E**, **F**, and **G**. The shaded boxes in Representation i indicate design partitions in a “tree” representation of the hierarchy. In Representation ii, the lower-level instances are represented inside the higher-level instances, and the partitions are illustrated with different colored shading. The top-level partition, called “Top”, automatically contains the top-level

entity in the design, and contains any logic not defined as part of another partition. The design file for the top level may be just a wrapper for the hierarchical instances below it, or it may contain its own logic. In this example, partition **B** contains the logic in instances **B**, **D**, and **E**. Entities **F** and **G** were first identified as separate partitions, and then merged together to create a partition **F-G**. The partition for the top-level entity **A**, called “Top”, includes the logic in one of its lower-level instances, **C**, because **C** was not defined as part of any other partition.

Figure 2-3. Partitions in a Hierarchical Design



You can create partition assignments to any design instance. The instance can be defined in HDL or schematic design, or come from a third-party synthesis tool as a VQM or EDIF netlist instance.

To take advantage of incremental compilation when source files change, create separate design files for each partition. If you define two different entities as separate partitions but they are in the same design file, you cannot maintain incremental compilation because the software would have to recompile both partitions if you changed either entity in the design file. Similarly, if two partitions rely on the same lower-level entity definition, changes in that lower-level affect both partitions.

The remainder of this section provides information to help you choose which design blocks you should assign as partitions.

Impact of Design Partitions on Design Optimization

The boundaries of your design partitions can impact the design's quality of results. Creating partitions prevents the compiler from performing logic optimizations across partition boundaries, which allows the software to synthesize and place each partition separately in an incremental flow. Therefore, consider partitioning guidelines to help reduce the effect of partition boundaries.

Whenever possible, register all inputs and outputs of each partition. This helps avoid any delay penalty on signals that cross partition boundaries and keeps each register-to-register timing path within one partition for optimization. In addition, minimize the number of paths that cross partition boundaries. If there are timing-critical paths that cross partition boundaries, rework the partitions to avoid these inter-partition paths. Including as many of the timing-critical connections as possible inside a partition allows you to effectively apply optimizations to that partition to improve timing, while leaving the rest of the design unchanged.

Avoid constant partition inputs and outputs, because to maintain incremental behavior, the software cannot use the constants to optimize logic on either side of the partition boundary. You can also merge two or more partitions to allow cross-boundary optimizations for paths that cross between the partitions, as long as the partitions have the same parent partition. Merging related logic from different hierarchy blocks into one partition can be useful if you cannot change the design hierarchy to accommodate partition assignments.

The Design Partition Planner can help you create good assignments, as described in [“Creating Design Partitions” on page 2–9](#). Refer to [“Partition Statistics Reports” on page 2–18](#), for information about the number of I/O connections and how many are unregistered or driven by a constant value. For information on timing reports and additional design guidelines, refer to [“Partition Timing Reports” on page 2–19](#) and [“Incremental Compilation Advisor” on page 2–19](#).

If critical timing paths cross partition boundaries, you can perform timing budgeting and make timing assignments to constrain the logic in each partition so that the entire timing path meets its requirements. In addition, because each partition is optimized independently during synthesis, you may have to perform resource allocation to ensure that each partition uses an appropriate number of device resources. If design partitions are compiled in separate Quartus II projects, there may be conflicts related to global routing resources for clock signals when the design is integrated into the top-level design. You can use the Global Signal logic option to specify which clocks should use global or regional routing, use the ALTCLK_CTRL megafunction to instantiate a clock control block and connect it appropriately in both the partitions being developed in separate Quartus II projects, or find the compiler-generated clock control node in your design and make clock control location assignments in the Assignment Editor.



For more partitioning guidelines and specific recommendations for fixing common design issues, as well as information on resource allocation, global signal usage, and timing budgeting, refer to the [Best Practices for Incremental Compilation Partitions and Floorplan Assignments](#) chapter in volume 1 of the *Quartus II Handbook*.

Design Partition Assignments Compared to Physical Placement Assignments

Design partitions for incremental compilation are logical partitions, which is different from physical placement assignments in the device floorplan. A logical design partition does not refer to a physical area of the device and does not directly control the placement of instances. A logical design partition sets up a virtual boundary between design hierarchies so that each is compiled separately, preventing logical optimizations from occurring between them. When the software compiles the design source code, the logic in each partition can be placed anywhere in the device unless you make additional placement assignments.

If you preserve the compilation results using a Post-Fit netlist, it is not necessary for you to back-annotate or make any location assignments for specific logic nodes. You should not use the incremental compilation and logic placement back-annotation features in the same Quartus II project. The incremental compilation feature does not use placement “assignments” to preserve placement results; it simply reuses the netlist database that includes the placement information.

You can assign design partitions to physical regions in the device floorplan using LogicLock region assignments. In the Quartus II software, LogicLock regions are used to constrain blocks of a design to a particular region of the device. Altera recommends using LogicLock regions for timing-critical design blocks that will change in subsequent compilations, or to improve the quality of results and avoid placement conflicts in some cases. Creating floorplan location assignments for design partitions using LogicLock regions is discussed in [“Creating a Design Floorplan With LogicLock Regions” on page 2-44](#).



For more information about when and why to create a design floorplan, refer to the [Best Practices for Incremental Compilation Partitions and Floorplan Assignments](#) chapter in volume 1 of the *Quartus II Handbook*.

Using Partitions With Third-Party Synthesis Tools

If you are using a third-party synthesis tool, set up your tool to create a separate VQM or EDIF netlist for each hierarchical partition. In the Quartus II software, assign the top-level entity from each netlist to be a design partition. The VQM or EDIF netlist file is treated as the source file for the partition in the Quartus II software.

Synopsys Synplify Pro/Premier and Mentor Graphics Precision RTL Plus

The Synplify Pro and Synplify Premier software include the MultiPoint synthesis feature to perform incremental synthesis for each design block assigned as a Compile Point in the user interface or a script. The Precision RTL Plus software includes an incremental synthesis feature that performs block-based synthesis based on Partition assignments in the source HDL code. These features provide automated block-based incremental synthesis flows and create different output netlist files for each block when set up for an Altera device.

Using incremental synthesis within your synthesis tool ensures that only those sections of a design that have been updated are resynthesized when the design is compiled, reducing synthesis run time and preserving the results for the unchanged blocks. You can change and resynthesize one section of a design without affecting other sections of the design.

- For more information about these incremental synthesis flows, refer to your tool vendor's documentation, or the appropriate chapter in volume 1 of the *Quartus II Handbook*: *Synopsys Synplify Support* or *Mentor Graphics Precision Synthesis Support*.

Other Synthesis Tools

You can also partition your design and create different netlist files manually with the basic Synplify software (non-Pro/Premier), the basic Precision RTL software (non-Plus), or any other supported synthesis tool by creating a separate project or implementation for each partition, including the top level. Set up each higher-level project to instantiate the lower-level VQM/EDIF netlists as black boxes. Synplify, Precision, and most synthesis tools automatically treat a design block as a black box if the logic definition is missing from the project. Each tool also includes options or attributes to specify that the design block should be treated as a black box, which you can use to avoid warnings about the missing logic.

Assessing Partition Quality

The Quartus II software provides various tools to assess the quality of your assigned design partitions. You can take advantage of these tools to assess your partition quality, and use the information to improve your design or assignments as required to achieve the best results.

- For more information about ensuring good partition quality, refer to the *Best Practices for Incremental Compilation Partitions and Floorplan Assignments* chapter in volume 1 of the *Quartus II Handbook*.

Partition Statistics Reports

After compilation, you can view statistics about design partitions in the Partition Merge Partition Statistics report, and on the **Statistics** tab in the **Design Partitions Properties** dialog box.

The Partition Merge Partition Statistics report lists statistics about each partition. The statistics for each partition (each row in the table) include the number of logic cells it contains, as well as the number of input and output pins it contains, and how many are registered or unconnected. This report is useful when optimizing your design partitions, ensuring that the partitions meet the guidelines presented in the *Best Practices for Incremental Compilation Partitions and Floorplan Assignments* chapter in volume 1 of the *Quartus II Handbook*.

Figure 2-4 shows the report window.

Figure 2-4. Partition Merge Partition Statistics Report

You can also view post-compilation statistics about the resource usage and port connections for a particular partition on the **Statistics** tab in the **Design Partition Properties** dialog box.

Partition Timing Reports

You can generate a Partition Timing Overview report and a Partition Timing Details report by clicking **Report Partitions** in the Tasks pane in the TimeQuest Timing Analyzer, or using the `report_partitions` Tcl command.

The Partition Timing Overview report shows the total number of failing paths for each partition and the worst-case slack for any path involving the partition.

The Partition Timing Details report shows the number of failing partition-to-partition paths and worst-case slack for partition-to-partition paths, to provide a more detailed breakdown of where the critical paths in the design are located with respect to design partitions.

Incremental Compilation Advisor

You can use the Incremental Compilation Advisor to check that your design follows Altera’s recommendations for creating design partitions and floorplan location assignments.

As shown in Figure 2-5, recommendations are split into **General Recommendations**, **Timing Recommendations**, and **Team-Based Design Recommendations** that apply to design flows in which partitions are compiled independently in separate Quartus II projects before being integrated into the top-level design. Each recommendation provides an explanation, describes the effect of the recommendation, and provides the action required to make a suggested change. In some cases, there is a link to the appropriate Quartus II settings page where you can make a suggested change to

assignments or settings. For some items, if your design does not follow the recommendation, the **Check Recommendations** operation creates a table that lists any nodes or paths in your design that could be improved. The relevant timing-independent recommendations for the design are also listed in the Design Partitions window and the LogicLock Regions window.

Figure 2-5. Incremental Compilation Advisor

The screenshot shows the Incremental Compilation Advisor window. The left pane contains a tree view of recommendations, with 'Register All Non-Global Ports' selected. The right pane displays the details for this recommendation, including a description, summary, and an action table listing unregistered ports.

Partition Name	Unregistered Port Name	Port Type	Node Driven By Port
1 des:des_inst2	des_blocks:desbl...	Input	des_blocks:desbl...
2 des:des_inst2	des_blocks:desbl...	Input	des_blocks:desbl...
3 des:des_inst2	des_blocks:desbl...	Input	des_blocks:desbl...
4 des:des_inst2	des_blocks:desbl...	Input	des_blocks:desbl...

To verify that your design follows the recommendations, go to the **Timing Independent Recommendations** page or the **Timing Dependent Recommendations** page, and then click **Check Recommendations**. For large designs, these operations can take a few minutes.

After you perform a check operation, symbols appear next to each recommendation to indicate whether the design or project setting follows the recommendations, or if some or all of the design or project settings do not follow the recommendations. Following these recommendations is not mandatory to use the incremental compilation feature. The recommendations are most important to ensure good results for timing-critical partitions.

For some items in the Advisor, if your design does not follow the recommendation, the **Check Recommendations** operation lists any parts of the design that could be improved. For example, if not all of the partition I/O ports follow the **Register All Non-Global Ports** recommendation, the advisor displays a list of unregistered ports with the partition name and the node name associated with the port.

When the advisor provides a list of nodes, you can right-click a node, and then click **Locate** to cross-probe to other Quartus II features, such as the RTL Viewer, Chip Planner, or the design source code in the text editor.



Opening a new TimeQuest report resets the Incremental Compilation Advisor results, so you must rerun the Check Recommendations process.

Specifying the Level of Results Preservation for Subsequent Compilations

As introduced in [“Incremental Compilation Summary”](#) on page 2-7 and [“Common Design Scenarios Using Incremental Compilation”](#) on page 2-10, the netlist type of each design partition allows you to specify the level of results preservation. The netlist type determines which type of netlist or source file the Partition Merge stage uses in the next incremental compilation.

When you choose to preserve a post-fit compilation netlist, the default level of Fitter preservation is the highest degree of placement and routing preservation supported by the device family. The advanced Fitter Preservation Level setting allows you to specify the amount of information that you want to preserve from the post-fit netlist file.

Netlist Type for Design Partitions

Before starting a new compilation, ensure that the appropriate netlist type is set for each partition to preserve the desired level of compilation results. [Table 2-2](#) describes the settings for the netlist type, explains the behavior of the Quartus II software for each setting, and provides guidance on when to use each setting.

Table 2-2. Partition Netlist Type Settings (Part 1 of 2)

Netlist Type	Quartus II Behavior for Partition During Compilation
Source File	<p>Always compiles the partition using the associated design source file(s). (1)</p> <p>Use this netlist type to recompile a partition from the source code using new synthesis or Fitter settings.</p>
Post-Synthesis	<p>Preserves post-synthesis results for the partition and reuses the post-synthesis netlist as long as the following conditions are true:</p> <ul style="list-style-type: none"> ■ A post-synthesis netlist is available from a previous synthesis. ■ No change that initiates an automatic resynthesis has been made to the partition since the previous synthesis. (2) For details, refer to “What Changes Initiate a Partition’s Automatic Resynthesis?” on page 2-24. <p>Compiles the partition from the source files if resynthesis is initiated or if a post-synthesis netlist is not available. (1)</p> <p>Use this netlist type to preserve the synthesis results unless you make design changes, but allow the Fitter to refit the partition using any new Fitter settings.</p>
Post-Fit	<p>Preserves post-fit results for the partition and reuses the post-fit netlist as long as the following conditions are true:</p> <ul style="list-style-type: none"> ■ A post-fit netlist is available from a previous fitting. ■ No change that initiates an automatic resynthesis has been made to the partition since the previous fitting. (2) For details, refer to “What Changes Initiate a Partition’s Automatic Resynthesis?” on page 2-24. <p>When a post-fit netlist is not available, the software reuses the post-synthesis netlist if it is available, or otherwise compiles from the source files. Compiles the partition from the source files if resynthesis is initiated. (1)</p> <p>The Fitter Preservation Level specifies what level of information is preserved from the post-fit netlist. For details, refer to “Fitter Preservation Level for Design Partitions” on page 2-22.</p> <p>Assignment changes, such as Fitter optimization settings, do not cause a partition set to Post-Fit to recompile.</p>

Table 2-2. Partition Netlist Type Settings (Part 2 of 2)

Netlist Type	Quartus II Behavior for Partition During Compilation
Empty	<p>Uses an empty placeholder netlist for the partition. The partition's port interface information is required during Analysis and Synthesis to connect the partition correctly to other logic and partitions in the design, and peripheral nodes in the source file including pins and PLLs are preserved to help connect the empty partition to the rest of the design and preserve timing of any lower-level non-empty partitions within empty partitions. If the source file is not available, you can create a wrapper file that defines the design block and specifies the input, output, and bidirectional ports. In Verilog HDL: a module declaration, and in VHDL: an entity and architecture declaration.</p> <p>You can use this netlist type to skip the compilation of a partition that is incomplete or missing from the top-level design. You can also set an empty partition if you want to compile only some partitions in the design, such as to optimize the placement of a timing-critical block such as an IP core before incorporating other design logic, or if the compilation time is large for one partition and you want to exclude it.</p> <p>If the project database includes a previously generated post-synthesis or post-fit netlist for an unchanged Empty partition, you can set the netlist type from Empty directly to Post-Synthesis or Post-Fit and the software reuses the previous netlist information without recompiling from the source files.</p>

Notes to Table 2-2:

- (1) If you turn on the **Rapid Recompile** option, the Quartus II software may not recompile the entire partition from the source code as described in this table; it will reuse compatible results if there have been only small changes to the logic in the partition. Refer to [“Incremental Capabilities Available When A Design Has No Partitions”](#) on page 2-2 for more information.
- (2) You can turn on the **Ignore changes in source files and strictly use the specified netlist, if available** option on the **Advanced** tab in the **Design Partitions Properties** dialog box to specify whether the Compiler should ignore source file changes when deciding whether to recompile the partition.

Fitter Preservation Level for Design Partitions

The default Fitter Preservation Level for partitions with a **Post-Fit** netlist type is the highest level of preservation available for the target device family and provides the most compilation time reduction.

You can change the advanced Fitter Preservation Level setting to provide more flexibility in the Fitter during placement and routing. You can set the Fitter Preservation Level on the **Advanced** tab in the **Design Partitions Properties** dialog box. [Table 2-3](#) describes the Fitter Preservation Level settings.

Table 2-3. Fitter Preservation Level Settings (Part 1 of 2)

Fitter Preservation Level	Quartus II Behavior for Partition During Compilation
Placement and Routing	<p>Preserves the design partition's netlist atoms and their placement and routing.</p> <p>This setting reduces compilation times compared to Placement only, but provides less flexibility to the router to make changes if there are changes in other parts of the design.</p> <p>By default, the Fitter preserves the usage of high-speed programmable power tiles contained within the selected partition, for devices that support high-speed and low-power tiles. You can turn off the Preserve high-speed tiles when preserving placement and routing option on the Advanced tab in the Design Partitions Properties dialog box.</p>

Table 2-3. Fitter Preservation Level Settings (Part 2 of 2)

Fitter Preservation Level	Quartus II Behavior for Partition During Compilation
Placement	Preserves the netlist atoms and their placement in the design partition. Re-routes the design partition and does not preserve high-speed power tile usage.
Netlist Only	<p>Preserves the netlist atoms of the design partition, but replaces and reroutes the design partition. A post-fit netlist with the atoms preserved can be different than the Post-Synthesis netlist because it contains Fitter optimizations; for example, Physical Synthesis changes made during a previous Fitting.</p> <p>You can use this setting to:</p> <ul style="list-style-type: none"> ■ Preserve Fitter optimizations but allow the software to perform placement and routing again. ■ Reapply certain Fitter optimizations that would otherwise be impossible when the placement is locked down. ■ Resolve resource conflicts between two imported partitions.

❓ For detailed information about how to set the **Netlist Type** and **Fitter Preservation Level** settings in the Quartus II software, refer to *Setting the Netlist Type and Fitter Preservation Level for Design Partitions* in Quartus II Help.

Where Are the Netlist Databases Saved?

The incremental compilation database folder (`\incremental_db`) includes all the netlist information from previous compilations. To avoid unnecessary recompilations, these database files must not be altered or deleted.

If you archive or reproduce the project in another location, you can use a Quartus II Archive File (`.qar`). Include the incremental compilation database files to preserve post-synthesis or post-fit compilation results. For details, refer to “Using Incremental Compilation With Quartus II Archive Files” on page 2-48.

To manually create a project archive that preserves compilation results without keeping the incremental compilation database, you can keep all source and settings files, and create and save a Quartus II Settings File (`.qxp`) for each partition in the design that will be integrated into the top-level design. Refer to “Exporting Design Partitions from Separate Quartus II Projects” on page 2-26 for more details about how to create a `.qxp` for a partition within your design.

Deleting Netlists

You can choose to abandon all levels of results preservation and remove all netlists that exist for a particular partition with the **Delete Netlists** command in the Design Partitions window. When you delete netlists for a partition, the partition is compiled using the associated design source file(s) in the next compilation. Resetting the netlist type for a partition to **Source** would have the same effect, though the netlists would not be permanently deleted and would be available for use in subsequent compilations. For an imported partition, the **Delete Netlists** command also optionally allows you to remove the imported `.qxp`.

What Changes Initiate a Partition's Automatic Resynthesis?

A partition is synthesized from its source files if there is no post-synthesis netlist available from a previous synthesis, or if the netlist type is set to **Source File**. Additionally, certain changes to a partition initiate an automatic resynthesis of the partition when the netlist type is **Post-Synthesis** or **Post-Fit**. The software resynthesizes the partition in these cases to ensure that the design description matches the post-place-and-route programming files. If you do not want this resynthesis to occur automatically, refer to [“Forcing Use of the Compilation Netlist When a Partition has Changed” on page 2-26](#).

The following list explains the changes that initiate a partition's automatic resynthesis when the netlist type is set to **Post-Synthesis** or **Post-Fit**:

- The device family setting has changed.
- Any dependent source design file has changed. Refer to [“Resynthesis Due to Source Code Changes” on page 2-25](#) for details.
- The partition boundary was changed by an addition, removal, or change to the port boundaries of a partition (for example, a new partition has been defined for a lower-level instance within this partition).
- A dependent source file was compiled into a different library (so it has a different `-library` argument).
- A dependent source file was added or removed; that is, the partition depends on a different set of source files.
- The partition's root instance has a different entity binding. In VHDL, an instance may be bound to a specific entity and architecture. If the target entity or architecture changes, it triggers resynthesis.
- The partition has different parameters on its root hierarchy or on an internal AHDL hierarchy (AHDL automatically inherits parameters from its parent hierarchies). This occurs if you modified the parameters on the hierarchy directly, or if you modified them indirectly by changing the parameters in a parent design hierarchy.
- You have moved the project and compiled database between a Windows and Linux system. Due to the differences in the way new line feeds are handled between the operating systems, the internal checksum algorithm may detect a design file change in this case.

The software reuses the post-synthesis results but re-fits the design if you change the device setting within the same device family. The software reuses the post-fitting netlist if you change only the device speed grade.

Synthesis and Fitter assignments, such as optimization settings, timing assignments, or Fitter location assignments including pin assignments, do not trigger automatic recompilation in the incremental compilation flow. To recompile a partition with new assignments, change the netlist type for that partition to one of the following:

- **Source File** to recompile with all new settings
- **Post-Synthesis** to recompile using existing synthesis results but new Fitter settings

- **Post-Fit** with the **Fitter Preservation Level** set to **Placement** to rerun routing using existing placement results, but new routing settings (such as delay chain settings)

You can use the LogicLock Origin location assignment to change or fine-tune the previous Fitter results from a Post-Fit netlist. For details about how you can affect placement with LogicLock regions, refer to [“Changing Partition Placement with LogicLock Changes”](#) on page 2-46.

Resynthesis Due to Source Code Changes

The Quartus II software uses an internal checksum algorithm to determine whether the contents of a source file have changed. Source files are the design description files used to create the design, and include Memory Initialization Files (.mif) as well as .qxp from exported partitions. When design files in a partition have dependencies on other files, changing one file may initiate an automatic recompilation of another file. The Partition Dependent Files table in the Analysis and Synthesis report lists the design files that contribute to each design partition. You can use this table to determine which partitions are recompiled when a specific file is changed.

For example, if a design has file **A.v** that contains entity **A**, **B.v** that contains entity **B**, and **C.v** that contains entity **C**, then the Partition Dependent Files table for the partition containing entity **A** lists file **A.v**, the table for the partition containing entity **B** lists file **B.v**, and the table for the partition containing entity **C** lists file **C.v**. Any dependencies are transitive, so if file **A.v** depends on **B.v**, and **B.v** depends on **C.v**, the entities in file **A.v** depend on files **B.v** and **C.v**. In this case, files **B.v** and **C.v** are listed in the report table as dependent files for the partition containing entity **A**.



If you turn on the **Rapid Recompile** option, the Quartus II software may not recompile the entire partition from the source code as described in this section; it will reuse compatible results if there have been only small changes to the logic in the partition. Refer to [“Incremental Capabilities Available When A Design Has No Partitions”](#) on page 2-2 for more information.

If you define module parameters in a higher-level module, the Quartus II software checks the parameter values when determining which partitions require resynthesis. If you change a parameter in a higher-level module that affects a lower-level module, the lower-level module is resynthesized. Parameter dependencies are tracked separately from source file dependencies; therefore, parameter definitions are not listed in the Partition Dependent Files list.

If a design contains common files, such as an **includes.v** file that is referenced in each entity by the command `include includes.v`, all partitions are dependent on this file. A change to **includes.v** causes the entire design to be recompiled. The VHDL statement `use work.all` also typically results in unnecessary recompilations, because it makes all entities in the work library visible in the current entity, which results in the current entity being dependent on all other entities in the design.

To avoid this type of problem, ensure that files common to all entities, such as a common include file, contain only the set of information that is truly common to all entities. Remove `use work.all` statements in your VHDL file or replace them by including only the specific design units needed for each entity.

Forcing Use of the Compilation Netlist When a Partition has Changed

Forcing the use of a post-compilation netlist when the contents of a source file has changed is recommended only for advanced users who understand when a partition must be recompiled. You might use this assignment, for example, if you are making source code changes but do not want to recompile the partition until you finish debugging a different partition, or if you are adding simple comments to the source file but you know the design logic itself is not being changed and you want to keep the previous compilation results.

To force the Fitter to use a previously generated netlist even when there are changes to the source files, right-click the partition in the Design Partitions window and then click **Design Partition Properties**. On the **Advanced** tab, turn on the **Ignore changes in source files and strictly use the specified netlist, if available** option.

Turning on this option can result in the generation of a functionally incorrect netlist when source design files change, because source file updates will not be recompiled. Use caution when setting this option.

Exporting Design Partitions from Separate Quartus II Projects

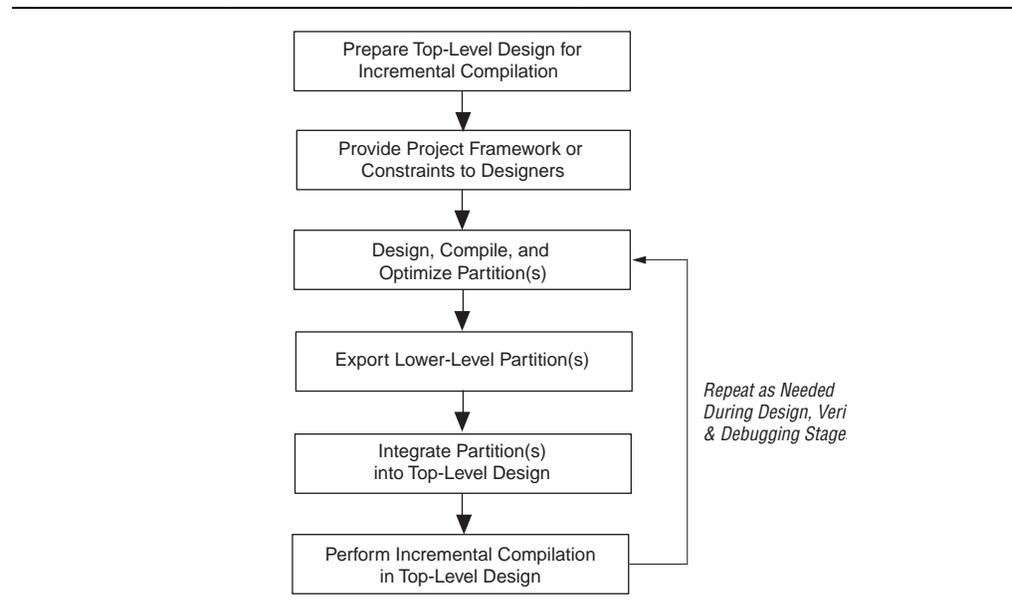
Partitions that are developed by other designers or team members in the same company or third-party IP providers can be exported as design partitions to a Quartus II Exported Partition File (**.qxp**), and then integrated into a top-level design. A **.qxp** is a binary file that contains compilation results describing the exported design partition and includes a post-synthesis netlist, a post-fit netlist, or both, and a set of assignments, sometimes including LogicLock placement constraints. The **.qxp** does not contain the source design files from the original Quartus II project.

To enable team-based development and third-party IP delivery, you can design and optimize partitions in separate copies of the top-level Quartus II project framework, or even in isolation. If the designers have access to the top-level project framework through a source control system, they can access project files as read-only and develop their partition within the source control system. If designers do not have access to a source control system, the project lead can provide the designer with a copy of the top-level project framework to use as they develop their partitions. The project lead also has the option to generate design partition scripts to manage resource and timing budgets in the top-level design when partitions are developed outside the top-level project framework.

The exported compilation results of completed partitions are given to the project lead, preferably using a source control system, who is then responsible for integrating them into the top-level design to obtain a fully functional design. This type of design flow is required only if partition designers want to optimize their placement and routing independently, and pass their design to the project lead to reuse placement and routing results. Otherwise, a project lead can integrate source HDL from several designers in a single Quartus II project, and use the standard incremental compilation flow described previously.

The diagram in [Figure 2-6](#) illustrates the team-based incremental compilation design flow using a methodology in which partitions are compiled in separate Quartus II projects before being integrated into the top-level design. This flow can be used when partitions are developed by other designers or IP providers.

Figure 2-6. Summary of Team-Based Incremental Compilation Flow



 You cannot export or import partitions that have been merged. For more information about merged partitions, refer to [“Deciding Which Design Blocks Should Be Design Partitions”](#) on page 2-14.

The topics in this section provide a description of the team-based design flow using exported partitions, describe how to generate a **.qxp** for a design partition, and explain how to integrate the **.qxp** into the top-level design:

There are some additional restrictions related to design flows using exported partitions, described in [“Incremental Compilation Restrictions”](#) on page 2-47.

Preparing the Top-Level Design

To prepare your design to incorporate exported partitions, first create the top-level project framework of the design to define the hierarchy for the subdesigns that will be implemented by other team members, designers, or IP providers.

In the top-level design, create project-wide settings, for example, device selection, global assignments for clocks and device I/O ports, and any global signal constraints to specify which signals can use global routing resources.

Next, create the appropriate design partition assignments and set the netlist type for each design partition that will be developed in a separate Quartus II project to **Empty**. Refer to “[Empty Partitions](#)” below for details. It may be necessary to constrain the location of partitions with LogicLock region assignments if they are timing-critical and are expected to change in future compilations, or if the designer or IP provider wants to place and route their design partition independently, to avoid location conflicts. For details, refer to “[Creating a Design Floorplan With LogicLock Regions](#)” on page 2-44.

Finally, provide the top-level project framework to the partition designers, preferably through a source control system. Refer to “[Project Management—Making the Top-Level Design Available to Other Designers](#)” on page 2-28 for more information.

Empty Partitions

You can use a design flow in which some partitions are set to an **Empty** netlist type to develop pieces of the design separately, and then integrate them into the top-level design at a later time. In a team-based design environment, you can set the netlist type to **Empty** for partitions in your design that will be developed by other designers or IP providers. The **Empty** setting directs the Compiler to skip the compilation of a partition and use an empty placeholder netlist for the partition.

When a netlist type is set to **Empty**, peripheral nodes including pins and PLLs are preserved and all other logic is removed. The peripheral nodes including pins help connect the empty partition to the design, and the PLLs help preserve timing of non-empty partitions within empty partitions.

When you set a design partition to **Empty**, a design file is required during Analysis and Synthesis to specify the port interface information so that it can connect the partition correctly to other logic and partitions in the design. If a partition is exported from another project, the **.qxp** contains this information. If there is no **.qxp** or design file to represent the design entity, you must create a wrapper file that defines the design block and specifies the input, output, and bidirectional ports. For example, in Verilog HDL, you should include a module declaration, and in VHDL, you should include an entity and architecture declaration.

Project Management—Making the Top-Level Design Available to Other Designers

In team-based incremental compilation flows, whenever possible, all designers or IP providers should work within the same top-level project framework. Using the same project framework among team members ensures that designers have the settings and constraints needed for their partition, and makes timing closure easier when integrating the partitions into the top-level design. If other designers do not have access to the top-level project framework, the Quartus II software provides tools for passing project information to partition designers.

Distributing the Top-Level Quartus II Project

There are several methods that the project lead can use to distribute the “skeleton” or top-level project framework to other partition designers or IP providers.

- If partition designers have access to the top-level project framework, the project will already include all the settings and constraints needed for the design. This framework should include PLLs and other interface logic if this information is important to optimize partitions.
 - If designers are part of the same design environment, they can check out the required project files from the same source control system. This is the recommended way to share a set of project files.
 - Otherwise, the project lead can provide a copy of the top-level project framework so that each design develops their partition within the same project framework.
- If a partition designer does not have access to the top-level project framework, the project lead can give the partition designer a Tcl script or other documentation to create the separate Quartus II project and all the assignments from the top-level design.

For details about project management scripts you can create with the Quartus II software, refer to “[Generating Design Partition Scripts](#)” on page 2-56.

If the partition designers provide the project lead with a post-synthesis .qxp and fitting is performed in the top-level design, integrating the design partitions should be quite easy. If you plan to develop a partition in a separate Quartus II project and integrate the optimized post-fitting results into the top-level design, use the following guidelines to improve the integration process:

- Ensure that a LogicLock region constrains the partition placement and uses only the resources allocated by the project lead.
- Ensure that you know which clocks should be allocated to global routing resources so that there are no resource conflicts in the top-level design.
 - Set the Global Signal assignment to **On** for the high fan-out signals that should be routed on global routing lines.
 - To avoid other signals being placed on global routing lines, turn off **Auto Global Clock and Auto Global Register Controls** under **More Settings** on the Fitter page in the **Settings** dialog box. Alternatively, you can set the Global Signal assignment to **Off** for signals that should not be placed on global routing lines.

Placement for LABs depends on whether the inputs to the logic cells within the LAB use a global clock. You may encounter problems if signals do not use global lines in the partition, but use global routing in the top-level design.

- Use the Virtual Pin assignment to indicate pins of a partition that do not drive pins in the top-level design. This is critical when a partition has more output ports than the number of pins available in the target device. Using virtual pins also helps optimize cross-partition paths for a complete design by enabling you to provide more information about the partition ports, such as location and timing assignments.

- When partitions are compiled independently without any information about each other, you might need to provide more information about the timing paths that may be affected by other partitions in the top-level design. You can apply location assignments for each pin to indicate the port location after incorporation in the top-level design. You can also apply timing assignments to the I/O ports of the partition to perform timing budgeting.

 For more information about resource balancing and timing allocation between partitions, refer to the *Best Practices for Incremental Compilation Partitions and Floorplan Assignments* chapter in volume 1 of the *Quartus II Handbook*.

Generating Design Partition Scripts

If IP providers or designers on a team want to optimize their design blocks independently and do not have access to a shared project framework, the project lead must perform some or all of the following tasks to ensure successful integration of the design blocks:

- Determine which assignments should be propagated from the top-level design to the partitions. This requires detailed knowledge of which assignments are required to set up low-level designs.
- Communicate the top-level assignments to the partitions. This requires detailed knowledge of Tcl or other scripting languages to efficiently communicate project constraints.
- Determine appropriate timing and location assignments that help overcome the limitations of team-based design. This requires examination of the logic in the partitions to determine appropriate timing constraints.
- Perform final timing closure and resource conflict avoidance in the top-level design. Because the partitions have no information about each other, meeting constraints at the lower levels does not guarantee they are met when integrated at the top-level. It then becomes the project lead's responsibility to resolve the issues, even though information about the partition implementation may not be available.

Design partition scripts automate the process of transferring the top-level project framework to partition designers in a flow where each design block is developed in separate Quartus II projects before being integrated into the top-level design. If the project lead cannot provide each designer with a copy of the top-level project framework, the Quartus II software provides an interface for managing resources and timing budgets in the top-level design. Design partition scripts make it easier for partition designers to implement the instructions from the project lead, and avoid conflicts between projects when integrating the partitions into the top-level design. This flow also helps to reduce the need to further optimize the designs after integration.

You can use options in the **Generate Design Partition Scripts** dialog box to choose which types of assignments you want to pass down and create in the partitions being developed in separate Quartus II projects.

For an example design scenario using design partition scripts, refer to “[Enabling Designers on a Team to Optimize Independently](#)” on page 2-39.

- ② For step-by-step information on how to generate design partition scripts, and a description of each option that can be included in design partition scripts, refer to *Generating Design Partition Scripts for Project Management*, and *Generate Design Partition Scripts Dialog Box* in Quartus II Help.

Exporting Partitions

When partition designers achieve the design requirements in their separate Quartus II projects, each designer can export their design as a partition so it can be integrated into the top-level design by the project lead. The **Export Design Partition** dialog box, available from the Project menu, allows designers to export a design partition to a Quartus II Exported Partition File (.qxp) with a post-synthesis netlist, a post-fit netlist, or both. The project lead then adds the .qxp to the top-level design to integrate the partition.

A designer developing a timing-critical partition or who wants to optimize their partition on their own would opt to export their completed partition with a post-fit netlist, allowing for the partition to more reliably meet timing requirements after integration. In this case, you must ensure that resources are allocated appropriately to avoid conflicts. If the placement and routing optimization can be performed in the top-level design, exporting a post-synthesis netlist allows the most flexibility in the top-level design and avoids potential placement or routing conflicts with other partitions.

When designing the partition logic to be exported into another project, you can add logic around the design block to be exported as a design partition. You can instantiate additional design components for the Quartus II project so that it matches the top-level design environment, especially in cases where you do not have access to the full top-level design project. For example, you can include a top-level PLL in the project, outside of the partition to be exported, so that you can optimize the design with information about the frequency multipliers, phase shifts, compensation delays, and any other PLL parameters. The software then captures timing and resource requirements more accurately while ensuring that the timing analysis in the partition is complete and accurate. You can export the partition for the top-level design without any auxiliary components that are instantiated outside the partition being exported.

If your design team uses makefiles and design partition scripts, the project lead can use the **make** command with the **master_makefile** command created by the scripts to export the partitions and create .qxp files. When a partition has been compiled and is ready to be integrated into the top-level design, you can export the partition with option on the **Export Design Partition** dialog box, available from the Project menu.

- ② For more information about how to export a design partition, refer to *Using a Team-Based Incremental Compilation Design Flow* in the Quartus II Help.

Viewing the Contents of a Quartus II Exported Partition File (.qxp)

The QXP report allows you to view a summary of the contents in a .qxp when you open the file in the Quartus II software. The .qxp is a binary file that contains compilation results so the file cannot be read in a text editor. The QXP report opens in the main Quartus II window and contains summary information including a list of the I/O ports, resource usage summary, and a list of the assignments used for the exported partition.

Integrating Partitions into the Top-Level Design

To integrate a partition developed in a separate Quartus II project into the top-level design, you can simply add the **.qxp** as a source file in your top-level design (just like a Verilog or VHDL source file). You can also use the **Import Design Partition** dialog box to import the partition, in certain situations, described in “[Advanced Importing Options](#)” on page 2-33.

The **.qxp** contains the design block exported from the partition and has the same name as the partition. When you instantiate the design block into a top-level design and include the **.qxp** as a source file, the software adds the exported netlist to the database for the top-level design. The **.qxp** port names are case sensitive if the original HDL of the partition was case sensitive.

When you use a **.qxp** as a source file in this way, you can choose whether you want the **.qxp** to be a partition in the top-level design. If you do not designate the **.qxp** instance as a partition, the software reuses just the post-synthesis compilation results from the **.qxp**, removes unconnected ports and unused logic just like a regular source file, and then performs placement and routing.

If you assigned the **.qxp** instance as a partition, you can set the netlist type in the Design Partitions Window to choose the level of results to preserve from the **.qxp**. To preserve the placement and routing results from the exported partition, set the netlist type to **Post-Fit** for the **.qxp** partition in the top-level design. If you assign the instance as a design partition, the partition boundary is preserved, as discussed in “[Impact of Design Partitions on Design Optimization](#)” on page 2-16.

Integrating Assignments from the .qxp

The Quartus II software filters assignments from **.qxp** files to include appropriate assignments in the top-level design. The assignments in the **.qxp** are treated like assignments made in an HDL source file, and are not listed in the Quartus II Settings File (**.qsf**) for the top-level design. Most assignments from the **.qxp** can be overridden by assignments in the top-level design.

The following subsections provide more details about specific assignment types:

Design Partition Assignments Within the Exported Partition

Design partition assignments defined within a separate Quartus II project are not added to the top-level design. All logic under the exported partition in the project hierarchy is treated as single instance in the **.qxp**.

Synopsys Design Constraint Files for the Quartus II TimeQuest Timing Analyzer

Timing assignments made for the Quartus II TimeQuest Timing Analyzer in a Synopsys Design Constraint File (**.sdc**) in the lower-level partition project are not added to the top-level design. Ensure that the top-level design includes all of the timing requirements for the entire project.



For recommendations about managing SDC constraints for the top-level design and independent lower-level partition projects, refer to the *Best Practices for Incremental Compilation Partitions and Floorplan Assignments* chapter in volume 1 of the *Quartus II Handbook*.

Global Assignments

The project lead should make all global project-wide assignments in the top-level design. Global assignments from the exported partition's project are not added to the top-level design. When it is possible for a particular constraint, the global assignment is converted to an instance-specific assignment for the exported design partition.

LogicLock Region Assignments

The project lead typically creates LogicLock region assignments in the top-level design for any lower-level partition designs where designer or IP providers plan to export post-fit information to be used in the top-level design, to help avoid placement conflicts between partitions. When you use the `.qxp` as a source file, LogicLock constraints from the exported partition are applied in the top-level design, but will not appear in your `.qsf` file or LogicLock Regions window for you to view or edit. The LogicLock region itself is not required to constrain the partition placement in the top-level design if the netlist type is set to **Post-Fit**, because the netlist contains all the placement information. For information on how to control LogicLock region assignments for exported partitions, refer to the [“Advanced Importing Options” on page 2-33](#).

Integrating Encrypted IP Cores from `.qxp` Files

Proper license information is required to compile encrypted IP cores. If an IP core is exported as a `.qxp` from another Quartus II project, the top-level designer instantiating the `.qxp` must have the correct license. The software requires a full license to generate an unrestricted programming file. If you do not have a license, but the IP in the `.qxp` was compiled with OpenCore Plus hardware evaluation support, you can generate an evaluation programming file without a license. If the IP supports OpenCore simulation only, you can fully compile the design and generate a simulation netlist, but you cannot create programming files unless you have a full license.

Advanced Importing Options

You can use advanced options in the **Import Design Partition** dialog box to integrate a partition developed in a separate Quartus II project into the top-level design. The import process adds more control than using the `.qxp` as a source file, and is useful only in the following circumstances:

- **If you want LogicLock regions in your top-level design (.qsf)**—If you have regions in your partitions that are not also in the top-level design, the regions will be added to your `.qsf` file during the import process.
- **If you want different settings or placement for different instantiations of the same entity**—You can control the setting import process with the advanced import options, and specify different settings for different instances of the same `.qxp` design block.

When you use the **Import Design Partition** dialog box to integrate a partition into the top-level design, the import process sets the partition's netlist type to **Imported** in the Design Partitions window.

After you compile the entire design, if you make changes to the place-and-route results (such as movement of an imported LogicLock region), use the **Post-Fit** netlist type on subsequent compilations. To discard an imported netlist and recompile from source code, you can compile the partition with the netlist type set to **Source File** and be sure to include the relevant source code in the top-level design. Refer to “[Netlist Type for Design Partitions](#)” on page 2-21 for details. The import process sets the partition’s Fitter Preservation Level to the setting with the highest degree of preservation supported by the imported netlist. For example, if a post-fit netlist is imported with placement information, the Fitter Preservation Level is set to **Placement**, but you can change it to the **Netlist Only** value. For more information about preserving previous compilation results, refer to “[Netlist Type for Design Partitions](#)” on page 2-21 and “[Fitter Preservation Level for Design Partitions](#)” on page 2-22.

When you import a partition from a **.qxp**, the **.qxp** itself is not part of the top-level design because the netlists from the file have been imported into the project database. Therefore if a new version of a **.qxp** is exported, the top-level designer must perform another import of the new **.qxp**.

When you import a partition into a top-level design with the **Import Design Partition** dialog box, the software imports relevant assignments from the partition into the top-level design, as described for the source file integration flow in “[Integrating Assignments from the .qxp](#)” on page 2-32. If required, you can change the way some assignments are imported, as described in the following subsections.

Importing LogicLock Assignments

LogicLock regions are set to a fixed size when imported. If you instantiate multiple instances of a subdesign in the top-level design, the imported LogicLock regions are set to a Floating location. Otherwise, they are set to a Fixed location. You can change the location of LogicLock regions after they are imported, or change them to a Floating location to allow the software to place each region but keep the relative locations of nodes within the region wherever possible. For details, refer to “[Changing Partition Placement with LogicLock Changes](#)” on page 2-46. To preserve changes made to a partition after compilation, use the **Post-Fit** netlist type.

The LogicLock Member State assignment is set to **Locked** to signify that it is a preserved region.

LogicLock back-annotation and node location data is not imported because the **.qxp** contains all of the relevant placement information. Altera strongly recommends that you do not add to or delete members from an imported LogicLock region.

Advanced Import Settings

The **Advanced Import Settings** dialog box allows you to disable assignment import and specify additional options that control how assignments and regions are integrated when importing a partition into a top-level design, including how to resolve assignment conflicts.

- ② For descriptions of the advanced import options available, refer to [Advanced Import Settings Dialog Box](#) in Quartus II Help.

Team-Based Design Optimization and Third-Party IP Delivery Scenarios

This section includes the following design flows with step-by-step descriptions when you have partitions being developed in separate Quartus II projects, or by a third-party IP provider.

- “Using an Exported Partition to Send to a Design Without Including Source Files” on page 2-35
- “Creating Precompiled Design Blocks (or Hard-Wired Macros) for Reuse” on page 2-36
- “Designing in a Team-Based Environment” on page 2-38
- “Enabling Designers on a Team to Optimize Independently” on page 2-39
- “Performing Design Iterations With Lower-Level Partitions” on page 2-42

Using an Exported Partition to Send to a Design Without Including Source Files

Scenario background: A designer wants to produce a design block and needs to send out their design, but to preserve their IP, they prefer to send a synthesized netlist instead of providing the HDL source code to the recipient.

Use this flow to package a full design as a single source file to send to an end customer or another design location.

As the sender in this scenario perform the following steps to export a design block:

1. Provide the device family name to the recipient. If you send placement information with the synthesized netlist, also provide the exact device selection so they can set up their project to match.
2. Create a black box wrapper file that defines the port interface for the design block and provide it to the recipient for instantiating the block as an empty partition in the top-level design.
3. Create a Quartus II project for the design block, and complete the design.
4. Export the level of hierarchy into a single **.qxp**. Following a successful compilation of the project, you can generate a **.qxp** from the GUI, the command-line, or with Tcl commands, as described in the following:
 - If you are using the Quartus II GUI, use the **Export Design Partition** command.
 - If you are using command-line executables, run **quartus_cdb** with the **--incremental_compilation_export** option.
 - If you are using Tcl commands, use the following command:
`execute_flow -incremental_compilation_export.`
5. Select the option to include just the **Post-synthesis netlist** if you do not have to send placement information. If the recipient wants to reproduce your exact Fitter results, you can select the **Post-fitting netlist** option, and optionally enable **Export routing**.
6. Provide the **.qxp** to the recipient. Note that you do not have to send any of your design source code.

As the recipient in this example, first create a Quartus II project for your top-level design and ensure that your project targets the same device (or at least the same device family if the `.qxp` does not include placement information), as specified by the IP designer sending the design block. Instantiate the design block using the port information provided, and then incorporate the design block into a top-level design.

Add the `.qxp` from the IP designer as a source file in your Quartus II project to replace any empty wrapper file. If you want to use just the post-synthesis information, you can choose whether you want the file to be a partition in the top-level design. To use the post-fit information from the `.qxp`, assign the instance as a design partition and set the netlist type to **Post-Fit**. Refer to “[Creating Design Partitions](#)” on page 2-9 and “[Netlist Type for Design Partitions](#)” on page 2-21.

Creating Precompiled Design Blocks (or Hard-Wired Macros) for Reuse

Scenario background: An IP provider wants to produce and sell an IP core for a component to be used in higher-level systems. The IP provider wants to optimize the placement of their block for maximum performance in a specific Altera device and then deliver the placement information to their end customer. To preserve their IP, they also prefer to send a compiled netlist instead of providing the HDL source code to their customer.

Use this design flow to create a precompiled IP block (sometimes known as a hard-wired macro) that can be instantiated in a top-level design. This flow provides the ability to export a design block with post-synthesis or placement (and, optionally, routing) information and to import any number of copies of this pre-compiled block into another design.

The customer first specifies which Altera device is being used for this project and provides the design specifications.

As the IP provider in this example, perform the following steps to export a preplaced IP core (or hard macro):

1. Create a black box wrapper file that defines the port interface for the IP core and provide the file to the customer to instantiate as an empty partition in the top-level design.
2. Create a Quartus II project for the IP core.
3. Create a LogicLock region for the design hierarchy to be exported.



Using a LogicLock region for the IP core allows the customer to create an empty placeholder region to reserve space for the IP in the design floorplan and ensures that there are no conflicts with the top-level design logic. Reserved space also helps ensure the IP core does not affect the timing performance of other logic in the top-level design. Additionally, with a LogicLock region, you can preserve placement either absolutely or relative to the origin of the associated region. This is important when a `.qxp` is imported for multiple partition hierarchies in the same project, because in this case, the location of at least one instance in the top-level design does not match the location used by the IP provider.

4. If required, add any logic (such as PLLs or other logic defined in the customer's top-level design) around the design hierarchy to be exported. If you do so, create a design partition for the design hierarchy that will be exported as an IP core.
5. Optimize the design and close timing to meet the design specifications.
6. Export the level of hierarchy for the IP core into a single **.qxp**.
7. Provide the **.qxp** to the customer. Note that you do not have to send any of your design source code to the customer; the design netlist and placement and routing information is contained within the **.qxp**.

As the customer in this example, incorporate the IP core in your design by performing the following steps:

1. Create a Quartus II project for the top-level design that targets the same device and instantiate a copy or multiple copies of the IP core. Use a black box wrapper file to define the port interface of the IP core.
 2. Perform Analysis and Elaboration to identify the design hierarchy.
 3. Create a design partition for each instance of the IP core (refer to [“Creating Design Partitions” on page 2-54](#)) with the netlist type set to **Empty** (refer to [“Netlist Type for Design Partitions” on page 2-21](#)).
 4. You can now continue work on your part of the design and accept the IP core from the IP provider when it is ready.
 5. Include the **.qxp** from the IP provider in your project to replace the empty wrapper-file for the IP instance. Or, if you are importing multiple copies of the design block and want to import relative placement, follow these additional steps:
 - a. Use the **Import** command to select each appropriate partition hierarchy. You can import a **.qxp** from the GUI, the command-line, or with Tcl commands:
 - If you are using the Quartus II GUI, use the **Import Design Partition** command.
 - If you are using command-line executables, run **quartus_cdb** with the `--incremental_compilation_import` option.
 - If you are using Tcl commands, use the following command:

```
execute_flow -incremental_compilation_import.
```
 - b. When you have multiple instances of the IP block, you can set the imported LogicLock regions to floating, or move them to a new location, and the software preserves the relative placement for each of the imported modules (relative to the origin of the LogicLock region). Routing information is preserved whenever possible. Refer to [“Changing Partition Placement with LogicLock Changes” on page 2-46](#)
-  The Fitter ignores relative placement assignments if the LogicLock region's location in the top-level design is not compatible with the locations exported in the **.qxp**.
6. You can control the level of results preservation with the **Netlist Type** setting. Refer to [“Netlist Type for Design Partitions” on page 2-21](#).



If the IP provider did not define a LogicLock region in the exported partition, the software preserves absolute placement locations and this leads to placement conflicts if the partition is imported for more than one instance.

Designing in a Team-Based Environment

Scenario background: A project includes several lower-level design blocks that are developed separately by different designers and instantiated exactly once in the top-level design.

This scenario describes how to use incremental compilation in a team-based design environment where each designer has access to the top-level project framework, but wants to optimize their design in a separate Quartus II project before integrating their design block into the top-level design.

As the project lead in this scenario, perform the following steps to prepare the top-level design:

1. Create a new Quartus II project to ultimately contain the full implementation of the entire design and include a "skeleton" or framework of the design that defines the hierarchy for the subdesigns implemented by separate designers. The top-level design implements the top-level entity in the design and instantiates wrapper files that represent each subdesign by defining only the port interfaces but not the implementation.
2. Make project-wide settings. Select the device, make global assignments such as device I/O ports, define the top-level timing constraints, and make any global signal allocation constraints to specify which signals can use global routing resources.
3. Make design partition assignments for each subdesign and set the netlist type for each design partition to be imported to **Empty** in the Design Partitions window.
4. Create LogicLock regions to create a design floorplan for each of the partitions that will be developed separately. This floorplan should consider the connectivity between partitions and estimates of the size of each partition based on any initial implementation numbers and knowledge of the design specifications.
5. Provide the top-level project framework to partition designers using one of the following procedures:
 - Allow access to the full project for all designers through a source control system. Each designer can check out the projects files as read-only and work on their blocks independently. This design flow provides each designer with the most information about the full design, which helps avoid resource conflicts and makes design integration easy.
 - Provide a copy of the top-level Quartus II project framework for each designer. You can use the **Copy Project** command on the Project menu or create a project archive.

As the designer of a lower-level design block in this scenario, design and optimize your partition in your copy of the top-level design, and then follow these steps when you have achieved the desired compilation results:

1. On the Project menu, click **Export Design Partition**.

2. In the **Export Design Partition** dialog box, choose the netlist(s) to export. You can export a Post-synthesis netlist if placement or performance preservation is not required, to provide the most flexibility for the Fitter in the top-level design. Select Post-fit netlist to preserve the placement and performance of the lower-level design block, and turn on **Export routing** to include the routing information, if required. One **.qxp** can include both post-synthesis and post-fitting netlists.
3. Provide the **.qxp** to the project lead.

Finally, as the project lead in this scenario, perform these steps to integrate the **.qxp** files received from designers of each partition:

1. Add the **.qxp** as a source file in the Quartus II project, to replace any empty wrapper file for the previously **Empty** partition.
2. Change the netlist type for the partition from **Empty** to the required level of results preservation.

Enabling Designers on a Team to Optimize Independently

Scenario background: A project consists of several lower-level design blocks that are developed separately by different designers who do not have access to a shared top-level project framework. This scenario is similar to the [“Creating Precompiled Design Blocks \(or Hard-Wired Macros\) for Reuse” on page 2–36](#) scenario, but assumes that there are several design blocks being developed independently (instead of just one IP block), and the project lead can provide some information about the design to the individual designers. If the designers have shared access to the top-level design, use the previous scenario [“Designing in a Team-Based Environment” on page 2–38](#).

This scenario describes how to use incremental compilation in a team-based design environment where designers or IP developers want to fully optimize the placement and routing of their design independently in a separate Quartus II project before sending the design to the project lead. In this scenario, the IP developers do not have access to the top-level project framework. This design flow requires more planning and careful resource allocation because design blocks are developed independently.

As the project lead in this scenario, perform the following steps to prepare the top-level design:

1. Create a new Quartus II project to ultimately contain the full implementation of the entire design and include a “skeleton” or framework of the design that defines the hierarchy for the subdesigns implemented by separate designers. The top-level design implements the top-level entity in the design and instantiates wrapper files that represent each subdesign by defining only the port interfaces but not the implementation.
2. Make project-wide settings. Select the device, make global assignments such as device I/O ports, define the top-level timing constraints, and make any global signal constraints to specify which signals can use global routing resources.
3. Make design partition assignments for each subdesign and set the netlist type for each design partition to be imported to **Empty** in the Design Partitions window.
4. Create LogicLock regions. This floorplan should consider the connectivity between partitions and estimates of the size of each partition based on any initial implementation numbers and knowledge of the design specifications.

5. Provide the constraints from the top-level design to partition designers using one of the following procedures:
 - Use design partition scripts to pass constraints and generate separate Quartus II projects. On the Project menu, use the **Generate Design Partition Scripts** command, or run the script generator from a Tcl or command prompt. Make changes to the default script options as required for your project. Altera recommends that you pass all the default constraints, including LogicLock regions, for all partitions and virtual pin location assignments. If partitions have not already been created by the other designers, use the partition script to set up the projects so that you can easily take advantage of makefiles. Provide each partition designer with the Tcl file to create their project with the appropriate constraints. If you are using makefiles, provide the makefile for each partition.
 - Use documentation or manually-created scripts to pass all constraints and assignments to each partition designer.

As the designer of a lower-level design block in this scenario, perform the appropriate set of steps to successfully export your design, whether the design team is using makefiles or exporting and importing the design manually.

If you are using makefiles with the design partition scripts, perform the following steps:

1. Use the **make** command and the makefile provided by the project lead to create a Quartus II project with all design constraints, and compile the project.
2. The information about which source file should be associated with which partition is not available to the software automatically, so you must specify this information in the makefile. You must specify the dependencies before the software rebuilds the project after the initial call to the makefile.
3. When you have achieved the desired compilation results and the design is ready to be imported into the top-level design, the project lead can use the `master_makefile` command to export this partition and create a **.qxp**, and then import it into the top-level design.

If you are not using makefiles, perform the following steps:

1. If you are using design partition scripts, source the Tcl script provided by the Project Lead to create a project with the required settings:
 - To source the Tcl script in the Quartus II software, on the Tools menu, click **Utility Windows** to open the Tcl console. Navigate to the script's directory, and type the following command: `source <filename>` ↵
 - To source the Tcl script at the system command prompt, type the following command: `quartus_cdb -t <filename>.tcl` ↵

2. If you are not using design partition scripts, create a new Quartus II project for the subdesign, and then apply the following settings and constraints to ensure successful integration:
 - Make LogicLock region assignments and global assignments (including clock settings) as specified by the project lead.
 - Make Virtual Pin assignments for ports which represent connections to core logic instead of external device pins in the top-level design.
 - Make floorplan location assignments to the Virtual Pins so they are placed in their corresponding regions as determined by the top-level design. This provides the Fitter with more information about the timing constraints between modules. Alternatively, you can apply timing I/O constraints to the paths that connect to virtual pins.
3. Proceed to compile and optimize the design as needed.
4. When you have achieved the desired compilation results, on the Project menu, click **Export Design Partition**.
5. In the **Export Design Partition** dialog box, choose the netlist(s) to export. You can export a Post-synthesis netlist instead if placement or performance preservation is not required, to provide the most flexibility for the Fitter in the top-level design. Select **Post-fit** to preserve the placement and performance of the lower-level design block, and turn on Export routing to include the routing information, if required. One **.qxp** can include both post-synthesis and post-fitting netlists.
6. Provide the **.qxp** to the project lead.

Finally, as the project lead in this scenario, perform the appropriate set of steps to import the **.qxp** files received from designers of each partition.

If you are using makefiles with the design partition scripts, perform the following steps:

1. Use the `master_makefile` command to export each partition and create **.qxp** files, and then import them into the top-level design.
2. The software does not have all the information about which source files should be associated with which partition, so you must specify this information in the makefile. The software cannot rebuild the project if source files change unless you specify the dependencies.

If you are not using makefiles, perform the following steps:

1. Add the **.qxp** as a source file in the Quartus II project, to replace any empty wrapper file for the previously Empty partition.
2. Change the netlist type for the partition from Empty to the required level of results preservation.

Resolving Assignment Conflicts During Integration

When integrating lower-level design blocks, the project lead may notice some assignment conflicts. This can occur, for example, if the lower-level design block designers changed their LogicLock regions to account for additional logic or placement constraints, or if the designers applied I/O port timing constraints that differ from constraints added to the top-level design by the project lead. The project

lead can address these conflicts by explicitly importing the partitions into the top-level design, and using options in the **Advanced Import Settings** dialog box, as described in “[Advanced Importing Options](#)” on page 2-33. After the project lead obtains the **.qxp** for each lower-level design block from the other designers, use the **Import Design Partition** command on the Project menu and specify the partition in the top-level design that is represented by the lower-level design block **.qxp**. Repeat this import process for each partition in the design. After you have imported each partition once, you can select all the design partitions and use the **Reimport using latest import files at previous locations** option to import all the files from their previous locations at one time. To address assignment conflicts, the project lead can take one or both of the following actions:

- Allow new assignments to be imported
- Allow existing assignments to be replaced or updated

When LogicLock region assignment conflicts occur, the project lead may take one of the following actions:

- Allow the imported region to replace the existing region
- Allow the imported region to update the existing region
- Skip assignment import for regions with conflicts

If the placement of different lower-level design blocks conflict, the project lead can also set the partition's **Fitter Preservation Level** to **Netlist Only**, which allows the software to re-perform placement and routing with the imported netlist.

Importing a Partition to be Instantiated Multiple Times

In this variation of the design scenario, one of the lower-level design blocks is instantiated more than once in the top-level design. The designer of the lower-level design block may want to compile and optimize the entity once under a partition, and then import the results as multiple partitions in the top-level design.

If you import multiple instances of a lower-level design block into the top-level design, the imported LogicLock regions are automatically set to Floating status.

If you resolve conflicts manually, you can use the import options and manual LogicLock assignments to specify the placement of each instance in the top-level design.

Performing Design Iterations With Lower-Level Partitions

Scenario background: A project consists of several lower-level subdesigns that have been exported from separate Quartus II projects and imported into the top-level design. In this example, integration at the top level has failed because the timing requirements are not met. The timing requirements might have been met in each individual lower-level project, but critical inter-partition paths in the top-level design are causing timing requirements to fail.

After trying various optimizations in the top-level design, the project lead determines that the design cannot meet the timing requirements given the current partition placements that were imported. The project lead decides to pass additional information to the lower-level partitions to improve the placement.

Use this flow if you re-optimize partitions exported from separate Quartus II projects by incorporating additional constraints from the integrated top-level design.

The best way to provide top-level design information to designers of lower-level partitions is to provide the complete top-level project framework using the following steps:

1. For all partitions other than the one(s) being optimized by a designer(s) in a separate Quartus II project(s), set the netlist type to **Post-Fit**.
2. Make the top-level design directory available in a shared source control system, if possible. Otherwise, copy the entire top-level design project directory (including database files), or create a project archive including the post-compilation database.
3. Provide each partition designer with a checked-out version or copy of the top-level design.
4. The partition designers recompile their designs within the new project framework that includes the rest of the design's placement and routing information as well top-level resource allocations and assignments, and optimize as needed.
5. When the results are satisfactory and the timing requirements are met, export the updated partition as a **.qxp**.

If this design flow is not possible, you can generate partition-specific scripts for individual designs to provide information about the top-level project framework with these steps:

1. In the top-level design, on the Project menu, click **Generate Design Partition Scripts**, or launch the script generator from Tcl or the command line.
2. If lower-level projects have already been created for each partition, you can turn off the **Create lower-level project if one does not exist** option.
3. Make additional changes to the default script options, as necessary. Altera recommends that you pass all the default constraints, including LogicLock regions, for all partitions and virtual pin location assignments. Altera also recommends that you add a maximum delay timing constraint for the virtual I/O connections in each partition.
4. The Quartus II software generates Tcl scripts for all partitions, but in this scenario, you would focus on the partitions that make up the cross-partition critical paths. The following assignments are important in the script:
 - Virtual pin assignments for module pins not connected to device I/O ports in the top-level design.
 - Location constraints for the virtual pins that reflect the initial top-level placement of the pin's source or destination. These help make the lower-level placement "aware" of its surroundings in the top-level design, leading to a greater chance of timing closure during integration at the top level.
 - INPUT_MAX_DELAY and OUTPUT_MAX_DELAY timing constraints on the paths to and from the I/O pins of the partition. These constrain the pins to optimize the timing paths to and from the pins.

5. The partition designers source the file provided by the project lead.
 - To source the Tcl script from the Quartus II GUI, on the Tools menu, click **Utility Windows** and open the Tcl console. Navigate to the script's directory, and type the following command: `source <filename>` ←
 - To source the Tcl script at the system command prompt, type the following command: `quartus_cdb -t <filename>.tcl` ←
6. The partition designers recompile their designs with the new project information or assignments and optimize as needed. When the results are satisfactory and the timing requirements are met, export the updated partition as a **.qxp**.

The project lead obtains the updated **.qxp** files from the partition designers and adds them to the top-level design. When a new **.qxp** is added to the files list, the software will detect the change in the "source file" and use the new **.qxp** results during the next compilation. If the project uses the advanced import flow, the project lead must perform another import of the new **.qxp**.

You can now analyze the design to determine whether the timing requirements have been achieved. Because the partitions were compiled with more information about connectivity at the top level, it is more likely that the inter-partition paths have improved placement which helps to meet the timing requirements.

Creating a Design Floorplan With LogicLock Regions

A floorplan represents the layout of the physical resources on the device. Creating a design floorplan, or floorplanning, describes the process of mapping the logical design hierarchy onto physical regions in the device floorplan. After you have partitioned the design, you can create floorplan location assignments for the design to improve the quality of results when using the incremental compilation design flow. Creating a design floorplan is not a requirement to use an incremental compilation flow, but it is recommended in certain cases. Floorplan location planning can be important for a design that uses incremental compilation for the following reasons:

- To avoid resource conflicts between partitions, predominantly when partitions are imported from another Quartus II project
- To ensure a good quality of results when recompiling individual timing-critical partitions

Design floorplan assignments prevent the situation in which the Fitter must place a partition in an area of the device where most resources are already used by other partitions. A physical region assignment provides a reasonable region to re-place logic after a change, so the Fitter does not have to scatter logic throughout the available space in the device.

Floorplan assignments are not required for non-critical partitions compiled as part of the top-level design. The logic for partitions that are not timing-critical (such as simple top-level glue logic) can be placed anywhere in the device on each recompilation, if that is best for your design.

The simplest way to create a floorplan for a partitioned design is to create one LogicLock region per partition (including the top-level partition). If you have a compilation result for a partitioned design with no LogicLock regions, you can use the Chip Planner with the Design Partition Planner to view the partition placement in the device floorplan. You can draw regions in the floorplan that match the general location and size of the logic in each partition. Or, initially, you can set each region with the default settings of **Auto** size and **Floating** location to allow the Quartus II software to determine the preliminary size and location for the regions. Then, after compilation, use the Fitter-determined size and origin location as a starting point for your design floorplan. Check the quality of results obtained for your floorplan location assignments and make changes to the regions as needed. Alternatively, you can perform synthesis, and then set the regions to the required size based on resource estimates. In this case, use your knowledge of the connections between partitions to place the regions in the floorplan.

Once you have created an initial floorplan, you can refine the region using tools in the Quartus II software. You can also use advanced techniques such as creating non-rectangular regions by merging LogicLock regions.

 For more information about when creating a design floorplan can be important, as well as guidelines for creating the floorplan, refer to the *Best Practices for Incremental Compilation Partitions and Floorplan Assignments* chapter in volume 1 of the *Quartus II Handbook*.

You can use the Incremental Compilation Advisor to check that your LogicLock regions meet Altera's guidelines, as described in "[Incremental Compilation Advisor](#)" on page 2-19.

Creating and Manipulating LogicLock Regions

Options in the **LogicLock Regions Properties** dialog box, available from the Assignments menu, allow you to enter specific sizing and location requirements for a region. You can also view and refine the size and location of LogicLock regions in the Quartus II Chip Planner. You can select a region in the graphical interface in the Chip Planner and use handles to move or resize the region.

Options in the **Layer Settings** panel in the Chip Planner allow you to create, delete, and modify tasks to determine which objects, including LogicLock regions and design partitions, to display in the Chip Planner.

 For more information about creating and viewing LogicLock regions in the LogicLock Regions window and Chip Planner, refer to *Creating and Manipulating LogicLock Regions* in Quartus II Help.

Changing Partition Placement with LogicLock Changes

When a partition is assigned to a LogicLock region as part of a design floorplan, you can modify the placement of a post-fit partition by moving the LogicLock region. As described in “[What Changes Initiate a Partition’s Automatic Resynthesis?](#)” on [page 2-24](#), most assignment changes do not initiate a recompilation of a partition if the netlist type specifies that Fitter results should be preserved. For example, changing a pin assignment does not initiate a recompilation; therefore, the design does not use the new pin assignment unless you change the netlist type to **Post-Synthesis** or **Source File**.

Similarly, if a partition’s placement is preserved, and the partition is assigned to a LogicLock region, the Fitter always reuses the corresponding LogicLock region size specified in the post-fit netlist. That is, changes to the LogicLock **Size** setting do not initiate refitting if a partition’s placement is preserved with the **Post-Fit** netlist type, or with **.qxp** that includes post-fit information.

However, you can use the LogicLock **Origin** location assignment to change or fine-tune the previous Fitter results. When you change the **Origin** setting for a region, the Fitter can move the region in the following manner, depending upon how the placement is preserved for that region's members:

- When you set a new region Origin, the Fitter uses the new origin and replaces the logic, preserving the relative placement of the member logic.
- When you set the region Origin to **Floating**, the following conditions apply:
 - If the region’s member placement is preserved with an imported partition, the Fitter chooses a new Origin and re-places the logic, preserving the relative placement of the member logic within the region.
 - If the region’s member placement is preserved with a **Post-Fit** netlist type, the Fitter does not change the Origin location, and reuses the previous placement results.

Taking Advantage of the Early Timing Estimator

When creating a floorplan you can take advantage of the Early Timing Estimator to enable quick compilations of the design while creating assignments. The Early Timing Estimator feature provides a timing estimate for a design without having to run a full compilation. You can use the Chip Planner to view the “placement estimate” created by this feature, identify critical paths by locating from the timing analyzer reports, and, if necessary, add or modify floorplan constraints. You can then rerun the Early Timing Estimator to quickly assess the impact of any floorplan location assignments or logic changes, enabling rapid iterations on design variants to help you find the best solution. This faster placement has an impact on the quality of results. If getting the best quality of results is important in a given design iteration, perform a full compilation with the Fitter instead of using the Early Timing Estimate feature.

Incremental Compilation Restrictions

This section documents the following restrictions and limitations that you may encounter when using incremental compilation, including interactions with other Quartus II features:

- “When Timing Performance May Not Be Preserved Exactly” on page 2-47
- “When Placement and Routing May Not Be Preserved Exactly” on page 2-47
- “Using Incremental Compilation With Quartus II Archive Files” on page 2-48
- “Formal Verification Support” on page 2-49
- “SignalProbe Pins and Engineering Change Orders” on page 2-49
- “SignalTap II Logic Analyzer in Exported Partitions” on page 2-49
- “External Logic Analyzer Interface in Exported Partitions” on page 2-50
- “Assignments Made in HDL Source Code in Exported Partitions” on page 2-50
- “Design Partition Script Limitations” on page 2-50
- “Restrictions on Megafunction Partitions” on page 2-52
- “Register Packing and Partition Boundaries” on page 2-53
- “I/O Register Packing” on page 2-53

When Timing Performance May Not Be Preserved Exactly

Timing performance might change slightly in a partition with placement and routing preserved when other partitions are incorporated or re-placed and routed. Timing changes are due to changes in parasitic loading or crosstalk introduced by the other (changed) partitions. These timing changes are very small, typically less than 30 ps on a timing path. Additional fan-out on routing lines when partitions are added can also degrade timing performance.

To ensure that a partition continues to meet its timing requirements when other partitions change, a very small timing margin might be required. The Fitter automatically works to achieve such margin when compiling any design, so you do not need to take any action.

When Placement and Routing May Not Be Preserved Exactly

The Fitter may have to refit affected nodes if the two nodes are assigned to the same location, due to imported netlists or empty partitions set to re-use a previous post-fit netlist. There are two cases in which routing information cannot be preserved exactly. First, when multiple partitions are imported, there might be routing conflicts because two lower-level blocks could be using the same routing wire, even if the floorplan assignments of the lower-level blocks do not overlap. These routing conflicts are automatically resolved by the Quartus II Fitter re-routing on the affected nets. Second, if an imported LogicLock region is moved in the top-level design, the relative placement of the nodes is preserved but the routing cannot be preserved, because the routing connectivity is not perfectly uniform throughout a device.

Using Incremental Compilation With Quartus II Archive Files

The post-synthesis and post-fitting netlist information for each design partition is stored in the project database, the **incremental_db** directory. When you archive a project, the database information is not included in the archive unless you include the compilation database in the **.qar** file.

If you want to re-use post-synthesis or post-fitting results, include the database files in the **Archive Project** dialog box so compilation results are preserved. Click **Advanced**, and choose a file set that includes the compilation database, or turn on **Incremental compilation database files** to create a Custom file set.

When you include the database, the file size of the **.qar** archive file may be significantly larger than an archive without the database.

The netlist information for imported partitions is already saved in the corresponding **.qxp**. Imported **.qxp** files are automatically saved in a subdirectory called **imported_partitions**, so you do not need to archive the project database to keep the results for imported partitions. When you restore a project archive, the partition is automatically reimported from the **.qxp** in this directory if it is available.

For new device families with advanced support, a version-compatible database might not be available. In this case, the archive will not include the compilation database. If you require the database files to reproduce the compilation results in the same Quartus II version, you can use the following command-line option to archive a full database:

```
quartus_sh --archive -use_file_set full_db [-revision <revision name>]  
<project name>
```

Limitations for HardCopy Compilation and Migration Flows

Incremental compilation within a single Quartus II project is supported for the base family in HardCopy migration flows for both the FPGA first and HardCopy first flows. Design partition assignments are migrated to the companion device. However, you can not make changes to the design after migration because the design would not match the compilation results for the base family. Therefore, you can perform incremental compilation on one device family, but cannot add new partitions or remove existing partitions after migration.

The Netlist Only preservation level is not supported for Post-fit netlists for FPGA or HardCopy ASIC device compilations when a migration device is specified (that is, for HardCopy ASIC device compilations with a FPGA migration device, or FPGA device compilations with a HardCopy ASIC migration device).

Exporting and importing partitions is not supported in HardCopy ASIC or FPGA device compilations when there is a migration device setting.

The Revision Compare feature requires that the HardCopy ASIC and FPGA netlists are the same. Therefore, all operations performed on one revision must also occur on the other revision. This is accomplished by logging all operations and replaying them on the other revision. Importing partitions does not support this requirement. You can often use **Empty** partitions to implement behavior similar to an exported partition flow, as long as you do not change any global assignments between compilations. All global assignments must be the same for all compiled partitions, so the assignments can be reproduced in the companion device after migration.

Formal Verification Support

You cannot use design partitions for incremental compilation if you are creating a netlist for a formal verification tool.

SignalProbe Pins and Engineering Change Orders

ECO and SignalProbe changes are performed only during ECO and SignalProbe compilations. Other compilation flows do not preserve these netlist changes.

When incremental compilation is turned on and your design contains one or more design partitions, partition boundaries are ignored while making ECO changes and SignalProbe signal settings. However, the presence of ECO and/or SignalProbe changes does not affect partition boundaries for incremental compilation. During subsequent compilations, ECO and SignalProbe changes are not preserved regardless of the **Netlist Type** or **Fitter Preservation Level** settings. To recover ECO changes and SignalProbe signals, you must use the Change Manager to re-apply the ECOs after compilation.

For partitions developed independently in separate Quartus II projects, the exported netlist includes all currently saved ECO changes and SignalProbe signals. If you make any ECO or SignalProbe changes that affect the interface to the lower-level partition, the software issues a warning message during the export process that this netlist does not work in the top-level design without modifying the top-level HDL code to reflect the lower-level change. After integrating the **.qxp** partition into the top-level design, the ECO changes will not appear in the Change Manager.



For more information about using the SignalProbe feature to debug your design, refer to the *Quick Design Debugging Using SignalProbe* chapter in volume 3 of the *Quartus II Handbook*. For more information about using the Chip Planner and the Resource Property Editor to make ECOs, refer to the *Engineering Change Management with the Chip Planner* chapter in volume 2 of the *Quartus II Handbook*.

SignalTap II Logic Analyzer in Exported Partitions

You can use the SignalTap II Embedded Logic Analyzer in any project that you can compile and program into an Altera device.

When incremental compilation is turned on, debugging logic is added to your design incrementally and you can tap post-fitting nodes and modify triggers and configuration without recompiling the full design. Use the appropriate filter in the Node Finder to find your node names. Use **SignalTap II: post-fitting** if the netlist type is Post-Fit to incrementally tap node names in the post-fit netlist database. Use **SignalTap II: pre-synthesis** if the netlist type is **Source File** to make connections to the source file (pre-synthesis) node names when you synthesize the partition from the source code.

If incremental compilation is turned off, the debugging logic is added to the design during Analysis and Elaboration, and you cannot tap post-fitting nodes or modify debug settings without fully compiling the design.

For design partitions that are being developed independently in separate Quartus II projects and contain the logic analyzer, when you export the partition, the Quartus II software automatically removes the SignalTap II logic analyzer and related SLD_HUB logic. You can tap any nodes in a Quartus II project, including nodes within **.qxp** partitions. Therefore, you can use the logic analyzer within the full top-level design to tap signals from the **.qxp** partition.

You can also instantiate the SignalTap II megafunction directly in your lower-level design (instead of using an **.stp** file) and export the entire design to the top level to include the logic analyzer in the top-level design.

-  For details about using the SignalTap II logic analyzer in an incremental design flow, refer to the *Design Debugging Using the SignalTap II Embedded Logic Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

External Logic Analyzer Interface in Exported Partitions

You can use the Logic Analyzer Interface in any project that you can compile and program into an Altera device. You cannot export a partition that uses the Logic Analyzer Interface. You must disable the Logic Analyzer Interface feature and recompile the design before you export the design as a partition.

-  For more information about the Logic Analyzer Interface, refer to the *In-System Debugging Using External Logic Analyzers* chapter in volume 3 of the *Quartus II Handbook*.

Assignments Made in HDL Source Code in Exported Partitions

Assignments made with I/O primitives or the `altera_attribute` HDL synthesis attribute in lower-level partitions are passed to the top-level design, but do not appear in the top-level **.qsf** file or Assignment Editor. These assignments are considered part of the source netlist files. You can override assignments made in these source files by changing the value with an assignment in the top-level design.

Design Partition Script Limitations

The Quartus II software has some additional limitations related to the design partition scripts described in “[Generating Design Partition Scripts](#)” on page 2–30.

Warnings About Extra Clocks Due to Design Partition Scripts

The generated scripts include applicable clock information for all clock signals in the top-level design. Some of those clocks may not exist in the lower-level projects, so you may see warning messages related to clocks that do not exist in the project. You can ignore these warnings or edit your constraints so the messages are not generated.

Synopsys Design Constraint Files for the TimeQuest Timing Analyzer in Design Partition Scripts

After you have compiled a design using TimeQuest constraints, and the timing assignments option is turned on in the scripts, a separate Tcl script is generated to create an `.sdc` file for each lower-level project. This script includes only clock constraints and minimum and maximum delay settings for the TimeQuest Timing Analyzer.



PLL settings and timing exceptions are not passed to lower-level designs in the scripts. For suggestions on managing SDC constraints between top-level and lower-level projects, refer to the *Best Practices for Incremental Compilation Partitions and Floorplan Assignments* chapter in volume 1 of the *Quartus II Handbook*.

Wildcard Support in Design Partition Scripts

When applying constraints with wildcards, note that wildcards are not analyzed across hierarchical boundaries. For example, an assignment could be made to these nodes: `Top|A:inst|B:inst|*`, where A and B are lower-level partitions, and hierarchy B is a child of A, that is B is instantiated in hierarchy A. This assignment is applied to modules A, B, and all children instances of B. However, the assignment `Top|A:inst|B:inst*` is applied to hierarchy A, but is not applied to the B instances because the single level of hierarchy represented by `B:inst*` is not expanded into multiple levels of hierarchy. To avoid this issue, ensure that you apply the wildcard to the hierarchical boundary if it should represent multiple levels of hierarchy.

When using the wildcard to represent a level of hierarchy, only single wildcards are supported. This means assignments such as `Top|A:inst|*|B:inst|*` are not supported. The Quartus II software issues a warning in these cases.

Derived Clocks and PLLs in Design Partition Scripts

If a clock in the top level is not directly connected to a pin of a lower-level partition, the lower-level partition does not receive assignments and constraints from the top-level pin in the design partition scripts.

This issue is of particular importance for clock pins that require timing constraints and clock group settings. Problems can occur if your design uses logic or inversion to derive a new clock from a clock input pin. Make appropriate timing assignments in your lower-level Quartus II project to ensure that clocks are not unconstrained.

If the lower-level design uses the top-level project framework from the project lead, the design will have all the required information about the clock and PLL settings. Otherwise, if you use a PLL in your top-level design and connect it to lower-level partitions, the lower-level partitions do not have information about the multiplication or phase shift factors in the PLL. Make appropriate timing assignments in your lower-level Quartus II project to ensure that clocks are not unconstrained or constrained with the incorrect frequency. Alternatively, you can manually duplicate the top-level derived clock logic or PLL in the lower-level design file to ensure that you have the correct multiplication or phase-shift factors, compensation delays and other PLL parameters for complete and accurate timing analysis. Create a design partition for the rest of the lower-level design logic for export to the top level. When the lower-level design is complete, export only the partition that contains the relevant logic.

Pin Assignments for GXB and LVDS Blocks in Design Partition Scripts

Pin assignments for high-speed GXB transceivers and hard LVDS blocks are not written in the scripts. You must add the pin assignments for these hard IP blocks in the lower-level projects manually.

Virtual Pin Timing Assignments in Design Partition Scripts

Design partition scripts use `INPUT_MAX_DELAY` and `OUTPUT_MAX_DELAY` assignments to specify inter-partition delays associated with input and output pins, which would not otherwise be visible to the project. These assignments require that the software specify the clock domain for the assignment and set this clock domain to " * ".

This clock domain assignment means that there may be some paths constrained and reported by the timing analysis engine that are not required.

To restrict which clock domains are included in these assignments, edit the generated scripts or change the assignments in your lower-level Quartus II project. In addition, because there is no known clock associated with the delay assignments, the software assumes the worst-case skew, which makes the paths seem more timing critical than they are in the top-level design. To make the paths appear less timing-critical, lower the delay values from the scripts. If required, enter negative numbers for input and output delay values.

Top-Level Ports that Feed Multiple Lower-Level Pins in Design Partition Scripts

When a single top-level I/O port drives multiple pins on a lower-level module, it unnecessarily restricts the quality of the synthesis and placement at the lower-level. This occurs because in the lower-level design, the software must maintain the hierarchical boundary and cannot use any information about pins being logically equivalent at the top level. In addition, because I/O constraints are passed from the top-level pin to each of the children, it is possible to have more pins in the lower level than at the top level. These pins use top-level I/O constraints and placement options that might make them impossible to place at the lower level. The software avoids this situation whenever possible, but it is best to avoid this design practice to avoid these potential problems. Restructure your design so that the single I/O port feeds the design partition boundary and the single connection is split into multiple signals within the lower-level partition.

Restrictions on Megafunction Partitions

The Quartus II software does not support partitions for megafunction instantiations. If you use the MegaWizard™ Plug-In Manager to customize a megafunction variation, the MegaWizard-generated wrapper file instantiates the megafunction. You can create a partition for the MegaWizard-generated megafunction custom variation wrapper file.

The Quartus II software does not support creating a partition for inferred megafunctions (that is, where the software infers a megafunction to implement logic in your design). If you have a module or entity for the logic that is inferred, you can create a partition for that hierarchy level in the design.

The Quartus II software does not support creating a partition for any Quartus II internal hierarchy that is dynamically generated during compilation to implement the contents of a megafunction.

Register Packing and Partition Boundaries

The Quartus II software performs register packing during compilation automatically. However, when incremental compilation is enabled, logic in different partitions cannot be packed together because partition boundaries prevent cross-boundary optimization. This restriction applies to all types of register packing, including I/O cells, DSP blocks, sequential logic, and unrelated logic. Similarly, logic from two partitions cannot be packed into the same ALM.

I/O Register Packing

Cross-partition register packing of I/O registers is allowed in certain cases where your input and output pins exist in the top-level hierarchy (and the Top partition), but the corresponding I/O registers exist in other partitions.

The following specific circumstances are required for input pin cross-partition register packing:

- The input pin feeds exactly one register.
- The path between the input pin and register includes only input ports of partitions that have one fan-out each.

The following specific circumstances are required for output register cross-partition register packing:

- The register feeds exactly one output pin.
- The output pin is fed by only one signal.
- The path between the register and output pin includes only output ports of partitions that have one fan-out each.

Output pins with an output enable signal cannot be packed into the device I/O cell if the output enable logic is part of a different partition from the output register. To allow register packing for output pins with an output enable signal, structure your HDL code or design partition assignments so that the register and tri-state logic are defined in the same partition.

Bidirectional pins are handled in the same way as output pins with an output enable signal. If the registers that need to be packed are in the same partition as the tri-state logic, you can perform register packing.

The restrictions on tri-state logic exist because the I/O atom (device primitive) is created as part of the partition that contains tri-state logic. If an I/O register and its tri-state logic are contained in the same partition, the register can always be packed with tri-state logic into the I/O atom. The same cross-partition register packing restrictions also apply to I/O atoms for input and output pins. The I/O atom must feed the I/O pin directly with exactly one signal. The path between the I/O atom and the I/O pin must include only ports of partitions that have one fan-out each.

- For more information and examples of cross-partition boundary I/O packing, refer to the *Best Practices for Incremental Compilation Partitions and Floorplan Assignments* chapter in volume 1 of the *Quartus II Handbook*.

Scripting Support

You can run procedures and make settings described in this chapter in a Tcl script or at a command-line prompt.

- For information about the `::quartus::incremental_compilation` Tcl package that contains a set of functions for manipulating design partitions and settings related to the incremental compilation feature, refer to *::quartus::incremental_compilation* in Quartus II Help.
- For scripting support information, including design examples and training, refer to the *Quartus II Software Scripting Support* page of the Altera website. For detailed Tcl scripting and command-line information, including design examples, refer to the *Tcl Scripting* and *Command-Line Scripting* chapters in volume 2 of the *Quartus II Handbook*.

Creating Design Partitions

To create a design partition, use the following Tcl command example:

```
set_instance_assignment -name PARTITION_HIERARCHY \  
<file name> -to <destination> -section_id <partition name>
```

- <file name>**—The name for the partition, which is auto-generated by the Quartus II software if you create the partition in the GUI or with the `set_partition` TCL command. The Quartus II software uses file name for the internally generated netlists files during incremental compilation. If you use Tcl to create your partitions, you must assign a file name that is unique across all partitions. For the top-level partition, the file name is ignored. To ensure the names are platform independent, file names must be unique and case-sensitive. For example, if a partition uses the file name `my_file`, no other partition can use the file name `MY_FILE`. For simplicity, Altera recommends that you base each file name on the corresponding instance name for the partition.
- <destination>**—The hierarchy name to make into a partition.
- <partition name>**—The user-specified partition name, which appears in the Design Partitions window in the Quartus II software. The partition name must be unique and less than 1024 characters. The name can consist only of alphanumeric characters, and the pipe (|), colon (:), and underscore (_) characters. Altera recommends enclosing the name in double quotation marks (" ").

- You can use the same value for file name and partition name in most cases, but be aware that the legal characters are different, for example, you cannot have the pipe " | " character in a filename. The Quartus software uses the hierarchy path as a default for the partition name, and a shortened MD5 hash as the default for the filename.

- For more information about hierarchical naming conventions, refer to the *Node-Naming Conventions in Quartus II Integrated Synthesis* section in the *Quartus II Integrated Synthesis* chapter in volume 1 of the *Quartus II Handbook*.

The Quartus II software stores all netlists in the `\incremental_db` compilation database directory.

Enabling or Disabling Design Partition Assignments During Compilation

To direct the Quartus II Compiler to enable or disable design partition assignments during compilation, use the following Tcl command:

```
set_global_assignment -name IGNORE_PARTITIONS <value> ←
```

- *OFF*—The Quartus II software recognizes the design partitions assignments set in the current Quartus II project and recompiles the partition in subsequent compilations depending on their netlist status.
- *ON*—The Quartus II software does not recognize design partitions assignments set in the current Quartus II project and performs a compilation without regard to partition boundaries or netlists.

Setting the Netlist Type

To set the partition netlist type, use the following Tcl command:

```
set_global_assignment -name PARTITION_NETLIST_TYPE <value> \  
-section_id <partition name>
```

- *SOURCE*
- *POST_SYNTH*
- *POST_FIT*
- *IMPORTED*
- *EMPTY*



For details about design partition properties, refer to [“Specifying the Level of Results Preservation for Subsequent Compilations”](#) on page 2-21.

Setting the Fitter Preservation Level for a Post-fit or Imported Netlist

To set the Fitter Preservation Level for a post-fit or imported netlist, use the following Tcl command:

```
set_global_assignment -name PARTITION_FITTER_PRESERVATION_LEVEL \
<value> -section_id <partition name>
```

- *NETLIST_ONLY*
- *PLACEMENT*
- *PLACEMENT_AND_ROUTING*



For details about design partition properties, refer to “[Specifying the Level of Results Preservation for Subsequent Compilations](#)” on page 2-21.

Preserving High-Speed Optimization

To preserve high-speed optimization for tiles contained within the selected partition, use the following Tcl command:

```
set_global_assignment -name PARTITION_PRESERVE_HIGH_SPEED_TILES
```

Specifying the Software Should Use the Specified Netlist and Ignore Source File Changes

To specify that the software should use the specified netlist and ignore source file changes, even if the source file has changed since the netlist was created, use the following command:

```
set_global_assignment -name PARTITION_IGNORE_SOURCE_FILE_CHANGES ON
-section_id "<partition name>".
```

Generating Design Partition Scripts

To generate design partition scripts, use the following script:

```
# load required package
load_package database_manager

# name and open the project
set project <project_path/project_name>
project_open $project

# generate the design partiion scripts
generate_bottom_up_scripts <options>

#close project
project_close
```



The options map to the same as those in the Quartus II software in the **Generate Design Partition Scripts** dialog box. For detailed information about each option, refer to [Generate Design Partition Scripts Dialog Box](#) in Quartus II Help.

To generate design partition scripts at the command-line, use the following command:

```
quartus_cdb <project name> --generate_bottom_up_scripts=on <options> ←
■ --include_makefiles_with_bottom_up_scripts <on|off>
```

- `--include_project_creation_in_bottom_up_scripts <on|off>`
- `--include_virtual_pins_in_bottom_up_scripts <on|off>`
- `--include_virtual_pin_timing_in_bottom_up_scripts <on|off>`
- `--bottom_up_scripts_virtual_pin_delay <delay in ns>`
- `--include_virtual_pin_locations_in_bottom_up_scripts <on|off>`
- `--include_logiclock_regions_in_bottom_up_scripts <on|off>`
- `--include_all_logiclock_regions_in_bottom_up_scripts <on|off>`
- `--include_global_signal_promotion_in_bottom_up_scripts <on|off>`
- `--include_pin_locations_in_bottom_up_scripts <on|off>`
- `--include_timing_assignments_in_bottom_up_scripts <on|off>`
- `--include_design_partitions_in_bottom_up_scripts <on|off>`
- `--remove_existing_regions_in_bottom_up_scripts <on|off>`
- `--disable_auto_global_promotion_in_bottom_up_scripts <on|off>`
- `--bottom_up_scripts_output_directory=<output directory>`

Exporting a Partition

To open a project and load the `::quartus::incremental_compilation` package before you use the Tcl commands to export a partition to a `.qxp` that contains both a post-synthesis and post-fit netlist, with routing, use the following script:

```
# load required package
package require ::quartus::incremental_compilation

# open project
project_open <project name>

# export partition to the .qxp and set preservation level
export_partition -partition <partition name>
-qxp <.qxp file name> -<options>

#close project
project_close

■ POST_FIT <on|off>
■ POST_SYNTH <on|off>
■ ROUTING <on|off>
```

Importing a Partition into the Top-Level Design

To import a `.qxp` into a top-level design, use the following script:

```
# load required packages
load_package incremental_compilation
load_package project
load_package flow
```

```
# open project
project_open <project name>

# make partition import file assignments
set_global_assignment -name PARTITION_IMPORT_FILE \
-entity <entity which instantiates imported .qxp> \
-section_id <imported partition name> \
<.qxp filename>

# execute import
execute_flow -incremental_compilation_import
```

Makefiles

For an example of how to use incremental compilation with a makefile as part of the team-based incremental compilation design flow, refer to the **read_me.txt** file that accompanies the `incr_comp` example located in the `/qdesigns/incr_comp_makefile` subdirectory.

- ② When using a team-based incremental compilation design flow, the **Generate Design Partition Scripts** dialog box can write makefiles that automatically export lower-level design partitions and import them into the top-level design whenever design files change. For more information about the **Generate Design Partition Scripts** dialog box, refer to *Generate Design Partition Scripts Dialog Box* in Quartus II Help.

Scripting and Command-Line Application Examples

This section provides scripting examples that cover some of the topics discussed in the main section of this chapter.

Reducing Opening a Project, Creating Design Partitions, and Performing an Initial Compilation

Example background: You open a project called `AB_project`, set up two design partitions, entities A and B, and then perform an initial complete compilation.

Example 2-1. `AB_project`

```
set project AB_project

package require ::quartus::flow
project_open $project

# Ensure that design partition assignments are not ignored
set_global_assignment -name IGNORE_PARTITIONS \ OFF

# Set up the partitions
set_instance_assignment -name PARTITION_HIERARCHY \
  incremental_db/A_inst -to A -section_id "Partition_A"
set_instance_assignment -name PARTITION_HIERARCHY \
  incremental_db/B_inst -to B -section_id "Partition_B"

# Set the netlist types to post-fit for subsequent
# compilations (all partitions are compiled during the
# initial compilation since there are no post-fit
# netlists)
set_global_assignment -name PARTITION_NETLIST_TYPE \
  POST_FIT -section_id "Partition_A"
set_global_assignment -name PARTITION_NETLIST_TYPE \
  POST_FIT -section_id "Partition_B"

# Run initial compilation:
export_assignments
execute_flow -full_compile

project_close
```

Reducing Compilation Time When Changing a Source File for One Partition—Command-Line Example

Example background: You have run the initial compilation shown in the example script in [Example 2-1](#). You have modified the HDL source file for partition A and want to recompile it.

Run the standard flow compilation command in your Tcl script:

```
execute_flow -full_compile
```

Or, type the following command at a system command prompt:

```
quartus_sh --flow compile AB_project
```

Assuming the source files for partition B do not depend on A, only A is recompiled. The placement of B and its timing performance is preserved, which also saves significant compilation time.

Optimizing the Placement for a Timing-Critical Partition

Example background: You have run the initial compilation shown in the example script under [Example 2-1](#). You would like to apply Fitter optimizations, such as physical synthesis, only to partition A. No changes have been made to the HDL files.

To ensure the previous compilation result for partition **B** is preserved, and to ensure that Fitter optimizations are applied to the post-synthesis netlist of partition **A**, set the netlist type of **B** to **Post-Fit** (which was already done in the initial compilation, but is repeated here for safety), and the netlist type of **A** to **Post-Synthesis**, as shown in the following example:

Example 2-2. AB_project (2)

```
set project AB_project

package require ::quartus::flow
project_open $project

# Turn on Physical Synthesis Optimization
set_global_assignment -name \
PHYSICAL_SYNTHESIS_REGISTER_RETIMING ON

# For A, set the netlist type to post-synthesis
set_global_assignment -name PARTITION_NETLIST_TYPE POST_SYNTH \
-section_id "Partition_A"

# For B, set the netlist type to post-fit
set_global_assignment -name PARTITION_NETLIST_TYPE POST_FIT \
-section_id "Partition_B"

# Run incremental compilation:
export_assignments
execute_flow -full_compile

project_close
```

Conclusion

With the Quartus II incremental compilation feature described in this chapter, you can preserve the results and performance of unchanged logic in your design as you make changes elsewhere. The various applications of incremental compilation enable you to improve your productivity while designing for high-density FPGAs.

Document Revision History

Table 2-4 shows the revision history for this document.

Table 2-4. Document Revision History

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Changed to new document template. ■ Reorganized Tcl scripting section. Added description for new feature: Ignore partitions assignments during compilation option. ■ Reorganized “Incremental Compilation Summary” on page 2-7 section.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Removed the explanation of the “bottom-up design flow” where designers work completely independently, and replaced with Altera’s recommendations for team-based environments where partitions are developed in the same top-level project framework, plus an explanation of the bottom-up process for including independent partitions from third-party IP designers. ■ Expanded the Merge command explanation to explain how it now accommodates cross-partition boundary optimizations. ■ Restructured Altera recommendations for when to use a floorplan. ■ Added “Viewing the Contents of a Quartus II Exported Partition File (.qxp)” on page 2-31 section. ■ Reorganized chapter to make design flow scenarios more visible; integrated into various sections rather than at the end of the chapter.
October 2009	9.1.0	<ul style="list-style-type: none"> ■ Redefined the bottom-up design flow as team-based and reorganized previous design flow examples to include steps on how to pass top-level design information to lower-level designers. ■ Moved SDC Constraints from Lower-Level Partitions section to the <i>Best Practices for Incremental Compilation Partitions and Floorplan Assignments</i> chapter in volume 1 of the <i>Quartus II Handbook</i>. ■ Reorganized the “Conclusion” on page 2-60 section. ■ Removed HardCopy APEX and HardCopy Stratix Devices section.
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Split up netlist types table ■ Moved “Team-Based Incremental Compilation Design Flow” into the “Including or Integrating partitions into the Top-Level Design” section. ■ Added new section “Including or Integrating Partitions into the Top-Level Design”. ■ Removed “Exporting a Lower-Level Partition that Uses a JTAG Feature” restriction ■ Other edits throughout chapter
November 2008	8.1.0	<ul style="list-style-type: none"> ■ Added new section “Importing SDC Constraints from Lower-Level Partitions” on page 2-44 ■ Removed the Incremental Synthesis Only option ■ Removed section “OpenCore Plus Feature for MegaCore Functions in Bottom-Up Flows” ■ Removed section “Compilation Time with Physical Synthesis Optimizations” ■ Added information about using a .qxp as a source design file without importing ■ Reorganized several sections ■ Updated Figure 2-10

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#)

 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter describes Quartus® II support for HardCopy® series devices.

Altera® HardCopy ASICs are the lowest risk, lowest total cost ASICs. The HardCopy system development methodology offers fast time-to-market, low risk, and with the Quartus II software, you can design with one set of RTL code and one IP set for both FPGA and ASIC implementations. This flow enables you to conduct true hardware/software co-design and completely prepare your system for production prior to ASIC design hand-off. Altera provides a turn-key process to convert your design to a HardCopy ASIC for production.

In this chapter, the term FPGA refers to a Stratix® II, Stratix III, or Stratix IV device, which is the prototype device for a HardCopy II, HardCopy III, or HardCopy IV device, respectively.

This chapter discusses the following topics:

- “HardCopy Development Flow” on page 3–2
- “HardCopy Companion Device Selection” on page 3–6
- “HardCopy Utilities” on page 3–7
- “HardCopy Device Resource Guide” on page 3–11
- “HardCopy Recommended Settings in the Quartus II Software” on page 3–12
- “HardCopy Design Readiness Check” on page 3–15
- “Performing ECOs with Quartus II Engineering Change Management with the Chip Planner” on page 3–20
- “Formal Verification of FPGA and HardCopy Revisions” on page 3–23



For more information about HardCopy series devices, refer to the respective HardCopy device handbook, which is available on the Literature page of the Altera website at www.altera.com.

HardCopy Series Design Benefits

Designing with HardCopy ASICs offers the following substantial benefits over other ASIC offerings:

- Seamless prototyping using an FPGA for at-speed system verification and system development reduces total project development time and cost
- Dependable conversion from an FPGA prototype to a HardCopy ASIC expands product planning options
- Unified design methodology for FPGA and HardCopy designs reduces the need for ASIC development software, two sets of intellectual property, and project risk
- System development methodology delivers lowest total cost

Quartus II Features for HardCopy Planning

With the Quartus II software, you can design a HardCopy ASIC using seamless FPGA prototyping. The Quartus II software provides the following expanded features, which are described in further detail later in this chapter, for HardCopy series device planning:

- **HardCopy Companion Device Assignment**—Identifies compatible HardCopy series devices for use with the FPGA prototype device currently selected.

This feature constrains the pins of your FPGA prototype, making it compatible with your HardCopy device. The feature also constrains the correct resources available for the HardCopy device, ensuring the compatibility of your FPGA design. You must compile the design targeting the HardCopy device to ensure that the design fits, routes, and meets timing requirements.
- **HardCopy Utilities**—The HardCopy Utilities menu provides a variety of functions to create or overwrite HardCopy companion revisions, set current revisions, and compare revisions for equivalency.
- **HardCopy Advisor**—The HardCopy Advisor guides you through the steps necessary to successfully submit a HardCopy design to the Altera HardCopy Design Center.
- **HardCopy Floorplan**—The Chip Planner can show a preliminary floorplan view of your HardCopy design's Fitter placement results.
- **HardCopy Device Preliminary Timing**—The TimeQuest Timing Analyzer performs a timing analysis of HardCopy devices based on preliminary timing models and Fitter placements. Final timing results for HardCopy devices are provided by the Altera HardCopy Design Center.
- **HardCopy Design Readiness Check**—The Quartus II software checks the project settings to ensure compliance with the HardCopy device settings, I/O, PLL, and RAM usage checks.
- **HardCopy Handoff Report**—The Quartus II software generates a handoff report containing information about the HardCopy design used by the Altera HardCopy Design Center in the design review process.
- **HardCopy Design Archiving**—The Quartus II software archives the HardCopy design project's files required to hand off the design to the Altera HardCopy Design Center.
- **Formal Verification**—Cadence Encounter Conformal software performs formal verification between the source RTL design files and post-compilation gate-level netlist from a HardCopy design.

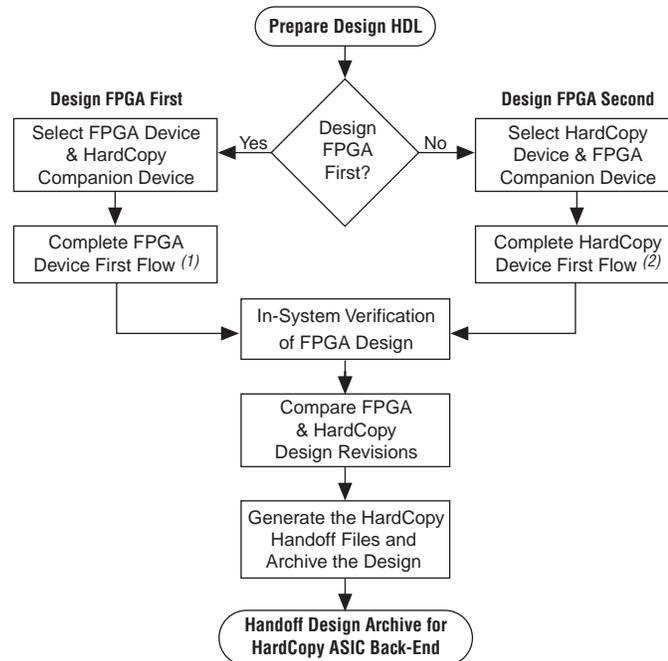
HardCopy Development Flow

In the Quartus II software, you design your FPGA and HardCopy companion device together in one Quartus II project using one of the following methods:

- Design the FPGA first for in-system verification, and then create a HardCopy companion device second.
- Design the HardCopy device first, and then create the FPGA companion device second for in-system verification.

Both of these flows are illustrated at a high level in [Figure 3-1](#). The features in the HardCopy Utilities menu help you complete your HardCopy design for submission to the Altera HardCopy Design Center for back-end implementation.

Figure 3-1. HardCopy Flow in Quartus II Software



Notes to Figure 3-1:

- (1) Refer to [Figure 3-2](#) on page 3-4 for an expanded description of this process.
- (2) Refer to [Figure 3-3](#) on page 3-6 for an expanded description of this process.



The FPGA first flow is the default flow and the rest of this chapter is based on this flow.

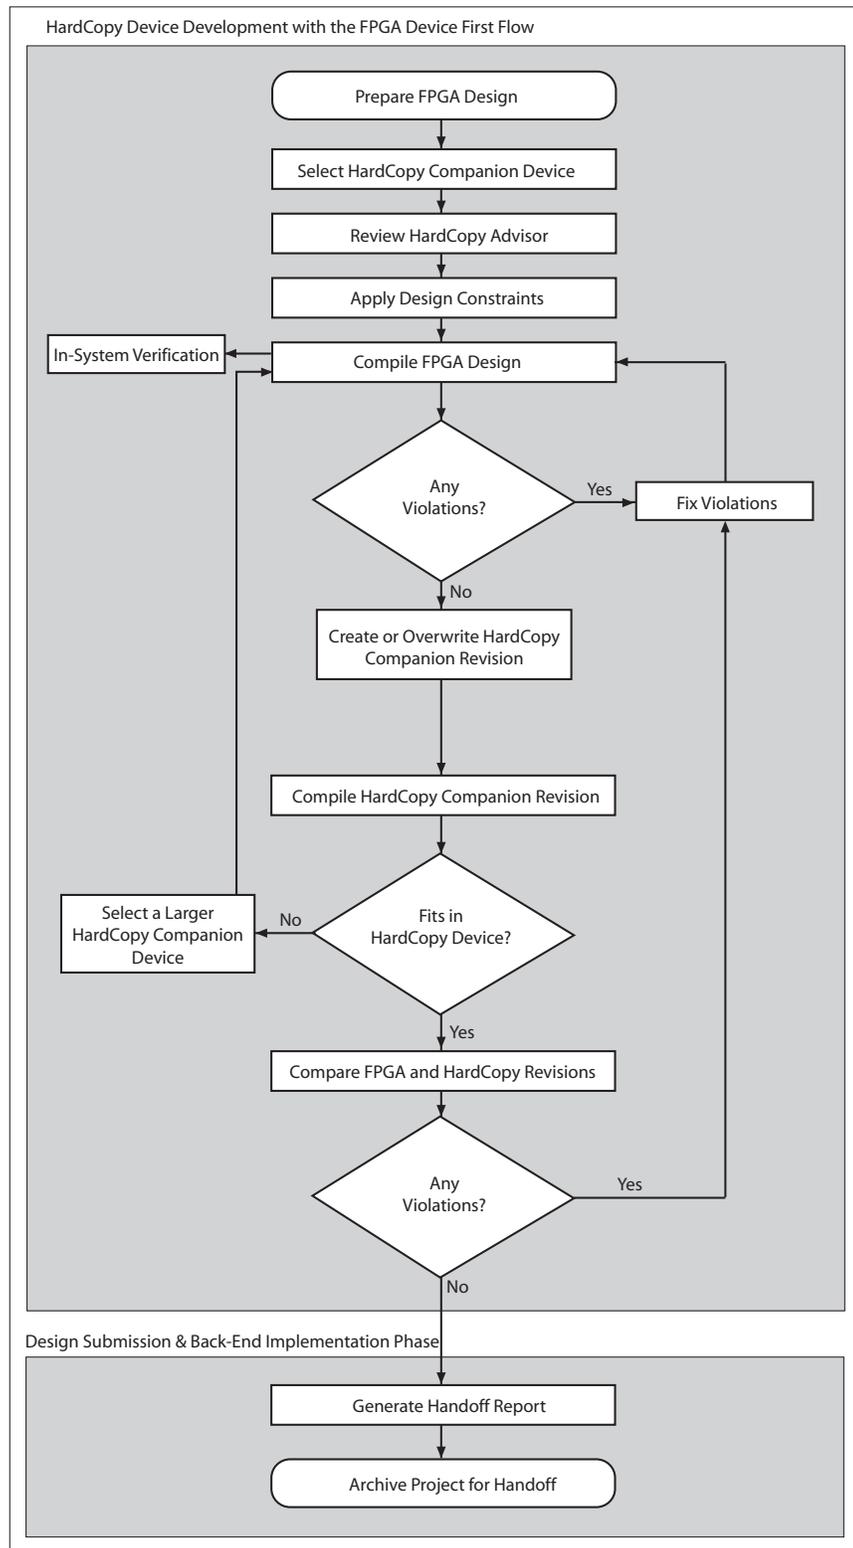
Designing the FPGA First

The HardCopy FPGA first flow development flow begins with seamless FPGA prototyping and is identical to the traditional FPGA design flow; plus a few additional tasks necessary to convert the design to the HardCopy companion device within the same project. To design your HardCopy device when selecting the FPGA companion device first, complete the following tasks:

- Specify an FPGA device and a HardCopy companion device
- Compile the FPGA design
- Create and compile the HardCopy companion revision
- Compare the HardCopy companion revision compilation to the FPGA device compilation

Figure 3-2 provides an overview of the development process for designing with an FPGA first and creating a HardCopy companion device second.

Figure 3-2. Designing FPGA Device First Flow



You must select a target FPGA device and a companion HardCopy device when compiling an FPGA design that you will migrate to a HardCopy device.

During the early stages of the design process, selecting the right HardCopy device may be difficult. The HardCopy Device Resource Guide can assist you in the selection process. After you select an FPGA and a HardCopy device, compile the FPGA and review the HardCopy Device Resource Guide to see if all resources are available in the targeted HardCopy device. If there are not enough resources available in the target HardCopy device, you must select a larger HardCopy device and recompile the FPGA design.

Once the FPGA and the HardCopy device selections have been finalized, perform the following tasks:

- Review the HardCopy Advisor for required and recommended tasks
- Enable the Design Assistant to run during compilation
- Add timing and location assignments
- Compile your FPGA design
- Create your HardCopy companion revision
- Compile your design for the HardCopy companion device
- Compare the HardCopy companion device compilation with the FPGA revision
- Generate a HardCopy handoff report
- Generate a HardCopy handoff archive
- Arrange for submission of your HardCopy Handoff Archive to the Altera HardCopy Design Center for back-end implementation

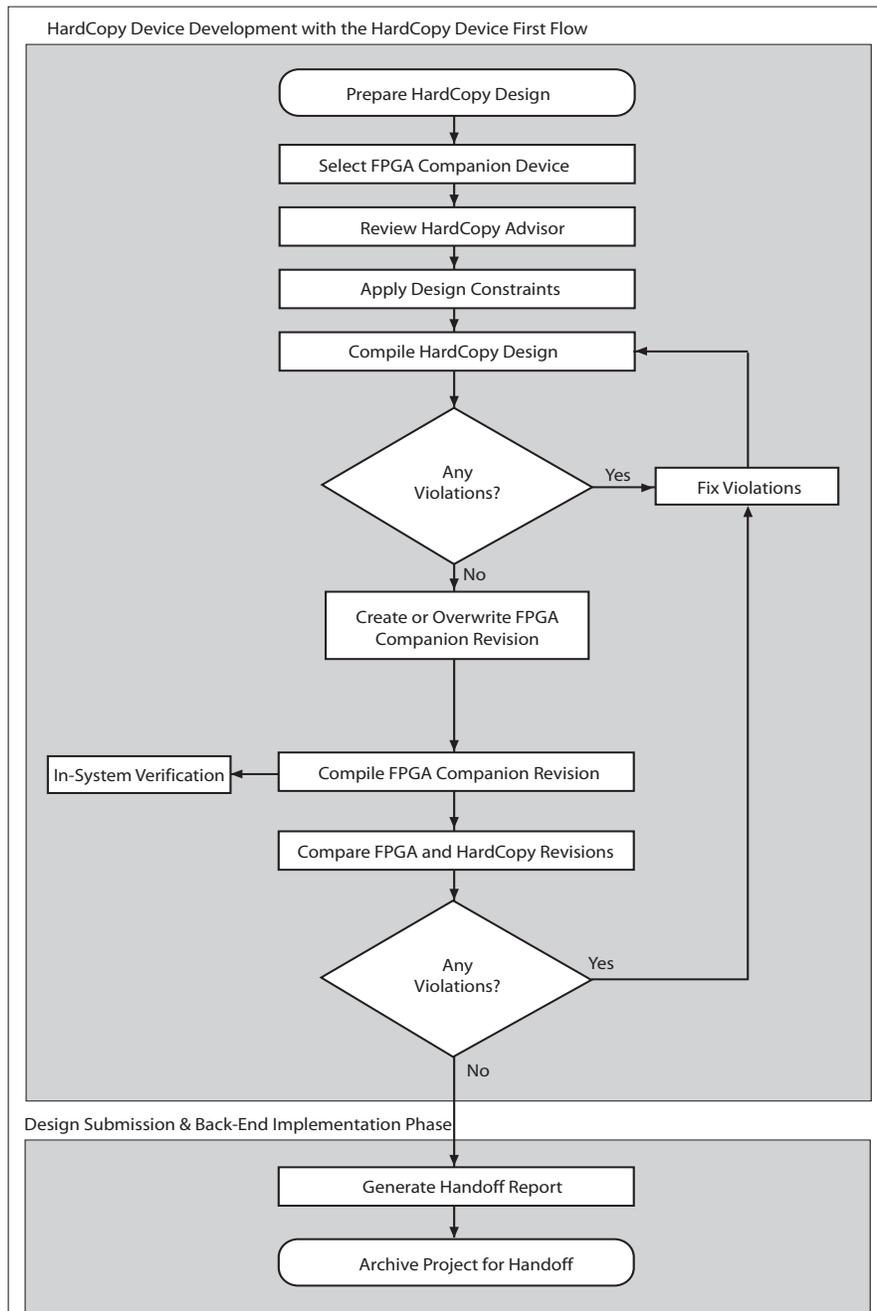


For more information about the overall design flow using the Quartus II software, refer to the *Introduction to the Quartus II Software* manual, which is available on the Literature page of the Altera website at www.altera.com.

Designing the HardCopy Device First

After you select an initial HardCopy ASIC device, you can design your HardCopy device first, and then create your FPGA prototype second. This approach is preferred when using the HardCopy device to achieve higher performance than the FPGA prototype, because you can see your potential maximum performance in the HardCopy device immediately during development, and you can create a slower performing FPGA prototype of the design for in-system verification. This design process is similar to the HardCopy FPGA first flow development flow, but you begin the design with a different initial device family instead. The remaining tasks to complete your design for both the FPGA and HardCopy devices roughly follow the same process (Figure 3-3). The HardCopy Advisor adjusts its list of tasks based on which device family you start with, FPGA or HardCopy, to help you complete the process seamlessly.

Figure 3-3. Designing HardCopy Device First Flow



HardCopy Companion Device Selection

In the Quartus II software, you can select a HardCopy companion device to ensure compatibility between the FPGA design and the HardCopy device's resources. To select your HardCopy companion device, on the Assignments menu, click **Device** and select a companion device from the **Companion device** list.

Selecting a HardCopy companion device for your FPGA prototype constrains the memory blocks, DSP blocks, and pin assignments, so that your design fits into the HardCopy device resources. Pin assignments are constrained in the FPGA design revision, so that the HardCopy device selected is pin-compatible. The Quartus II software also constrains the FPGA design revision so that identical device resources are targeted in both the FPGA and the HardCopy ASIC.

You can also specify your HardCopy companion device using the following tool command language (Tcl) command:

```
set_global_assignment -name\
DEVICE_TECHNOLOGY_MIGRATION_LIST <HardCopy Device Part Number>
```

For example, to select the HC230F1020 device as your HardCopy companion device for the EP2S130F1020C4 FPGA, use the following the Tcl command:

```
set_global_assignment -name\
DEVICE_TECHNOLOGY_MIGRATION_LIST HC230F1020C
```

HardCopy Utilities

The HardCopy Utilities menu contains the main functions you use to develop your HardCopy design and FPGA prototype companion revision. To access this menu, on the Project menu, click **HardCopy Utilities**. From the HardCopy Utilities menu, you can perform the following tasks:

- Create or overwrite HardCopy companion revisions
- Specify the current HardCopy companion revision
- Compare the companion revisions for functional equivalence
- Generate a HardCopy Handoff report for design reviews
- Archive HardCopy handoff files for submission to the Altera HardCopy Design Center
- Start the HardCopy Design Readiness Check
- Track your design progress using the HardCopy Advisor

Each **HardCopy Utilities** feature is summarized in [Table 3-1](#). The process for using each of these features is explained in the following sections.

Table 3-1. HardCopy Utilities Menu Options (Part 1 of 2)

Menu	Description	Applicable Design Revision	Restrictions
Create/Overwrite HardCopy Companion Revision	Creates a new companion revision or overwrites an existing companion revision for your FPGA and HardCopy design	FPGA prototype design and HardCopy companion revision	<ul style="list-style-type: none"> ■ The Auto device selected by the Fitter option must be turned off ■ An FPGA device and a HardCopy companion device must be set
Set Current HardCopy Companion Revision	Specifies the companion revision to associate with the current design revision	FPGA prototype design and HardCopy companion revision	A companion revision must already exist

Table 3-1. HardCopy Utilities Menu Options (Part 2 of 2)

Menu	Description	Applicable Design Revision	Restrictions
Compare HardCopy Companion Revisions	Compares the FPGA design revision with the HardCopy companion design revision and generates a report	FPGA prototype design and HardCopy companion revision	Both revisions must be compiled
Generate HardCopy Handoff Report	Generates a report containing important design information files and messages generated by the Quartus II Compiler	FPGA prototype design and HardCopy companion revision	<ul style="list-style-type: none"> ■ Both revisions must be compiled ■ The Compare HardCopy Companion Revisions command must be successfully run
Archive HardCopy Handoff Files	Generates a Quartus II Archive File (.qar) specifically for submitting the design to the Altera HardCopy Design Center	HardCopy companion revision	<ul style="list-style-type: none"> ■ Both revisions must be compiled ■ The Compare HardCopy Companion Revisions command must be successfully run ■ The Generate HardCopy Handoff Report command must be successfully run
Start HardCopy Design Readiness Check	Generates a reports with the design's settings, I/O check, PLL, and RAM usage checks	FPGA prototype design and HardCopy companion revision	None
HardCopy Advisor	Opens the HardCopy Advisor, which walks you through the process of creating a HardCopy project	FPGA prototype design and HardCopy companion revision	None

Companion Revisions

You can create multiple design revisions for both the FPGA and the HardCopy device. For example, if your initial FPGA revision is named *top* and the corresponding HardCopy revision is named *top_hc*, you could create another FPGA revision, named *top_fpga*, and the corresponding HardCopy revision would be named *top_fpga_hc*. The Quartus II software creates specific HardCopy design revisions of the project in conjunction with the primary project revisions. These parallel design revisions for HardCopy devices are called companion revisions.



Although you can create multiple design revisions, Altera recommends that you maintain only one FPGA revision once you create the HardCopy companion revision.

Once you have successfully compiled your FPGA prototype, you can create and compile the HardCopy companion revision of your design.

You can associate only one FPGA revision to one HardCopy companion revision. If you create more than one revision or companion revision, set the current companion for the revision you are working on.



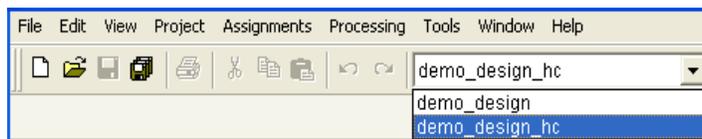
For more information about creating or setting a companion revision in the Quartus II software, refer to *Migrating a Design to a HardCopy or FPGA Device* in Quartus II Help.

Compiling the HardCopy Companion Revision

You can compile your HardCopy design with preliminary timing information in the Quartus II software. The timing constraints for the HardCopy companion revision can be the same as the FPGA design used to create the revision. The Quartus II software contains preliminary timing models for HardCopy devices and you can gauge the degree of performance improvement you can achieve in the HardCopy device compared to the FPGA. Altera verifies that the HardCopy companion device timing requirements are met in the Altera HardCopy Design Center.

After you create your HardCopy companion revision from your compiled FPGA design, select the companion revision in the Quartus II software design revision pull-down list (Figure 3-4) or from the **Revisions** list. Compile the HardCopy companion revision. After you compile your design in the Quartus II software, you can perform a comparison check of the HardCopy companion revision to the FPGA prototype revision.

Figure 3-4. Changing Current Revision



Comparing HardCopy and FPGA Companion Revisions

Altera uses the companion revisions in a single Quartus II project to maintain compatibility between the FPGA and HardCopy ASIC. This methodology allows you to design with one set of RTL code that is used in both the FPGA and HardCopy ASIC, guaranteeing functional equivalency.

When making changes to companion revisions, use the **Compare HardCopy Companion Revisions** command to ensure that your design matches your HardCopy design functionality and compilation settings.

The Comparison Revision Summary in the Compilation report identifies where assignments were changed between revisions or if there is a change in the logic resource count due to different compilation settings.

- ④ For more information about comparing companion revisions in the Quartus II software, refer to *Migrating a Design to a HardCopy or FPGA Device* in Quartus II Help.

Generating a HardCopy Handoff Report

To submit a design to the Altera HardCopy Design Center, you must generate a HardCopy Handoff report, which provides important information about the design that you want the Altera HardCopy Design Center to review.

After you generate the HardCopy Handoff report, you can archive the design using the **Archive HardCopy Handoff Files** command, which is described in “[Archiving HardCopy Handoff Files](#)”.

- ② For more information about the **Generate HardCopy Handoff Report** command in the Quartus II software, refer to *Generate HardCopy Handoff Report Command* in Quartus II Help.

Archiving HardCopy Handoff Files

The last step in the HardCopy design methodology is to archive the HardCopy project for submission to the Altera HardCopy Design Center for HardCopy back-end implementation. The **Archive HardCopy Handoff** command creates a unique **.qar** file, which is different than the standard Quartus II project archive utility generates. This archive contains only the necessary data from the Quartus II project required to implement the design in the Altera HardCopy Design Center.

- ② For more information about the **Archive HardCopy Handoff Files** command in the Quartus II software, refer to *Archive HardCopy Handoff Files Command* in Quartus II Help.

HardCopy Advisor

The HardCopy Advisor provides a list of tasks to help guide you through the development of your FPGA prototype and your HardCopy design. The following tasks highlight the checkpoints that the HardCopy Advisor reviews, including the major checkpoints in the design process, but they do not include show every step in the process for completing your FPGA and HardCopy designs:

1. Select an FPGA device.
2. Select a HardCopy device.
3. Turn on the Design Assistant.
4. Set up timing constraints.
5. Check for incompatible assignments.
6. Compile and check the FPGA design.
7. Create or overwrite the companion revision.
8. Compile and check the HardCopy companion results.
9. Compare companion revisions.
10. Generate a Handoff report.
11. Archive handoff files and send them to Altera.

When your design uses an FPGA as your starting point, Altera recommends following the HardCopy Advisor guidelines for your FPGA until you complete the prototype revision.

When the FPGA design is complete, create and switch to your HardCopy companion revision. Follow the HardCopy Advisor steps for that revision until completion, and then submit the design to Altera for the HardCopy back-end implementation process.

- ② For more information about the HardCopy Advisor in the Quartus II software, refer to *About the HardCopy Advisor* in Quartus II Help.

HardCopy Device Resource Guide

The HardCopy Device Resource Guide compares the resources required to successfully compile a design with the resources available in the various HardCopy devices. The report rates each HardCopy device and each device resource according to how well it fits the design. The Quartus II software generates the HardCopy Device Resource Guide for all designs successfully compiled for FPGA devices. This guide is found in the Fitter folder of the Compilation report. Refer to [Table 3-2](#) for an explanation of the color codes.

Use this report to identify potential HardCopy device candidates for your design. The HardCopy and FPGA device package must be compatible. A logic resource usage greater than 100% or a ratio greater than 1:1 in any category indicates that the design probably will not fit in that specific HardCopy device.

Table 3-2. HardCopy Device Resource Guide Color Legend

Color	Package Resource ⁽¹⁾	Device Resources
Green (High)	The design can map to the HardCopy package and has been fitted with target device migration enabled in the HardCopy Companion Device dialog box.	The resource quantity is within the range of the HardCopy device and the design can likely map if all other resources also fit. You still must compile the HardCopy revision to ensure the design is able to route and close timing.
Orange (Medium)	The design can map to the HardCopy package. However, the design has not been fitted with the target device migration enabled in the HardCopy Companion Device dialog box.	The resource quantity is within the range of the HardCopy device. However, the resource is at risk of exceeding the range for the HardCopy package. If your target HardCopy device falls in this category, compile your design targeting the HardCopy device as soon as possible to check if the design fits and is able to route and migrate all other resources. You might have to select a larger device.
Red (None)	The design cannot map to the HardCopy package.	The resource quantity exceeds the range of the HardCopy device. The design cannot migrate to this HardCopy device.

Note to Table 3-2:

- (1) The package resource is constrained by the FPGA for which the design was compiled. Only vertical migration devices within the same package are able to migrate to HardCopy devices.

The HardCopy architecture consists of an array of fine-grained HCells, which are used to build logic equivalent to FPGA adaptive logic modules (ALMs) and digital signal processing (DSP) blocks. The DSP blocks in HardCopy devices match the functionality of the FPGA DSP blocks, though timing of these blocks is different than the FPGA DSP blocks because they are constructed of HCell macros. The memory blocks in HardCopy devices are equivalent to the FPGA memory blocks. Preliminary timing reports of the HardCopy device are available in the Quartus II software. Final timing results of the HardCopy device are provided by the Altera HardCopy Design Center after the HardCopy back-end implementation process is complete.



For more information about the HardCopy device resources, refer to the respective HardCopy series device handbook, which is available on the Literature page of the Altera website at www.altera.com.

The report example in [Figure 3-5](#) shows the resource comparisons for a design compiled for an EP2S130F1020 device. Based on the report, the HC230F1020 device in the 1,020-pin FineLine BGA package is an appropriate HardCopy device. If the HC230F1020 device is not specified as a migration target during the compilation, its package and migration compatibility is rated medium (orange). The migration compatibilities of the other HardCopy devices are rated none (red), because the package types are incompatible with the FPGA device. The 1,020-pin FBGA HC240 device is rated none (red) because it is only compatible with the EP2S180F1020 device.

[Figure 3-5](#) shows the report after the (unchanged) design was recompiled with the HardCopy HC230F1020 device specified as a migration target. Now the HC230F1020 device package and migration compatibility is rated high (green).

Figure 3-5. HardCopy Device Resource Guide with Target Migration Enabled

HardCopy Device Resource Guide									
Color Legend: -- Green: -- Package Resource: The HardCopy device package can be migrated from the selected FPGA device package, and the design has been fitted with the target device migration enabled.									
Resource	Stratix II EP2S130	HC210w*	HC210	HC220	HC220	HC230	HC240	HC240	
1	Migration Compatibility	None	None	None	None	High	None	None	
2	Primary Migration Constraint	Package	Package	Package	Package		Package	Package	
3	Package	FBGA - 1020	FBGA - 484	FBGA - 484	FBGA - 672	FBGA - 780	FBGA - 1020	FBGA - 1020	FBGA - 1508

HardCopy Recommended Settings in the Quartus II Software

The HardCopy development flow involves additional planning and preparation in the Quartus II software when compared to a standard FPGA design. Additional planning and preparation is required because you are developing your design for implementation in two devices: a prototype of your design/system in an FPGA and a companion revision in a HardCopy device for production. Additional settings and constraints are required to make the FPGA design compatible with the HardCopy device, and in some cases, you must remove certain settings in the design. This section explains the additional settings and constraints necessary for your design to be successful in both FPGA and HardCopy ASIC devices.

Limit DSP and RAM to HardCopy Device Resources

The **Limit DSP & RAM to HardCopy device resources** option in the **Device** dialog box maintains compatibility between the FPGA and HardCopy devices by ensuring that your design does not use resources in the FPGA device that are not available in the selected HardCopy device or vice versa.

- For more information about the **Limit DSP & RAM to HardCopy device resources** option in the Quartus II software, refer to *Device Dialog Box* in Quartus II Help.

Enabling Design Assistant to Run During Compile

You must use the Design Assistant in the Quartus II software to check all HardCopy designs for design rule violations before submitting the designs to the Altera HardCopy Design Center. Additionally, you must fix all critical and high-level errors.

-  Altera recommends turning on the Design Assistant to run automatically during each compilation so that you can review the violations you must fix or waive.
-  For more information about the Design Assistant and its rules in the Quartus II software, refer to *About the Design Assistant* in Quartus II Help.

Timing Settings

The TimeQuest Timing Analyzer is a complete static timing analysis tool that you use as a sign-off tool for Altera FPGAs and HardCopy ASICs. The TimeQuest Timing Analyzer guides the Fitter and analyzes timing results after compilation and is the required timing analysis tool for all designs.

-  For more information about the TimeQuest Timing Analyzer, refer to the *TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook* and *About TimeQuest Timing Analysis* in Quartus II Help.

Constraints for Clock Effect Characteristics

The `create_clock` and `create_generated_clock` commands create ideal clocks, but do not account for board effects. To account for clock effect characteristics, you can use the `set_clock_latency` and `set_clock_uncertainty` commands.

-  For more information about how to use these commands, refer to the *TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

You can use the `derive_clock_uncertainty` command to automatically derive the clock uncertainties in your `.sdc` file. This command is useful when you are unsure of the clock uncertainties. The calculated clock uncertainty values are based on I/O buffer, static phase errors (SPE) and jitter in the PLLs, clock networks, and core noise.

The `derive_clock_uncertainty` command applies interclock, intraclock, and I/O interface uncertainties. This command automatically calculates and applies setup and hold clock uncertainties for each clock-to-clock transfer found in your design.

To determine I/O interface uncertainty, you must create a virtual clock, then assign delays to the input/output ports with the `set_input_delay` and `set_output_delay` commands for that virtual clock.

-  These uncertainties are applied in addition to those you specified with the `set_clock_uncertainty` command. However, if a clock uncertainty assignment for a source and destination pair is already defined, the new one is ignored. In this case, you can use either the `-overwrite` command to overwrite the previous clock uncertainty command, or manually remove them with the `remove_clock_uncertainty` command.

The following syntax is for the `derive_clock_uncertainty` command:

```
derive_clock_uncertainty [-h | -help] [-long_help] [-add]
\[-overwrite]
```

The arguments for this command are listed in [Table 3-3](#):

Table 3-3. Arguments for derive_clock_uncertainty

Option	Description
-h -help	Short help
-long_help	Long help with examples and possible return values
-add	Adds results user-defined clock uncertainty assignments
-overwrite	Overwrites previously performed clock uncertainty assignments

When the `derive_clock_uncertainty` command is used, a `PLLJ_PLLSPE_INFO.txt` file is automatically generated in the project directory. This file lists the names of the PLLs, as well as their jitter and SPE values in the design. This text file can be used by the `HCII_DTW_CU_Calculator`.

-  For more information about the `derive_clock_uncertainty` command, refer to the *TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.
-  Altera strongly recommends that you use the `derive_clock_uncertainty` command in the HardCopy revision. The Altera HardCopy Design Center does not accept designs that do not have clock uncertainty constraints by either using the `derive_clock_uncertainty` command or the HardCopy II Clock Uncertainty Calculator, and then using the `set_clock_uncertainty` command.
-  For more information about how to use the HardCopy II Clock Uncertainty Calculator, refer to the *HardCopy II Clock Uncertainty Calculator User Guide*.

Quartus II Software Features Supported for HardCopy Designs

The Quartus II software supports optimization features for HardCopy prototype development, including the following features discussed in this section:

- “Physical Synthesis Optimization”
- “LogicLock Regions” on page 3-15
- “PowerPlay Power Analyzer” on page 3-15
- “Incremental Compilation” on page 3-15

Physical Synthesis Optimization

The physical synthesis optimizations performed in the FPGA device are passed to the HardCopy companion revision for placement and timing closure. When designing with a HardCopy device first, you can enable physical synthesis optimizations for the HardCopy device, and these post-fit optimizations are passed to the FPGA revision.

The **Effort level** on the **Physical Synthesis Optimizations** page of the **Settings** dialog box for HardCopy III and HardCopy IV devices must be **Fast** because the performance gain achieved compared to the compilation time is very limited.

-  For more information about setting physical synthesis optimizations for the FPGA revision of the designs in the Quartus II software, refer to *Setting up and Running the Fitter* in Quartus II Help.

LogicLock Regions

LogicLock regions are flexible floorplan location constraints that help you place logic on the target device. You can use LogicLock regions in FPGA designs targeted to HardCopy devices, which are also passed onto the HardCopy companion revision. LogicLock regions in HardCopy devices cannot have their size set to **Auto**. Although floating LogicLock regions are supported, Altera recommends that you do not use floating LogicLock regions for HardCopy devices, because floating LogicLock regions may affect the design's ability to meet timing closure. HardCopy LogicLock regions must be manually sized and placed in the floorplan. When LogicLock regions are created in a HardCopy device, they start with width and height dimensions set to (1,1), and the origin coordinates for placement are at **X1_Y1** in the lower left corner of the floorplan. You must adjust the size and location of the LogicLock regions that you create in the HardCopy device before compiling the design.

-  For more information about using LogicLock regions, refer to the *Analyzing and Optimizing the Design Floorplan* chapter in volume 2 of the *Quartus II Handbook*.

PowerPlay Power Analyzer

You can initially perform power estimation and analysis of your HardCopy and FPGA devices using the PowerPlay Early Power Estimator. You can then use the PowerPlay Power Analyzer for a more accurate estimation of your device's power consumption.

-  For more information about using the PowerPlay Power Analyzer, refer to the *PowerPlay Power Analysis* chapter in volume 3 of the *Quartus II Handbook*.

Incremental Compilation

Quartus II incremental compilation within a single Quartus II project is supported for the base family in HardCopy migration flows for both the FPGA first and HardCopy first flows. Exporting and importing partitions is not supported in HardCopy ASIC or FPGA device compilations when there is a migration device setting.

-  For more information about using Quartus II incremental compilation, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook* and *About Incremental Compilation* in Quartus II Help.

HardCopy Design Readiness Check

The HardCopy Design Readiness Check (HCDRC) is available as one of the processing steps in the default compilation of either the FPGA first or the HardCopy first flow. This feature checks issues that must be addressed prior to handing off the HardCopy design to the Altera HardCopy Design Center for the HardCopy back-end process. This is different from the user-driven approach in the HardCopy Advisor, in which you must manually open the advisor to check for any violations.

The checks performed in the HCDRC for the Quartus II software include I/Os, PLL, RAM, ALTX, and settings checks (global, instance, and operating settings).

Turning the HardCopy Design Readiness Check On and Off

The HCDRC is turned on by default, but can be turned on or off using the following .qsf file assignments:

```
set_global_assignment -name \ FLOW_HARDCOPY_DESIGN_READINESS_CHECK ON
```

```
set_global_assignment -name \ FLOW_HARDCOPY_DESIGN_READINESS_CHECK OFF
```

You can also turn on the HCDRC in the **More Compilation Process Settings** dialog box.

Setting Check

The Setting Check report lists the results of the setting checks from the Handoff report. The Setting Check report consists of the following sections.

Summary

The Summary section displays the number of settings that do not meet recommendations. One of the following messages is displayed:

```
<number> global setting(s) do not meet recommendation. Please review the recommendation and do appropriate correction as it may affect the result of the migration to HardCopy.
```

or

```
<number> instance setting(s) do not meet recommendation. Please review the recommendation and do appropriate correction as it may affect the result of the migration to HardCopy.
```

Global Setting

The Global Setting section displays recommendations for global settings only. Global settings with values other than the recommended values are highlighted in red.

Instance Setting

The Instance Setting section is identical to the Global Setting section, but checks only for instances assignments.

Operating Setting

The Operating Setting section displays checks related to the recommended operating settings for the FPGA and the HardCopy device.

This check is primarily applicable to Stratix III devices used as prototype FPGAs because HardCopy III devices only support 0.9 V core voltage, whereas Stratix III devices support both 1.1 V and 0.9 V core voltage.

The Setting Check reports also include checking for illegal assignments in the HardCopy design flow. The illegal assignments checks are shown in [Example 3-1](#).

Example 3-1. Illegal Assignment Checks

```
USE_CHECKERED_PATTERN_AS_UNINITIALIZED_RAM_CONTENT ON (1)
SIGNAL_PROBE_ENABLE ON|OFF
SIGNAL_PROBE_SOURCE ON|OFF (2)
```

Notes to Example 3-1:

- (1) Refer to the section “RAM Usage Check” on page 3-18.
 - (2) SignalProbe is not supported in HardCopy ASICs.
-

I/O Check

The HCDRC I/O check ensures that you have assigned location assignments for the pins, I/O standard, current strength assignment, output pin load assignment, termination assignments, and also checks for any unconnected pins. The tool issues a warning if you have not specified the assignment for the I/O check.

For example, for missing I/O standard assignments, the HCDRC issues the following warning:

```
5 pin(s) have no explicit I/O Standard assignments provided in the setting
file and default values are being used. Please add a specific I/O Standard
assignment for these pins.
```

Input Pin Placement for Global and Regional Clock

Due to the difference in the interconnect delays between the FPGA and HardCopy device, the use of non-primary clock inputs as clock inputs in a design can cause timing closure to be a problem when migrating the FPGA to the HardCopy device. The Input Pin Placement for Global and Regional Clock check informs you of the problem before finalizing the pin location, so that any clock inputs can be moved to the primary clock input.

This check lists all the pins that drive the global or regional clock but are not placed in a dedicated clock pad. All pins are required to have manual location assignments. Pins that are missing location assignments are listed in the Missing Pin Location Assignment report.

The following message appears in the message panel during compilation and also appears in the I/O Check Summary:

```
<number> pin(s) drives global or regional clock, but is not placed in a
dedicated clock pin position. Clock insertion delay will be different
between FPGA and HardCopy companion revisions because of differences in
local routing interconnect delays.
```

PLL Usage Check

The PLL Usage Check report lists PLL usage requirements and violations checks.

PLL Real-Time Reconfigurable Check

This check highlights the PLLs that do not have PLL reconfiguration. PLL reconfiguration allows fine tuning of the PLLs in the design after manufacturing.

The following message appears in the message panel during compilation and also appears in the Logic Check Summary:

```
<number> PLL(s) don't have real time reconfiguration. It is highly recommended that each PLL to have PLL reconfiguration for designs migrating to HardCopy.
```

PLL elements that do not have PLL reconfiguration are listed in a table.

PLL Clock Outputs Driving Multiple Clock Network Types Check

This check is derived from the Design Assistant rule check for HardCopy (H102) and lists all PLL instances in the current design that have clock outputs driving multiple clock network types.

The following message is displayed if the tool detects violations of this type:

```
Found <number> PLL(s) with clock outputs that drives multiple clock network types.
```

PLL with No Compensation Mode Check

This check lists all PLLs that are in No Compensation operating mode. This setting is not recommended for a design migrating to a HardCopy device because of differences in the clock networks and the clock delays between the FPGA and HardCopy device.

The following warning message appears during compilation when a PLL is in a No Compensation mode:

```
<number> PLL(s) is operating in a "No compensation" mode.
```

PLL with Normal or Source Synchronous Mode Feeding Output Pin Check

When a PLL is directly feeding an output pin, it must be set to Zero Delay Buffer operating mode. However, if a PLL is set either in Normal Compensation mode or Source Synchronous mode, a warning message is issued during compilation.

The following warning message appears during the runtime of HC Ready:

```
<number> PLL(s) is in normal or source synchronous mode that is not fully compensated because it feeds an output pin -- only PLLs in zero delay buffer mode can fully compensate output pins.
```

RAM Usage Check

HardCopy series devices do not support initialized RAM blocks upon power-up. However, you can use the ALTMEM_INIT megafunction to initialize the RAMs of a HardCopy series device in your design with the content of a ROM.



For more information about the ALTMEM_INIT megafunction, refer to the [RAM Initializer \(ALTMEM_INIT\) Megafunction User Guide](#).

In HardCopy series devices, RAM blocks power up uninitialized. During the RAM Usage check, the HCDRC tool checks for RAMs that are initialized using a Memory Initialization File (.mif). Any RAM with a .mif file is listed in a table with the following compilation warning message:

```
<number> RAM(s) have Memory Initialization File (MIF). HardCopy devices do not allow initialized RAM. Please ensure that no RAM is initialized by a MIF file.
```

Initialized Memory Dependency Testing

The Assembler module of the Compiler allows you to write an FPGA programming file with an initialized checkerboard pattern for memory contents of M4K memory blocks for the FPGA revision. You should not use this option in a FPGA revision used to migrate to the HardCopy revision because it creates irreconcilable revision differences between the FPGA and HardCopy designs because the HardCopy handoff cannot physically have any initialized memory content. Use this option only on a parallel copy of your compiled FPGA design that you want to test on your board.

To create a programming file with an initialized checkerboard pattern, perform the following steps:

1. Compile your completed FPGA design revision to use for prototype testing. This is the revision you should eventually use to create your HardCopy companion revision.
2. Create the HardCopy companion revision.
3. Compile, compare, and generate the hand off archive files for your design.
4. Switch back to your FPGA revision, and on the Project menu, click **Revisions**, and then double click <<**new revision**>> in the **Revisions** table.
5. In the **Create Revision** dialog box, type a revision name in the **Revision name** box and turn on **Copy database** and **Set as current revision**. This step copies your FPGA revision and sets the new revision as the current open revision in the Quartus II software.
6. On the Assignments menu, click **Settings**, and then click **Assembler** in the **Category** list. Turn on **Use checkered pattern as uninitialized RAM content** on the **Assembler** page, or add the following line to the revision .qsf file:

```
set_global_assignment -name  
USE_CHECKERED_PATTERN_AS_UNINITIALIZED_RAM_CONTENT ON
```
7. Run the Assembler in the FPGA revision to generate a new programming file for your FPGA.
8. Test the new programming file in your prototype environment to determine if your design has a dependency for FPGA RAM contents initialized with zeros after power-up and configuration.

Because the checkerboard pattern is used for testing, the patterns written into the RAM blocks for the new programming file may not detect all cases of zero-initialized RAM content dependencies. Some designs may detect only one bit as zero (for example, the LSB of a memory word), so this method may not detect all cases. This checkerboard pattern test will detect a case when a full RAM word line is expected as zeros at startup.

ALTGX Usage Check

Beginning in the Quartus II software version 10.0, the ALTGX Usage check of the HCDRC performs checks on ALTGX instance usage for designs targeting Stratix IV GX and HardCopy IV GX devices.

The HCDRC tool checks all the ALTGX instances that are initialized in the design for connectivity with the ALTGX_RECONFIG instance. For ALTGX instances that do not connect to an ALTGX_RECONFIG instance, the following warning message appears with the respective instance HSSI_CMU atom name:

```
ALTGX megafunctions do not have ALTGX_RECONFIG megafunctions connected.  
Altera recommends connecting ALTGX_RECONFIG megafunction to each ALTGX  
megafunction when migrating your designs to HardCopy devices.
```

Performing ECOs with Quartus II Engineering Change Management with the Chip Planner

As designs grow larger in density, analyzing designs for performance, routing congestion, logic placement, and executing Engineering Change Orders (ECOs) becomes critical. In addition to design analysis, you can use various bottom-up and top-down flows to implement and manage your design. This process may become difficult to manage, because ECOs are often implemented as last minute changes to your design.

With the Altera Chip Planner, you can shorten the design cycle time significantly. When changes are made to your design as ECOs, you do not have to perform a full compilation in the Quartus II software. Instead, you can make changes directly to the post place-and-route netlist, generate a new programming file, test the revised design by performing a gate-level simulation and timing analysis, and then verify the fix on the system. When the fix has been verified on the FPGA, switch to the HardCopy revision, apply the same ECOs, run timing analysis and the Assembler, compare the revisions, and then run the HardCopy Netlist Writer for design submission.

There are three types of migration scenarios:

- One-to-one changes, which are changes that can be implemented on both architectures—FPGA and HardCopy.
- Changes that must be implemented differently on the two architectures to achieve the same result.
- Changes that cannot be implemented on both architectures.

The following sections outline the methods for migrating each of these types of changes.

Migrating One-to-One Changes

One-to-one changes are implemented using identical commands in both architectures. In general, such changes include those that affect only I/O cells or PLL cells. Some examples of one-to-one changes include creating, deleting, or moving pins, changing pin or PLL properties, or changing pin connectivity (provided the source and destination of the connectivity changes are I/Os or PLLs). These types of changes can be implemented identically on both architectures.

The following list is a partial list of one-to-one change examples:

- I/O creation, deletion, and moves
- I/O property changes (for example, I/O standards, delay chain settings, and so forth)
- PLL property changes
- Connectivity changes between non-LCELL_COMB atoms (for example, PLL to I/O, DSP to I/O, and so forth)

If such changes are exported to Tcl, directly reapplying the generated Tcl script (with a minor text edit) on the companion revision implements the appropriate changes as described in the following steps:

1. Export the changes from the Change Manager to Tcl.
2. Open the generated Tcl script and change the `project_open <project> -revision <revision>` to refer line to the appropriate companion revision.
3. Apply the Tcl script to the companion revision.

Migrating Changes that Must Be Implemented Differently

Some changes must be implemented differently on the two architectures, such as changes affecting the logic of the design. Examples include LUTMASK changes, LC_COMB/HSADDER creation and deletion, connectivity changes not described in the previous section, and different PLL settings for the FPGA and the HardCopy revisions.

 For more information about how to use different PLL settings for the FPGA and HardCopy devices, refer to *AN 432: Using Different PLL Settings Between Stratix II and HardCopy II Devices*.

Table 3-4 summarizes suggested implementation of various changes.

Table 3-4. Implementation Suggestions for Changes that Must Be Implemented Differently

Change Type	Suggested Implementation
LUTMASK changes	Because a single FPGA atom can require multiple HardCopy atoms to implement, you may need to change multiple HardCopy atoms to implement the change, including adding or modifying connectivity.
Make/Delete LC_COMB	If you are using an FPGA LC_COMB in extended mode (7-LUT) or are using a SHARE chain, you must create multiple atoms to implement the same logic functions in the HardCopy device. Additionally, the placement of the LC_COMB cell has no meaning in the companion revision because the underlying resources are different.
Make/Delete LC_FF	Basic creation and deletion is the same on both architectures; however, similar to LC_COMB creation and deletion, the location of an LC_FF in a HardCopy and FPGA revision do not translate.
Editing logic connectivity	Because a LCELL_COMB atom might have to be broken up into several HardCopy LCELL_COMB atoms, the source or destination ports for connectivity changes might have to be analyzed to properly implement the change in the companion revision.

Changes that Cannot be Migrated

A small set of changes are incompatible and cannot be implemented in both architectures. The best example of this incompatibility occurs when moving logic in a design; because the logic fabric is different between the two architectures, locations in the FPGA are not compatible in HardCopy, and vice versa.

Overall Migration Flow

This section outlines the migration flow and the suggested procedure for implementing changes in both revisions to ensure a successful revision comparison such that the design can be submitted to the Altera HardCopy Design Center.

Preparing the Revisions

The general procedure for migrating changes between devices is the same, whether migrating from the FPGA to HardCopy device or vice versa. The steps are described below:

1. Compile the design on the initial device.
2. Migrate the design from the initial device to the target device in the companion revision.
3. Compile the companion revision.
4. Run the **Compare HardCopy Companion Revisions** command. Both revisions should pass the revision comparison.

If testing identifies problems requiring ECO changes, equivalent changes can be applied to both FPGA and HardCopy revisions, as described in the following section.

Applying ECO Changes

The general flow for applying equivalent changes in companion revisions is described below:

1. Make changes in one revision using the Chip Planner tools (Chip Planner, Resource Property Editor, and Change Manager), and then, to verify and export these changes, follow these steps:
 - a. Make changes using a Chip Planner tool.
 - b. Perform a netlist check using the **Check and Save All Netlist Changes** command.
 - c. Verify correctness using timing analysis, simulation, and prototyping (FPGA only). If more changes are required, repeat steps **a** and **b**.
 - d. Export change records from the Change Manager to Tcl scripts, or **.csv** or **.txt** file formats. This exported file is used to assist in making the equivalent changes in the companion revision.
2. Open the companion revision in the Quartus II software.
3. Using the exported file, manually reapply the changes using a Chip Planner tool. As stated previously, some changes can be reapplied directly to the companion revision (either manually or by applying the Tcl commands), while others require some modifications.

4. Run the **Compare HardCopy Revisions** command. The revisions should match.
5. Verify the correctness of all changes, which may require running timing analysis.
6. Run the **HardCopy Assembler** command and the **HardCopy Netlist Writer** command for design submission along with handoff files.

The following command is the Tcl command for running the HardCopy Assembler:

```
execute_module -tool asm -args "--read_settings_files=off --  
write_settings_files=off"
```

The following command is the Tcl command for the HardCopy Netlist Writer:

```
execute_module -tool cdb \  
-args "--generate_hardcopy_files"
```



For more information about using the Chip Planner, refer to the *Quartus II Engineering Change Management with the Chip Planner* chapter in volume 2 of the *Quartus II Handbook*.

Formal Verification of FPGA and HardCopy Revisions

Third-party formal verification software, Cadence Encounter Conformal verification software, is used for several FPGA and HardCopy families.

The formal verification flow for HardCopy ASIC designs is a two-step process. First, run formal verification on the FPGA netlist to ensure that the FPGA netlist matches the RTL. Second, use the **Compare HardCopy Revisions** command in the Quartus II software to ensure that the HardCopy implementation matches the FPGA.



While this flow is enabled, performing formal verification is not necessary due to the one-to-one mapping of logic between the FPGA prototype and the HardCopy ASIC.

To use the Conformal software with the Quartus II software project for your FPGA design revision, you must enable the EDA Netlist Writer so it can generate the necessary netlist and command files required to run the Conformal software. To automatically run the EDA Netlist Writer during the compilation of your FPGA revision, perform the following steps:

1. On the Assignments menu, click **Settings**. The **Settings** dialog box appears.
2. In the **Category** list, under **EDA Tool Settings**, select **Formal Verification**, and then in the **Tool name** list, select **Conformal LEC**.
3. Compile your FPGA and HardCopy design revisions.

The Quartus II EDA Netlist Writer produces the netlist for the FPGA when run on that revision. You can compare your FPGA post-compilation netlist to your RTL source code using the scripts generated by the EDA Netlist Writer.

After both the FPGA and HardCopy revisions have been compiled, you can run the **Compare HardCopy Revisions** command to ensure that the HardCopy implementation matches the FPGA.

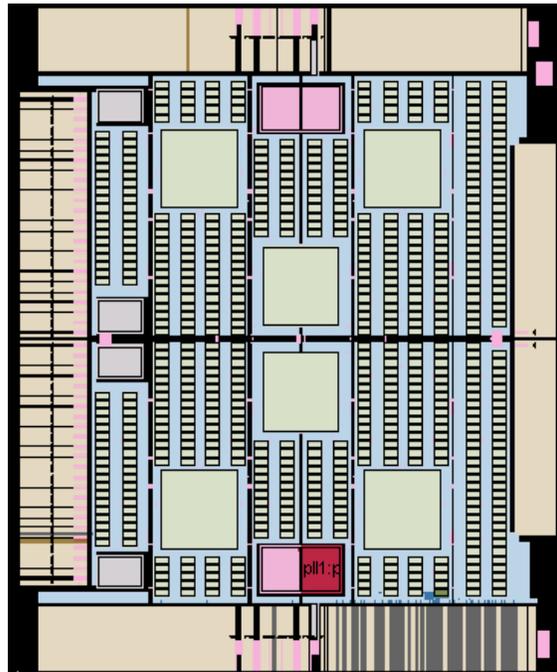
 For more information about using the Cadence Encounter Conformal verification software, refer to the *Cadence Encounter Conformal Support* chapter in volume 3 of the *Quartus II Handbook*.

HardCopy Floorplan View

The Quartus II software displays the floorplan and placement of your HardCopy companion revision. This floorplan shows the preliminary placement and connectivity of all I/O pins, PLLs, memory blocks, HCell macros, and DSP HCell macros. Congestion mapping of routing connections can be viewed using the **Layers Setting** dialog box (from the View menu of the Chip Planner). This is useful in analyzing densely packed areas of your floorplan that can reduce the peak performance of your design. The Altera HardCopy Design Center verifies final HCell macro timing and placement to guarantee that timing closure is achieved.

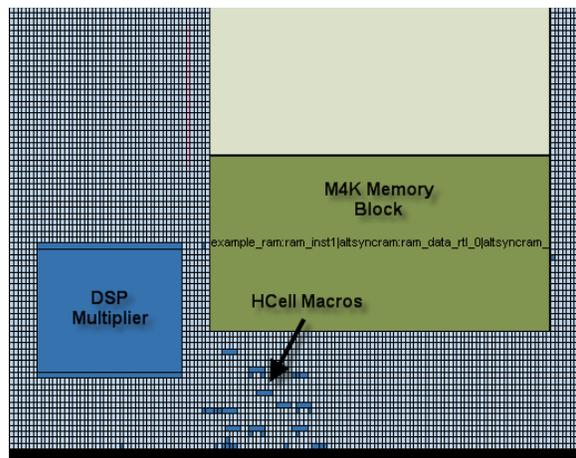
Figure 3-6 shows an example of the HC230F1020 device floorplan.

Figure 3-6. HC230F1020 Device Floorplan



In this small example design, the logic is placed near the bottom edge. You can see the placement of a DSP block constructed of HCell macros, various logic HCell macros, and an M4K memory block. A labeled close-up view of this region is shown in [Figure 3-7](#).

Figure 3-7. Close-Up View of Floorplan



The Altera HardCopy Design Center performs final placement and timing closure on your HardCopy design based on the timing constraints provided in the FPGA design.

 For more information about the Altera HardCopy Design Center process, refer to the respective HardCopy series device handbook, which is available on the Literature page of the Altera website at www.altera.com.

Document Revision History

[Table 3-5](#) shows the revision history for this chapter.

Table 3-5. Document Revision History (Part 1 of 2)

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Edited the “Timing Settings” on page 3-14 section to remove support for the Classic Timing Analyzer ■ Changed to new document template ■ Editorial changes
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Added new section “ALTGX Usage Check” on page 3-21 ■ Updated “LogicLock Regions” on page 3-16 to reflect updated companion revision support ■ Updated “Incremental Compilation” on page 3-16 to reflect updated companion revision support ■ Linked sections throughout the chapter to Quartus II Help ■ Removed “Referenced Documents”

Table 3-5. Document Revision History (Part 2 of 2)

Date	Version	Changes
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Removed HardCopy Stratix legacy support information ■ Updated “Physical Synthesis Optimization” on page 3-15 ■ Updated “Quartus II Software Features Supported for HardCopy Designs” on page 3-15 ■ Updated “Referenced Documents” ■ Updated the tables and figures for HardCopy Series devices
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Updated “RAM Usage Check” on page 3-19 ■ Updated “Referenced Documents”
November 2008	8.1.0	<ul style="list-style-type: none"> ■ Added HardCopy IV E support information ■ Added notes for Initialized Memory Dependency testing ■ Changed page size to 8.5” × 11”
May 2008	8.0.0	<ul style="list-style-type: none"> ■ Updated “RAM Usage Check” on page 3-19 ■ Updated “Referenced Documents”

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

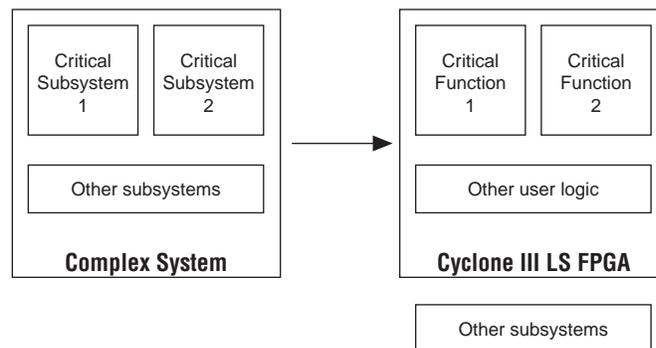
 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter contains rules and guidelines for creating a floorplan with the design separation flow, and assumes familiarity with the Quartus® II incremental compilation flow and floorplanning with the LogicLock™ feature.

The basic principle of a secure and reliable system is that critical subsystems in the design have physical and functional independence. Systems with redundancy require physical independence to ensure fault isolation—that a failure or corruption of any single subsystem will not adversely affect any other part of the system. Furthermore, if errors occur, physical independence simplifies analysis by allowing developers to evaluate each subsystem separately.

Traditionally, systems that require redundancy implement critical IP structures using multiple devices. The Quartus II design separation flow, used in Cyclone® III LS devices, provides the ability to design physically independent structures on a single device. This functionality allows system designers to achieve a higher level of integration on a single FPGA, and alleviates increasingly strict Size Weight and Power (SWaP) requirements. Figure 4-1 illustrates this concept.

Figure 4-1. Achieving Higher Level Integration on a Single Cyclone III LS Device



The Quartus II design separation flow introduces the constraints necessary to create secured regions and floorplan a secured system. When implemented in Cyclone III LS devices, a secured region provides physical independence through controlled routing and a boundary of unused resources. By restricting routing resources and providing a physical guard band of unused logic array blocks (LABs), faults or unintended signals originating in one secured region are prevented from adversely affecting other design blocks on the device.



The Quartus II design separation flow features require specific licensing in addition to licensing the Quartus II software. For further details, contact your local Altera sales representative or Altera distributor.

© 2010 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, HARDCOPY, MAX, MEGACORE, NIOS, QUARTUS and STRATIX are Reg. U.S. Pat. & Tm. Off. and/or trademarks of Altera Corporation in the U.S. and other countries. All other trademarks and service marks are the property of their respective holders as described at www.altera.com/common/legal.html. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

The Quartus II design separation flow incorporates additional LogicLock and floorplanning features into the incremental compilation flow. The following three chapters in the *Quartus II Handbook* serve as companion references to this chapter:

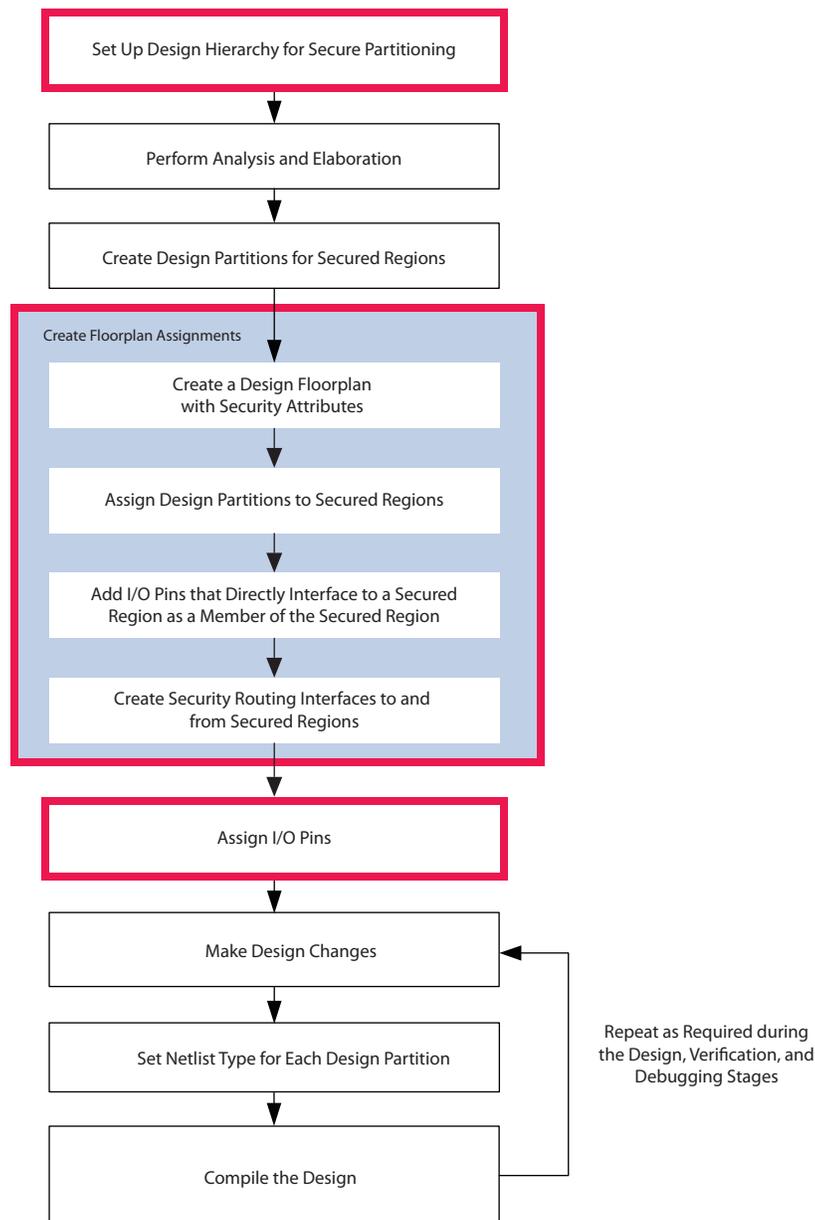
- *Quartus II Incremental Compilation for Hierarchical and Team-Based Design*—Describes the Quartus II incremental compilation flow
- *Best Practices for Incremental Compilation Partitions and Floorplan Assignments*—Contains guidelines for using the incremental compilation flow and creating a design floorplan
- *Analyzing and Optimizing the Design Floorplan*—Describes various attributes associated with LogicLock location constraints and introduces the Chip Planner for creating and modifying a floorplan

Design Flow Overview

The design separation flow is based on the incremental compilation flow. You begin with an incremental compilation design flow and then apply design separation constraints to each design partition that you want to physically isolate from the rest of the design. This section provides an overview of the design separation flow steps.

Figure 4-2 shows a flow chart of the design separation flow. Red boxes in the flow chart highlight steps that are specific to the design separation flow, while the remaining boxes in the flow chart are common to both the design separation and incremental compilation flows. A brief description is given for each step in the flow chart below and serves as a quick-start guide for the design separation flow.

Figure 4-2. Design Separation Compilation Flow



1. **Set up design hierarchy for secured partitioning**—Prepare your design for implementation of the design separation flow, by setting up your design hierarchy for secured partitioning along logical hierarchical boundaries. If necessary, create wrapper files to create logical boundaries in the design hierarchy to support the design entities that you must separate from the remainder of the design.
2. **Perform analysis and elaboration**—Run analysis and elaboration to identify the hierarchy in your design.
3. **Create design partitions for secured regions**—For each design entity that requires physical independence, create a logical design partition for each design entity. Partition logic using guidelines from the incremental compilation flow.

Refer to “[Creating Design Partitions for the Design Separation Flow](#)” on page 4-4 for more information.

4. **Create a design floorplan with security attributes**—After creating design partitions, create LogicLock location assignments and a floorplan, at minimum, for all the entities to be secured in your design. Use the security attributes in the LogicLock Regions window to specify the security level of each LogicLock region. When you apply this attribute, fencing regions are automatically created in your floorplan to isolate the secured LogicLock regions. Refer to “[Creating a Design Floorplan with Secured Regions](#)” on page 4-6 for more information.
5. **Assign design partitions to secured regions**—Assign design partitions to secured LogicLock regions to separate them from each other and from all other hierarchy blocks. Refer to “[Using Secured Regions](#)” on page 4-9 for more information.
6. **Add I/O pins that directly interface with a secured region as a member of the secured region**—If a secured region interfaces with one or more I/O pins, make the I/O pins members of the secured region. If a secured region has I/O pins as members, that region must overlap the I/O pads. Refer to “[Adding I/O Pins as Members of Secured Regions](#)” on page 4-9 for more information.
7. **Create security routing interfaces to and from secured regions**—Create security routing interfaces by applying the security routing interface attribute to LogicLock regions.

Only routing resources can be used within a security routing interface; no logic can be placed. Each security routing interface must abut one or two secured regions. After you create an interface region for each signal or group of signals entering or exiting a secured region, assign the signals to the appropriate routing interfaces.

For signals routing between secured regions with different security attributes or between a secured region and an unsecured region, you must lower the security attribute for the signal exiting the stricter security region. Refer to “[Making Signal Security Assignments](#)” on page 4-19 for more information.

8. **Assign I/O pins**—After creating secured regions and security routing interfaces, if the secured regions contain I/O pins as members, assign I/O pins to meet design separation flow requirements. For example, I/O banks cannot be shared between secured regions. If a secured region contains I/O pins as members, the entire I/O bank is usable only by the secured region that sinks or sources the I/O pin. Refer to “[Assigning I/O Pins](#)” on page 4-25 for more information.
9. **Make design changes, set the netlist type for each design partition, and compile the design**—After making the necessary I/O pin assignments, you complete the design separation flow-specific steps, and you can start the iterative process of making design changes, setting the netlist type for each design partition, and then compiling your design until you achieve a floorplan that meets your design requirements.

The design separation flow-specific steps, step 1 and steps 4 through 8, are described in further detail in subsequent sections in this chapter.

Creating Design Partitions for the Design Separation Flow

After setting up your design to support secured partitioning and running analysis and elaboration, you can create design partitions.

Each secured region floorplan assignment uses a single design partition in the incremental compilation flow to identify the functional elements belonging to a secured region. Design partition assignments are made along entity boundaries in the hierarchy of your RTL design.

Because only a single design partition may be used in a secured region, you must plan your design entities such that logic that requires physical isolation from the rest of the design is packed into a single design entity. Additionally, you should create wrapper files where necessary to reorganize your hierarchy, so that all your secured regions are contained within a single entity or module in your RTL. The incremental compilation feature allows functional independence of each design partition because it disables netlist optimizations across partition boundaries.

Most of the rules, guidelines, and tools for creating design partitions used in the incremental compilation flow are applicable in the design separation flow. You can use the Incremental Compilation Advisor, the Design Partition Planner, and the Chip Planner features in the Quartus II software to help you create design partition assignments.

When creating design partitions, the following considerations are important:

- Register the inputs and outputs of a design partition to avoid cross-boundary logic optimizations and to maintain timing performance along the signal path.
- Minimize the number of I/O paths that cross partition boundaries to keep logic paths within a single partition for optimization. Minimizing the number of cross-boundary I/O paths makes partitions more independent for both logic and placement optimization.
- Avoid logic that requires cross-boundary logic optimizations.



For more details about guidelines for creating design partitions, refer to the *Best Practices for Incremental Compilation Partitions and Floorplan Assignments* chapter in volume 1 of the *Quartus II Handbook*.

When creating your design in the design separation flow, you must be aware of some restrictions and special considerations that differ from the incremental compilation flow. These considerations are discussed in the following “[Merging PLL Resources](#)” and “[Avoiding Multiple Design Partitions With a Secured Region](#)” sections.

Merging PLL Resources

In the Quartus II incremental compilation flow without design separation constraints, the Fitter can use the same PLL resource on the device when multiple design partitions instantiate a PLL with the same parameters. This resource merge occurs even if optimization across design partitions is required. When the design separation flow is enabled and a design contains one or more secured regions, PLL merging across design partitions is disabled, which helps to maintain the physical separation between design partitions. PLL merging is disabled for the entire design, even if LogicLock regions in a Cyclone III LS design contain no security attributes. For partitions that require shared PLL resources, the PLL must be instantiated outside of the design partitions.

Avoiding Multiple Design Partitions With a Secured Region

Multiple design partitions, including child partitions and multi-hierarchy partitions, are not allowed in a secured region. Each secured region, which you designate after creating design partitions, must contain only a single design partition.

Child partitions are design partitions created from a subentity of an existing design partition and would potentially create multiple design partitions in a secured region, so they are not allowed in the design separation flow.

Multi-hierarchy partitions are created by merging multiple design partitions from different branches of the hierarchy. These partitions are merged into a single netlist during elaboration to allow cross-boundary optimizations during synthesis and fitting, and result in a single incremental result for each multi-hierarchy partition. Multi-hierarchy partitions function similarly as single-hierarchy partitions, but must contain hierarchies from a common parent partition and are not allowed in the design separation flow.

Creating a Design Floorplan with Secured Regions

After creating design partitions, you can create a design floorplan with secured regions with the Chip Planner and security attributes in the LogicLock Regions window.

The Quartus II software uses LogicLock location assignments to map logic in your design hierarchy to physical resources on the device. The Chip Planner provides a visual floorplan of the entire device and allows you to move and resize your LogicLock location constraints on the floorplan of the device. The design separation flow adds a security attribute constraint to each LogicLock region to further constrain routing to achieve physical isolation between LogicLock regions. Signals that require connectivity between two secured regions or between a secured region and unsecured logic are assigned to a special LogicLock region known as a security routing interface. A security routing interface is a controlled region that limits the routing of the contained signals to only the one or two LogicLock regions that this region abuts.

To create fault isolation between secured regions, the design separation flow selectively shuts off routing around the periphery of a secured region. Because signal connectivity at the boundary of the secured region is unused, any faults that occur within the secured region are prevented from adversely affecting neighboring regions. Fault isolation, when using the design separation flow, is possible because no physical connection exists to propagate the fault outside of the region.

Cyclone III LS devices use a MultiTrack interconnect architecture consisting of row and column interconnects that span fixed distances to achieve signal connectivity between LABs. In the horizontal direction, row interconnects use wire resources that span 1 LAB, 4 LABs, and 24 LABs. These row-routing resources are direct link interconnects, R4 interconnects, and R24 interconnects, respectively. In the vertical direction, routing resources span distances of 1 LAB, 4 LABs, and 16 LABs. These column routing resources are register chain interconnects, C4 interconnects, and C16 interconnects, respectively. In the design separation flow, LogicLock region routing wires (C4, C16, R4, and R16) that cross outside the border of a boundary are turned off. Each secured region uses an unused boundary (or a fence) of LABs to guard against the faults from wire resources spanning a length of one-LAB (direct link

and register chain routing resources) from affecting a neighboring region.

The rules and guidelines for floorplanning in the design separation flow are similar to those in a typical compilation flow. However, there are some special considerations for the relative placement of secured regions in your design floorplan. Because each secured region is a keep-out region for routing resources from other LogicLock regions, ensure that a routing path with valid communication interfaces exists between secured regions. Furthermore, the routing path (encapsulated in a security routing interface) should not follow a circuitous path and must be simple enough to meet your timing requirements.

A Fitter-generated floorplan is not possible while a security attribute is applied to a LogicLock region; that is, the size attribute cannot be **Auto**, and the state attribute cannot be **Floating** for any LogicLock region in a secured design.



You can use a Fitter-generated floorplan, created without security attributes, as a starting point to create a final floorplan for the design separation flow.

To use a Fitter-generated floorplan as an initial floorplan, apply **Reserved** attributes to LogicLock regions that must be physically isolated from the rest of the design. A Fitter-generated floorplan with **Reserved** attributes generates non-overlapping LogicLock regions. You can modify the initial floorplan by adjusting the relative placement for each secured region, taking into account the connectivity requirements for each region.

Subsequent sections further detail the rules and guidelines for floorplanning that are specific to the design separation flow.



For more information about using the Chip Planner settings and options, refer to the *Analyzing and Optimizing the Design Floorplan* chapter in volume 2 of the *Quartus II Handbook*.

Using Security Attributes

The **Security Attributes** column in the LogicLock Regions window and the **Security** tab in the **LogicLock Regions Properties** dialog box are available when your version of the Quartus II software is licensed specifically for the design separation feature. Setting the **Security attribute** applies a constraint to a LogicLock region, making the region either a secured region or a security routing interface, from where signals enter or exit a secured region.

The **Signals** list is populated after analysis and synthesis with the inputs and outputs of secured regions. Columns in the **Signals** list describe the **Security Level**, the security routing interface the signal is assigned to, and whether the signal is an output or input to the region.

The design separation flow security features are highlighted in the LogicLock Regions window and the **LogicLock Regions Properties** dialog box shown in [Figure 4-3](#) and [Figure 4-4](#), respectively.

Figure 4-3. Security Attribute Column Available in the Design Separation Flow

Region Name	Size	Width	Height	State	Origin	Reserved	Security Attribute
CPU_to_top	Fixed	8	7	Locked	X96_Y2	On	Security routing interface
CPU_wrap_2_Pattern_gen	Fixed	18	3	Locked	X77_Y59	On	Security routing interface
Eth1_to_Eth2	Fixed	23	9	Locked	X24_Y32	On	Security routing interface
Eth2_to_Pattern_check	Fixed	17	8	Locked	X59_Y7	On	Security routing interface
Ethernet_1_wrap:Ethernet_1_wrap...	Fixed	Custom Shape	Custom Shape	Locked	Custom Shape	On	1
Ethernet_2_wrap:Ethernet_2_wrap...	Fixed	59	25	Locked	X0_Y7	On	1
Pattern_checker_wrap:Pattern_ch...	Fixed	19	8	Locked	X76_Y7	On	1
Pattern_gen_to_Ether1	Fixed	1	6	Locked	X75_Y63	On	Security routing interface

Figure 4-4. Security Tab Available in the Design Separation Flow

Specify the security settings of the LogicLock region and its signals. A classified region must either contain exactly one design partition, or pins, or both, as its members. Signals of a region inherit the region's security level, unless manually lowered. Signals going out of or into a classified region must do so via a security routing interface. A security routing interface is a user-created region placed adjacently to at least one classified region, and has one or more signals assigned to it as members.

Security attribute: 1

Signal	Security Level	Security Routing Interface	Type
Ethernet_1_wrap:Ethernet_1_wrap_1 Ethernet_1 control port readdata[0]	Auto	CPU to Ethernet1	Region input
Ethernet_1_wrap:Ethernet_1_wrap_1 Ethernet_1 control port readdata[10]	Auto	CPU to Ethernet1	Region input
Ethernet_1_wrap:Ethernet_1_wrap_1 Ethernet_1 control port readdata[11]	Auto	CPU to Ethernet1	Region input
Ethernet_1_wrap:Ethernet_1_wrap_1 Ethernet_1 control port readdata[12]	Auto	CPU to Ethernet1	Region input
Ethernet_1_wrap:Ethernet_1_wrap_1 Ethernet_1 control port readdata[13]	Auto	CPU to Ethernet1	Region input
Ethernet_1_wrap:Ethernet_1_wrap_1 Ethernet_1 control port readdata[14]	Auto	CPU to Ethernet1	Region input
Ethernet_1_wrap:Ethernet_1_wrap_1 Ethernet_1 control port readdata[15]	Auto	CPU to Ethernet1	Region input
Ethernet_1_wrap:Ethernet_1_wrap_1 Ethernet_1 control port readdata[16]	Auto	CPU to Ethernet1	Region input
Ethernet_1_wrap:Ethernet_1_wrap_1 Ethernet_1 control port readdata[17]	Auto	CPU to Ethernet1	Region input
Ethernet_1_wrap:Ethernet_1_wrap_1 Ethernet_1 control port readdata[18]	Auto	CPU to Ethernet1	Region input
Ethernet_1_wrap:Ethernet_1_wrap_1 Ethernet_1 control port readdata[19]	Auto	CPU to Ethernet1	Region input
Ethernet_1_wrap:Ethernet_1_wrap_1 Ethernet_1 control port readdata[1]	Auto	CPU to Ethernet1	Region input
Ethernet_1_wrap:Ethernet_1_wrap_1 Ethernet_1 control port readdata[20]	Auto	CPU to Ethernet1	Region input
Ethernet_1_wrap:Ethernet_1_wrap_1 Ethernet_1 control port readdata[21]	Auto	CPU to Ethernet1	Region input
Ethernet_1_wrap:Ethernet_1_wrap_1 Ethernet_1 control port readdata[22]	Auto	CPU to Ethernet1	Region input
Ethernet_1_wrap:Ethernet_1_wrap_1 Ethernet_1 control port readdata[23]	Auto	CPU to Ethernet1	Region input
Ethernet_1_wrap:Ethernet_1_wrap_1 Ethernet_1 control port readdata[24]	Auto	CPU to Ethernet1	Region input
Ethernet_1_wrap:Ethernet_1_wrap_1 Ethernet_1 control port readdata[25]	Auto	CPU to Ethernet1	Region input

Note: List is automatically populated after Analysis & Synthesis and Partition Merge.

Table 4-1 lists a summary of the Security Attributes available for the design separation flow.

Table 4-1. Security Attributes for LogicLock Regions (Part 1 of 2)

Security Attribute	Description
Unsecured	Removes the constraint for physical isolation.
1	Creates a secured region. Physically isolates the LogicLock region by restricting routing resources from leaving the region. Creates a one-lab width border of unused LABs around the LogicLock region. Applying this attribute to a LogicLock region sets the global assignment LL_REGION_SECURITY_LEVEL 1 for the LogicLock region.

Table 4-1. Security Attributes for LogicLock Regions (Part 2 of 2)

Security Attribute	Description
2	<p>Creates a secured region. Security attribute 2 typically represents a stricter level of fault isolation than security attribute 1. For Cyclone III LS devices, implementation of security attribute 2 is the same as security attribute 1; however, this may not be true in subsequent architectures supporting the design separation flow. When selected for the Cyclone III LS family, the Quartus II software creates a one-lab width border of unused LABs around the LogicLock region.</p> <p>Applying this attribute to a LogicLock region sets the global assignment <code>LL_REGION_SECURITY_LEVEL 2</code> for the LogicLock region.</p>
Security Routing Interface	<p>Creates a routing interface for signals entering or exiting a secured region. Only routing resources (no logic) may be used within a security routing interface.</p> <p>Applying this attribute to a LogicLock region sets the global assignment <code>LL_SECURITY_ROUTING_INTERFACE ON</code> for the LogicLock region.</p>

Using Secured Regions

When you apply a secured region attribute (1 or 2) to an existing LogicLock region, the LogicLock region must have a fixed size with a locked origin. Each secured region must have a minimum size of eight-LABs in both the horizontal and vertical dimensions. A region smaller than 8×8 LABs may be non-routable when using the design separation flow.

Child regions are not allowed when creating a secured region because a secured region contains only a single partition. In the non-secured compilation flow, child regions are used primarily to ensure that logic in a child partition is physically contained inside the LogicLock region of the parent partition.

Adding I/O Pins as Members of Secured Regions

A secured region must contain all physical device resources required to complete compilation. I/O pads that are members of a secured region must be contained within the boundaries of the secured region that sources or sinks it. That is, a secured region must overlap the I/O pads that are members of the region. If the logic in the secured region instantiates a PLL or a clock block, those physical device resources must also be overlapped by the region.

You can add I/O pins as members of a secured region using the **LogicLock Region Properties** dialog box.

Using Security Routing Interfaces

A LogicLock region with the security routing interface security attribute creates a routing channel for signals to and from a secured region. No logic may be placed in a security routing interface. Each security routing interface can connect two secured regions, or a secured region with one or more unsecured regions. If you are connecting two secured regions, a fencing region is automatically placed around the interface region. You can assign each signal entering or exiting a secured region to a security routing interface on the **Security** tab in the **LogicLock Regions Properties** dialog box.

For information about assigning signals to a security routing interface, refer to [“Making Signal Security Assignments” on page 4-19](#).

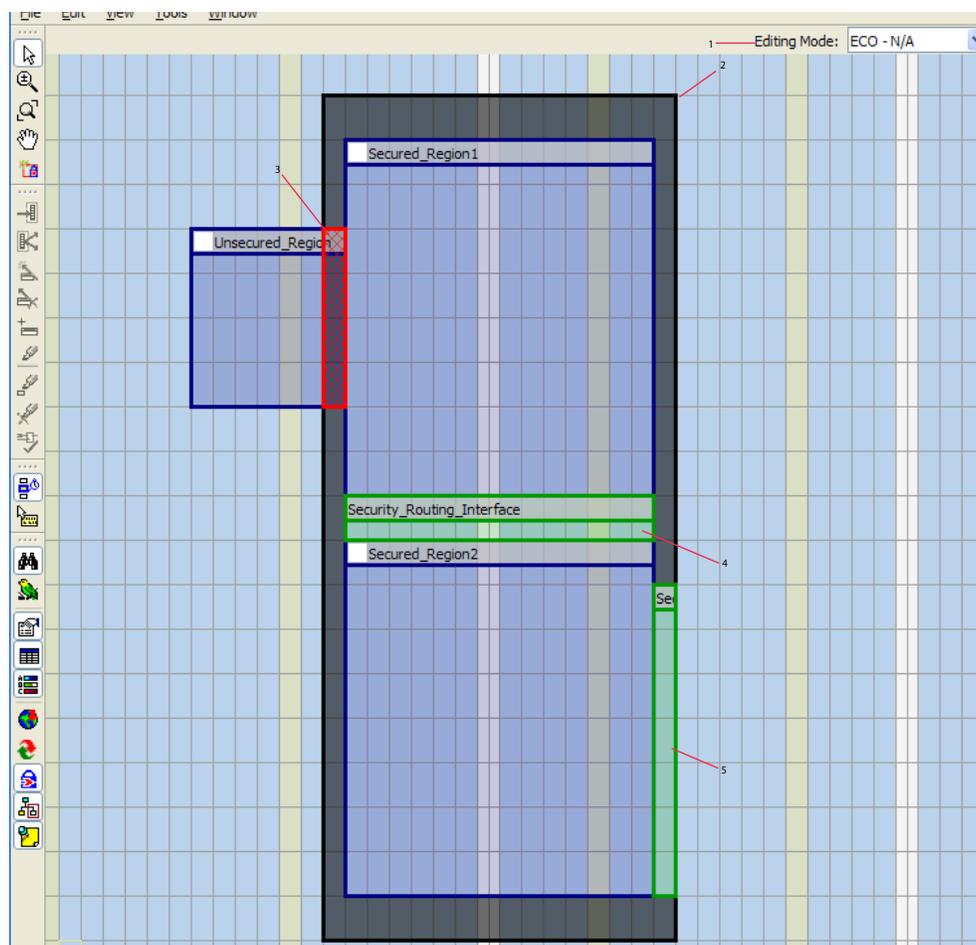
For information about the number of signals that can be contained in a security routing interface, refer to “Routing Restrictions” on page 4-26.

Making Design Separation Flow Location Assignments in the Chip Planner

The Chip Planner allows you to visually modify the size and location of LogicLock regions. This section describes the attributes of LogicLock regions within the context of the design separation flow.

When the design separation flow is enabled, the fencing region around each secured region in the Chip Planner is shaded grey. Security routing interfaces are shaded green. Illegal placements that violate secured region boundaries are highlighted in red at the location where the violation occurs. Figure 4-5 shows the LogicLock regions with security attributes in the Chip Planner.

Figure 4-5. LogicLock Regions With Security Attributes



Notes to Figure 4-5:

- (1) Floorplan Editing Mode task.
- (2) Unused fence around a secured region
- (3) Security violation, created by a LogicLock region placement within a fencing region of a secured region
- (4) Security routing interface region connecting two secured regions
- (5) Security routing interface region connecting secured region and unsecured logic

Understanding Fencing Regions

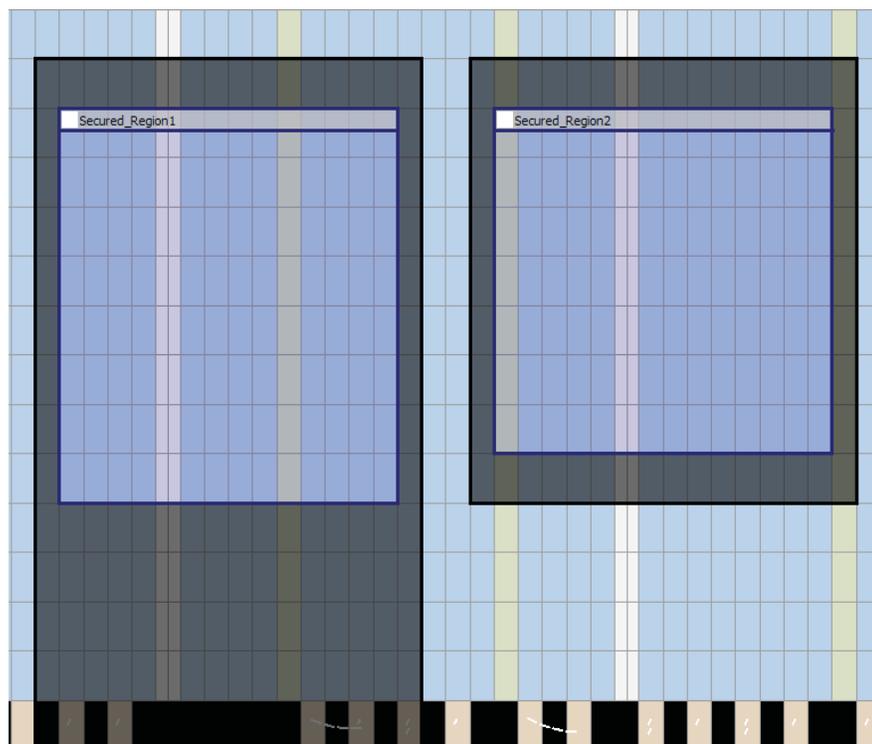
The Quartus II software automatically adds a fencing region, which is a border of unused LABs, when you apply security attribute 1 or security attribute 2 to a LogicLock region. No logic may be placed into a fencing region. The Fitter does not use any routing wires that exit the fence boundary of a secured region. Because direct drive and carry chains can be used at the edge of a secured region, the fencing region prevents signals driven on one length one wires (in the horizontal and vertical directions) from exiting the secured region.

The fencing region around a secured region is generally one-LAB horizontally and one-LAB vertically. There are two regions that require special fencing regions:

- Vertical I/O regions
- Areas around the configuration engine

I/O banks along the top and bottom of the chip use only vertical routing wires to and from the I/O Elements (IOEs). The heavy use of C4 wires from IOEs creates a four-LAB fence between the vertical I/O banks and a secured region. Secured regions requiring a connection to I/O in the top or bottom banks of the device optimally use resources if you add the I/O signals as members of the secured region and overlap the corresponding I/O pads in the floorplan. In [Figure 4-6](#), **Secured_Region2** is five LABs away from the bottom of the device and **Secured-Region1** is four LABs away from the bottom of the device.

Figure 4-6. Vertical Fencing Near I/O Banks



A configuration engine is a hard IP block that manages the configuration of the device. Additionally, the configuration engine routes the control signals for the CRC detection circuit and the internal oscillator into the core logic on the device. In the

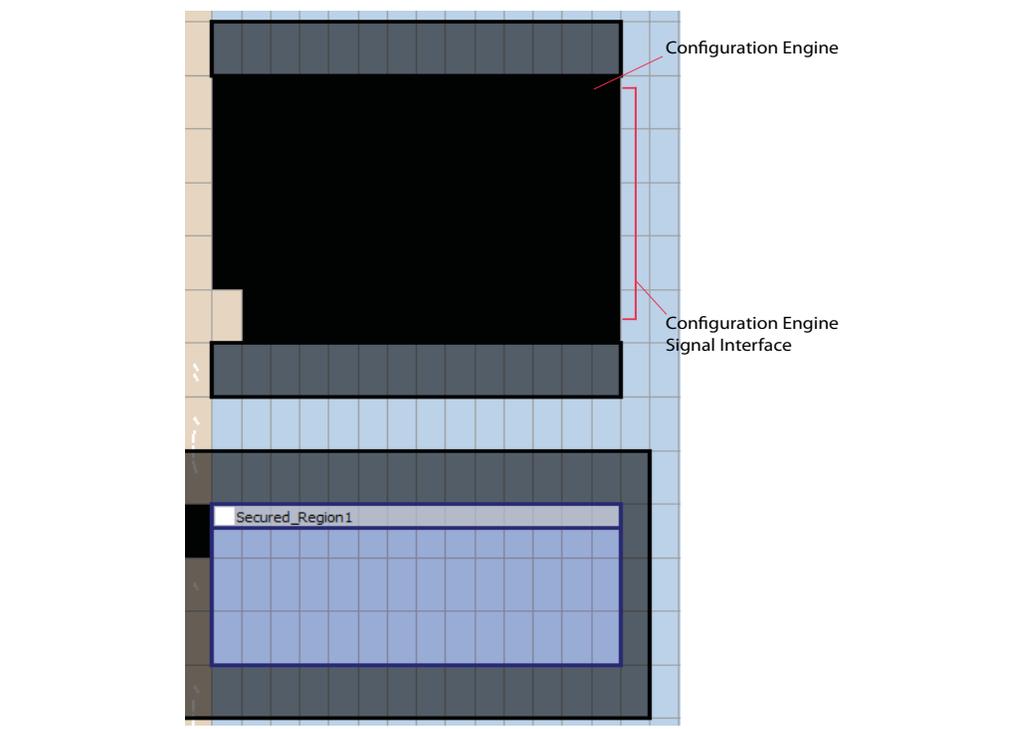
design separation flow, a one-LAB fence is automatically added around the configuration engine whenever a secured region occupies the same LAB column as the configuration engine. The configuration engine is a region notched out of the left side in the middle of the device.

All control signals to and from the configuration engine route from its right edge. If you use an instantiated WYSIWYG that uses any control signals to and from the configuration engine, the signals must either interface with unsecured logic or they must interface with a secured region through a security routing interface.

 If your design routes signals to and from the configuration engine, placing a secured region that directly abuts the configuration engine signal interface (along the right side of the configuration engine) causes a Fitter error.

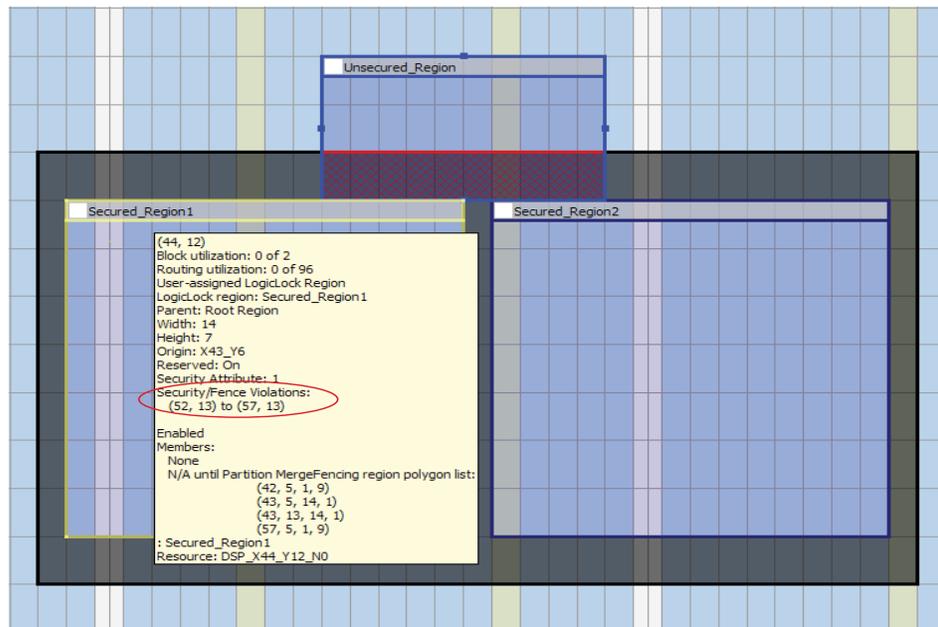
Figure 4-7 shows a configuration engine with a fencing region in the floorplan.

Figure 4-7. Configuration Engine



Fencing regions between two secured regions are allowed to overlap. That is, two adjacent secured regions can be separated by a one-LAB fence. The Chip Planner issues a security warning violation if a LogicLock region is placed within the boundary of a secured region. Security violations are highlighted in red and the tooltip of a secured region indicates the locations of all security violations. You may receive an error if you try to compile a design with a security violation. Figure 4-8 shows two regions with overlapping fences and a security violation from an unsecured region.

Figure 4–8. Overlapping Regions



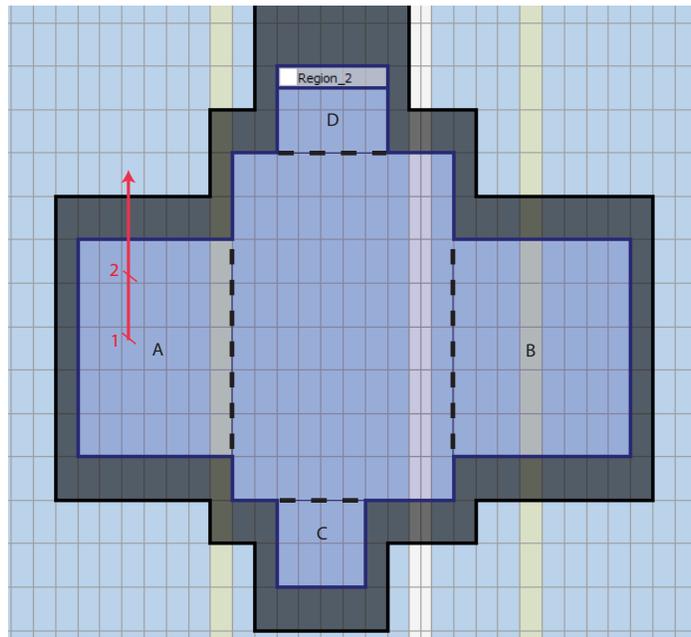
Creating Non-Rectangular Regions

You can create non-rectangular regions by creating multiple rectangular regions and then merging them.

- For more information about creating non-rectangular regions in the Chip Planner in the Quartus II software, refer to *Creating and Manipulating LogicLock Regions* in Quartus II Help.

Non-rectangular LogicLock regions in the design separation flow make circuitous routes more likely. As such, non-rectangular regions can have an adverse affect on performance when used with the design separation flow.

If a secured non-rectangular region contains a subregion that is less than 8×8 LABs, the chances of a non-routable situation occurring increases. Subregions that deterministically require the use of certain routing resources may not fit successfully if a violation of the secured region is occurs. As a general guideline, each subregion should be 8×8 LABs or larger, to ensure that routing resources with a length four LABs are readily available. In Figure 4–9, each subregion of Region 2 (labeled A, B, C, and D) are less than 8×8 LABs in dimension. These subregions can potentially cause a no-fit situation. Depending on the placement and connectivity of LABs, certain routes may be difficult to achieve. For example, the Fitter would not be able to route a connection from LAB 1 to LAB 2 in region A directly. While another path may be possible, a series of hops that do not leave the LogicLock region may not be available and may not satisfy the timing requirements of the route.

Figure 4-9. Non-Rectangular LogicLock Regions

Guidelines for the Relative Placement of Secured LogicLock Regions

Because each secured region is a keep-out region for placement and routing of any logic that is not a member of the secured region, you should be aware of the guidelines in this section as you lay out your floorplan. Placement that does not account for the connectivity requirements between LogicLock regions may cause poor performance or a non-routable design. The guidelines for floorplanning when using the design separation flow include:

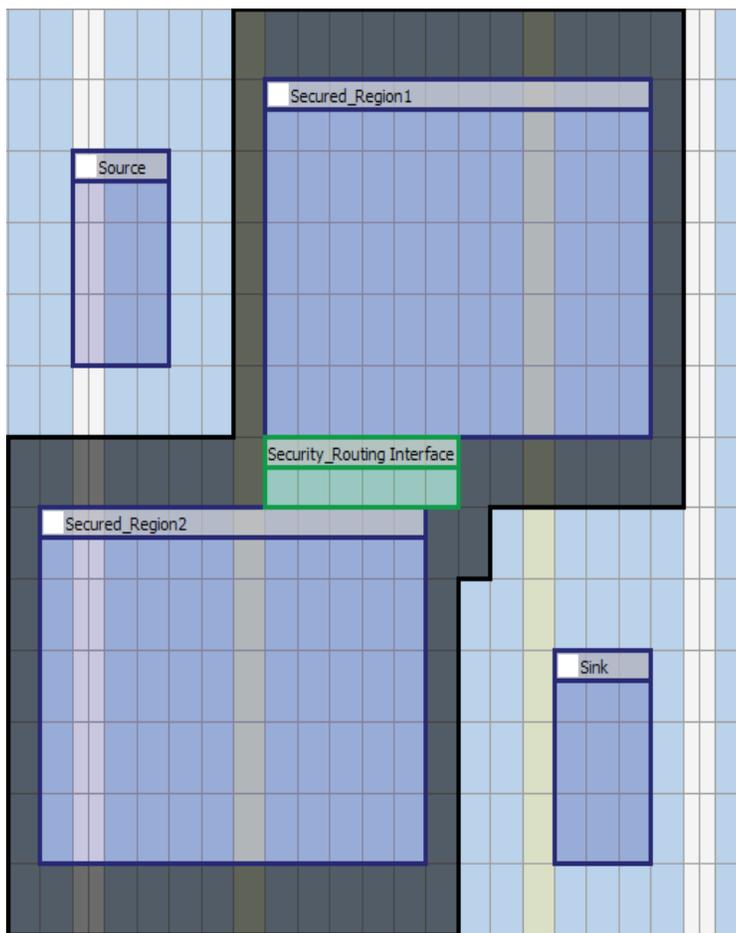
- Create a complete floorplan, including location assignments for unsecured logic.
- Create a non-circuitous route between secured regions requiring a routing region. Generally, routing regions between secured regions should be rectangular.
- Create security routing interfaces between secured regions that do not intersect with other routing regions; secured regions and their routing edges must fit on a single plane. A secured region must overlap any physical resources (such as I/Os, PLLs, and CLKCTRL) that are instantiated by the design partition contained in the secured region.
- Abut the secured region to the edge of the device whenever possible.

Creating a Complete Floorplan

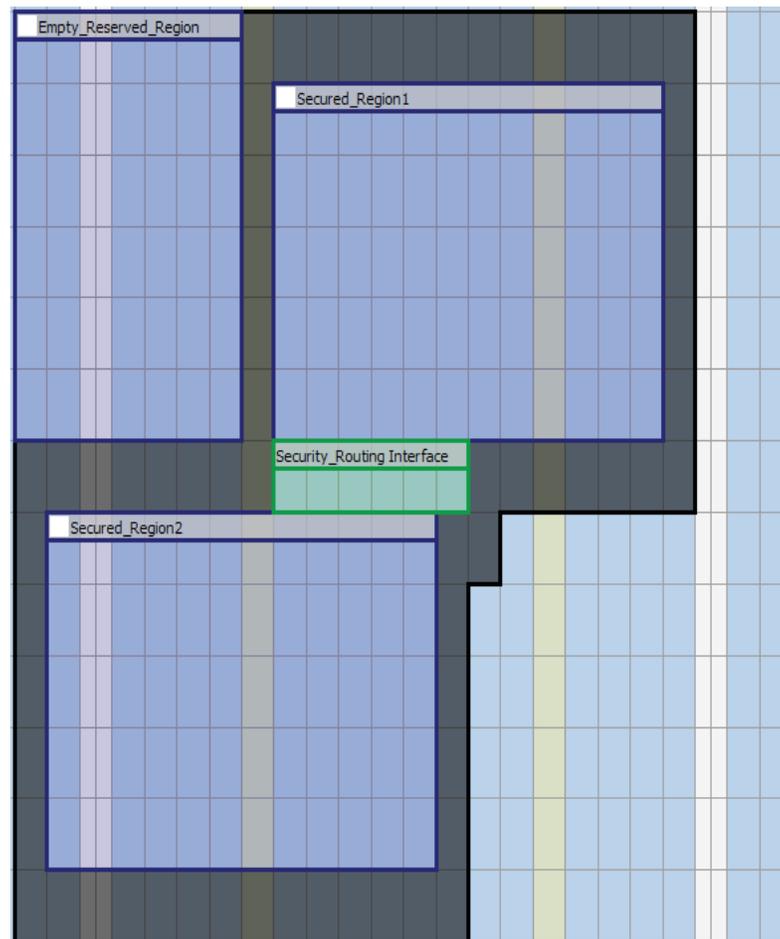
You should allocate a region for all logic in your design. If you have a large secured region that divides the device into multiple disjointed regions, and you have unsecured logic that is not floorplanned, the design may not be routable.

If an unsecured partition does not contain any location assignments, the placement algorithms may make logic assignments on any unallocated space on the device. In the floorplan shown in [Figure 4-10](#), the source and sink registers do not have a valid path through the device, because all routing channels are occupied by Secured Region 1 and Secured Region 2.

Figure 4-10. Non-Routable Placement Example



If a complete floorplan is not possible for all partitions in your design, you can use empty LogicLock regions with the **Reserved** attribute to prevent the Fitter from placing any logic in a region that can potentially cause a no-fit. For the example provided in Figure 4-10, an empty region can be placed in the upper-left corner of the device to prevent any logic that has not been floorplanned from being placed there, which is then shown in Figure 4-11.

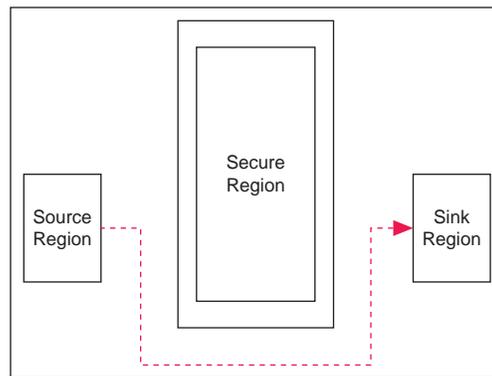
Figure 4-11. Empty Reserved Region Preventing Fitter From Placing Logic

Ensuring Routability Between Regions

The Quartus II software cannot create auto-generated location constraints for any region with a security attribute. If you use a Fitter-generated placement as a starting point for a floorplan with security attributes, an optimal floorplan in a design without separation may not work in the same design. In a floorplan without secured regions, only the placement of logic is restricted. All routing resources on the device are available for the Fitter, and may be routed through a region. Secured regions reserve all routing resources within the LogicLock boundary to the design partition contained in the region.

Having a circuitous route between two regions degrades performance and may cause a non-routable design. Modify any regions that have signal connectivity and must route around a secured region to achieve a connection. [Figure 4-12](#) shows a floorplan that does not contain disjointed parts. However, the source region must route around a secured region to connect to the sink region.

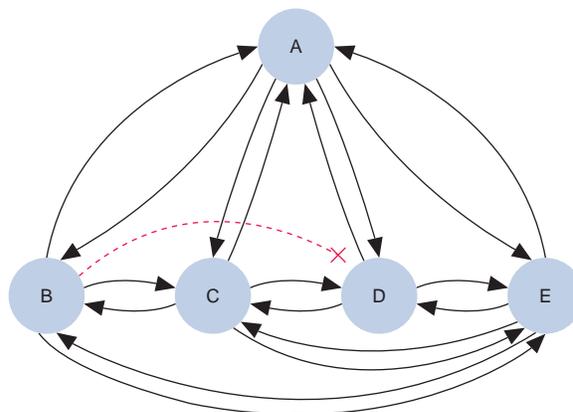
Figure 4-12. Relative Placement of Regions Containing a Circuitous Path



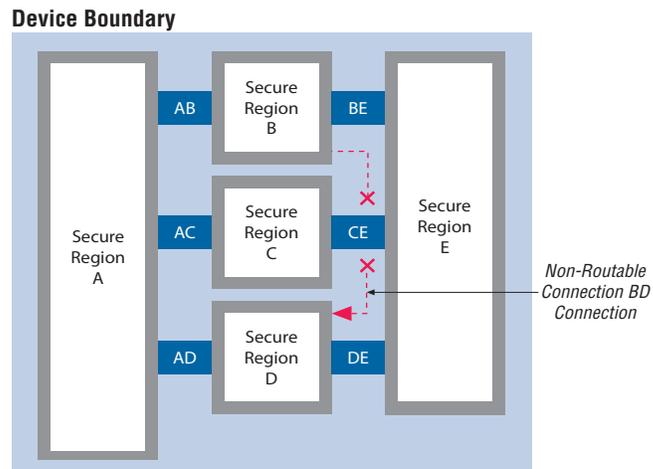
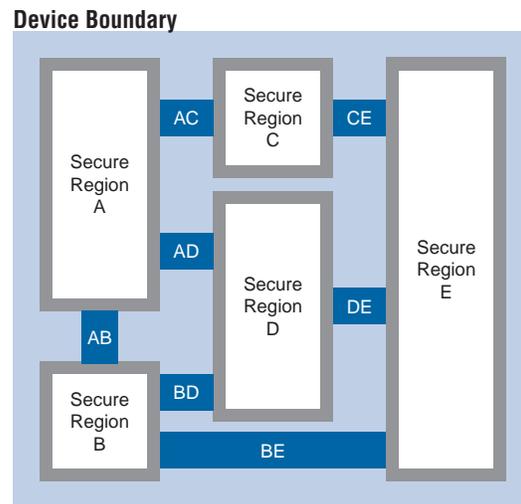
Ensuring Planarity

A fence is automatically created around a security routing interface connecting two secured regions. Because no other routing resources may pass through a security routing interface connecting two secured regions, you should model all secured regions as nodes in a routing graph and all security routing interfaces as the edges, and all nodes and their edges must fit on a planar graph (that is, none of the edges can intersect). If you have five or more secured regions on the device, and each secured region contains signals that fan-out to multiple secured regions, a planar floorplan may not be possible. Figure 4-13 shows a routing graph with five nodes. A complete graph having each pair of distinct vertices connected by an edge is not possible without having any of the edges cross. If the topology of your floorplan contains such a non-routable arrangement, your design hierarchy must be rearranged to collapse related design partitions into a single design partition.

Figure 4-13. Non-Planar Routing Graph: Connection BD Not Possible

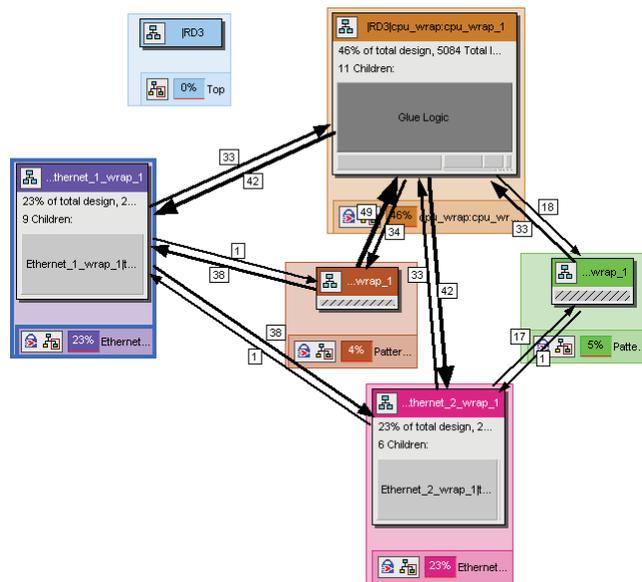


If your secured regions and security routing interfaces can be modeled as a planar graph, but have a high degree of connectivity between the components, you may have to rearrange the shape, size, or location of the secured regions to generate a routable floorplan. For instance, the hypothetical floorplan shown in Figure 4-14 does not have a valid routing path BD (between region B and region D). The modified floorplan in Figure 4-15 shows how all of the required connections can be achieved on a planar surface.

Figure 4-14. Floorplan with Non-Routable Connection BD**Figure 4-15. Floorplan Arranged to Accommodate Connection BD**

You can use the Design Partition Planner for a visual representation of the connectivity between design partitions. This tool helps you determine if the secured regions in the design can be arranged on a planar floorplan. [Figure 4-16](#) shows the Design Partition Planner.

Figure 4-16. Design Partition Planner



Placing Physical Resources

All physical resources that are required by the secured region must be contained inside the boundary of the secured region, including I/O pins connected to the secured region and primitives that have been instantiated within the secured region, such as PLLs and clock control blocks.

Making Signal Security Assignments

Each signal that enters or exits a secured region must have a security level attribute and be explicitly assigned to a security routing interface. The security level for each signal is automatically assigned a default value and matches the secured region that is the source of the signal. Possible security levels of a signal include: **Auto**, **Unsecured**, **1**, and **2**. An assignment of **Auto** sets the default security level for the signal.

A signal with a security attribute may connect to a region with an equivalent or higher security level. For example, a signal with a security level of **Unsecured** can drive logic in a region set to **Unsecured**, **1**, or **2** and a signal with a security level of **1** can drive logic in a region set to **1** or **2**. A signal originating from a secured region may not drive logic in a region with a lower security level. If you have a signal from a higher security level that must drive logic in a lower security level, you can direct the Fitter to honor the connection by explicitly lowering the security level of the signal.

At most, each security routing interface connects two regions. If a signal fans out to multiple regions, assign the signal to multiple security routing interface regions; one interface region per destination.

You can assign signals to security routing interfaces and the security level of signals with the **Security** tab in the **LogicLock Region Properties** dialog box, as shown in [Figure 4-4](#).

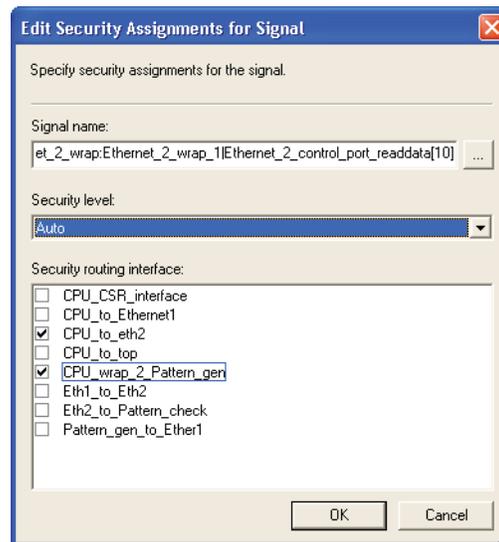
To assign a signal to a security routing interface, follow these steps:

1. On the **Security** tab of the **LogicLock Regions Properties** dialog box, select a signal name in the **Signals** list, and then click **Edit**. The **Edit Security Assignments for Signal** dialog box appears, as shown in [Figure 4-17](#).

 Alternatively, you can select multiple names in the **Signal** list by pressing the **Ctrl** key, clicking multiple names, and then clicking **Edit**.

The **Signals** list is populated with the names of signals entering and exiting the secured region after analysis and synthesis and a partition merge have been run successfully.

Figure 4-17. Edit Security Assignments for Signal Dialog Box



2. If necessary, lower the security level of the signal by specifying the **Security level**.
3. Select the security routing interface for signal(s) assignment. Signals that fan-out or fan-in to multiple regions can be assigned to multiple security routing interfaces.

Understanding Signal Names

The list of signals entering and exiting a secured region are signal names from the post-map netlist. Signal outputs from a secured region are derived from the output port name, as specified in the top-level RTL entity contained in the secured region. Signal inputs to a secured region are derived from the name of the output port name that feeds the secured region. In the design separation flow, output port names are preserved through the compilation process, and are used as an alias for the logic or register that feed them.

The post-map region output signals listed in the signal list coincide with the signal name in the post-fit netlist. However, combinational signal names from unsecured or unpartitioned logic that feed a secured region may change through the compilation process. Many of the RTL signals are optimized during the process of synthesis and

place-and-route. Frequently, RTL signal names may not appear in the post-fit netlist after optimization. For example, the compilation process can add tildes (~) to nets that fan out from a node, making it difficult to decipher which signal nets they actually represent. Use the post-compilation filter in the node finder to add additional signals to a security routing interface. When possible, use registered signals as inputs into a secured region, and register the output signals from a secured design partition.

Working with Global Signals

Global signals are low-skew routing lines that drive throughout the device. Global signals do not require an interface region to drive into a secured region. In Cyclone III LS devices, there are 20 global routing resources for use with high fan-out signals, such as clocks or control signals. Each global signal is accessed by a clock control block, which are located on the periphery boundary of the device. Each clock control block can be driven directly by external clock pins, PLL outputs, or a signal generated from internal logic.



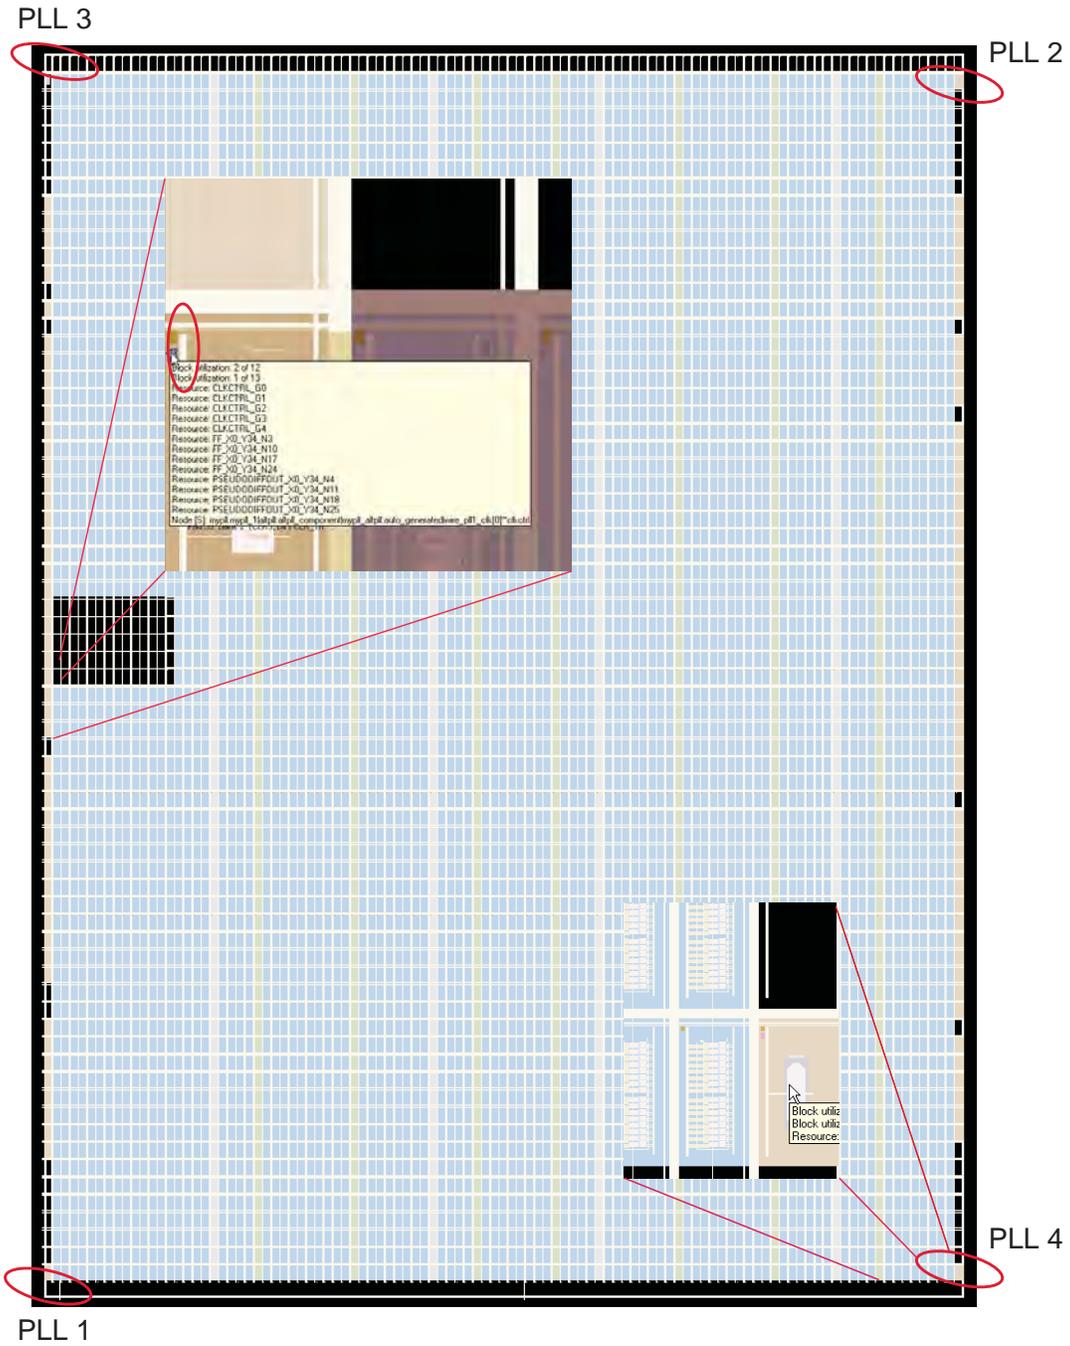
For more information about the clock networks in Cyclone III LS devices, refer to the *Clock Networks and PLLs in Cyclone III LS Devices* chapter in volume 1 of the *Cyclone III Device Handbook*.

In a compilation flow without security assignments, signals with a high fan-out (such as clock pins and control signals) are automatically promoted to use global clock resources. In the design separation flow, automatic global promotion is not turned on. Signals with high fan-out requiring global routing resources must be manually promoted to drive a clock control block.

Signals cannot be promoted onto a global routing resource through a global signal assignment from within a secured region. The Fitter only allows a clock promotion assignment to a signal if the signal is in an unsecured region. If you have a signal inside of a secured region that must use a global routing resource, you must first route the signal outside of the secured region before applying a global promotion assignment. The signal must be assigned to a security routing interface and the security level of the signal must be lowered.

For a global promotion assignment to be honored, there must be an available clock control block that is not overlapped by a secured region, and an available routing path to the clock control block. There are five clock control blocks located on each side of the device, along the horizontal and vertical axes that run through the center of the device. [Figure 4-18](#) shows the location of the clock control blocks and the PLLs for a 3CLS70 device in the Chip Planner floorplan.

Figure 4-18. PLL and Clock Control Block Location on a EPC3SL70 Device



PLLs and clock control blocks can be manually instantiated in the design partition of a secured region using the ALTPLL and ALTCLKCTRL megafunctions, respectively. Instantiation of the ALTCLKCTRL megafunction within a secured partition forces the global promotion of the signal driving the clock control block. To generate a valid placement when you instantiate PLLs or a clock control block, the secured region containing the physical resource must overlap a free PLL, a free clock control block, or both.

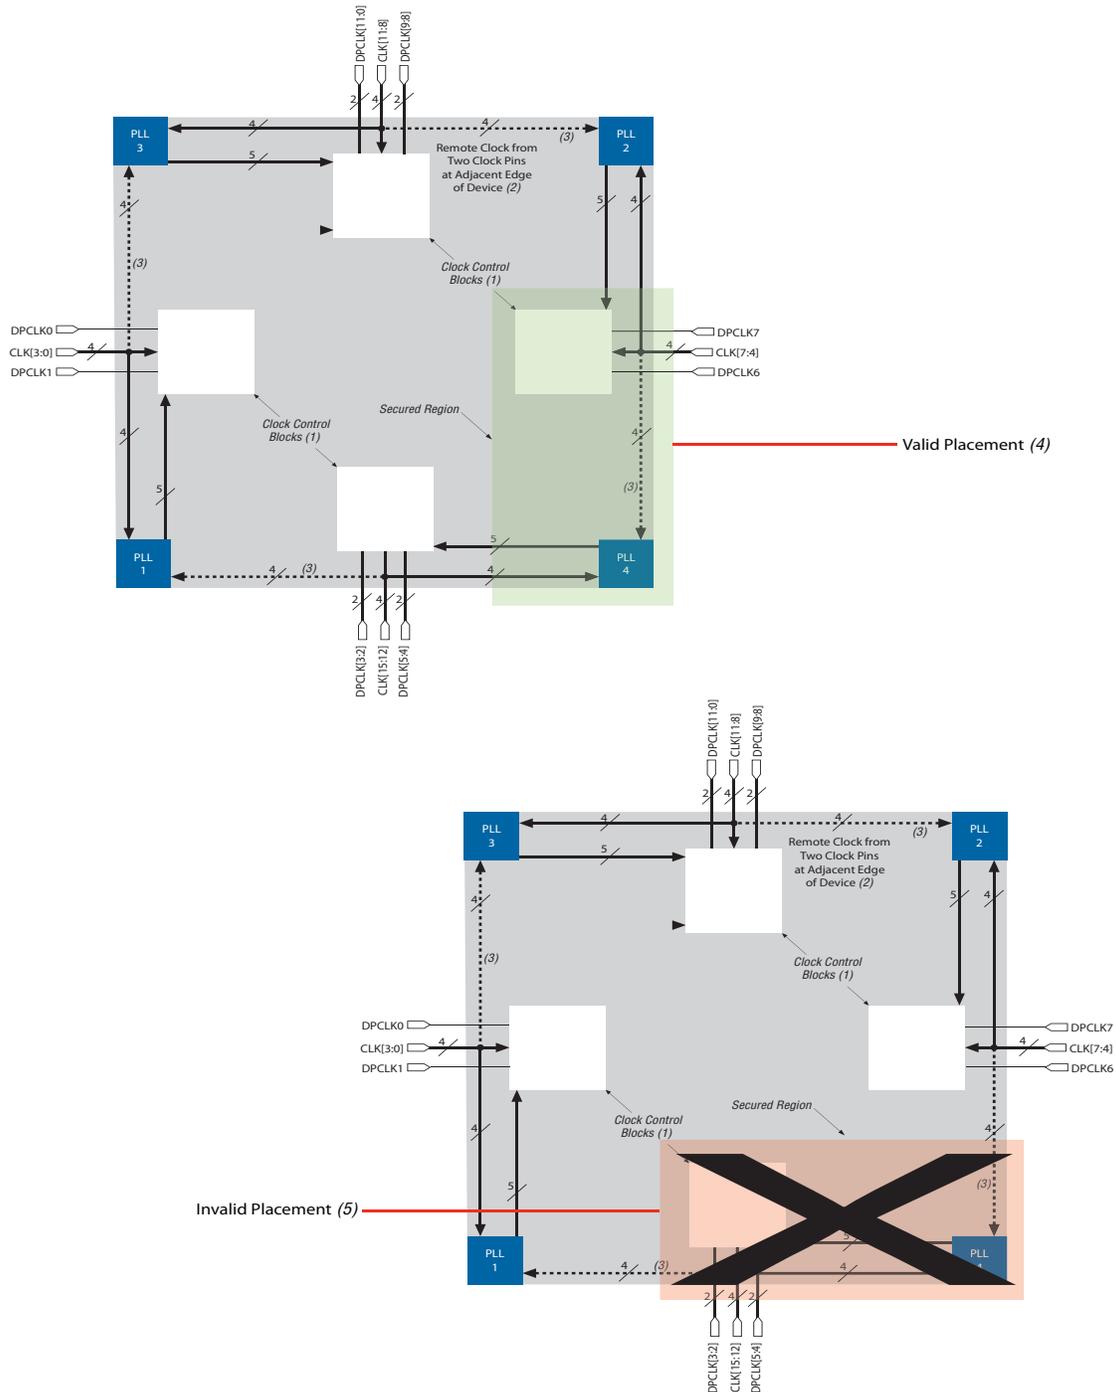
There are certain restrictions you should be aware of when you instantiate a PLL within a secured region. Secured regions with a PLL that are fed by an external clock pin must contain the PLL and a valid clock pin that can drive the PLL. Each PLL has a set of dedicated clock control blocks that it can access, located to the right (clockwise) of the PLL in the device floorplan.

Because automatic promotion of signals onto a global resource is not allowed, a PLL and the clock control block it drives must not be located in the same secured region. If your design has a PLL inside of a secured region, you must assign the PLL output to a security routing interface and then lower the security level of the PLL output.

The clock control block associated with the PLL must not be covered by any secured region. There are two sets of dedicated clock pins that can drive a PLL input. The pads for the clock input pins are co-located with the clock control blocks. If you use the clock input pin that is co-located with the clock control block associated with the PLL, the clock pin cannot be added as a member of the secured region. Instead, you must either assign the clock pin to a security routing interface that is connected with the secured region, or you can apply the `LL_IGNORE_IO_PIN_SECURITY_CONSTRAINT` assignment to relax the fitter restriction on the clock input pin.

For more information about the `LL_IGNORE_IO_PIN_SECURITY_CONSTRAINT` assignment, refer to [“Assigning I/O Pins” on page 4-25](#).

[Figure 4-19](#) shows examples of valid placement and invalid placement of secured regions that instantiate PLLs, if the `LL_IGNORE_IO_PIN_SECURITY_CONSTRAINT` assignment has not been applied.

Figure 4-19. Location of Valid and Invalid PLL, Clock Pin, and Clock Control Block Placement in a Cyclone III LS Device**Notes to Figure 4-19:**

- (1) There are five clock control blocks on each side.
- (2) Remote clocks cannot be used to feed the PLLs.
- (3) Dedicated clock paths can feed into this PLL. However, these are not fully-compensated paths.
- (4) This secured region contains a PLL that is fed by an external clock pin, whose outputs drive the clock control block through an unsecured region.
- (5) This secured region contains a PLL whose output drives an clock control block within the same secured region. This placement is invalid.

Assigning I/O Pins

After ensuring that signals that enter or exit a secured region contain a security level attribute and are explicitly assigned to a security routing interface, you must also ensure that I/O pin assignments adhere to design separation flow guidelines. Consider the following three rules, in addition to the typical pin assignment rules, when assigning I/O pins with the design separation flow enabled:

- I/O pins that are connected to a secured region must be assigned as a member of that secured region or assigned to a security routing interface region that abuts the secured region.
- Secured regions with I/O pins as members cannot share the I/O banks with any other region.
- I/O pins associated with different secured regions or security levels may not use adjacent pins.

I/O pins may be added as members of a secured region, typically when directly connected to the secured region. To add I/O pins as members of a secured region, in the **LogicLock Regions Properties** dialog box, on the **General** tab, click **Add node**. If an I/O pin is a member of a secured region, the I/O pad must be physically contained within the region, and the secured region must overlap the I/O resource.

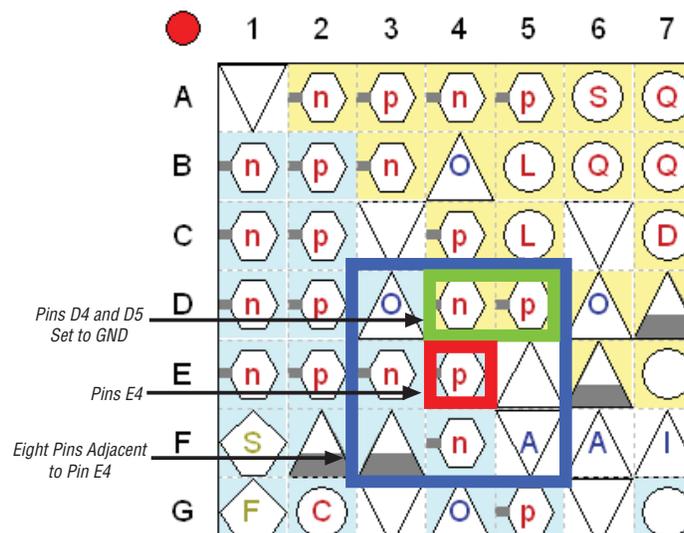
If you do not add the I/O pin as a member of the secured region, you must assign the I/O signal to a security routing interface that abuts the secured region. This security routing interface must connect the secured region to the root region or another unsecured region. Explicitly lower the security level of any output signals from the secured region that are connected to I/O pins.



I/O signals that are routed out to unsecured logic are no longer guaranteed to be physically isolated from other signals in the design.

Each I/O pin is adjacent to eight other pins: four along the horizontal and vertical axes, and four in the two diagonal axes, as shown in [Figure 4-20](#).

Figure 4-20. Pin Adjacency



Pins from different I/O banks may not share an adjacent I/O pin if one of the I/O banks contains pins that are members of a secured region. User I/O pins that are adjacent to a signal in a secured region, which belong to a different I/O bank than the secured signal, should be assigned to GND in the Quartus II software. For example, in [Figure 4-20](#), pin E4 is assigned a signal from a secured region, and I/O banks 1 and 7 belong to different LogicLock regions. Pins D4 and D5 are assigned to GND to ensure that no signal adjacencies exist between the I/O banks.

As a general rule, all unused I/O pins should be assigned to GND in the Quartus II software and assigned to a ground plane on the PCB. By default, the Quartus II software assigns unused pins to GND. You can configure this option in the **Unused Pins** page of the **Device and Pin Options** dialog box.

If you must relax a particular I/O restriction for specific signals to meet your design requirements, you may use the `LL_IGNORE_IO_PIN_SECURITY_CONSTRAINT` assignment, which is used to bypass normal I/O pin checks for a specific signal. For example, you can apply this assignment to a clock pin assigned to one of the dedicated clock inputs.



Apply the `LL_IGNORE_IO_PIN_SECURITY_CONSTRAINT` assignment in the Quartus Settings File (`.qsf`) located within the project directory of the active design. Note that there is a single `.qsf` per project revision.

To disable the I/O signal rule check for the specified pin name in the `.qsf`, add the assignment line:

```
set_instance_assignment -name LL_IGNORE_IO_PIN_SECURITY_CONSTRAINT ON -to  
<pin_name>
```



For more information about the pinouts and pin adjacencies for Cyclone III LS devices, refer to the [Cyclone III Device Pin-Out](#) tables. For more information and guidance about I/O assignments, refer to the [Cyclone III Device Family Pin Connection Guidelines](#) for Cyclone III LS devices and the [I/O Management](#) chapter in volume 2 of the *Quartus II Handbook*.

Making Post Compilation Edits

Engineering Change Orders (ECOs) and the **Rapid Recompile** feature make incremental changes to routing in a post-fit netlist. ECOs are small changes made to the functionality of a design after the design has been fully compiled. A design is fully compiled when synthesis and place-and-route are completed.

Any ECOs that do not affect routing, such as changing the LUT mask on an ALM, are supported. ECOs that affect routing or make incremental changes to the routing in a post-fit netlist are not permitted in the design separation flow.



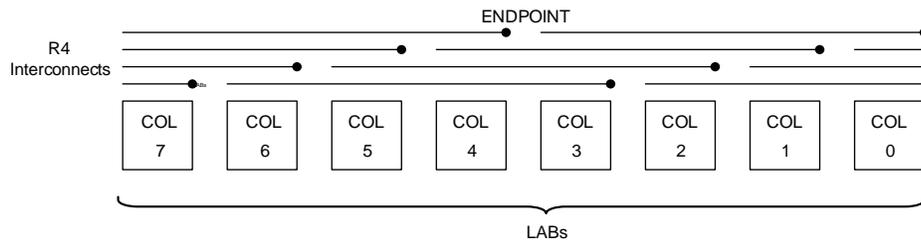
For more information about **Rapid Recompile** option in the Quartus II software, refer to [Incremental Compilation Page \(Settings Dialog Box\)](#) in Quartus II Help.

Routing Restrictions

During the overall planning of your design, you should be aware of specific design separation flow routing restrictions, especially during the floorplanning stages. These routing restrictions are discussed in this section.

Column and row interconnect routing resources on Cyclone III LS devices are staggered, with a group of routing elements starting at each LAB location. Each routing element is driven by the LAB location where the wire starts and can reach any LAB destination along the length of the routing element. Figure 4-21 shows a set of staggered R4 interconnects.

Figure 4-21. Staggered R4 Interconnects



The Fitter disables routing wires near the edge of a secured region, where routing is confined within the region. Figure 4-22 shows the Chip Planner displaying used routing elements in a design with secured regions, using options in the **Layer Settings** dialog box and using the background color map I/O banks, with only the **Global Routing** and **Used Resources** options turned on.

Figure 4-22. Chip Planner View of Used Resources

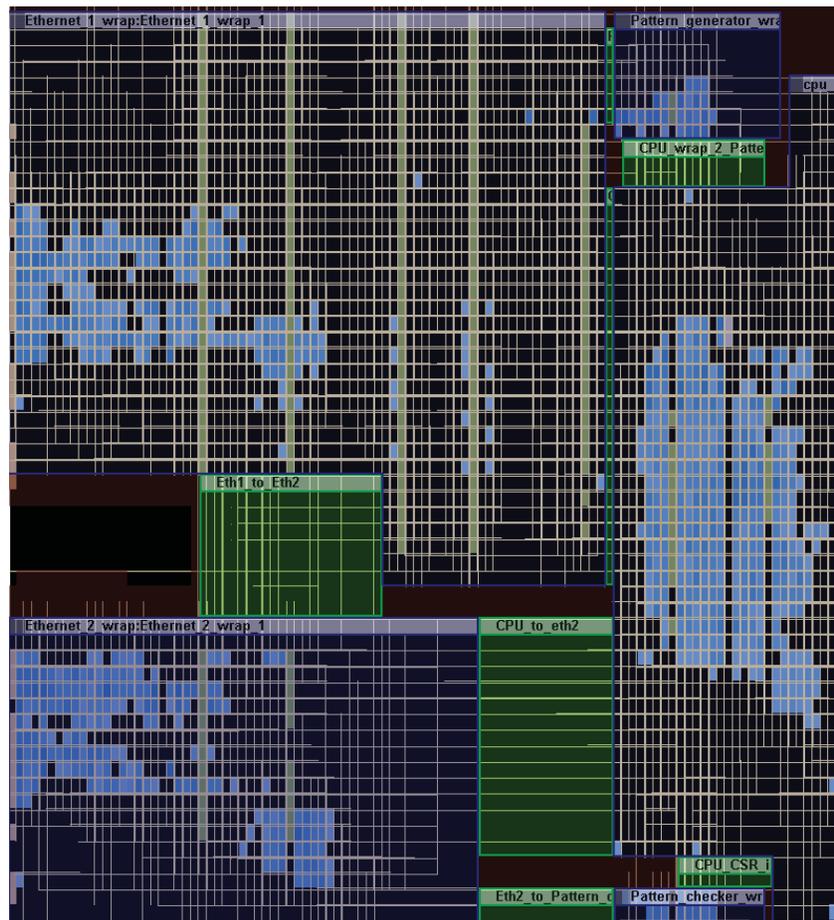


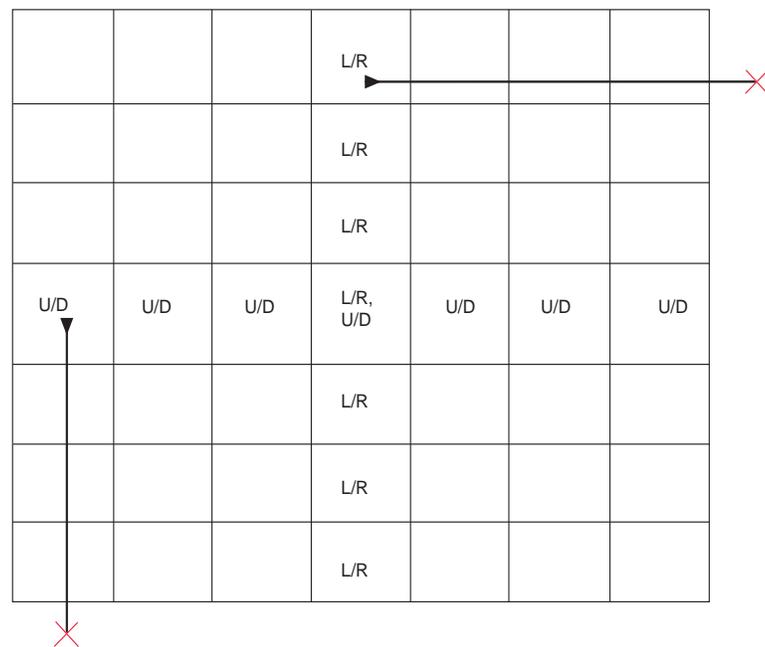
Figure 4-22 shows that no routing resources reach outside of LogicLock region boundaries, except for global routing signals and signals through interface regions.

Long wires are often unusable in secured regions because their length extends beyond the border of the region. If a secured region abuts the device boundary, you can often attain an increase in routability, because all of the routing interconnects that start inside the region and drive toward the edge of the device can be used.

I/O pads along the top and bottom of the device can only use column interconnects to drive into the device fabric. The shortest routing element from the I/O to core logic is a C4 routing wire. I/O pads on the left and right sides of the device can use both C4 and R4 routing elements to reach their LAB destinations. Because column I/Os are restricted to using C4 interconnects going into the device, a four-LAB fence is created around secured regions when the boundary of the secured region is within four-LABs of the top and bottom I/O pads.

Secured regions should be sized at a minimum of 8×8 LABs. If a region is smaller than 8×8 LABs, a connection between two LABs that violates the secured region boundary may occur. For example, in Figure 4-23, any elements along the middle axis of the 7×7 LAB array cannot use any C4 or R4 routing elements, because a C4 routing element would reach outside the secured region.

Figure 4-23. 7x7 LAB Array



Number of Signals in Routing Interfaces

In Cyclone III LS devices, every LAB location has 68 routing elements (R4) driving horizontally in each direction and 48 routing elements (C4) driving vertically in each direction. The number of connections that can be directly driven by an individual LAB is 17 connections in the horizontal direction and 12 in the vertical direction. To guarantee routability, Altera recommends that you have a routing interface height of at least one-LAB for every 17 signals routing either left or right, and a routing interface width of one-LAB for every 12 signals routing either up or down.

Figure 4-24, Table 4-2, and Table 4-3 illustrate this concept. Figure 4-24 shows three secured regions with two security routing regions; one routing signals horizontally and the other routing signals vertically. Table 4-2 and Table 4-3 list the maximum and the recommended number of signals crossing each security region.

In Figure 4-24, H_{AB} is both the smaller of the height of the region and the height of the routing interface. The minimum W_{AB} is one-LAB. W_{BC} is both the smaller width of the region and the width of the routing interface. The minimum H_{BC} is one-LAB. Changing W_{AB} or H_{BC} does not affect the values in Table 4-3.

Figure 4-24. Signals Crossing a Routing Interface

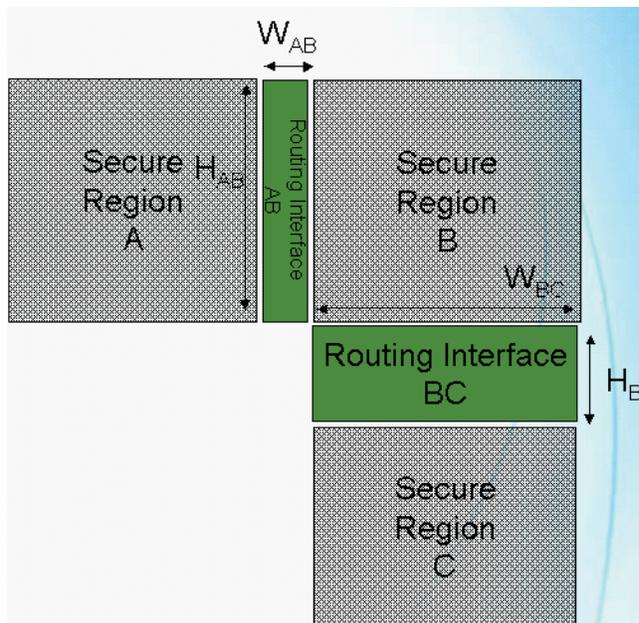


Table 4-2. Maximum Number of Signals Assigned in an Interface Region

To	From		
	A	B	C
A	—	$68 \times H_{AB}$	—
B	$68 \times H_{AB}$	—	$48 \times W_{BC}$
C	—	$48 \times W_{BC}$	—

Table 4-3. Recommended Number of Signals to Ensure Routability

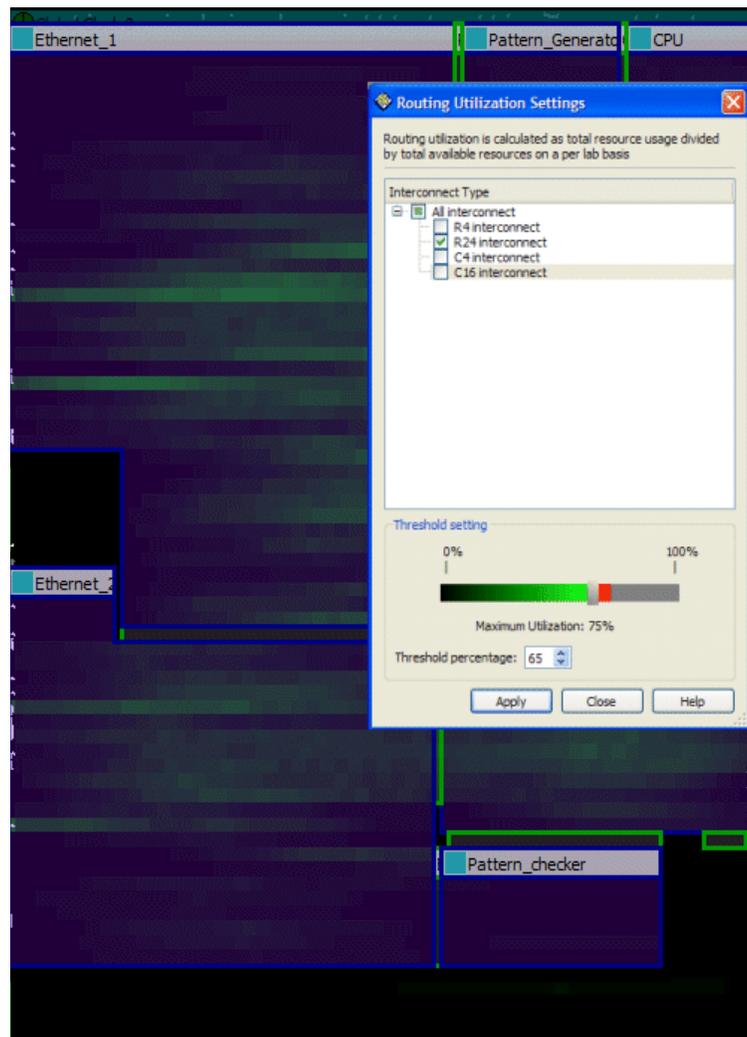
To	From		
	A	B	C
A	—	$17 \times H_{AB}$	—
B	$17 \times H_{AB}$	—	$12 \times W_{BC}$
C	—	$12 \times W_{BC}$	—

As a general guideline, keep the security routing interface channel width between the two connecting secured regions as short as possible and the depth of the channel as wide as possible. The channel width is the number of LABs that a security routing interface abuts and the depth of the channel is the number of LABs a signal passes as it goes through the routing channel.

In [Figure 4-25](#), an optimal security interface for routing AB would have a channel width equal to the height of secured region A (H_{AB}) and a channel depth of one-LAB (W_{AB}). Having a wide channel with a short depth increases the number of routing resources available between two secured regions.

You can use the **Routing Congestion** task in the Chip Planner for a visual representation of the routing utilization between secured regions. Routing resources are filtered by type. Utilization of each routing resource type is highlighted on a color gradient over the range that you specify. This tool can help you adjust region sizes and security routing interface channel widths to help you achieve an optimal floorplan. A design with the **Routing Congestion** task in the Chip Planner and R24 routing utilization is shown in [Figure 4-25](#).

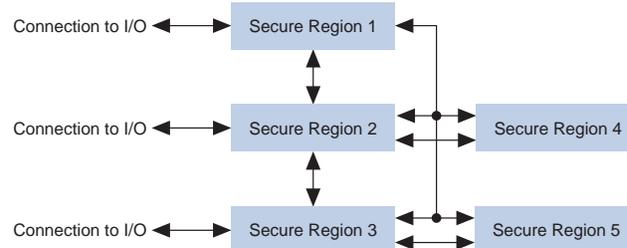
Figure 4-25. Routing Congestion



Application Example: Modifying a Fitter-Generated Floorplan for the Design Separation Flow

In this application example, the design contains five partitions that must be packed into secured regions. Figure 4-26 shows a block diagram of the design, the entities of the design, and the connectivity between the five secured partitions.

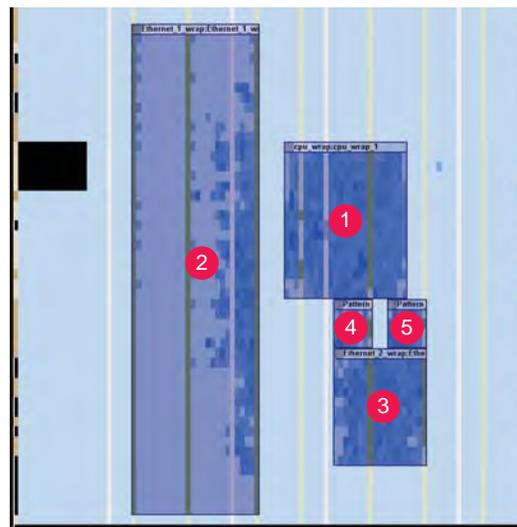
Figure 4-26. Connectivity between Five Secured Partitions



The following steps outline a recommended design flow for creating a floorplan for this design:

1. Create a LogicLock region for each partition that must be packed into a secured region.
2. Set each LogicLock region with the **Auto**, **Floating**, **On**, and **Unsecured** attributes for the **Size**, **State**, **Reserved**, and **Security Attributes** columns, respectively. Running an initial placement with these settings generates non-overlapping LogicLock regions that can be used as an initial floorplan.
3. On the Processing menu, point to **Start** and click **Start Early Timing Estimate** to run an initial place and route. The initial place and route approximates the size of each region and the general placement of the LogicLock regions relative to other LogicLock regions to achieve timing closure. The floorplan generated by the early timing estimate is shown in Figure 4-27.

Figure 4-27. Initial Fitter-Generated Floorplan



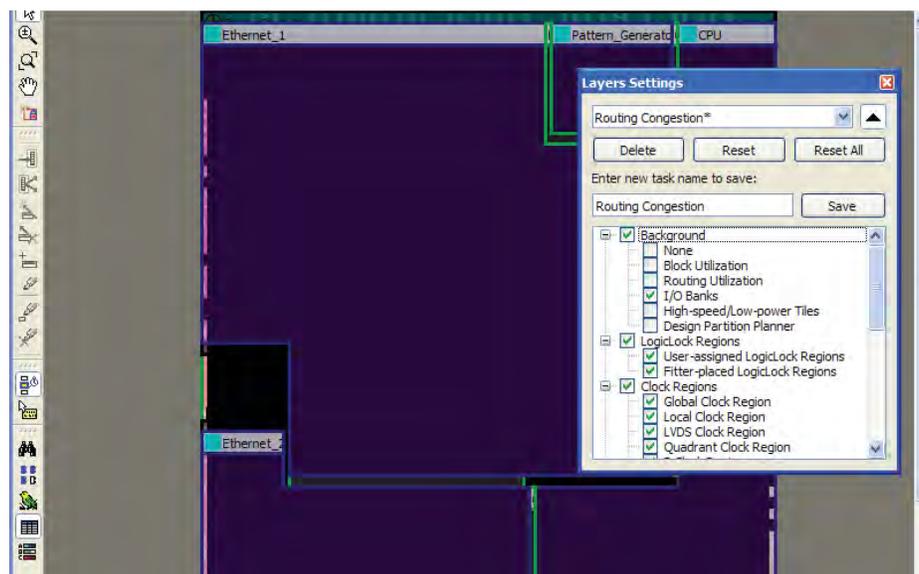
4. In the LogicLock Regions window, select the LogicLock regions, right-click, and then click **Set Size and Origin to Previous Fitter Results**.
5. Use the Design Partition Planner to view the connectivity between the different regions. You can experiment with the relative placement of the blocks by dragging and dropping each design partition. The wire bundles between design partitions help you to determine a placement that has non-overlapping routing channels.



You must also consider the connectivity to the I/O banks when arranging your floorplan. You can toggle the display of the connections between the partitions and the I/O banks within the Design Partition Planner to help you properly allocate I/O resources, as well as avoid conflicts between I/O connections and inter-partition signals. To display routing between partitions and the I/O banks, turn on **Display connections to I/O banks** in the **Bundle Configuration** dialog box.

6. Set each LogicLock region to the desired security attribute.
7. In the Chip Planner, adjust the size and placement of each LogicLock region using the relative placement you created with the Design Partition Planner. Note the following considerations when modifying the floorplan:
 - The floorplan must be complete. If there is unsecured logic that is non-contiguous due to the placement of a secured region, use an empty reserved LogicLock region to prevent a non-routable placement.
 - Each secured region must be a minimum of 8×8 LABs.
 - Each region that has I/O pins added as members of the LogicLock region should overlap the I/O bank to which it is connected. You can use the I/O bank background color map to visualize the boundaries between the I/O banks (Figure 4-28).
 - All global resources (such as clock pins and PLLs) that are required by unsecured logic must not be covered by a secured region.

Figure 4-28. I/O Banks Layers Setting for Viewing Connectivity of LogicLock Regions to I/O Banks



8. Create security routing interfaces between each of the secured regions. Assign all signals entering or exiting a region to a security routing interface.

The final floorplan result for this application example is shown in [Figure 4-29](#).

Figure 4-29. Final Floorplan



Report Panels

After the Fitter successfully places and routes a design with secured regions, the Quartus II software generates security reports. Use the security reports to review the secured regions, their associated routing interfaces, all inputs and outputs from each secured region, and the I/O bank usage for each secured region. The security reports are located in the **Fitter** section of the Compilation reports.

Secured LogicLock Region Summary

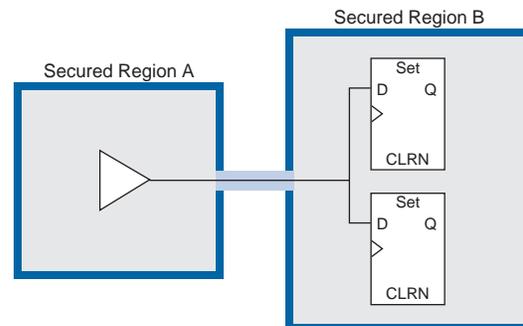
This report provides a summary of all secured regions in your design. [Table 4-4](#) describes each column in the Secured LogicLock Region Summary report.

Table 4-4. Secured LogicLock Region Report (Part 1 of 2)

Column Name	Description
Secured LogicLock Region	Lists all secured LogicLock regions in the design.
Security Attribute	Lists the security attribute (unsecured, 1, 2, or security routing interface) of the LogicLock region.
Partition Assigned	Lists the design partition assigned to the secured region.

Table 4-4. Secured LogicLock Region Report (Part 2 of 2)

Column Name	Description
Number of Input Signals (Total Fanout)	Lists the number of inputs and fan-outs into a region. The input counts the number of unique drivers that feed a secured region. The fan-out counts the total number of unique destinations being fed by the input signals into the secured region. Figure 4-30 illustrates input signals and fan-outs to a region.
Number of Output Signals (Total Fanout)	Lists the number of outputs and fan-outs from a region. The output counts the number of unique drivers sourcing a signal from the secured region. The fan-out counts the total number of unique destinations fed by the output signal.

Figure 4-30. Input Signals and Fan-Outs to a Region

Secured Region A - Number of Output Signals (Total Fanout) : 1
 Secured Region B - Number of Input Signals (Total Fanout) : 1

Security Routing Interfaces

This report summarizes the security routing interfaces. [Table 4-5](#) describes each column in the Security Routing Interfaces report.

Table 4-5. Security Routing Interface Report

Column Name	Description
Interface Name	Lists all security routing interfaces in the design.
Abutting Region A	First region that the security routing interface abuts (touches the border of the secured region).
Abutting Region B	Second region that the security routing interface abuts (touches the border of the secured region).
Number of Signals A to B (Total Fanout in B)	Lists the number of signal connections between region A and region B. The counts are shown as signals and fan-outs. Signals list the number of unique drivers from region A. Fan-out lists the number of unique destinations in region B that are fed by region A.
Number of Signals B to A (Total Fanout in A)	Lists the number of signal connections between region B and region A. The counts are shown as signals and fan-outs. Signals list the number of unique drivers from region B. Fan-out lists the number of unique destinations in region A that are fed by region B.

Secured LogicLock Region Inputs and Outputs

This set of reports provides a detailed list of every signal that enters or exits a secured region. There is one report per secured region.

Security I/O Bank Usage

This report displays the secured LogicLock region associated with each I/O bank, lists the number of pins within each region, and lists the number of pins in use. Table 4-6 describes each column in the Secured LogicLock Region Inputs and Outputs report.

Table 4-6. Secured LogicLock Region Input and Output Report

Column Name	Description
I/O Bank	Lists all available I/O banks on the device.
Associated Region	An I/O bank becomes associated with a secured LogicLock region if any portion of the I/O bank is covered by the region. If no secured region covers an I/O bank, "Unsecured Logic" is displayed, and all pins of the I/O bank are available for unsecured use.
Pin Locations Used / Pin Locations Covered by Region	Displays the ratio of pins with a signal assignment in the I/O bank to the number of possible I/O pin assignments.

Quartus Settings File Syntax

This section contains the syntax description for each Quartus Settings File (.qsf) assignment in the design separation flow.

LL_SECURITY_ROUTING_INTERFACE

This command changes a LogicLock region assignment to a security routing interface.

Type: Boolean; (ON/OFF - Defaults to OFF)

Syntax:

```
set_global_assignment -name LL_SECURITY_ROUTING_INTERFACE <value> \
-section_id <section_identifier>LL_REGION_SECURITY_LEVEL
```

LL_REGION_SECURITY_LEVEL

This command identifies the security level of a LogicLock region.

Type: Enumeration—defaults to UNSECURED

- 1
- 2
- UNSECURED

Syntax:

```
set_global_assignment -name LL_REGION_SECURITY_LEVEL <value> \
-section_id <section_identifier>
```

LL_MEMBER_OF_SECURITY_ROUTING_INTERFACE

This command assigns an I/O pin from a secured region to a security routing interface. Both <value> and <section_id> denote the name of the routing interface region. <to> specifies the name of the signal.

Type: String

Syntax:

```
set_instance_assignment -name \ LL_MEMBER_OF_SECURITY_ROUTING_INTERFACE
<value> -to <to> \
-section_id <section_id>
```

LL_SIGNAL_SECURITY_LEVEL

This command sets the security level of a signal. The default value is the security level of the region that generates the signal. This assignment may be used only to lower a security level.

Type: Enumeration

- UNSECURED
- 1
- 2

Syntax:

```
set_instance_assignment -name LL_SIGNAL_SECURITY_LEVEL <value> \
-to <to> -section_id <section_id>
```

Document Revision History

Table 4-7 shows the revision history for this chapter.

Table 4-7. Document Revision History

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Modified the former “Avoiding Child Partitions” section into the new “Avoiding Multiple Design Partitions With a Secured Region” on page 4-6 section and added information about multi-hierarchy partitions ■ Updated the “Using Secured Regions” on page 4-9 section ■ Updated the “Making Design Separation Flow Location Assignments in the Chip Planner” on page 4-10 section ■ Updated the “Creating a Complete Floorplan” on page 4-14 section ■ Updated the “Working with Global Signals” on page 4-21 and “Assigning I/O Pins” on page 4-25 sections with information about the LL_IGNORE_IO_PIN_SECURITY_CONSTRAINT assignment ■ Added the “Making Post Compilation Edits” on page 4-26 section ■ Updated the “Number of Signals in Routing Interfaces” on page 4-28 section ■ Added feature licensing information ■ Updated figures and overall editorial update ■ Template update
July 2010	10.0.0	Initial release. Content originated from <i>AN 567: Quartus II Design Separation Flow</i> .

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.



Altera's Qsys system integration tool is now available as beta for evaluation in the Quartus® II software subscription edition version 10.1. Altera does not recommend using the beta release of Qsys in the Quartus II software version 10.1 for designs that are close to completion and are meeting design requirements. Before using Qsys, review the *Quartus II Software Version 10.1 Release Notes* and *AN 632: SOPC Builder to Qsys Migration Guidelines* for known issues and limitations. To submit general feedback or technical support on the beta release of Qsys, submit a service request through mysupport.altera.com. Alternatively, to submit general feedback, click **Feedback** on the Quartus II software Help menu.

This section provides information about Qsys. Qsys is a powerful system integration tool which is included as part of the Quartus II software. Qsys automates the task of capturing of integrating customized HDL components, which may include IP cores, verification IP, and other design modules. You can use Qsys to integrate your own components with the components that Altera® or third-party developers provide. In some cases, you can implement an entire design using components from the Qsys component library.

This section includes the following chapters:

- [Chapter 5, Creating a System with Qsys](#)

This chapter describes the Qsys system integration tool.

- [Chapter 6, Creating Qsys Components](#)

This chapter describes Qsys components and how to define Qsys components using a GUI.

- [Chapter 7, Qsys Interconnect](#)

This chapter discusses the Qsys interconnect, a high-bandwidth structure for connecting components that use Avalon® interfaces.

- [Chapter 8, Component Interface Tcl Reference](#)

This chapter describes an alternative method for defining Qsys components by declaring their properties and behaviors in a Hardware Component Description File (`_hw.tcl`). It also provides a reference for the Tool Command Language (Tcl) commands that describe Qsys components.

Qsys Beta



Altera's Qsys system integration tool is now available as beta for evaluation in the Quartus® II software subscription edition version 10.1. Altera does not recommend using the beta release of Qsys in the Quartus II software version 10.1 for designs that are close to completion and are meeting design requirements. Before using Qsys, review the *Quartus II Software Version 10.1 Release Notes* and *AN 632: SOPC Builder to Qsys Migration Guidelines* for known issues and limitations. To submit general feedback or technical support on the beta release of Qsys, submit a service request through mysupport.altera.com. Alternatively, to submit general feedback, click **Feedback** on the Quartus II software Help menu.

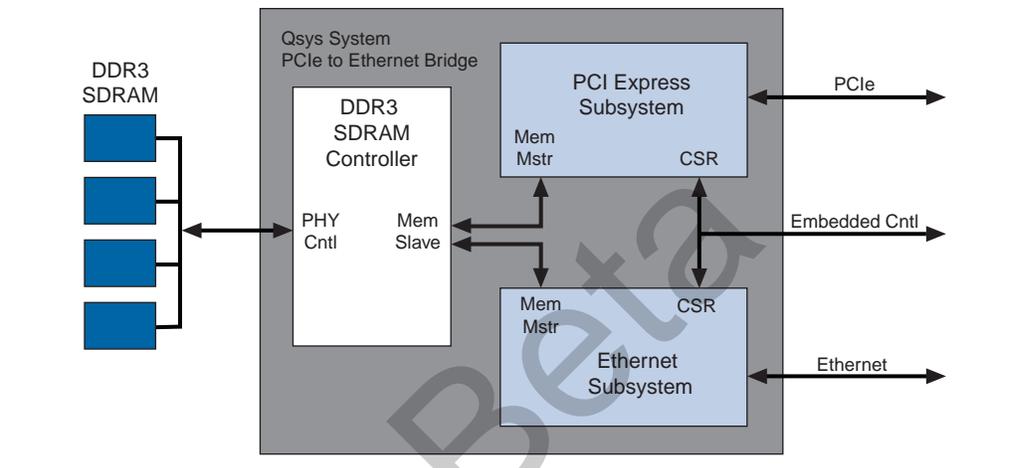
Qsys is a powerful system integration tool which is included as part of the Quartus® II software. Qsys captures system level hardware designs at a relatively high level of abstraction and also automates the task of defining and integrating customized HDL components, which may include IP cores, verification IP, and other design modules. Qsys facilitates design reuse by packaging and making available your custom components and systems. Qsys integrates your custom components with Altera® and third-party developer components. In some cases, you can implement an entire design using components from the Altera component library. During system generation, Qsys automatically creates high-performance interconnect logic from the connectivity options you specify, eliminating the error-prone and time-consuming task of writing HDL to specify the system-level connections.

Qsys provides the following advantages for hardware system design:

- Automates the process of customizing and integrating components
- Supports modular system design
- Supports visualization of large systems
- Supports optimization of interconnect fabric and pipelining within the system
- Fully integrated with the Quartus II software

Qsys supports hierarchical system design. You can include any Qsys system that exports an interface as a component in another Qsys system. [Figure 5-1](#) shows the top level of a Qsys example design that implements a PCI Express to Ethernet bridge. This example combines separate PCI Express and Ethernet subsystems with Altera's DDR3 SDRAM Controller with UniPHY IP core. Different members of the design team could develop the various subsystems simultaneously, decreasing time-to-market for the complete design. For a detailed discussion of this example design refer to ["Example Hierarchical System"](#) on page 5-12.

Figure 5-1. Top-Level Block Diagram for a PCI Express to Ethernet Bridge



Hierarchical system design in Qsys offers the following advantages:

- Enables team-based, modular design by dividing large designs into subsystems
- Enables design reuse by allowing you to use any Qsys system as a component
- Enables scalability by allowing you to instantiate multiple instances of a Qsys system

Qsys supports enhanced component parameterization, allowing you to design for maximum efficiency and utility. For example, you can specify parameters that can be fed to a program that generates the RTL at run time, allowing unlimited parameterization.

This chapter introduces Qsys, including the following sections:

- ["Qsys GUI"](#) on page 5-2
- ["Qsys Design Flow"](#) on page 5-8
- ["Example Hierarchical System"](#) on page 5-12

Qsys GUI

You can launch Qsys from the Quartus II software GUI or from a command prompt. To start Qsys from the Quartus II software, on the File menu click **New**. In the New dialog box, click **Qsys System File**. To start Qsys at the command line, type `qsys-edit` ↵.

- ① For more information about the Qsys GUI, refer to [About Qsys](#) in Quartus II Help.

Qsys Component Library

The Qsys component library contains all the components found on the component search path that you specify, whether or not they are included in a Quartus II project. These components include Altera-provided IP cores, third-party IP cores, and custom IP cores that you provide. Qsys components are listed in the component library and can be used in designs if they have exported interfaces. The component library also includes all the components that are used the Qsys interconnect.

Altera recommends that you use standard Avalon interfaces in your component designs. By using these standard interfaces, you can create components that interoperate with the components in the Qsys component library. In addition, you can take advantage of bus functional models (BFMs), monitors, and other verification IP when verifying your design. However, Qsys allows you to design with any interface that your design requires. If a set of signals cannot adhere to the *Avalon Interface Specifications*, you can encapsulate any arbitrary collection of signals as a conduit interface. You can connect conduit interfaces inside of Qsys or export them for connection outside of the immediate module.

 For more information all interface types, refer to the *Avalon Interface Specifications*. For more information about BFMS, refer to the *Avalon Verification IP Suite User Guide*.

Altera and third-party developers provide ready-to-use Qsys components. The component library has many different types of components including all of the following:

- Microprocessors, such as the Nios II processor
- DSP IP cores, such as the FIR Compiler II
- Interface protocols, such as the PCI Express Compiler IP core
- Transceiver PHYs, such as the XAUI PHY IP core
- Memory controllers, such as the RLDRAM II Controller with UniPHY
- Avalon Streaming (Avalon-ST) components, such as the Avalon-ST Multiplexer IP core
- Qsys interconnect components, such as the Qsys master router which decodes addresses

These components are installed automatically with the Quartus II software, and are available in the Qsys component library.

Integrating Custom Components

You can use the following steps to integrate you custom components into a Qsys system:

1. Determine the interfaces that interact with your custom component.
2. Create the component logic using either Verilog HDL or VHDL.
3. Use the Qsys component editor to define the `_hw.tcl` file.
4. Instantiate the component in the system.

Once you have created a Qsys component, you can use the component in other Qsys systems and share the component with other design teams.

- For instructions on developing a custom Qsys component, details about the file structure of a component, or the component editor, refer to the *Creating Qsys Components* chapter in volume 1 of the *Quartus II Handbook*.

Integrating Third-Party Components

You can also use Qsys components created by third-party IP developers. Altera awards the Qsys Certified label to IP cores that are fully supported in Qsys. These cores support Avalon interfaces and may include timing and placement constraints, software drivers, simulation models, and reference designs.

To find Qsys third-party components that you can use in Qsys systems, follow these steps:

1. On the Tools menu in Qsys, click **Download Components**.
2. On the **Intellectual Property & Reference Designs** web page, type **Qsys Certified** in the **Search** box labeled **Search IP and Reference Designs products by their descriptions**.

- You can also do advanced searches by clicking on **Launch the Altera Product Selector Guide** from the **Altera Product Selector** web page and then clicking the **IP Selector** tab.

Adding System Contents

The **System Contents** tab displays the components that you have added to your system.

Adding Components

To add a component to your system, click on the component in the **Component Library** and then click the **Add** button. A parameter editor appears allowing you to customize the component.

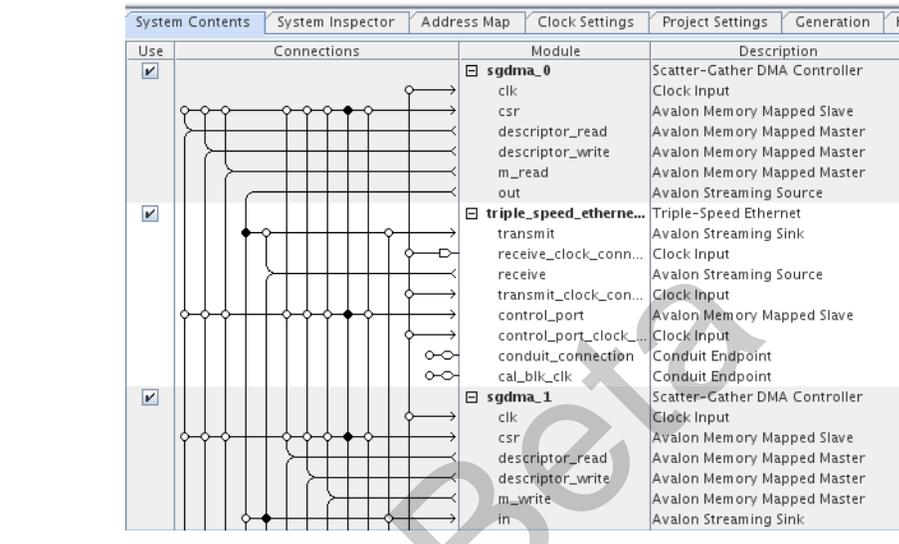
- You can type some or all of the component's name in the **Component Library** search box to display all components including that string.

Connecting Components

You specify connections at the interface level; individual signals within connected interfaces are connected automatically. You connect interfaces of compatible types and opposing directions. For example, you can connect an Avalon Memory-Mapped (Avalon-MM) master interface to an Avalon-MM slave interface, and an interrupt sender interface to an interrupt receiver interface. To see the possible connections for an interface, click the **System Contents** tab and then click the interface name. Hover

your cursor in the **Connections** column. To see the connectivity matrix where open circles represent possible connections and black circles indicate connections that you have made. To make a connection, click on the open circle at the intersection of the two interface names. Clicking a second time removes the connection. [Figure 5-2](#) illustrates the connectivity matrix.

Figure 5-2. Connections Column



? For more information, refer to [Connecting Qsys Components](#) in Quartus II Help.

Filtering Components

You can use the **Filters** dialog box to filter the display of your system in the **System Contents** tab. You can filter the display of your system by interface type, instance name, interface type, or using custom tags. For example, you can use filtering to view only instances that include an Avalon-MM interface or instances that are connected to a particular Nios II processor.

? For more information, refer to the [Filters Dialog Box](#) in Quartus II Help.

Using the System Inspector

The **System Inspector** tab displays the underlying model of your complete system. The **System Inspector** provides comprehensive details about your system such as the following information:

- The connections between all signals.
- The signal names of all signals included in exported interfaces.
- The internal connections of Qsys subsystems that are included as components.

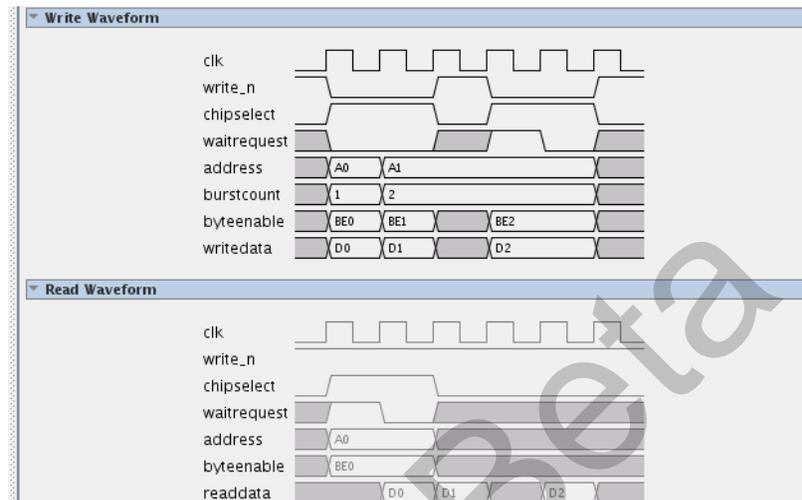


In contrast, the System Contents tab, displays only the exported interfaces of Qsys subsystems included as components.

- The global parameter settings that you specified on the **Project Settings** tab.

You can use the **System Inspector** tab to review and change component parameters and to review interface timing. For example, [Figure 5-3](#) shows the timing for the Avalon-MM DMA write master for the PCI Express-to-Ethernet system illustrated in [Figure 5-8 on page 5-14](#).

Figure 5-3. Avalon-MM Write Master Timing Waveforms Available on Project Settings Tab



- ☞ To display the timing for an interface, expand the component to display its interfaces and click on the interface name.

Defining the Address Map

The **Address Map** tab provides a table including all the Avalon-MM slaves in your design and the address range that each connected Avalon-MM master uses to address that slave. The table shows the slaves on the left and masters along the top, with the address span of the connection shown in each cell. A blank cell implies that there is no connection between that master and slave.

Follow these steps to change or create a connection between master and slave components:

1. In Qsys, click the **Address Map** tab.
2. Locate the table cell that represents the connection between the Avalon-MM master and slave component pair.
3. Either type in a base address or update the current base address in the cell.

- ☞ You can design a system where two Avalon-MM masters see Avalon-MM slave at different addresses. If you use this feature, the **Base** and **End** address columns of the **System Contents** tab are labeled *mixed* rather than providing the address range.

Specifying Clock Settings

You can use the **Clock Settings** tab to define the clocks in your system. The **Clock Settings** tab defines the **Name**, **Source**, and frequency (**MHz**) of each clock. Clicking the **Add** button adds a new clock. To change the default values, click in the appropriate column, backspace, and type the new value.

- ❓ For more information, refer to the [Adding Components to a Qsys System \(To define clock domains in a system\)](#) in Quartus II Help.

Specifying Project Settings

You can use the **Project Settings** tab to view and change the properties of your Qsys system. [Table 5-1](#) describes system-level parameters available on the **Project Settings** tab.

Table 5-1. Project Settings Parameters

Parameter Name	Description
Device Family	Specifies the Altera device family. If your final design targets a HardCopy [®] series device, specify that device here.
Clock Crossing Adapter Type	Specifies the default implementation for automatically inserted clock crossing adapters. The following choices are available: <ul style="list-style-type: none"> ■ Handshake—This adapter uses a simple hand-shaking protocol to propagate transfer control signals and responses across the clock boundary. This methodology uses fewer hardware resources because each transfer is safely propagated to the target domain before the next transfer can begin. The Handshake adapter is appropriate for systems with low throughput requirements. ■ FIFO—This adapter uses dual-clock FIFOs for synchronization. The latency of the FIFO-based adapter is a couple of clock cycles more than the handshaking clock crossing component, but the FIFO-based adapter can sustain higher throughput because it can support multiple transactions in flight at any given time. The FIFO-based clock crossers require more resources. The FIFO adapter is appropriate for memory mapped transfers requiring high throughput across clock domains. ■ Auto—if you select Auto, Qsys specifies the FIFO adapter for bursting links and the Handshake adapter for all other links.
Max Additional Latency	Specifies the maximum number of pipeline stages on each command and response path that Qsys may insert to increase the f_{MAX} at the expense of additional latency. You can specify between 0–4 pipeline stages, where 0 means that the interconnect has a combinatorial datapath. This setting is per-Qsys system or subsystem, meaning that each subsystem can have a different setting. Note that this additional latency is for both the command and response directions for the two Qsys systems even if you combine them into a single Quartus II project.
Global Reset	When this option is on, Qsys creates a global reset bus for your design. Turn this option off if you want control over reset connectivity. If you turn this option off, you must avoid reset loops and system lockup. For example, resetting an Avalon-MM slave while it is responding to a burst read request results in system lockup. The default setting for this option is on.
Use SOPC Builder Port Naming	If you turn this option on, Qsys uses the SOPC Builder naming conventions for port names and exports the interfaces that were exported in the original SOPC Builder system. This option is only for backwards compatibility for legacy designs created in SOPC Builder.

System Generation

You specify the files you want to generate on the **Generation** tab. You can generate simulation models, HDL files for Quartus II synthesis, or a Block Symbol File (**.bsf**) for schematic design. By default, Qsys places these output files in a subdirectory of your project directory. To change the default behavior, click on the **Output Path** directory, specify a new directory.

You must add the Quartus II IP File (**.qip**) file to your Quartus II project. The **.qip** file is stored in the synthesis directory. It lists the files necessary for Quartus II compilation. The **.qip** file includes references to the following information:

- HDL files used in the Qsys system
- TimeQuest Timing Analyzer Synopsys Design Constraint (**.sdc**) files
- Component definition files for archiving purposes

❓ For more information about adding files to your Quartus II project, refer to *Managing Files in a Project* in Quartus II Help.

Viewing the HDL Example

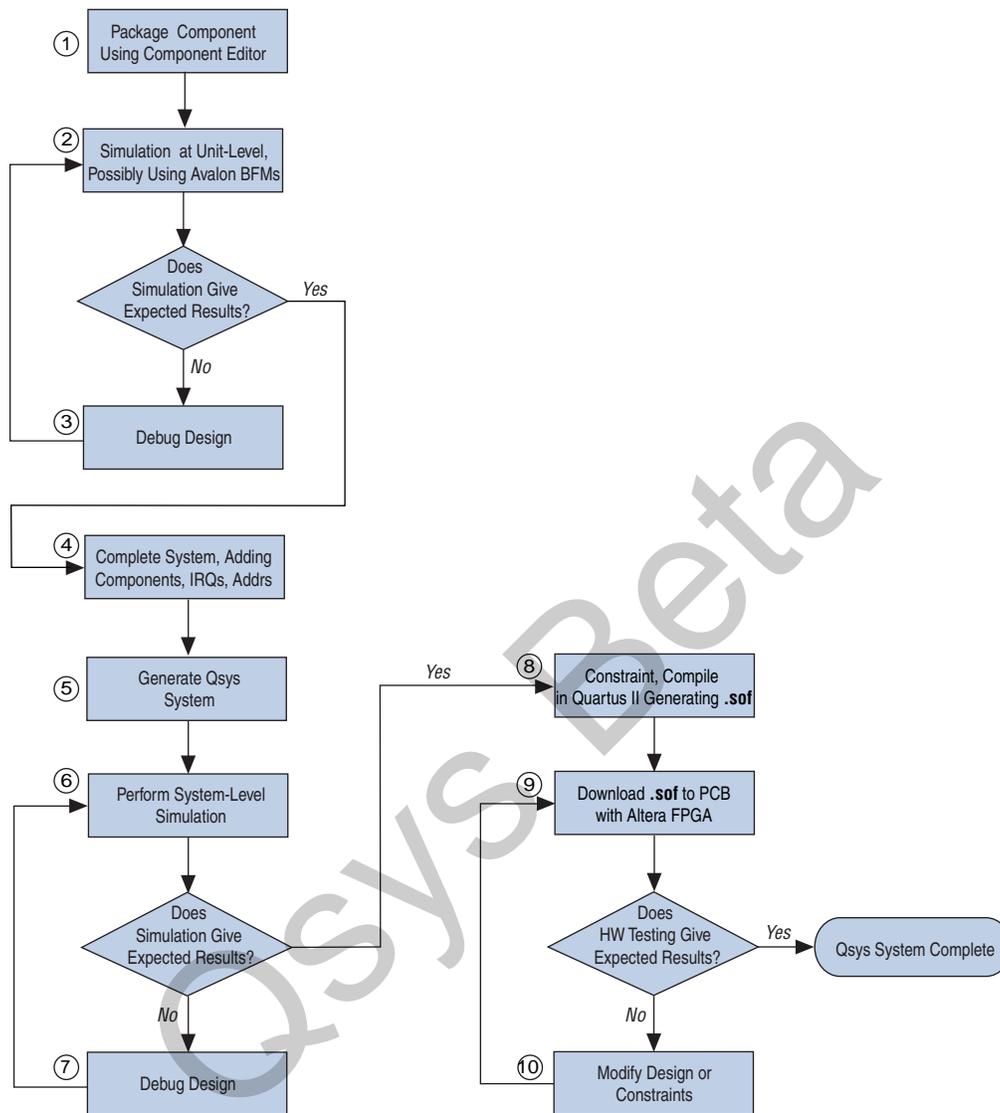
The **HDL Example** tab provides the top-level HDL definition of your Qsys system in either Verilog HDL or VHDL. This tab also displays VHDL component declarations. You can copy this example and paste it into a top-level HDL file that instantiates the Qsys system, if the system is not the top-level module in your Quartus II project.

Qsys Design Flow

Figure 5-4 illustrates an example bottom-up design flow in Qsys which starts with component design. As this flow diagram illustrates, the typical design flow includes the following high-level steps:

1. Package your component for Qsys using the Component Editor.
2. Simulate at the unit-level, possibly incorporating Avalon BFMs to verify the system.
3. Complete the Qsys design by adding other components, specifying interrupts, clocks, resets, and addresses.
4. Generate the Qsys system.
5. Perform system level simulation.
6. Constrain and compile the design.
7. Download the design to an Altera device.
8. Test in hardware.

Figure 5-4. Complete Qsys Design Flow



 In the alternative top-down valid design flow, you begin by designing the Qsys system and then define and instantiate custom Qsys component. This approach clarifies the system requirements earlier in the design process.

Designs targeting HardCopy devices require specific design constraints. Consequently, if you are targeting a HardCopy series device, you must verify you design for the HardCopy companion device.

Follow these guidelines to verify your design for both devices:

1. In the Quartus II **Device** dialog box, select both the FPGA and the appropriate HardCopy companion device.
2. In Step 8 of the design flow shown in [Figure 5-4](#), compile for both the FPGA and HardCopy device.
3. After Step 10 of the design flow shown in [Figure 5-4](#), if FPGA passes all functional simulation and hardware verification tests, generate the HardCopy handoff archive and send this archive to the HardCopy Design Center for the backend flow and tapeout.

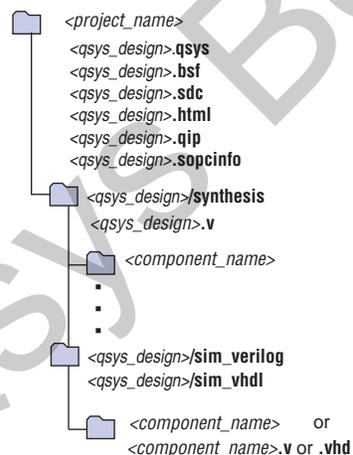
For more information about designing for HardCopy devices, refer to [About Designing HardCopy Devices](#) in Quartus II Help.

Generating Output Files

Qsys system generation creates the interconnect between components. In addition to the files for simulation and synthesis, Qsys creates a **.bsf** and an HTML datasheet.

[Figure 5-5](#) illustrates the directory structure for the output files.

Figure 5-5. Qsys Generated Files Directory Structure



[Table 5-2](#) describes the files that Qsys generates. Each time you generate your system, Qsys overwrites these files. If you have additional constraints, such as board-level timing constraints, Altera recommends that you create a separate Synopsys Design Constraints File (**.sdc**) and include that file in your Quartus II project.

Table 5-2. Qsys Generated Files (Part 1 of 2)

File Name or Directory Name	Description
<code><system_name>.qsys</code>	Qsys System File (.qsys). The .qsys file contains a list of system components, connections and their parameterizations.
<code><system_name>.bsf</code>	A Block Symbol File (.bsf) representation of the top-level Qsys system for use in Quartus II Block Diagram Files (.bdf).

Table 5-2. Qsys Generated Files (Part 2 of 2)

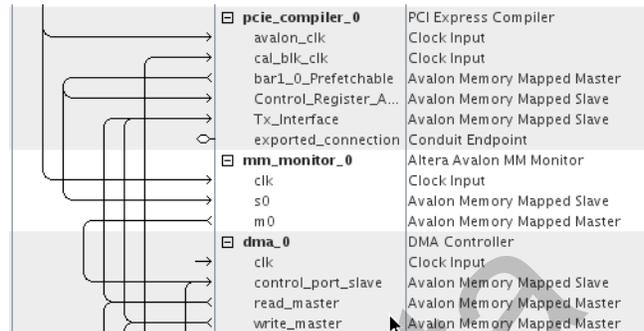
File Name or Directory Name	Description
<system_name>.html	This is a datasheet for the system which provides a system overview including the following information: <ul style="list-style-type: none"> ■ All external connections for the system ■ A memory map showing the address of each Avalon-MM slave with respect to each Avalon-MM master to which it is connected ■ All parameter assignments for each component
<system_name>.sopcinfo	Qsys information file (.sopcinfo) that describes all of the components and connections in your system. This file is a complete system description, and is used by downstream tools such as the Nios II tool chain. It also describes the parameterization of each component in the system; consequently, you can parse its contents to get requirements when developing software drivers for Qsys components.
<system_name>.sdc	Synopsys Design Constraints File (.sdc). Provides constraints for timing analysis for Timequest. This file only includes constraints for the Qsys system.
synthesis	This directory includes the files that the Quartus II software uses to synthesize your design. These files are over-written each time you generate your system.
<system_name>.qip	This file lists the Quartus II software needed to compile your design. You must add the .qip file to your Quartus II project.
<system_name>.v or <system_name>.vhd <component_name>.v or <component_name>.vhd	An HDL file for the top-level Qsys system and for each component in the system. The files under the <system_name>/ synthesis directory are used for synthesis.
sim_verilog	This directory includes the files to simulate your design.
<system_name>.v or <system_name>.vhd <component_name>.v or <component_name>.vhd	An HDL file for the top-level Qsys system and for each component in the system. The files under the <system_name>/ simulation directory are used for simulation.

Simulating a Qsys System

To simulate your Qsys system, first use Qsys to create a simulation model. Enable **Create Verilog simulation model** and click **Generate** on the **Generation** tab. This command creates Verilog simulation model files in the specified **Output Directory**, along with a ModelSim® example simulation script that compiles the system design files in the correct order and sets up the ModelSim simulation environment. You can use the generated ModelSim script **mti_setup.tcl** as an example for your testbench and simulation environment, or to create scripts for the VCS and NC-Sim simulators.

You can add monitors to Avalon-MM and Avalon-ST interfaces in your system to verify protocol correctness and test coverage with a SystemVerilog simulator. [Figure 5-6](#) demonstrates the use of monitors. It places an Avalon-MM monitor between the previously connected `pcie_compiler` `bar1_0_Prefetchable` Avalon-MM master interface and the `dma_0` `control_port_slave` Avalon-MM slave interface.

Figure 5-6. Inserting an Avalon-MM Monitor between Avalon-MM Master and Slave Interfaces



In a similar manner, you can insert an Avalon-ST monitor between Avalon-ST source and sink interfaces. You can also add BFM components that drive the external interfaces of your system components.

- For more information about using BFMs and system monitors, including tutorials demonstrating sample systems, refer to the [Avalon Verification IP Suite User Guide](#).

To simulate a Nios II processor running software code, you must generate a Nios II memory initialization file in the Nios II Software Build Tools (Nios II SBT) for Eclipse. Right-click on the application project in Eclipse, point to **Make Targets**, and choose **Build**, then select the `mem_init_install` and click **Build**. This Tcl file provides an example of a script that compiles the design files and sets up the ModelSim simulation environment.

- In the Quartus II software version 10.1, you can simulate application code for one Nios II processor in your system at a time, using one `info.meminit` file.
- For more information about the Nios II SBT for Eclipse, refer to [Getting Started with the Graphical User Interface](#) in the *Nios II Software Developer's Handbook*. For more information on the status of simulation in the beta release of Qsys in the Quartus II version 10.1, refer to the [Quartus II Release Notes](#).

Example Hierarchical System

[Figure 5-7](#) shows the details of the PCI Express example subsystem which is also illustrated at a very high level in [Figure 5-1 on page 5-2](#). In this example system, a software application running on the root complex processor programs the DMA controller. The DMA controller's Avalon-MM read and write master interfaces initiate transfers to and from the DDR3 memory and to the PCI Express Avalon-MM Tx data port. The system exports the DMA master interfaces through an Avalon-MM pipeline bridge. As the figure illustrates, all three masters connect to a single slave interface.

During system generation, Qsys automatically inserts arbitration logic to control access to this slave interface. By default, the arbiter provides equal access to all requesting masters; however, you can weight the arbitration by changing the number of arbitration shares for the requesting masters. A second Avalon-MM pipeline bridge allows the control and status interfaces to be connected internally and also exported.

For more information, refer to “Arbitration” in the *Qsys Interconnect* chapter in volume 1 of the *Quartus II Handbook*.

Figure 5-7. PCI Express Subsystem Block Diagram

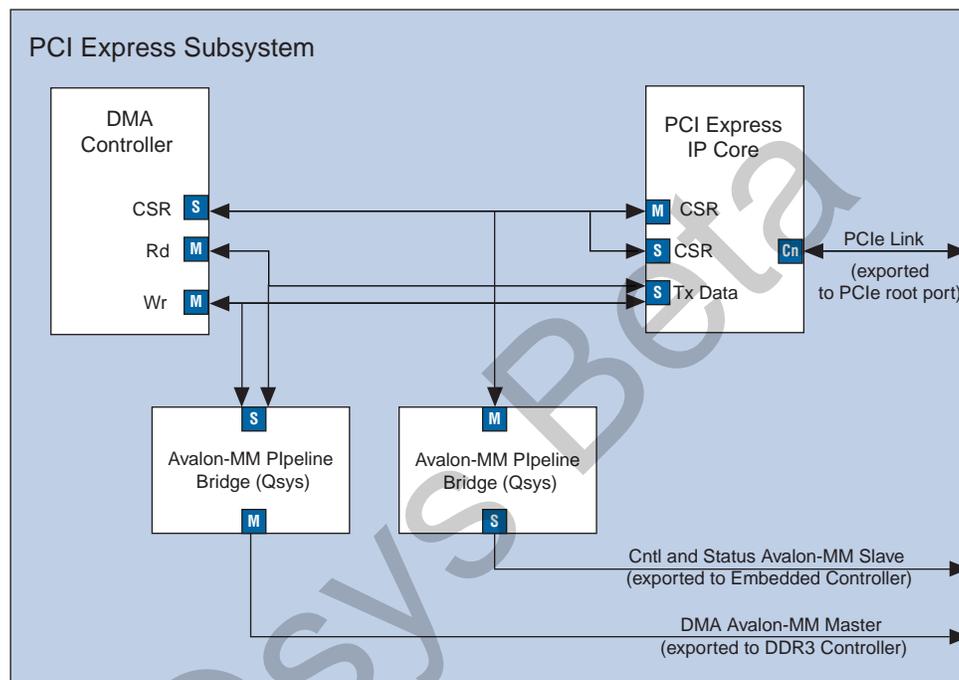


Figure 5-8 shows the Qsys representation of the PCI Express subsystem.

Figure 5-8. Qsys Representation of the PCI Express Subsystem

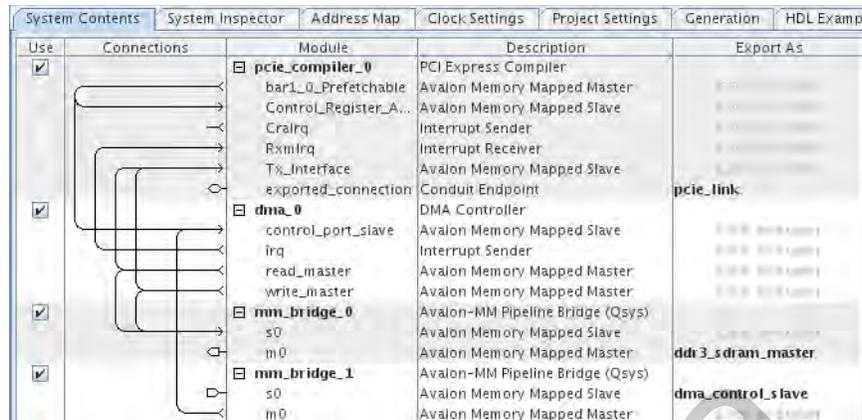


Figure 5-9 shows the details of the Ethernet example subsystem from Figure 5-1. In this subsystem, the transmit (TX) DMA receives data from the DDR3 memory and writes it to the Altera Triple-Speed Ethernet IP core using an Avalon-ST source interface. The receive (RX) DMA accepts data from the Triple-Speed Ethernet IP core on its Avalon-ST sink interface and writes it to DDR3 memory.

The read and write masters of both Scatter-Gather DMA controllers and the Triple-Speed Ethernet IP core connect to the DDR3 memory through an Avalon-MM pipeline bridge. This Ethernet example subsystem exports all three control and status interfaces through an Avalon-MM pipeline bridge which connects to a controller outside of the Qsys system.

Figure 5-9. Scatter-Gather DMA-to-Ethernet Example Subsystem

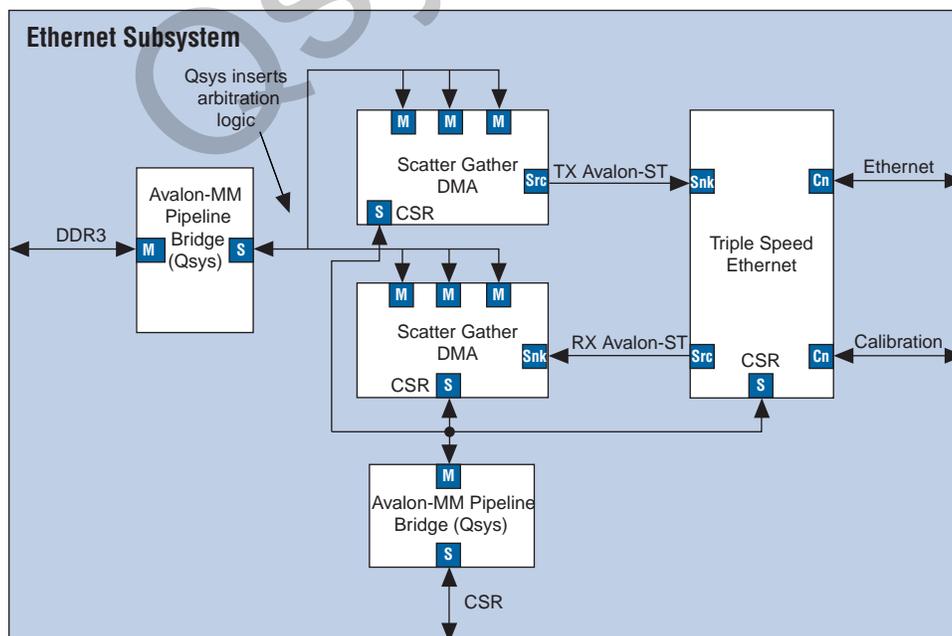
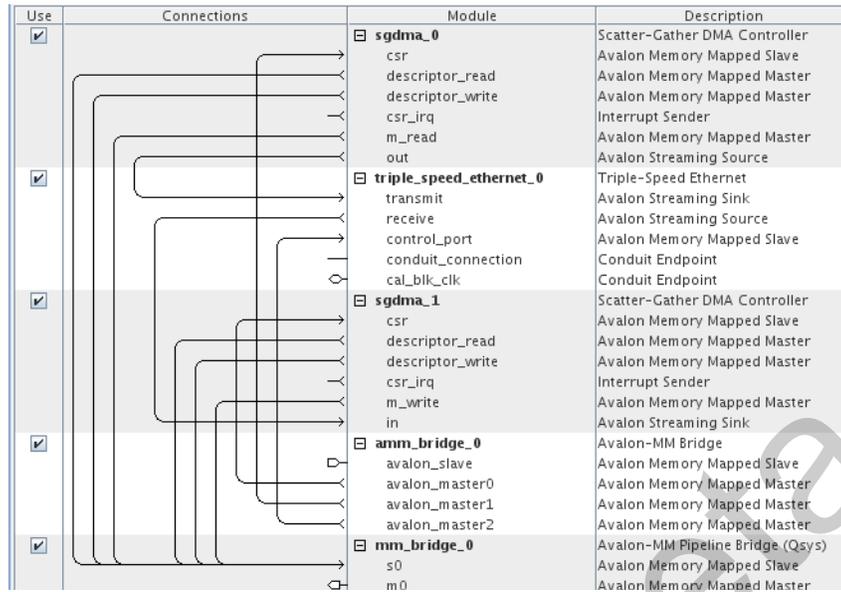


Figure 5-10 shows the Qsys representation of the Ethernet subsystem.

Figure 5-10. Qsys Representation of the Ethernet Subsystem



This example system includes two clock domains. The PCI Express and Ethernet subsystems run at 125 MHz. The DDR3 SDRAM controller runs at 200 MHz. Qsys automatically inserts clock crossing logic to synchronize the DDR3 SDRAM Controller with the with the PCI Express and Ethernet subsystems. Figure 5-11 shows the top level of the example system.

Figure 5-11. PCI Express-to-Ethernet Bridge Example System

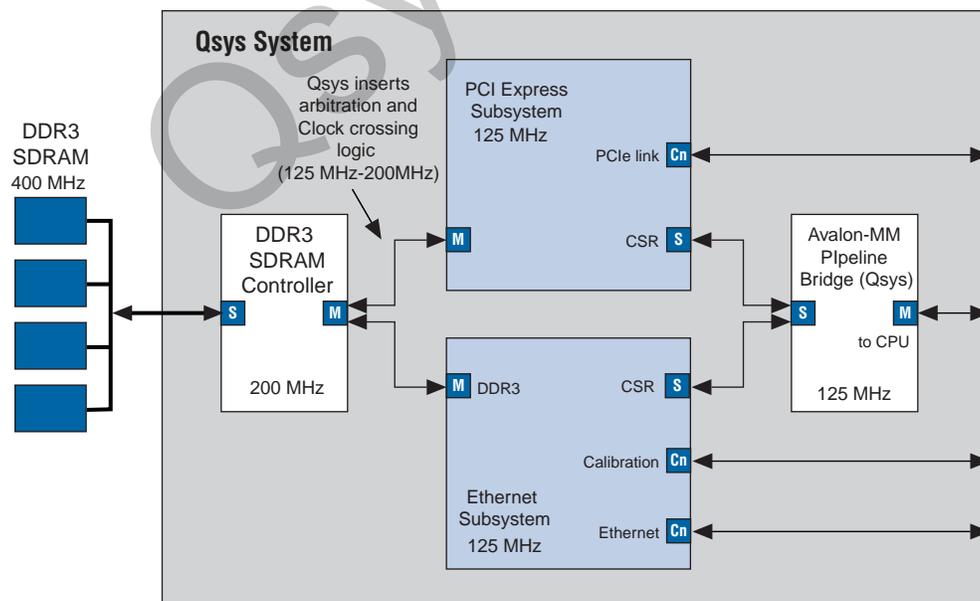


Figure 5-12 shows the Qsys representation of the complete design.

Figure 5-12. Qsys Representation of the Complete PCI Express to Ethernet Bridge

Use	Connections	Module	Description	Export As
<input checked="" type="checkbox"/>		uniphy_ddr3_0	DDR3 SDRAM Controller with UniPHY (New)	
		avalon_slave_0	Avalon Memory Mapped Slave	ddr3_sdr3ram
		memory_phy	Conduit	
		Other	Conduit	
		PLL_sharing	Conduit	
<input checked="" type="checkbox"/>		ethernet_dma_subsystem	ethernet_dma_subsystem	dma_calibration
		cal_blk_if	Conduit Endpoint	
		csr_if	Avalon Memory Mapped Slave	
		dma_if	Avalon Memory Mapped Master	
		ethernet_if	Conduit Endpoint	ethernet
<input checked="" type="checkbox"/>		pcie_subsystem	pcie_subsystem	pcie_link
		pcie_link	Conduit Endpoint	
		ddr3_sdr3ram_master	Avalon Memory Mapped Master	
		dma_control_slave	Avalon Memory Mapped Slave	
<input checked="" type="checkbox"/>		mm_bridge_0	Avalon-MM Pipeline Bridge (Qsys)	embedded_control
		s0	Avalon Memory Mapped Slave	
		m0	Avalon Memory Mapped Master	

Using Pipeline Bridges

The PCI Express to Ethernet bridge example system uses several pipeline bridges. These bridges must be configured to accommodate the address range of all of connected components, including the components in the originating subsystem and the components in the next higher level of the system hierarchy. As the name suggests, the pipeline bridge inserts a pipeline stage between the connected components. Altera recommends registering signals at the subsystem interface level for the following reasons:

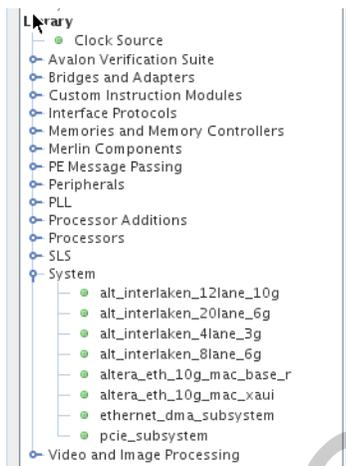
- Decreases the amount of combinational logic that must be completed in one cycle making it easier to close timing.
- Raises the potential frequency, or f_{MAX} , of your design at the expense of an additional cycle of latency which might adversely affect system throughput.
- Incremental compilation can get better f_{MAX} results if the subsystem boundary is registered.

For more information about incremental compilation, refer to *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* in volume 1 of the *Quartus II Handbook*.

Creating Hierarchical Components

Any Qsys system that exports an interface is available for use in other Qsys systems. Figure 5-13 shows the component library including the PCI Express and Ethernet subsystems as components in the component library for the PCI Express to Ethernet bridge example system. To include this Qsys component in other designs, you can add it to the component library or include the directory for this component in component search path for Qsys.

Figure 5-13. Qsys Component Library



- For more information about your IP search path, refer to the “Component Search Path” section in the [Creating Qsys Components](#) chapter in volume 1 of the *Quartus II Handbook*.
- For more information about adding components to your library, refer to the “Installing Additional Components” section in the [Creating Qsys Components](#) chapter in volume 1 of the *Quartus II Handbook*.

Document Revision History

Table 5-3 shows the revision history for this document.

Table 5-3. Document Revision History

Date	Version	Changes
December 2010	10.1.0	Initial release.

- For previous versions of the Quartus II Handbook, refer to the [Quartus II Handbook Archive](#).
- Take an [online survey](#) to provide feedback about this handbook chapter.

Qsys Beta



Altera's Qsys system integration tool is now available as beta for evaluation in the Quartus® II software subscription edition version 10.1. Altera does not recommend using the beta release of Qsys in the Quartus II software version 10.1 for designs that are close to completion and are meeting design requirements. Before using Qsys, review the *Quartus II Software Version 10.1 Release Notes* and *AN 632: SOPC Builder to Qsys Migration Guidelines* for known issues and limitations. To submit general feedback or technical support on the beta release of Qsys, submit a service request through mysupport.altera.com. Alternatively, to submit general feedback, click **Feedback** on the Quartus II software Help menu.

A Qsys *component* is a hardware design block available within Qsys that can be instantiated in a Qsys system. You can use Altera-provided or define custom Qsys components as hierarchical building blocks in creating Qsys systems. This chapter describes the structure of Qsys components, with an emphasis on the using the component editor to create and edit the Hardware Component Description Files (`_hw.tcl`) that describes a component to Qsys.

This chapter includes the following major sections:

- “Qsys Components” on page 6-1
- “Component Editor” on page 6-9

Qsys Components

A Qsys component includes the following elements:

- The HDL description of the component's hardware.
- A description of the interface to the component hardware, such as the names and types of I/O signals.
- A description of the parameters that determine the operation of the component.
- A parameter editor for customizing an instance of the component in Qsys.
- Scripts and other information Qsys requires to generate the HDL files for the component and integrate the component into the Qsys system.
- Other component-related information, such as references to software drivers, necessary for development steps downstream of Qsys.

Component Providers

Qsys components are provided by multiple sources, including the following:

- Altera provides a great variety of components automatically installed with the Quartus® II software.
- You can use the Qsys component editor to define your own custom Qsys components.
- Third-party IP developers provide Qsys-compliant components. You can display a list of third-party components by clicking **Download Components** on the Tools menu. Then, on the [Intellectual Property & Reference Designs](#) web page, type [Qsys Certified](#) in the Search box labeled **Search for IP and Reference Designs products by their descriptions**.
- Altera® development kits which are listed on the [All Development Kits](#) web page.

Component Interfaces

You can design Qsys components with any combination of the following Avalon interface types:

- Avalon Memory-Mapped (Avalon-MM)—for Avalon-MM master and slaves that communicate using read and write commands.
- Avalon Streaming (Avalon-ST)—for point-to-point connections between Avalon-ST sources and sinks that stream data.
- Tristate conduits—for a tristate conduit controller in your Qsys system to tristate devices on your PCB.
- Interrupts—for point-to-point connections between interrupt senders that generate interrupts and interrupt receivers that service interrupts.
- Clocks—for point-to-point connections between clock sources and clock sinks.
- Resets—for point-to-point connections between reset sources and reset sinks.
- Conduits—for point-to-point connections between conduit interfaces. You can use the conduit interface type to define an arbitrary collection of signals that do not fit into any of the other Avalon interface categories.

A single component can use as many of these interface types as it requires. For example, a component might provide an Avalon-ST source port for high-throughput data, in addition to an Avalon-MM slave for control. All components must include the clock interface type.

Component Types

You can build more flexibility into your components by writing a generation callback routine which generates your HDL programmatically. The following sections describe the different component types.

Static HDL Components

A static `_hw.tcl` file defines the top-level HDL file and associated component files. The HDL that describes a static component is created by the component author and is not changed by users of the component. HDL parameters are available when instantiating the component.

Generated HDL Components

Alternatively, you can also define a component whose HDL is generated based on the value of its declared parameters. These components use a custom callback to generate the HDL for each instance of the component.

For example, you could write a custom callback to include a control and status interface based on the value of a status interface parameter. The callback overcomes a limitation of HDL languages which do not allow runtime parameters.

Composed HDL Components

Composed components are constructed from combinations of other components. You can use a compose callback to connect and parameterize a composed component. Composed components can static or generated.

 For more information about defining your own generation or compose callback procedure, refer to the “Generation Callback” and “Compose Callback” sections in the *Component Interface Tcl Reference* chapter in volume 1 of the *Quartus II Handbook*.

Component Structure

Components are defined with a `_hw.tcl` file, a text file written in the Tcl scripting language that describes the component to Qsys. You can author an `_hw.tcl` file by creating a text file manually or using the component editor. This section describes the structure of `_hw.tcl` components and how they are stored.

Component Description File (`_hw.tcl`)

Component files include the following elements:

- A component description file, which is a Tcl file with file name of the form `<entity name>_hw.tcl`.
- SystemVerilog, Verilog HDL, or VHDL files that define the custom component.

The `_hw.tcl` file defines everything that Qsys requires about the name and location of component design files, including files for simulation and constraint files.

The component editor simplifies the process of creating components and automatically saves components in the `_hw.tcl` format. You can use these Tcl files as a template for editing components by hand. When you edit a previously saved `_hw.tcl` file, Qsys automatically backs up the earlier version as `_hw.tcl~`.

For more information about `_hw.tcl` file details, refer to the *Component Interface Tcl Reference* chapter in volume 1 of the *Quartus II Handbook*.

Component File Organization

A typical component uses the following directory structure. The precise names of the directories are not significant.

- `<component_directory>/`
 - `<hdl>/`—a directory that contains the component HDL design files and the `_hw.tcl` file.
 - `<component_name>_hw.tcl`—the component description file.
 - `<component_name>.v` or `.vhd`—the HDL file that contains the top-level module.
 - `<component_name>_sw.tcl`—the software driver configuration file. This file specifies the paths for the `.c` and `.h` files associated with the component.
 - `<software>/`—a directory that contains software drivers or libraries related to the component, if any. Altera recommends that the software directory be a subdirectory of the directory that contains the `_hw.tcl` file.

 For information on writing a device driver or software package suitable for use with the Nios II processor, refer to the *Hardware Abstraction Layer* section of the *Nios II Software Developer's Handbook*. The *Nios II Software Build Tool Reference* chapter of the *Nios II Software Developer's Handbook* describes the commands you can use in the Tcl script.

Component Versioning

You can create and maintain multiple versions of the same component using one of the following options:

- Define the module property `version` in your `_hw.tcl` file.
- If multiple versions of the component are defined in your component libraries, you can add a different version of a component by right-clicking on the component and selecting **Add version** `<version_number>`.
- You can create an `.ipx` file in the same directory as your Qsys project to control the search path for your project.

Component Search Path

Qsys searches for component files each time you open the tool. Qsys locates and displays the list of available components in the **Component Library**. Qsys searches the directories in the IP search path for the following component file types:

- Hardware Component Description Files (`_hw.tcl`) files. Each `_hw.tcl` file defines a single component.
- IP Index (`.ipx`) files. Each file indexes a collection of available components, or a reference to other directories to search. In general, `.ipx` files facilitate faster startup for Qsys and other tools because fewer files need to be read and analyzed.

Qsys searches the directories recursively, while others only to a specific depth. In the following list of search locations, a recursive descent is annotated by `**`. The `*` signifies any file. When a directory is recursively searched, the search stops at any directory containing a `_hw.tcl` or `.ipx` file; subdirectories are not searched.

- \$\$PROJECT_DIR/*
- \$\$PROJECT_DIR/ip/**/*
- \$QUARTUS_INSTALLDIR/./ip/**/*

Complete the following steps to extend the default search path by specifying additional directories:

1. On the Tools menu click **Options**.
2. In the **Category** list, click **IP Search Path**.
3. Click **Add**.
4. Browse to locate additional directories and click **Open** to add them to your search path.

 These additional paths apply to all projects; that is, the paths are global to the current version of Qsys. The search path is ultimately defined by the file, `<$QUARTUS_INSTALLDIR>/sopc_builder/bin/root_components.ipx`.

Adding Components to the Library

Use one of the following methods to add components to the Component Library:

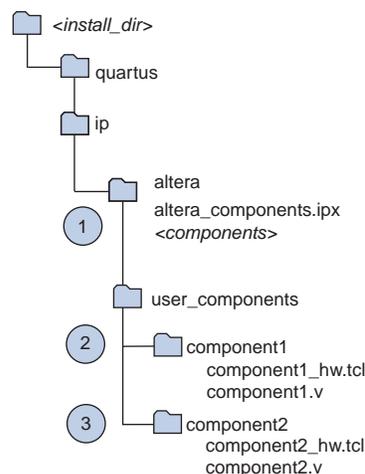
Copy to the IP Root Directory

The simplest method to add a new component is to copy your components into the standard IP directory provided by Altera. This approach is useful in the following situations:

- You want to associate your components with a specific release of the Quartus II software
- You want to and have the same components available across multiple projects

Figure 6-1 illustrates this approach.

Figure 6-1. User Library Included In Subdirectory of \$IP_ROOTDIR



In [Figure 6-1](#), the circled numbers identify three steps of the algorithm that Qsys follows during initialization. These steps are explained in the following paragraphs.

1. Qsys recursively searches the `<install_dir>/ip/` directory by default. It finds the file in the `altera` subdirectory, which tells it about all of the Altera components. **altera_library.ipx** includes listings for all components found in its subdirectories. The recursive search stops when Qsys finds this **.ipx** file.
2. As part of its recursive search, Qsys also looks in the adjacent **user_components** directory. Qsys finds the **component1** directory, which contains **component1_hw.tcl**. When Qsys finds that component, the recursive search stops so that no components in subdirectories of **component1** are found.
3. Qsys then searches in the adjacent **component2** directory, which includes **component2_hw.tcl**. If Qsys finds that component, the recursive search stops.

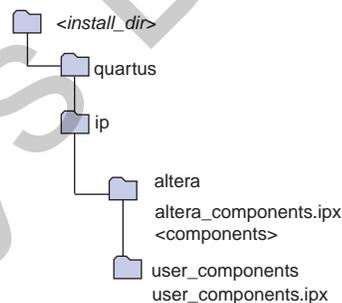


If you save your `_hw.tcl` file in the `<install_dir>/ip/` directory, Qsys finds your `_hw.tcl` file and stops. Qsys does not conduct the search just described.

Reference Components in an .ipx File

Alternatively, you can specify the search path in a **user_components.ipx** file under `<install_dir>/ip` path. This method allows you to store components in a location that is not linked to your Quartus II installation and to add a location that is independent of the default search path. [Figure 6-2](#) illustrates this approach.

Figure 6-2. Specifying A User .ipx directory



The **user_components.ipx** file includes a single line of code redirecting Qsys to the location of the user library. [Example 6-1](#) shows the code for this redirection.

Example 6-1. Redirect to User Library

```

<library>
  <path path="<user_lib_dir>/user_ip/**/*" />

/<library>
  
```



For both of these approaches, if you install a new version of the Quartus II software, you must also repeat the steps to include your components.

You can verify that components are available and also decrease the time it takes to launch Qsys by using the utilities, `ip-catalog` and `ip-make-ipx` commands. The following sections describe these commands.

ip-catalog

This commands displays the a catalog of available components in either plain text or XML format.

Usage

```
ip-catalog --project-dir[=<directory>] --name[=<value>]  
--verbose[=<true/false>] --xml[=<true/false>] --help
```

Options

- `--project-dir[=<directory>]`. Optional. Components can be found in certain locations relative to the project, if any. By default, the current directory, `.'` is used. To exclude any project directory, use `..`.
- `--name[=<value>]`. Optional. This argument provides a pattern to filter the names of the components found. To show all components, use a `*` or `'`. By default, all components are shown. The argument is not case sensitive.
- `--verbose[=<true/false>]`. Optional. When true, reports the progress of the command.
- `--xml[=<true/false>]`. Optional. When true, prints the output in XML format instead of a line- and colon-delimited format.
- `--help`. Shows help for the `ip-catalog` command.

ip-make-ipx

This command creates an index file for the directory specified. It returns a 0 for successful completion and a non-zero value for failure.

Usage

```
ip-make-ipx --source-directory[=<directory>] --output[=<file>]  
--relative-vars[=<value>] --thorough-descent  
--message-before[=<value>] --message-after[=<value>] --help
```

Options

- `--source-directory=<directory>`. Optional. The directory to index. The default directory is `."`. You can also provide a comma separated list of directories.
- `--output[=<file>]`. Optional. The name of the file to generate. The default name is `./components.ipx`.
- `--relative-vars[=<value>]`. Optional. Causes the output file to include references relative to the specified variable or variables where possible. You can specify multiple variables as a comma-separated list.
- `--thorough-descent[=<true/false>]`. Optional. If set, a component or `.ipx` file in a directory does not prevent subdirectories from being searched.
- `--message-before[=<value>]`. Optional. A message to print to stdout when indexing begins
- `--message-after[=<value>]`. Optional. A message to print to stdout when indexing completes
- `--help`. Show help for this command

Understanding IPX File Syntax

An `.ipx` file is an XML file whose top-level element is `<library>` with a `<path>` subelements are `<path>` and `<component>`.

A `<path>` element contains a single attribute, also called `path` and may reference a directory with a wildcard, `(*)`, or reference a single file. Two asterisks designate any number of subdirectories. A single asterisk designates a match to a single file or directory. In searching down the designated path, the following three types of files are identified:

- `.ipx`—additional index files
- `_hw.tcl`—Qsys component definitions
- `_sw.tcl`—Nios II board support package (BSP) software component definitions

A `<component>` element contains several attributes to define a component. If you provide all the required details for each component in an `.ipx` file, the start-up time for Qsys is less than if Qsys must discover the files in a directory. [Example 6-2](#) shows two `<component>` elements. Note that the paths for file names are specified relative to the `.ipx` file.

Example 6-2. Component Elements

```
<library>
  <component
    name="A Qsys Component"
    displayName="Qsys FIR Filter Component"
    version="2.1"
    file="./components/qsys_filters/fir_hw.tcl"
  />
  <component
    name="rgb2cmyk_component"
    displayName="RGB2CMYK Converter(Color Conversion Category!)"
    version="0.9"
    file="./components/qsys_converters/color/rgb2cmyk_hw.tcl"
  />
</library>
```

Component Editor

The Qsys component editor is a GUI that allows you to define a component and its parameter editor GUI. You use the component editor to do the following:

- Specify the SystemVerilog, Verilog HDL, or VHDL files that describe the modules in your component, and simulation and constraint files
- Conversely, create an HDL template for a component by first defining its interface using the **HDL Files** tab of the component editor.
- Specify the signals for each of the component's interfaces, and define the behavior of each interface signal.
- Specify relationships between interfaces, such as determining which clock interface is used by a slave interface.
- Declare any parameters that alter the component structure or functionality, and define a user interface to let users parameterize instances of the component.

After you define your component in the component editor the component is available in the component library. The following sections explain how to use the component editor.

Component Hardware Structure

The component editor allows you to define components with one or more interfaces. For a description of the available interface types refer to ["Component Interfaces" on page 6-2](#). You can specify exported interfaces which appear at the top-level of the Qsys system. You can connect exported interfaces to devices on the PCB or to other Qsys subsystems in hierarchical designs.

You can also use the component editor to generate an early version of the `_hw.tcl` file and then manually edit this file to complete the component definition.

Starting the Component Editor

To start the component editor in Qsys, on the **File** menu, click **New Component**. When the component editor starts, the **Introduction** tab describes how to use the component editor.

 Each tab in the component editor provides on-screen information that describes how to use the tab. Click the triangle labeled **About** at the top-left of each tab to view these instructions. You can also refer to *Component Editor (Qsys)* in Quartus II Help for additional information about the component editor.

HDL Files Tab

The **HDL Files** tab allows you to create an Qsys component from existing Verilog HDL or VHDL files, or to create an HDL template in either Verilog HDL or VHDL for a Qsys component by first specifying its interfaces. The following sections describe both the bottom-up and top-down approaches to component design.

Bottom-Up Component Design

You can use the **HDL Files** tab to specify Verilog HDL or VHDL files that describe the component logic. Files are provided to downstream tools such as the Quartus II software and ModelSim® in the same order as they appear in the **HDL Files** table.

You can also use the component editor to define the interface to components outside the Qsys system. In this case, you do not provide HDL files. Instead, you use the component editor to interactively define the hardware interface.

After you specify an HDL file, the Quartus II Analysis and Elaboration analyzes signals and parameters declared for all modules in the top-level file. After successful analysis, the component editor **Signals** tab lists all design modules in the **Top Level Module** list. If your HDL contains more than one module, you must select the appropriate top-level module from the **Top Level Module** list.

All files are managed in a single table, with options for **Synth** and **Sim**. You can select the **Top** option to select the top-level file for synthesis. When the top-level module is changed, the component editor performs best-effort signal matching against the existing port definitions. If a port is absent from the module, it is removed from the port list. You can use the up and down arrows to specify the HDL file analysis order.

By default, all files are added with both **Synth** and **Sim** options turned on. To add a simulation-only file, turn off the **Synth** option for that file. Simulation files are passed to ModelSim for simulation. To add a synthesis-only file, turn off the **Sim** file option.



The component editor determines the signals on the component when only the top-level module or entity is added to the table, but all of the files required for the component must be added for the component to compile in Quartus II software or work in simulation.

Top-Down Component Design

The **Create HDL Template** button on the **HDL Files** tab allows you to create an HDL template for a component if you have not provided a HDL description for it. Clicking the **Create HDL Template** button shows you the component HDL and lets you choose between Verilog HDL and VHDL. Altera recommends that you define your signals, interfaces, parameters and basic component information, including the component name, before creating the HDL template by clicking **Save**. The component editor writes `<component_name>.v` or `<component_name>.vhd` to your project directory.

After you have created the component's HDL code, you can add other files that are required to define your component, including the `_hw.tcl` file, and synthesis and simulation files using the **Add** button on the **HDL Files** tab.

Signals Tab

You use the **Signals** tab to specify the purpose of each signal on the top-level component module. If you specified a file on the **HDL Files** tab, the signals on the top-level module appear on the **Signals** tab.

The **Interface** list also allows creation of a new interface so that you can assign a signal to a different interface without first switching to the **Interfaces** tab. Each signal must belong to an interface and be assigned a legal signal type for that interface. In addition to Avalon-MM and Avalon-ST interfaces, components may have clock, interrupt, reset, tristate conduit and conduit interfaces.

Naming Signals for Automatic Type and Interface Recognition

The component editor recognizes signal types and interfaces based on the names of signals in the source HDL file, if they conform to the following naming conventions:

Signal associated with a specific interface—`<interface type>_<interface name>_<signal type>[_n]`

For any value of `<interface_name>` the component editor automatically creates an interface by that name, if necessary, and assigns the signal to it. The `<signal_type>` must match one of the valid signal types for the type of interface. Refer to the [Avalon Interface Specifications](#) for the signal types available for each interface type. You can append `_n` to indicate an active-low signal. [Table 6-1](#) lists the valid values for `<interface_type>`.

Table 6-1. Valid Values for <Interface Type> (Part 1 of 2)

Value	Meaning
avs	Avalon-MM slave
avm	Avalon-MM master
aso	Avalon-ST source
asi	Avalon-ST sink
cso	Clock output
csi	Clock input
coe	Conduit
inr	Interrupt receiver
ins	Interrupt sender

Table 6-1. Valid Values for <Interface Type> (Part 2 of 2)

Value	Meaning
ncm	Nios II custom instruction master
ncs	Nios II custom instruction slave
rsi	Reset sink
rso	Reset source
tcm	Tristate conduit master
tcs	Tristate conduit slave

Example 6-3 shows a Verilog HDL module declaration with signal names that infer two Avalon-MM slaves.

Example 6-3. Verilog HDL Module With Automatically Recognized Signal Names

```

module my_slave_irq_component (

    csi_clockreset_clk; // clock interface
    csi_clockreset_reset_n; //reset clock interface

    avs_sl_address; //s1 slave interface
    avs_sl_read; //s1 slave interface
    avs_sl_write; //s1 slave interface
    avs_sl_writedata; //s1 slave interface
    avs_sl_readdata; //s1 slave interface
    ins_irq0_irq; //irq0 interrupt sender interface
);

input csi_clockreset_clk;
input csi_clockreset_reset_n;
input [7:0] avs_sl_address;
input avs_sl_read;
input avs_sl_write;
input [31:0] avs_sl_writedata;
output wire[31:0] avs_sl_readdata;
output wire ins_irq0_irq;

/* Insert your logic here */

endmodule

```

Templates for Interfaces to External Logic

You can use the **Create HDL Template** command to generate an HDL template for the component. Then, you connect these signals outside of the Qsys system. If your component uses an Avalon interface to interface outside of the Qsys system, you can use the Templates menu in the component editor to add typical interface signals to your signal list. There are templates for the following interfaces:

- Avalon-MM Slave
- Avalon-MM Slave with Interrupt
- Avalon-MM Master
- Avalon-MM Master with Interrupt

- Avalon-ST Source
- Avalon-ST Sink
- Nios Custom Instruction Slave

After adding a typical Avalon interface using a template, you can add or delete signals to customize the interface.

Interfaces Tab

The **Interfaces** tab allows you to configure the interfaces on your component and specify a name for each interface. The interface name identifies the interface and appears in the Qsys connection panel. The interface name is also used to uniquely identify any signals that are ports on the top-level Qsys system.

The **Interfaces** tab allows you to configure the type and properties of each interface. For example, an Avalon-MM slave interface has timing parameters that you must set appropriately. The **Interfaces** tab displays waveforms that illustrate the timing that you specify. If you update the timing parameters, the waveforms automatically update to illustrate the new timing. The waveforms are available for the following interface types:

- Avalon-MM
- Avalon-ST
- Interrupts

HDL Parameters Tab

You specify the parameters that users of your component can set to configure your component on the **HDL Parameters** tab. The **Parameters** table included on this tab displays Verilog HDL parameters or VHDL generics that you declared in the top-level HDL module. Using the **Parameters** table, you can specify the following information about each parameter:

- Default value
- Whether or not it is user-editable
- Type
- Group
- Tool tip

Click **Preview the GUI** at any time to see how the component GUI appears.

The following rules apply to HDL parameters exposed via the component parameter editor:

- Editable parameters cannot contain computed expressions.
- If a parameter $\langle n \rangle$ defines the width of a signal, the signal width must be of the form $\langle n-1 \rangle:0$.

- When a VHDL component is used in a Verilog HDL Qsys system, or a Verilog HDL component is used in a VHDL Qsys system, numeric parameters must be 32-bit decimal integers. When passing other numeric parameter types, unpredictable results occur. (The interconnect fabric is written in Verilog HDL and SystemVerilog.)

 Refer to *Component Interface Tcl Reference* chapter in volume 1 of the *Quartus II Handbook* for detailed information about creating and displaying parameters using Tcl scripts.

Library Info

The **Library Info** tab allows you to specify the following information about your component:

- **Name**—Specifies the component name. When you save your component, the component editor saves your component to the string that you specified concatenated to the `_hw.tcl` suffix, for example, `my_component_hw.tcl`
- **Display Name**—Specifies the user-visible name for this component in Qsys.
- **Version**—Specifies the version number of the component.
- **Group**—Specifies which group in Qsys displays your component in the list of available components. If you enter a previously unused group name, Qsys creates a new group by that name.
- **Description**—Allows you to describe the component.
- **Created By**—Allows you to specify the author of the component.
- **Icon**—Allows you to place an image in the title bar of your component, in place of the MegaCore logo. The icon can be a `.jpg`, `.gif`, or `.png` file. The directory for the icon is relative to the directory that contains the `_hw.tcl` file.
- **Documentation**—Allows you to specify multiple documents that pertain to your component. You can use this property to specify a file on the internet or in your company's file system. The specified file can be in either `.html` or `.pdf` format. To specify an internet file, begin your path with `http://`, for example: `http://mydomain.com/datasheets/my_memory_controller.html`. To specify a file in your company's file system, you begin your path with `file:///` for Linux and `file://` for Windows, for example: `file:///company_server/datasheets/my_memory_controller.pdf`. For handwritten `_hw.tcl` files, you can specify documentation using the `add_documentation_link` Tcl command. Shows how to specify documentation that is included in the component directory.

Example 6-4. Documentation Link for Documentation Stored with Component HDL Files

```
set_module_property DATASHEET_URL
"file://[get_module_property MODULE_DIRECTORY]Modular_SGDMA_Dispatcher_Core_UG.pdf"
```

 For more information refer to the `add_documentation_link` command in the *Component Tcl Interface Reference*.

Saving a Component

You can save the component by clicking **Finish** on any of the tabs, or by clicking **Save** on the File menu. Based on the settings you specify in the component editor, the component editor creates a component description file with the file name `<class-name>_hw.tcl`. The component editor saves the file in the same directory as the HDL file that describes the component's hardware interface. If you did not specify an HDL file, you can save the component description file to any location you choose.

You can relocate component files later. For example, you could move component files into a subdirectory and store it in a central network location so that other users can instantiate the component in their systems. The `_hw.tcl` file contains relative paths to the other files, so if you move the `_hw.tcl` file you should move all the HDL and other files associated with it.

 Altera recommends that you store `_hw.tcl` files for a project in the `ip/<class-name>` directory for the project. You should store the HDL and other files in the same directory as the `_hw.tcl` file.

Editing a Component

After you save a component and exit the component editor, you can edit it in Qsys. To edit a component, right-click it in the list of available components on the **System Contents** tab and click **Edit Component**. The component editor appears.

 You cannot edit components that were created outside of the component editor, such as Altera-provided components.

If you edit the HDL for a component and change the interface to the top-level module, you need to edit the component to reflect the changes you made to the HDL.

Registering Software Assignments

You can use Tcl commands to create software assignments. You can register any software assignment that you want, as arbitrary key-value pairs. [Example 6-5](#) shows a typical Tcl API script:

Example 6-5. Typical Software Assignment with Tcl API Scripting

```
set_module_assignment name value  
set_interface_assignment name value
```

The assignments are added to the Qsys information file (`.sopcinfo`), available for use for downstream components.

 For more information about these software assignments, refer to the [Publishing Component Information to Embedded Software](#) chapter in the Nios II software Developer's Handbook.

Component Parameterization

To edit component instance parameters, select a component in the **System Contents** tab of Qsys and click **Edit**.

Document Revision History

Table 6-2 shows the revision history for this document.

Table 6-2. Document Revision History

Date	Version	Changes
December 2010	10.1.0	Initial release.

-  For previous versions of the Quartus II Handbook, refer to the [Quartus II Handbook Archive](#).
-  Take an [online survey](#) to provide feedback about this handbook chapter.

Qsys Beta



Altera's Qsys system integration tool is now available as beta for evaluation in the Quartus® II software subscription edition version 10.1. Altera does not recommend using the beta release of Qsys in the Quartus II software version 10.1 for designs that are close to completion and are meeting design requirements. Before using Qsys, review the *Quartus II Software Version 10.1 Release Notes* and *AN 632: SOPC Builder to Qsys Migration Guidelines* for known issues and limitations. To submit general feedback or technical support on the beta release of Qsys, submit a service request through mysupport.altera.com. Alternatively, to submit general feedback, click **Feedback** on the Quartus II software Help menu.

The Qsys interconnect is a high-bandwidth structure for connecting components that use Avalon® interfaces. This chapter describes the Qsys interconnect. The interconnect uses algorithmic transformations to insert interconnect components in implementing the Qsys system. This chapter also provides brief descriptions of the Qsys interconnect components that implement the interconnect. All Qsys interconnect components are available to be used in your own designs. The Qsys interconnect connects the following Avalon interface types:

- Avalon-ST—connects Avalon-ST sources and sinks that stream unidirectional data.
- Avalon-MM—connects Avalon-MM master and slaves that communicate using read and write commands.
- Tristate conduits— connects tristate conduit controllers in the FPGA to tristate devices on the PCB using a three-signal encoding of tristate information.
- Interrupts—connects interrupt senders and the interrupt receivers of the component that service them.
- Clocks—connects clock sources and clock sinks.
- Resets—connects reset sources and reset sinks.
- Conduits—connects point-to-point conduit interfaces. You can use the conduit interface type to define an arbitrary collection of signals that does not fit into any of the other Avalon interface categories.



For more information about the Avalon interfaces, refer to the *Avalon Interface Specifications*.

For Avalon-ST interfaces, Qsys provides adapters that allow flexibility in creating point-to-point connections. For example, the Avalon-ST data format adapter allows you to connect streaming interfaces of different widths.

For Avalon-MM interfaces, the implementation of the Qsys interconnect is based on a network-on-chip architecture. Transactions between masters and slaves are encapsulated in packets and transmitted on a network that carries the packets between masters and slaves. The master command network transports read and write command packets from master interfaces to slave interfaces. The slave response network transports read response packets from slave interfaces to master interfaces.

This chapter includes the following sections:

- “Avalon-MM Interface Components” on page 7-2
- “Avalon-ST Interfaces” on page 7-18
- “Tristate Conduit Components” on page 7-21
- “Interrupt Interfaces” on page 7-27
- “Clock Interfaces” on page 7-29
- “Reset Interfaces” on page 7-29
- “Conduits” on page 7-30

Avalon-MM Interface Components

Qsys interconnect for memory-mapped interfaces connects Avalon-MM master and slave interfaces. It supports the following items:

- Any number of master and slave components. The master-to-slave relationship can be one-to-one, one-to-many, many-to-one, or many-to-many.
- Master and slaves of different data widths.
- Components operating in different clock domains.
- Components with different interface properties and signals. Qsys can adapt the component interfaces so that interfaces with the following types differences can be connected:
 - Interfaces that use active-high and active-low signalling
 - Interfaces with different burst characteristics
 - Interfaces with different latencies
 - Interfaces with different port signatures

Figure 7-1 is a simplified representation of the Qsys interconnect for an Avalon-MM system with multiple masters. As this figure illustrates, the underlying implementation of the master and slave connections uses a network topology. When you generate a Qsys system, Qsys implements the interconnect connectivity that you specified, replacing the point-to-point connections you created in the **Connections** column with a network topology.

Figure 7-1. Qsys interconnect—Example System

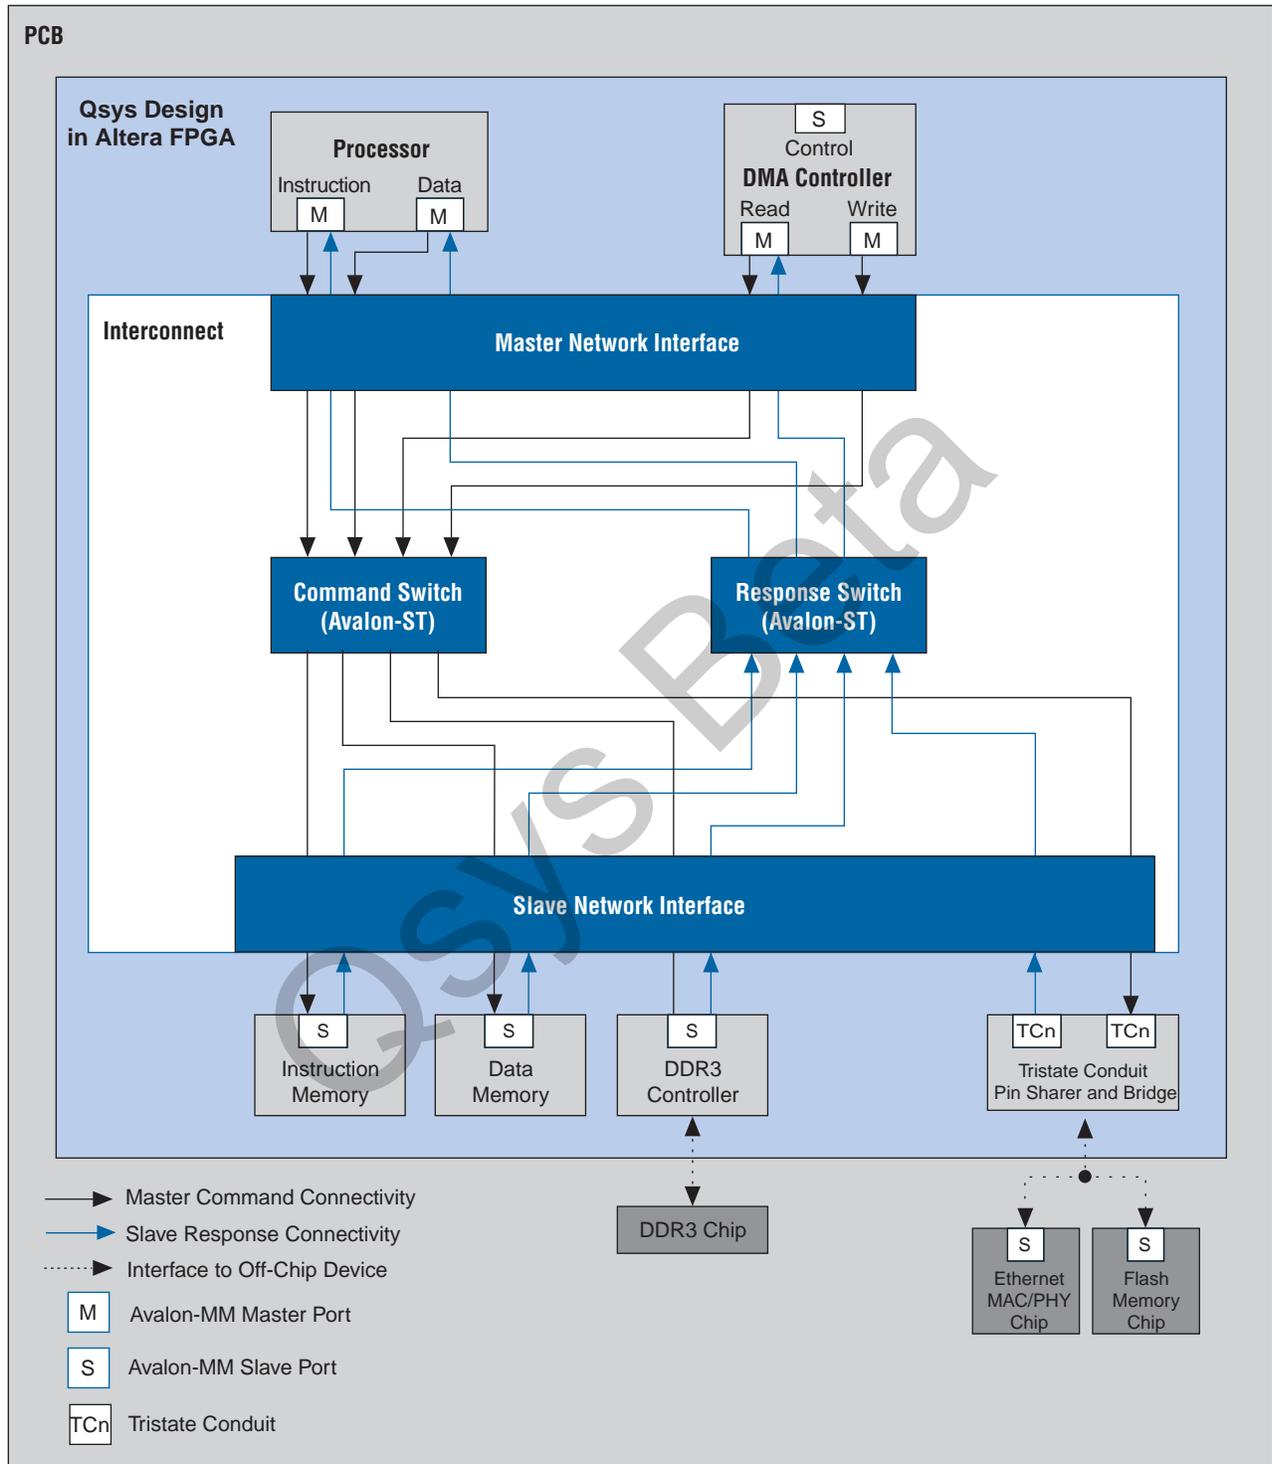


Figure 7-2 shows the format of the Qsys packet that encapsulates the Avalon-MM master commands and Avalon-MM slave responses.

Figure 7-2. Qsys Packet Format

Address	Transaction type	Data	Byte enable	Source ID	Destination ID	Byte count	Burstwrap	Protection
---------	------------------	------	-------------	-----------	----------------	------------	-----------	------------

Table 7-1 describes the fields of Qsys packet.

Table 7-1. Qsys Packet Format

Field	Description
Address	Specifies the byte address for the lowest byte in the current cycle.
Transaction_type	Indicates the transaction type. Table 7-2 lists the 5 transaction types.
Data	For command packets, carries the data to be written. For read response packets, carries the data that has been read.
Byteenable	Specifies which symbol of the data are valid. The following values are legal for Avalon-MM master and slaves transferring 32-bit data: <ul style="list-style-type: none"> ■ 1111 writes full 32 bits ■ 0011 writes lower 2 bytes ■ 1100 writes upper 2 bytes ■ 0001 writes byte 0 only ■ 0010 writes byte 1 only ■ 0100 writes byte 2 only ■ 1000 writes byte 3 only
Source_ID	The ID of the master or slave that initiated the command or response.
Destination_ID	The ID master or slave to which the command or response is directed.
Burstwrap	The burstwrap value specifies the wrapping behavior of the current burst. The burstwrap value is of the form $2^{<n>-1}$. The following types are defined: <ul style="list-style-type: none"> ■ Variable wrap—Variable wrap bursts can wrap at any integer power of 2 value. When the burst reaches the wrap boundary, it wraps back to the previous burst boundary so that only the low order bits are used for addressing. For example, a burst starting at address 0x1C, with a burst wrap boundary of 32 bytes and a burst size of 20 bytes, would write to addresses 0x1C, 0x0, 0x4, 0x8, and 0xC. For a burst wrap boundary of size $<m>$, $Burstwrap = <m> - 1$, or for this case $Burstwrap = (32 - 1) = 31$ which is $2^5 - 1$. ■ Sequential—Sequential bursts increment the address for each transfer in the burst. For sequential bursts, the <code>Burstwrap</code> field is set to all 1s. For example, with a 6-bit <code>Burstwrap</code> field, the value for a sequential burst is 6'b111111 or 63, which is $2^6 - 1$. <p>In version 10.1 of the Quartus II software, adaptation logic sets a hardwired value for the burstwrap field, according the declared master burst properties. For example, for a master which declares sequential bursting, the burstwrap field is set to all 1-bits. Similarly, masters that declare linewrap burst have their burstwrap field set to the appropriate constant value.</p>
Protection	Access level protection. When 0, the packet has normal access. When 1, the packet has privileged access. For Avalon-MM interfaces, this field maps directly to the privileged access signal, which allows an Avalon-MM master to write to an on-chip memory ROM instance.

Table 7-2 lists the transaction type encodings.

Table 7-2. Transaction Types

Bit	Name	Definition
0	PKT_TRANS_READ	When asserted, indicates a read transaction.
1	PKT_TRANS_COMPRESSED_READ	For read transactions, specifies whether or not the read command can be expressed in a single cycle, that is whether or not it has all byteenables asserted on every cycle.
2	PKT_TRANS_WRITE :	When asserted, indicates a write transaction.
3	PKT_TRANS_POSTED	When asserted, no response is required.
4	PKT_TRANS_LOCK	When asserted, indicates arbitration is locked. Applies to write packets.

The fields of the Qsys packet format are variable length to minimize the resources used. However, if the majority of components in a design have a single data width, for example 32 bits, and a single component has a data width of 64 bits, Qsys inserts a width adapter to accommodate 64-bit transfers.

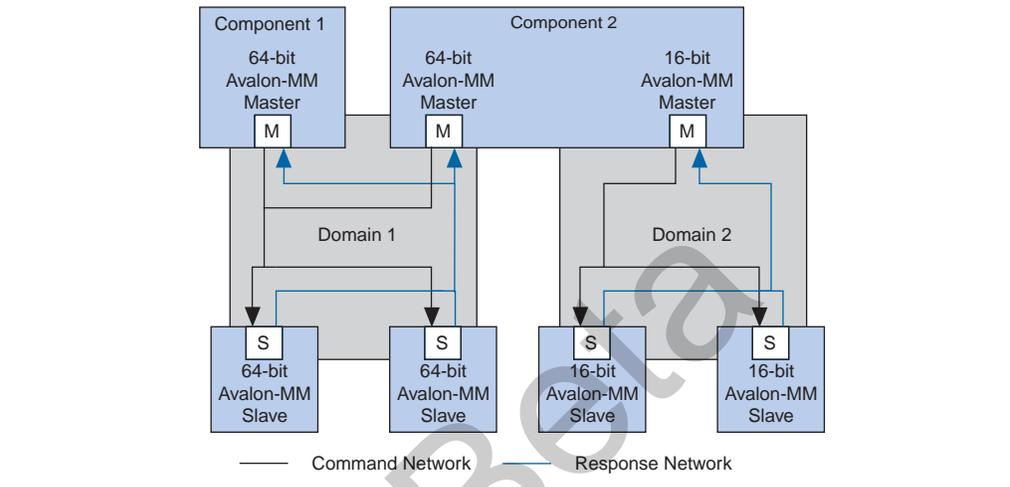
Component Interconnect Domains

A group of connected Avalon-MM masters and slaves is called an *interconnect domain*. The components in a single interconnect domain share the same packet format. The following two examples illustrate this point.

Using Two Separate Domains

Figure 7-3 illustrates the use of two separate domains. The first domain includes two, 64-bit masters connected to two, 64-bit slaves. The second domain includes one, 16-bit master connected to two, 16-bit slaves. Because the interfaces in Domain 1 and Domain 2 do not share any connections, Qsys can optimize the packet format for the two separate domains. In this example, the first domain uses a 64-bit data width and the second domain uses 16-bit data.

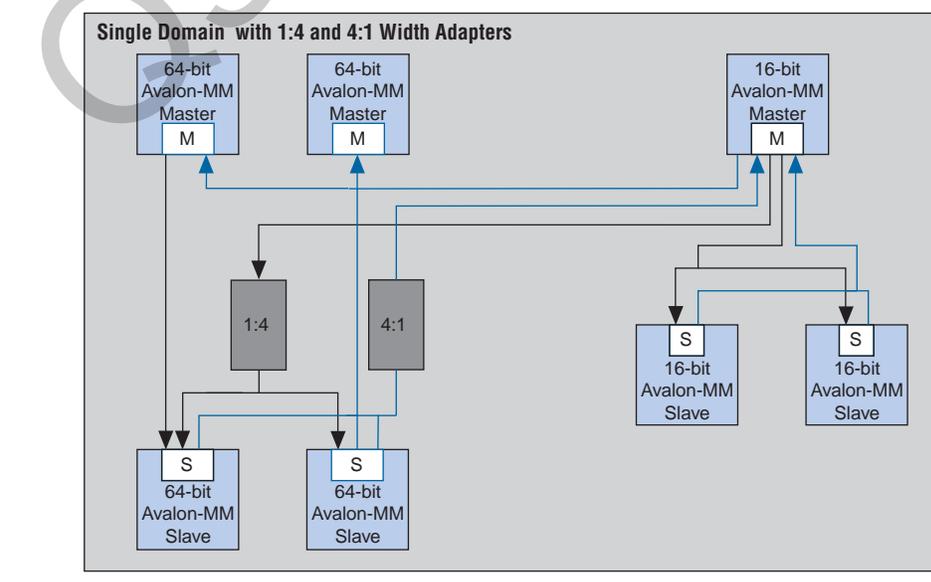
Figure 7-3. Two Domains



Using One Domain with Width Adaptation

Figure 7-4 illustrates a Qsys system that includes two, 64-bit masters that access two, 64-bit slaves. It also includes one, 16-bit Master, accessing two, 16-bit slaves and one, 64-bit slave. Because one of the masters connects to all of the slaves, Qsys creates a single domain with two packet formats: one with 64-bit data and one with 16-bit data. A width adapter manages accesses between the 16-bit master and 64-bit slaves.

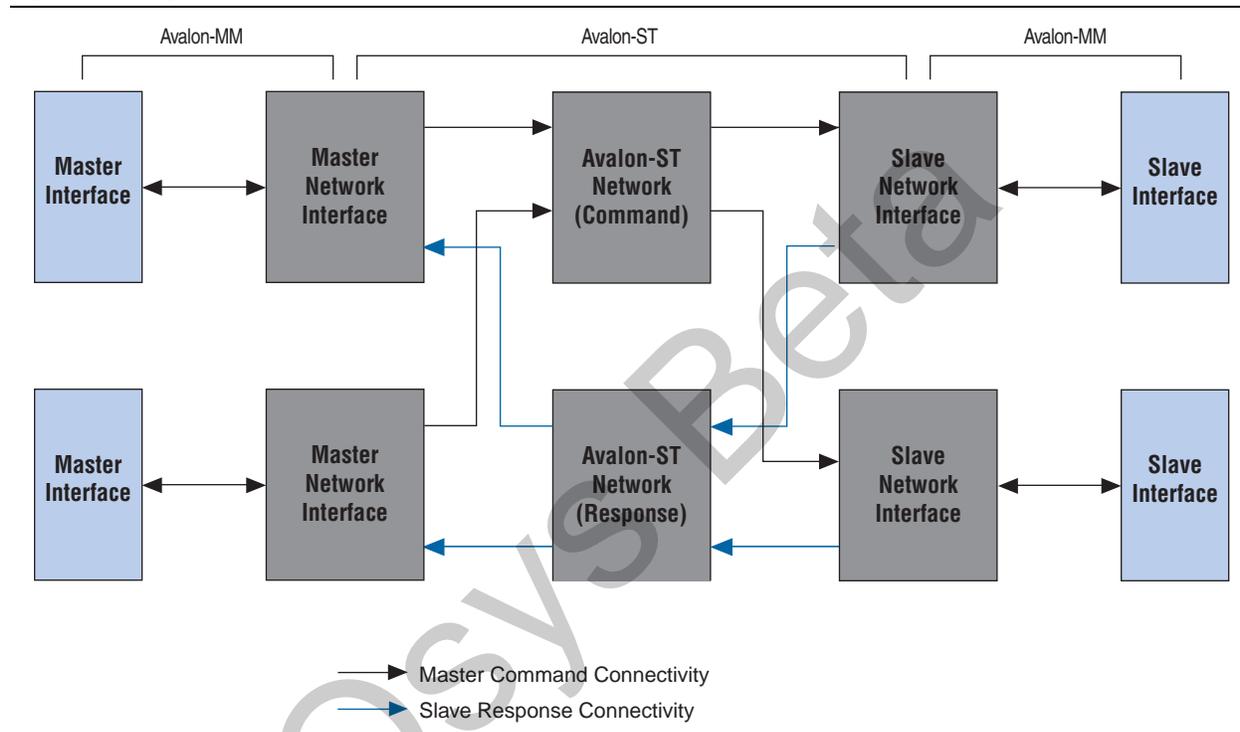
Figure 7-4. One Domain with 1:4 and 1:4 Width Adapters



Qsys Transformations

Figure 7-5 provides a more detailed view of the transformation that occurs when you generate a Qsys system with Avalon-MM master and slave components. As this figure illustrates, the Avalon-MM master and slave components connect to network interface modules that encapsulate the transaction in Avalon ST packets. The Avalon-MM interfaces have no information about the encapsulation or the function of the layer transporting the packets and simply operate in accordance with Avalon-MM protocol, using the read and write signals and transfers as defined in the *Avalon Interface Specifications*.

Figure 7-5. Qsys Transform from Avalon-MM to Avalon-ST



Master Command and Slave Response Networks

Many Qsys components implement the Qsys interconnect and network interfaces represented by the Avalon-ST Network (Command) and Avalon-ST Network (Response) blocks in Figure 7-5. All of these Qsys components are provided by Altera and included in the Component Library available in Qsys. They are available for you to be used stand-alone in your designs. For example, you may want to include the Avalon-ST pipeline stage in your datapath to pipeline a streaming connection, thus increasing the clock frequency of your design.

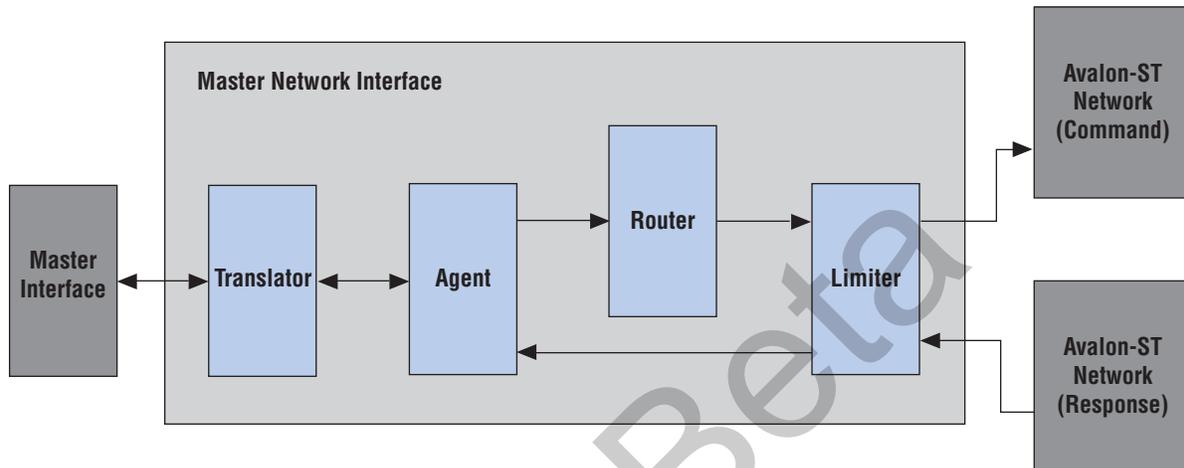
The subsequent sections describe the components that are part of the Avalon-ST master command and Avalon-ST slave response network, including the following components:

- Merlin Master Translator
- Merlin Master Agent
- Merlin Router

- Merlin Traffic Limiter
- Merlin Slave Translator
- Merlin Slave Agent

Figure 7-6 provides a block diagram for the Master command network showing the Merlin master translator, agent, router and limiter.

Figure 7-6. Qsys Components in the Master Command Network



Merlin Master Translator

The Merlin master translator interfaces to an Avalon-MM master component. It converts the Avalon-MM master interface to a simpler representation that the Qsys network uses. It performs the following functions:

- Translates active low signalling to active high signalling
- Inserts wait states to prevent an Avalon-MM master from reading invalid data
- Translates word and symbol addresses
- Translates word and symbol burst counts
- Handles burst count timing and sequencing
- Removes unnecessary address bits

Merlin Master Agent

The agent translates Avalon-MM master transactions into Qsys command packets and translates the Qsys Avalon-MM slave response packets into Avalon-MM responses.

Merlin Router

The router routes command packets from the master to the slave and response packets from the slave to the master. For master command packets, the router uses the Avalon-MM address to set the `Destination_ID` and Avalon-ST channel. For the slave response packet, the router uses the `Destination_ID` to set the Avalon-ST channel. The demultiplexers use the Avalon-ST channel to route the packet to the correct destination.

Merlin Traffic Limiter

The limiter ensures the responses arrive in order. It prevents any command from being sent if the response could conflict with the response for a command that has already been issued. By guaranteeing in-order responses, the limiter simplifies the response network.

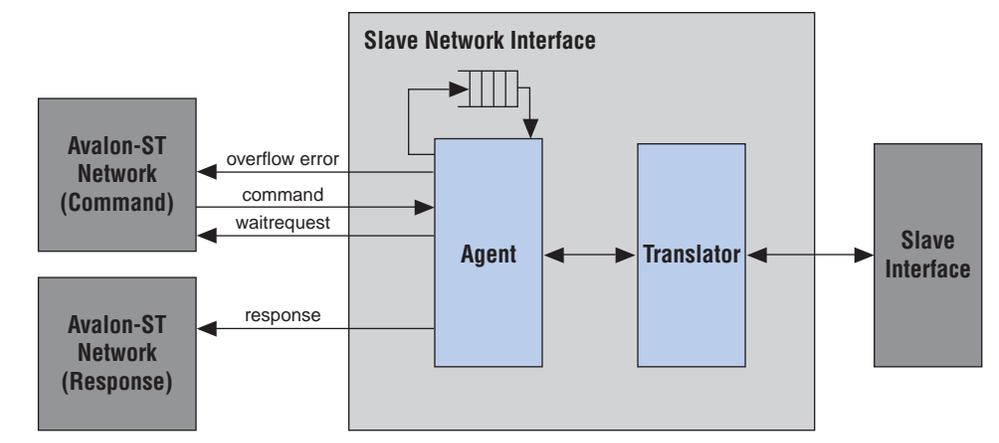
Merlin Slave Translator

The Merlin slave translator interfaces to an Avalon-MM slave component as [Figure 7-7](#) illustrates. It converts the Avalon-MM slave interface to a simplified representation that the Qsys network uses. An Avalon-MM Merlin slave translator performs the following functions:

- Drives the `begintransfer`, `beginbursttransfer`, and `writebyteenable` signals
- Supports Avalon-MM slaves that operate using fixed timing and or slaves that use the `readdatavalid` signal to identify valid data
- Translates the read, write, and chipselect signals into the representation that the Avalon-ST slave response network uses
- Converts active low signals to active high signals
- Translates word and symbol addresses and burstcounts
- Handles burstcount timing and sequencing
- Removes unnecessary address bits

[Figure 7-7](#) shows the Qsys components that comprise the slave response network.

Figure 7-7. Qsys Components in the Slave Response Network



Merlin Slave Agent

The agent accepts command packets, and issues the resulting transactions to the Avalon interface. For pipelined slaves, an Avalon-ST FIFO stores information about pending transactions. The size of this FIFO is the maximum number of pending responses that you specify when creating the slave component.

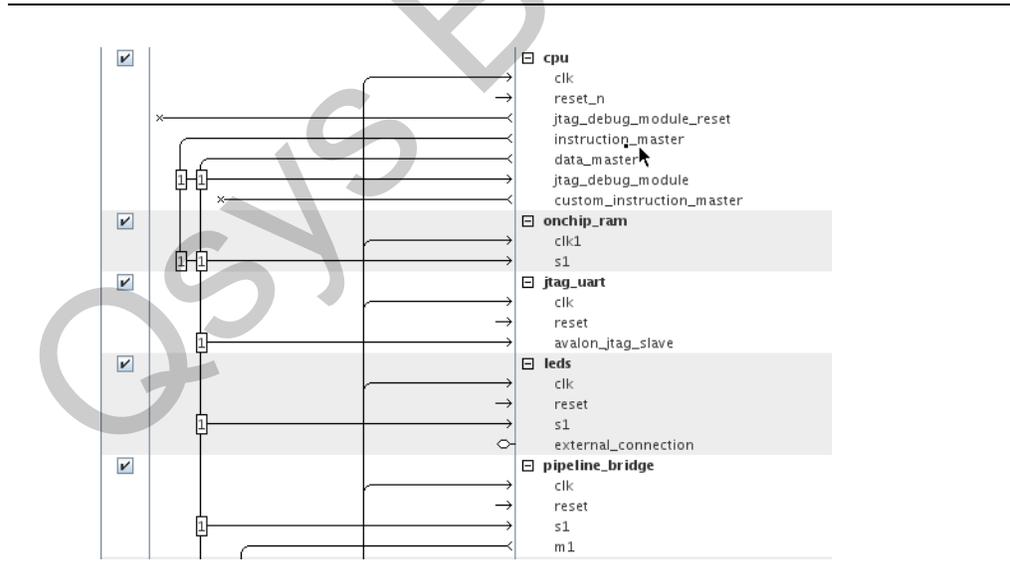
The agent also backpressures the Avalon-MM master command interface when the FIFO is full if the slave component includes the `waitrequest` signal.

Arbitration

When multiple masters contend for access to a slave, Qsys automatically inserts arbitration which grants access in fairness-based, round-robin order. In a fairness-based arbitration scheme, each master has an integer value of transfer *shares* with respect to a slave. One share represents permission to perform one transfer. The default arbitration scheme is equal share round-robin granting equal, sequential access to all requesting masters. You can change the arbitration scheme to weighted round robin by specifying a relative number of arbitration shares to the masters that access a particular slave. To display arbitration settings, on the View menu, click **Show Arbitration**.

Figure 7-8 illustrates the arbitration shares.

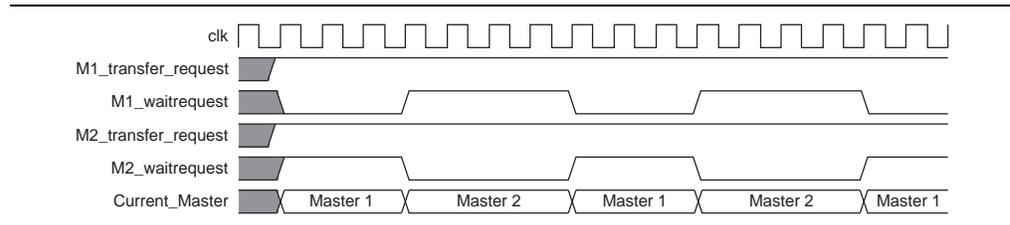
Figure 7-8. Arbitration Settings on the System Contents Tab



Arbitration Examples

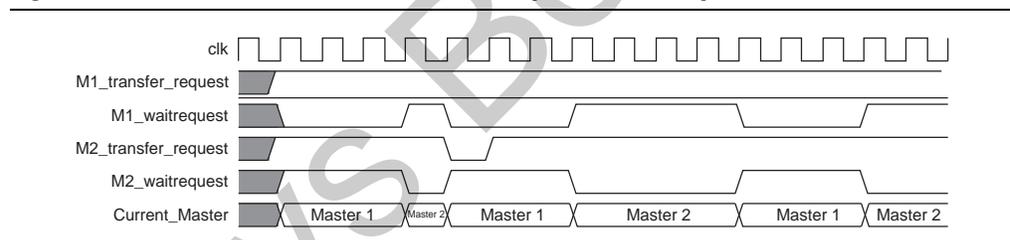
Figure 7-9 illustrates the timing for two Avalon-MM masters continuously accessing a single Avalon-MM slave to perform back-to-back transfers. Master 1 has three shares and Master 2 has four shares. The Merlin arbiter grants Master 1 access for three transfers, then Master 2 for four transfers. This cycle repeats indefinitely.

Figure 7-9. Arbitration of Continuous Transfer Requests from Two Masters



If a master stops requesting transfers before it exhausts its shares, it forfeits all of its remaining shares, and the Merlin arbiter grants access to another requesting master as Figure 7-10 illustrates. After completing one transfer, Master 2 stops requesting for one clock cycle. As a result, the arbiter grants access back to Master 1, which gets three shares.

Figure 7-10. Arbitration of Two Masters with a Gap in Transfer Requests

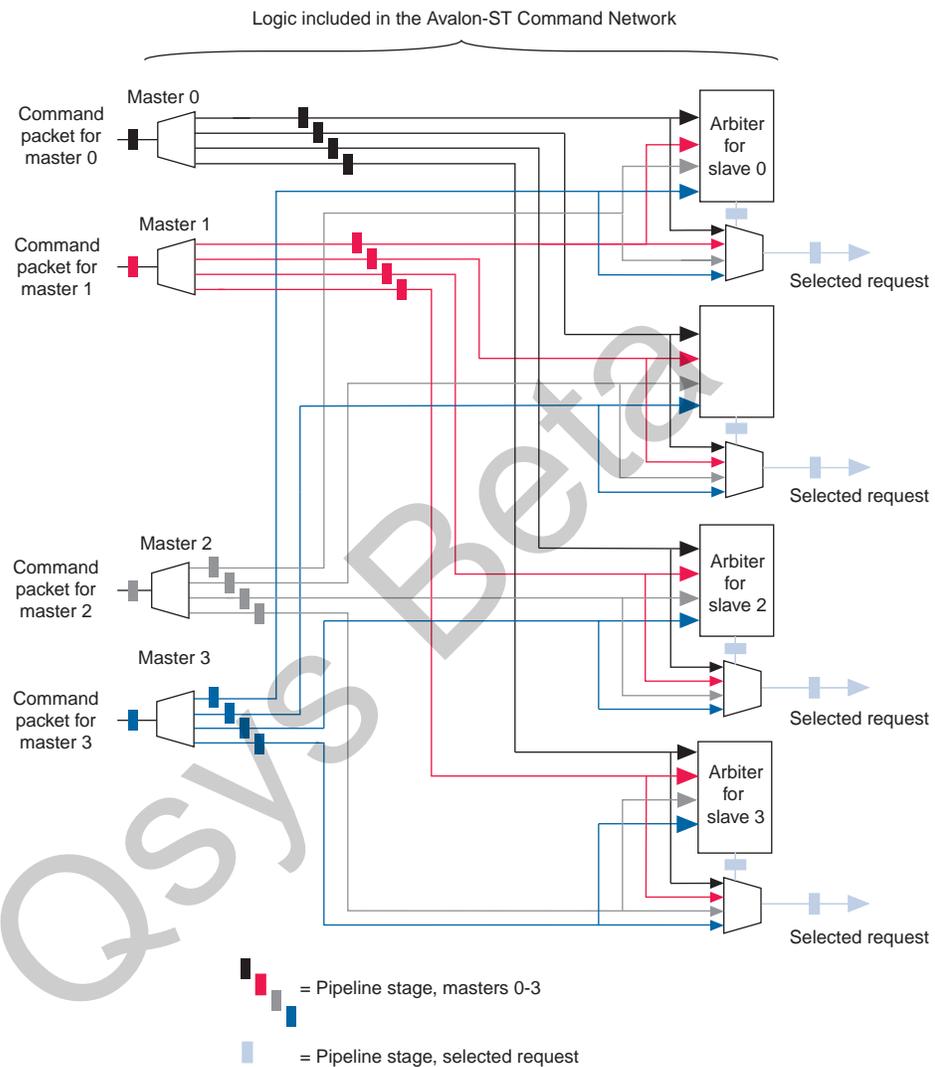


Merlin Arbiter

The input to the Merlin arbiter is the Avalon-MM master command packet for all masters requesting access to the a particular slave. The arbiter outputs the channel number for the selected master. This channel number controls the output of a multiplexer that selects slave device. Figure 7-11 illustrates this logic.

In [Figure 7-11](#), four Avalon-MM masters connect to four Avalon-MM slaves. In each cycle, an arbiter positioned in front of each Avalon-MM slave, selects among the requesting the Avalon-MM masters.

Figure 7-11. Arbitration Logic



If you specified a **Max Additional Latency** parameter greater than zero on the Qsys **Project Settings** tab, the output of the arbiter is registered. Registering this output reduces the amount of combinational logic between the master and fabric, increasing the f_{MAX} of the system.

For more information about the **Max Additional Latency** parameter refer to the “Project Settings” section in the *Creating a System with Qsys* chapter in volume 1 of the *Quartus II Handbook*.

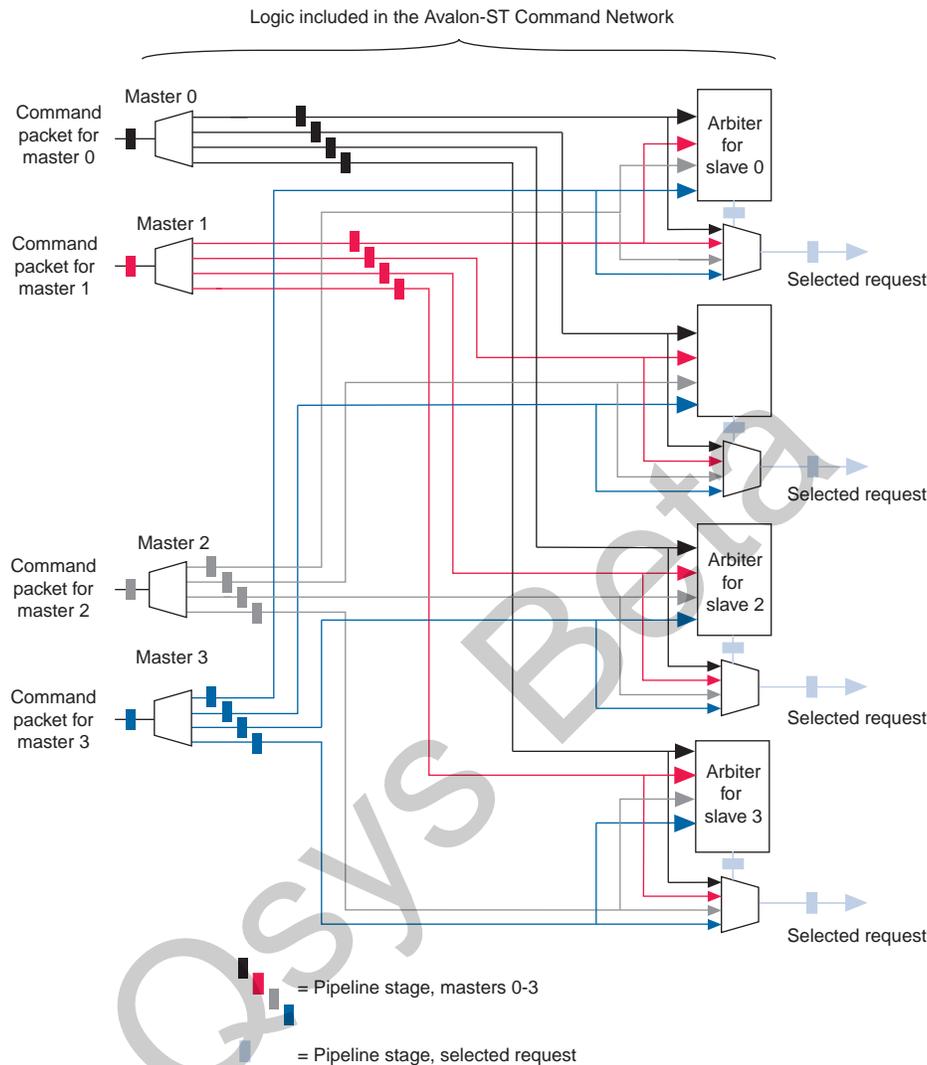
Interconnect Pipelining

If you set the **Max Additional Latency** parameter to a value greater than 0 on the **Project Settings** tab, Qsys automatically inserts Avalon-ST pipeline stages when you generate your design. The pipeline stages increase the f_{MAX} of your design by reducing the combinational logic depth. The cost is additional latency and logic.

Figure 7-12 shows the placement of up to four potential pipeline stages inserted by Qsys in the following locations:

- Before the input to the demultiplexer
- At the output of the multiplexer
- Between the arbiter and the multiplexer
- At the outputs of the demultiplexer
- Please add an additional paragraph:

 The insertion of pipeline stages depends upon the existence of certain interconnect components. For example, in a single-slave system, no multiplexer exists; therefore multiplexer pipelining does not occur. In an extreme case, of a single-master to single-slave system, no pipelining occurs, regardless of the value of **Max Additional Latency**.

Figure 7-12. Pipeline Placement in Arbitration Logic

Additional Qsys Interconnect Components

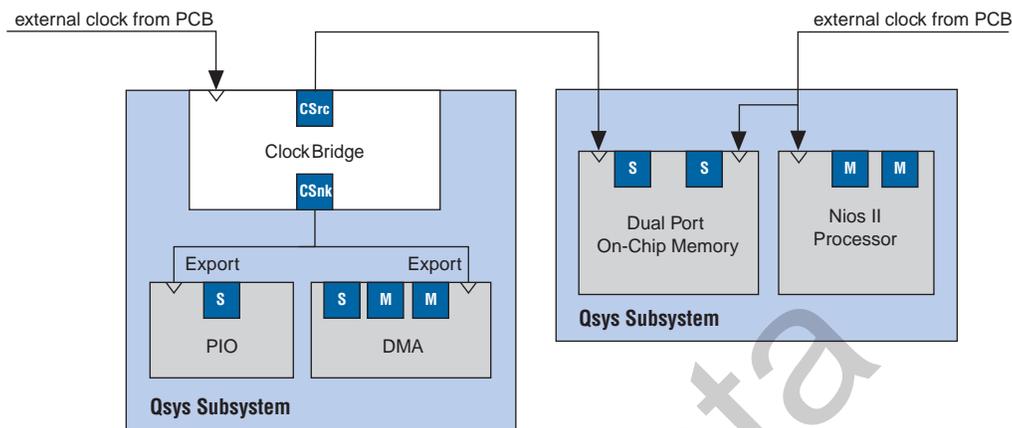
The following sections describe additional components used by the Qsys interconnect. All of these components are in the Qsys Component Library for use in your designs.

- “Clock Bridge” on page 7-15
- “Avalon-MM Clock Crossing Bridge (Qsys)” on page 7-15
- “Avalon-MM Pipeline Bridge (Qsys)” on page 7-15
- “Merlin Width Adapter” on page 7-16

Clock Bridge

The clock bridge allows you to route clocks between Qsys subsystems. You can use this bridge to connect a single clock source to the input clocks of multiple Qsys subsystems. Figure 7-13 illustrates the use of this bridge.

Figure 7-13. Clock Bridge



Avalon-MM Clock Crossing Bridge (Qsys)

The Avalon-MM clock crossing bridge transfers Avalon-MM commands and responses between asynchronous clock domains. It uses asynchronous FIFOs to implement the clock crossing logic. The Avalon-MM clock crossing bridge has a number of parameters, including parameters to control the depth of the synchronization FIFO in both the master and slave clock domains.



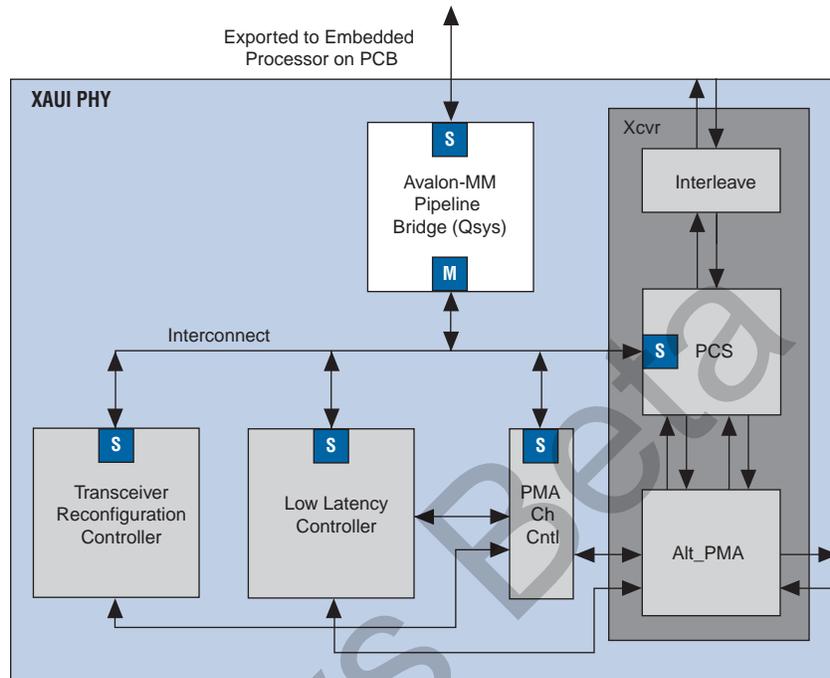
The Avalon-MM clock crossing bridge (Qsys) core is implemented to work with the Qsys interconnect. The legacy Avalon-MM clock crossing bridge core is available for SOPC Builder systems. If you port an SOPC Builder design that includes the Avalon-MM clock crossing bridge to Qsys, Qsys automatically changes the older version to the Qsys version.

Avalon-MM Pipeline Bridge (Qsys)

The Avalon-MM Pipeline Bridge inserts a register stage in the Avalon-MM command and response paths. It accepts commands on its Avalon-MM slave port and propagates them to its Avalon-MM master port. It provides separate parameters to turn on pipelining in the command and response networks.

Because you can turn the pipelining feature of this bridge off, you can also use the Avalon-MM bridge to export a single Avalon-MM slave interface that can be used to control multiple Avalon-MM slave devices. In this configuration, it transfers commands received on its Avalon-MM slave interface to its Avalon-MM master port. [Figure 7-14](#) illustrates its use.

Figure 7-14. Avalon Bridge



Because the Avalon-MM slave interface is exported to the pins of the device, having a single Avalon-MM slave port, rather than separate ports for each Avalon-MM slave device, reduces the pin count of the FPGA.

The Avalon-MM pipeline bridge (Qsys) is implemented to work with the Qsys interconnect. The older Avalon-MM pipeline bridge is available for SOPC Builder systems. If you upgrade from SOPC Builder to Qsys, Qsys automatically replaces the bridge.

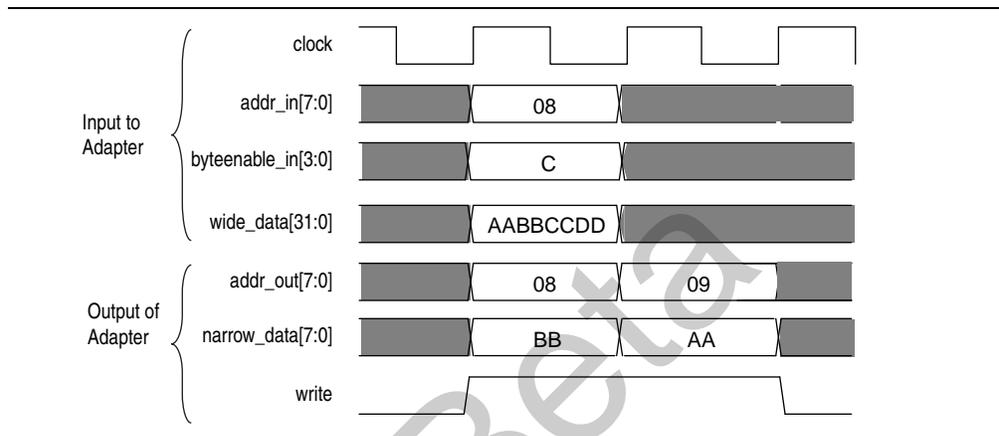
Merlin Width Adapter

The Merlin width adapter converts between Avalon-MM master and slaves with different data and byteenable widths. This adapter is used in the Avalon-ST domain and operates with information contained in the packet format illustrated [Figure 7-2 on page 7-4](#). It accepts packets on its sink interface with one data width and produces output packets on its source interface with a different data width. The ratio of the wider data width to the narrower width must be a power of two, such as 4:1, 8:1, and 16:1. This adapter assumes that the field ordering of the input and output packets is the same, with the only difference being the width of the data and accompanying byte enable signals.

When the width adapter converts from a wide data to a narrow data, the narrower data is transmitted over several beats. The first output beat contains the lowest addressed segment of the input data and byte enables. Figure 7-15 illustrates the timing for a 4:1 width adapter.

When the width adapter converts from narrow data to wide data, each input beat's data and byte enables are copied to the appropriate segment of the wider output data and byte enables signals.

Figure 7-15. Width Adapter Timing for a 4:1 Adapter



Burst Transfers

Avalon-MM burst transactions grant a master uninterrupted access to an Avalon-MM slave for a specified number of transfers. The master specifies the number of transfers when it initiates the burst using the `burstcount` signal. Once a burst begins between a master-slave pair, arbiter logic is locked until the burst completes. For burst masters, the size of the burst is the number of cycles that the master has access to the slave, and the selected arbitration shares have no effect.

Merlin Burst Adapter

The Qsys interconnect uses the Qsys Merlin burst adapter to accommodate the burst capabilities of each interface in the system, including interfaces that do not support burst transfers. The maximum burst length for each interface is a property of the component interface and is independent of other interfaces in the system. Therefore, a particular master might be capable of initiating a burst longer than a slave's maximum supported burst length. In this case, the burst adapter translates the large master burst into smaller bursts, or into individual slave transfers if the slave does not support bursting. Until the master completes the burst, the arbiter logic prevents other masters from accessing the target slave. For example, if a master initiates a burst of 16 transfers to a slave with maximum burst length of 8, the burst adapter initiates 2 bursts of length 8 to the slave.

Avalon-MM masters always issue addresses that are aligned to the size of the transfer. However, in some cases, when a narrow-to-wide width adaptation is used, the resulting address may be unaligned. In the case of unaligned addresses, the burst adapter issues the maximum possible sized bursts, with appropriate byte enables, to bring the burst-in-progress up to an aligned slave address. Then, it completes the burst on aligned addresses.

Burst Types

The burst adapter supports variable wrap or sequential burst types to accommodate the different properties of the Avalon-MM masters. Refer to [Table 7-1 on page 7-4](#) for definitions of these burst types. Some bursting masters can issue more than one burst type.

Avalon-ST Interfaces

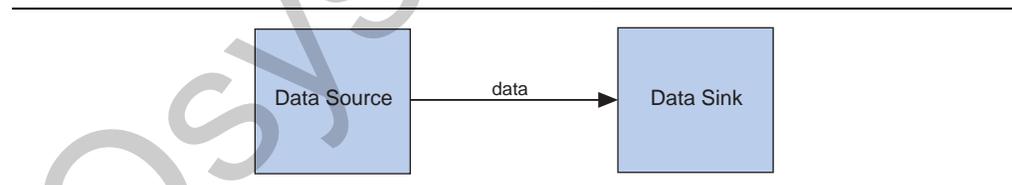
The interconnect for Avalon-ST connects high-bandwidth, low-latency components that use the Avalon-ST interface. This interconnect creates datapaths for unidirectional traffic including multichannel streams, packets, and DSP data. The Avalon-ST interconnect is flexible and can be used to implement on-chip interfaces for industry standard telecommunications and data communications cores, such as Ethernet, Interlaken, and video. In all cases, you can define bus widths, packets, and error conditions.

You specify how Avalon-ST source and sink ports connect in Qsys. If your source and sink interfaces have different properties, selecting **Insert Avalon-ST adapters** on the XXXX menu Qsys inserts the necessary adapters which are visible in the **System Contents** tab.

Avalon-ST Examples

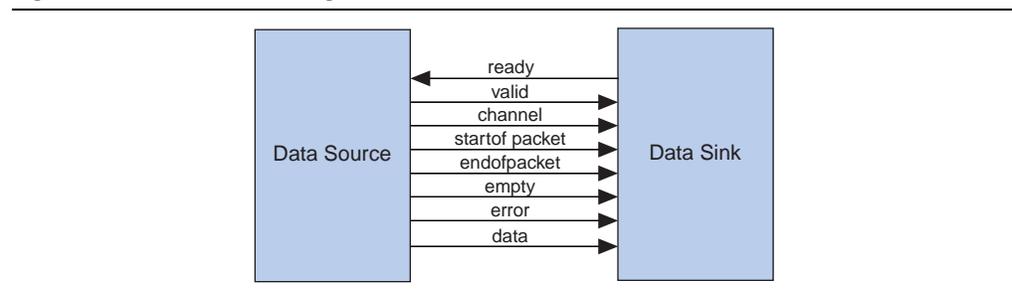
[Figure 7-16](#) illustrates the simplest system example with an Avalon-ST connection between the source and sink. This source-sink pair includes only the data signal. The sink must be able to receive data as soon as the source interface comes out of reset.

Figure 7-16. Interconnect for a Simple Avalon Streaming Source-Sink Pair



[Figure 7-17](#) illustrates a more extensive interface that includes signals indicating the start and end of packets, channel numbers, error conditions, and back pressure.

Figure 7-17. Avalon Streaming Interface for Packet Data



All data transfers using Avalon-ST interconnect occur synchronously to the rising edge of the associated clock interface. Throughput and frequency of a system depends on the components and how they are connected.

 For details about the Avalon-ST interface protocol, refer to the *Avalon Interface Specification*.

Avalon-ST Components

The Qsys **Component Library** includes a number of Avalon-ST components that you can use to create datapaths, including datapaths whose input and output streams have different properties. Generated systems that include Avalon-MM master and slave components may also use these Avalon-ST components because the generation process creates an interconnect whose structure resembles a network topology as “Qsys Transformations” on page 7-7 describes. The following sections introduce the Avalon-ST components.

Avalon-ST Handshake Clock Crosser

The Avalon-ST handshake clock crossing adapter connects streams that operate at different frequencies. This adapter uses a simple hand-shaking protocol to propagate transfer control signals and responses across the clock boundary and responses in the other direction. This methodology uses fewer FPGA resources because each transfer is safely propagated to the target domain before the next transfer can begin. The Avalon-ST handshake clock crosser is appropriate for lithotripsy connections because the handshake incurs at least four cycles of round-trip latency for every read command, limiting throughput.

You can use the parameter editor for the Avalon-ST handshake clock crosser to specify parameter values. Among the parameters that you can specify are the data width, whether or not to include packet support, and synchronizer depths.

Avalon-ST Pipeline Stage

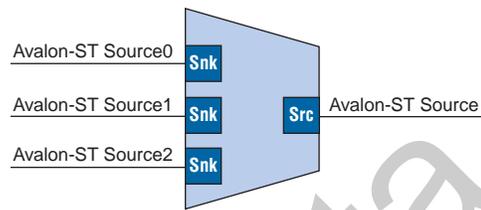
The Avalon-ST pipeline stage optionally inserts a single pipeline (register) stage in the Avalon-ST command and response datapaths. It receives data on its Avalon-ST sink interfaces and drives it unchanged on its Avalon-ST source interface.

 The Qsys Component Library also includes an Avalon-MM pipeline bridge whose data interfaces use the Avalon-MM protocol, rather than the Avalon-ST protocol.

Merlin Multiplexer

The Merlin multiplexer accepts data on its Avalon-ST sink interface and multiplexes the data for transmission on its Avalon-ST source interface. You can parameterize the multiplexer to append channel information on the source to indicate which sink is driving the source data. The multiplexer includes internal arbitration logic which selects between inputs using a round-robin arbitration algorithm. [Figure 7-18](#) illustrates the Avalon-ST multiplexer. Among the parameters that you can specify are the option to use packet scheduling, which guarantees that the multiplexer only changes inputs at the end of a packet.

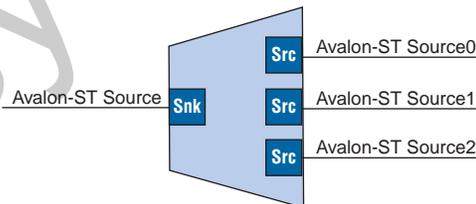
Figure 7-18. Merlin Multiplexer



Merlin Demultiplexer

The Merlin demultiplexer accepts channelized data on its sink interfaces, and transmits the data on one of its source interfaces. The channel bits of the source stream indicate which port the drives the output data. [Figure 7-19](#) illustrates the Merlin multiplexer. Among the parameters that you can specify are the number of output ports and the width of the channel signal.

Figure 7-19. Avalon-ST Demultiplexer



Avalon-ST and Avalon-MM Interfaces

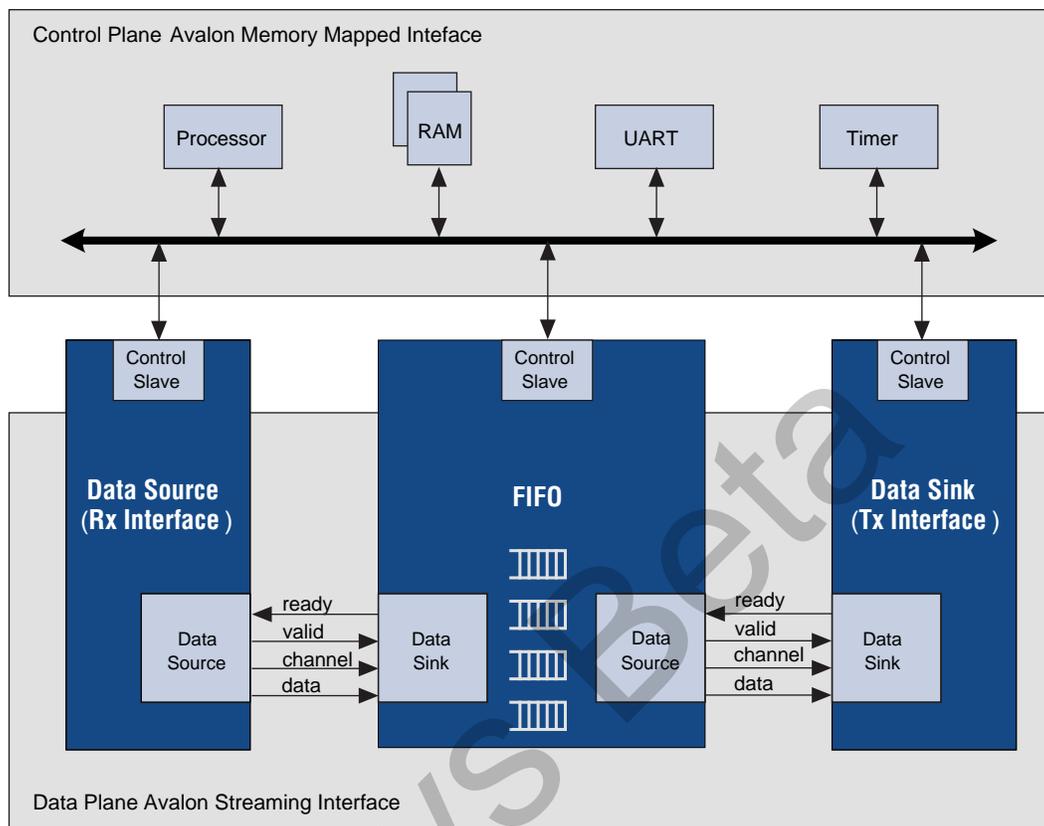
The Avalon-ST and Avalon-MM interfaces are complementary. High bandwidth components with streaming data typically use Avalon-ST interfaces for the high throughput datapath. These components can also use Avalon-MM connection interfaces to provide an access point for control. In contrast to the Avalon-MM interconnect, which can be used to create a wide variety of topologies, the Avalon-ST interconnect fabric always creates a point-to-point between a single data source and data sink, as [Figure 7-20](#) illustrates.

There are two connection pairs in this figure:

- The data source in the Rx Interface transfers data to the data sink in the FIFO.
- The data source in the FIFO transfers data to the Tx Interface data sink.

In [Figure 7-20](#), the Avalon-MM interface allows a processor to access the data source, FIFO or data sink to provide system control.

Figure 7-20. Use of the Avalon-MM Avalon-ST Interfaces



Tristate Conduit Components

The tristate conduit interface type allows you to design Qsys subsystems that connect to tristate devices on your PCB. The following three components implement the tristate conduit functionality:

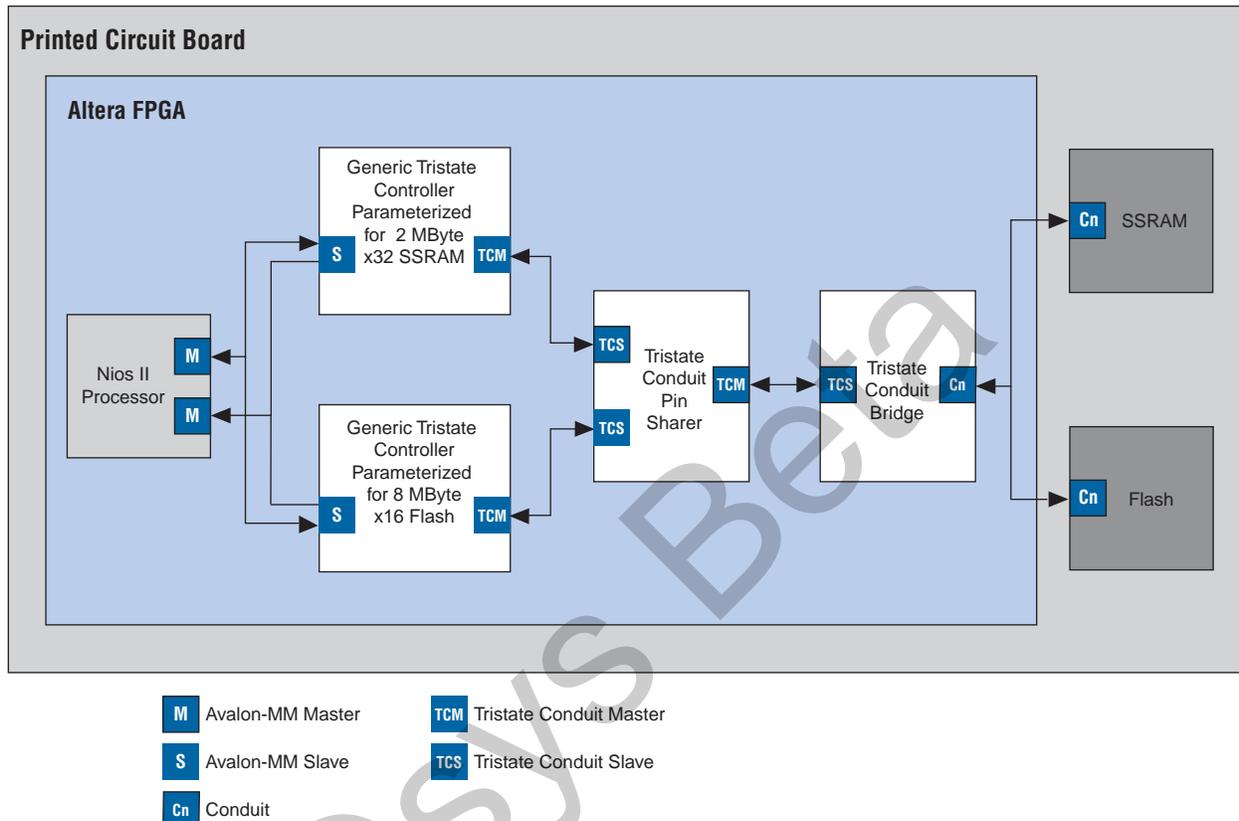
- [Generic Tristate Controller](#)
- [Tristate Conduit Pin Sharer](#)
- [Tristate Conduit Bridge](#)

You can use these components to implement pin sharing, convert between unidirectional and bidirectional signals, and create tristate controllers for devices whose interfaces can be described using the Avalon-MM signal types.

 For more information about the Avalon-MM signal types, refer to the *Avalon Memory-Mapped Slave Interfaces* chapter in the [Avalon Interface Specifications](#).

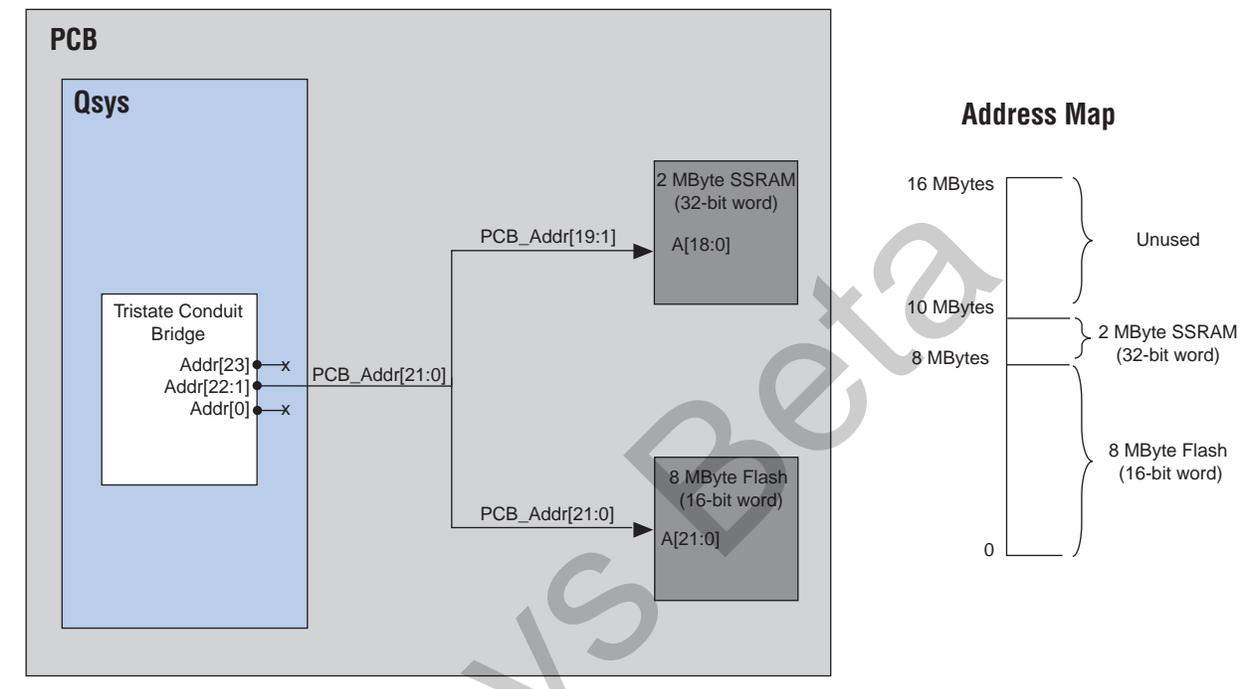
Figure 7-21 illustrates the typical use of these components. This figure includes two generic tristate conduit controllers. The first is customized to control a flash memory. The second is customized to control an off-chip SSRAM. The tristate conduit pin sharer multiplexes between these two controllers, and the tristate conduit bridge converts between an onchip encoding of tristate signals and true bidirectional signals. The tristate conduit bridge converts between an onchip encoding of tristate signals and true bidirectional signals.

Figure 7-21. Tristate Conduit System to Control Off-Chip SRAM and Flash Devices



By default, the tristate conduit pin sharer and tristate conduit bridge presents byte addresses. Each address location in many memory devices contains more than one byte of data. In the example presented in Figure 7-21, the flash device operates on 16-bit words and must ignore the least-significant bit of the Avalon-MM address. The SSRAM memory operates on 32-bit words and must ignore the two, low-order memory bits. Because neither device requires a byte address, `addr[0]` is not routed on the PCB. Figure 7-22 shows `addr[0]` as a unconnected.

Figure 7-22. Address Connections from Qsys System to PCB



In this example design, the flash device responds to address range 0 MBytes–8 MBytes-1. The SSRAM responds to address range 8 MBytes–10 MBytes-1. The PCB schematic for the PCB connects `addr[20:2]` to `addr[18:0]` of the SSRAM device because the SSRAM responds to 32-bit word address. The 8 MByte flash device accesses 16-bit words; consequently, the schematic does not connect `addr[0]`. Chipselect signals select between the two devices.

 If you create a custom tristate conduit master with word-aligned addresses, the tristate conduit pin sharer does nothing to change or align the address signals.

Figure 7-23 illustrates this example system in Qsys.

Figure 7-23. Tristate Conduit System in Qsys

Connections	Module	Description	Export As
	<input type="checkbox"/> nios_2_qsys_0 data_master instruction_master d_irq jtag_debug_module	Nios II Processor Avalon Memory Mapped Master Avalon Memory Mapped Master Interrupt Receiver Avalon Memory Mapped Slave	 Click to export Click to export Click to export Click to export
	<input type="checkbox"/> flash_controller universal_avalon_slave_0 tristate_conduit_master_0	Generic Tristate Controller Avalon Memory Mapped Slave Tristate Conduit Master	 Click to export Click to export
	<input type="checkbox"/> SDRAM_controller universal_avalon_slave_0 tristate_conduit_master_0	Generic Tristate Controller Avalon Memory Mapped Slave Tristate Conduit Master	 Click to export Click to export
	<input type="checkbox"/> tristate_conduit_pin_sharer ps_tcm tcs_0 tcs_1	Tristate Conduit Pin Sharer Tristate Conduit Master Tristate Conduit Slave Tristate Conduit Slave	 Click to export Click to export Click to export
	<input type="checkbox"/> tristate_conduit_bridge out tcs	Tristate Conduit Bridge Conduit Tristate Conduit Slave	 out Click to export
	<input type="checkbox"/> multiplexer_0	Avalon-ST Multiplexer	Click to export

Generic Tristate Controller

The generic tristate controller provides a template for a controller that you can parameterize to reflect the behavior of an off-chip device. The generic tristate controller has many parameters that you can use to customize this component such as the following examples:

- The width of the address and data signals
- The read and write wait times
- The bus turnaround time



In calculating delays, the generic tristate controller chooses the larger of the bus turnaround time and read latency. Turnaround time is measured from the time that a command is accepted, not from the time that the previous read returned data.

- The data hold time

The generic tristate controller always includes the following interfaces:

- Avalon-MM slave interface—This interface connects to an Avalon-MM master, such as a Nios II processor.
- Tristate conduit master—This interface usually connects to the tristate conduit slave interface of the tristate conduit pin sharer.
- Clock sink—The component's clock reference. This interface must be connected to a clock source.
- Resets sink—This interface connects to a reset source interface.

In addition, the generic tristate controller includes optional reset source and interrupt sender interfaces.

To work correctly with the Nios II SBT, the controller must include appropriate module and interface assignments. You can include these assignments by adding assignments using the **Module Assignments** and **Avalon Connection Point Assignments** section of the parameter editor. Downstream embedded software tools use these assignments.

For more information about these configuration names, refer to the *Publishing Component Information to Embedded Software* chapter in the *Nios II Software Developer's Handbook*.

Altera provides preset values for all configuration names for many commonly used devices. Figure 7-24 illustrates the **Module Assignments** and **Avalon Connection Point Assignments** for the Flash Memory Interface (CFI).

Figure 7-24. Module and Avalon Connection Point Assignments

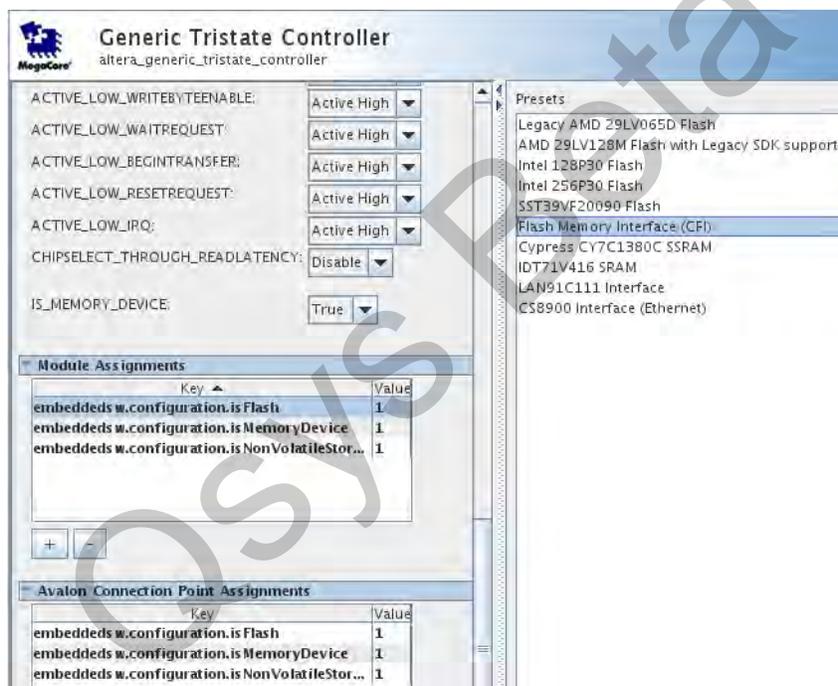


Table 7-3 lists configuration names that you can use to identify your components to downstream embedded software tools.

Table 7-3. Configuration Names

Assignment	Value
embeddedsw.configuration.isNonVolatileStorage	0 or 1
embeddedsw.configuration.isPrintableDevice	0 or 1
embeddedsw.configuration.isMemoryDevice	0 or 1
embeddedsw.configuration.isFlash	0 or 1
embeddedsw.configuration.isEthernetMacDevice	0 or 1

Tristate Conduit Pin Sharer

The tristate conduit pin sharer multiplexes between the signals of the connected tristate controllers. You connect all signals from the tristate controllers to the tristate conduit pin sharer and use the parameter editor to specify the signals that are shared. The parameter editor includes the **Originating Signal** name of the connected signals and a **Shared Signal Name** column for you to type the shared name as [Figure 7-25](#) illustrates.

Figure 7-25. Specifying Shared Signals Using the Tristate Conduit Pin Sharer



If the widths of shared signals differ, the signals are aligned on their 0th bit and the higher-order pins are driven to 0 whenever the smaller signal has control of the bus. Unshared signals always propagate through the pin sharer. The tristate conduit pin sharer uses the round-robin arbiter that is described in [“Arbitration” on page 7-10](#) to select between tristate conduit controllers.

 All tristate conduit components connected to a given pin sharer must be in the same clock domain.

Tristate Conduit Bridge

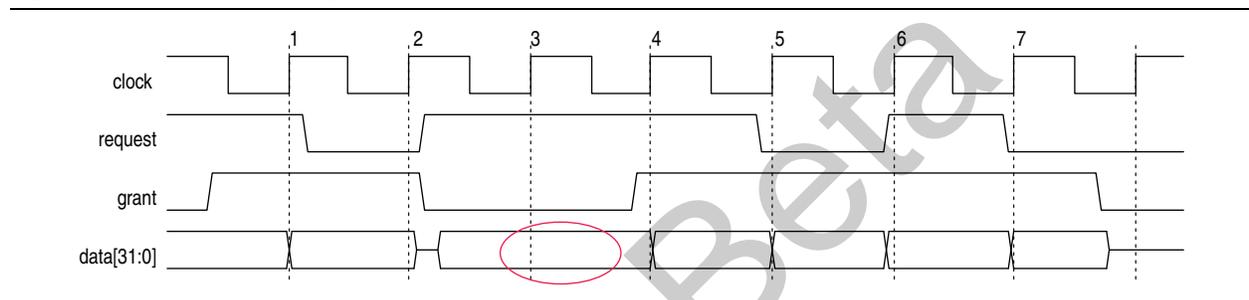
The tristate conduit bridge is the final component on the edge of the Qsys system. It instantiates bidirectional signals for each tristate triplet while passing all other signals straight through the component. The tristate conduit bridge registers all outgoing and incoming signals, which adds two cycles of latency for a read request. You must account for this additional pipelining when designing a custom controller. During reset, all outputs are placed in a high-impedance state; outputs are enabled in the first clock cycle after reset is deasserted. The Quartus II software labels these output signals bidirectional.

Timing

Figure 7–26 illustrates the arbitration timing. As this figure illustrates, a device can drive valid data in the granted cycle. Figure 7–26 shows the following sequence of events:

1. In cycle one, the arbiter grants a request. The granted device drives valid data in cycles one and two.
2. In cycle 4, the arbiter grants a request. The granted device drives valid data in cycles 4 and 5.
3. In cycle 6, the arbiter grants a request. The granted device drives valid data in cycles 6 and 7.
4. Cycle 3 is the only cycle that does not contain valid data.

Figure 7–26. Arbitration Timing



Interrupt Interfaces

In systems with interrupt sender interfaces, the Qsys interconnect includes several components to implement interrupt handling. Qsys handles individual, single-bit interrupt requests (IRQs). In the event that multiple senders assert their IRQs simultaneously, the receiver logic (typically under software control) determines which IRQ has highest priority, then responds appropriately.

Using individual requests, the interrupt logic can handle up to 32 IRQ inputs connected to each interrupt receiver. With this logic, the interrupt sender connected to `interrupt_receiver_0` is the highest priority with sequential receivers being successively lower priority. You can redefine the priority of interrupt senders by instantiating the Merlin IRQ mapper component. For more information refer to the “Merlin IRQ Mapper” on page 7–28.

Assigning IRQs in Qsys

You assign IRQ connections on the **System Contents** tab of Qsys. After adding all components to the system, you connect interrupt senders and receivers. You can use the **IRQ** column to specify an IRQ number with respect to each receiver or specify not to connect the IRQ.

- ❓ For more information, refer to *Connecting Qsys Components* in Quartus II Help.

Qsys uses the following four components to implement interrupt handling:

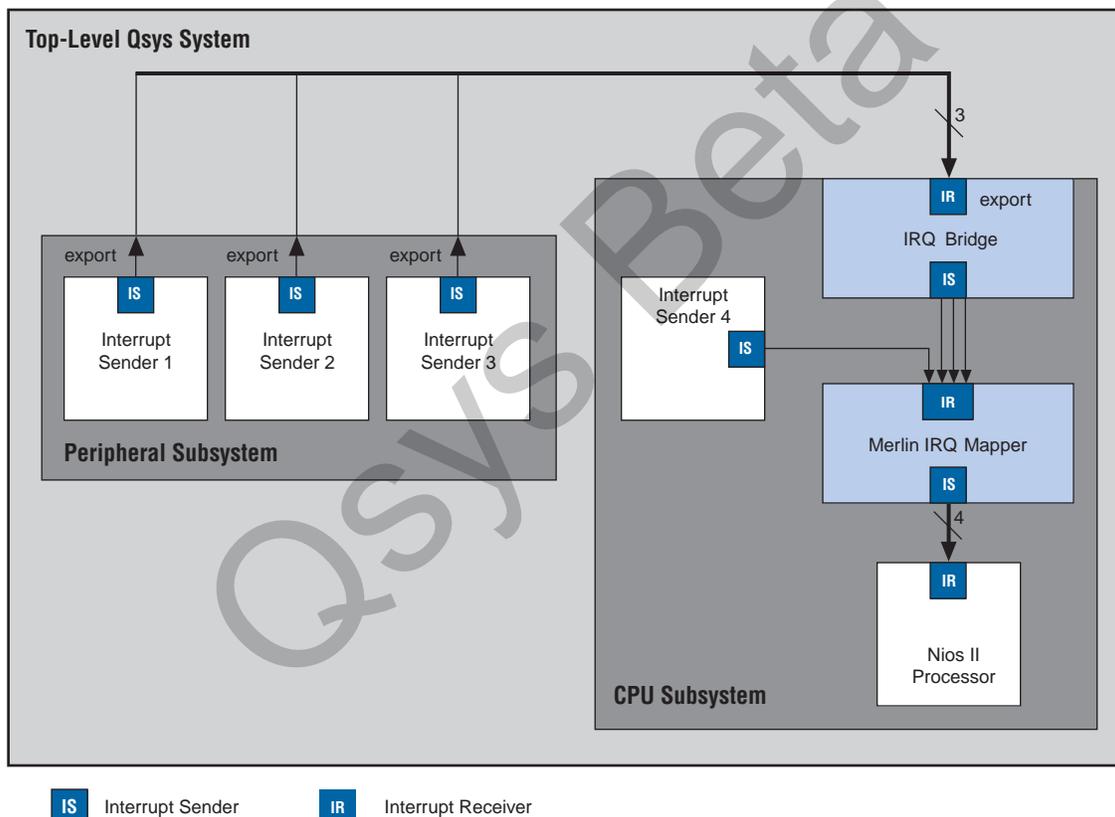
- IRQ Bridge
- Merlin IRQ Mapper
- Merlin IRQ Clock Crosser

The following sections describe these components.

IRQ Bridge

The IRQ bridge allows you to route interrupt wires between Qsys subsystems. In [Figure 7-27](#), the Peripheral Subsystem has three interrupt senders that are exported to the top level of the subsystem. These interrupts are routed to the Merlin IRQ receiver bridge in the CPU Subsystem.

Figure 7-27. Qsys IRQ Bridge Application



Merlin IRQ Mapper

The Merlin IRQ mapper converts individual interrupt wires into a bus. In addition, you can use the IRQ mapper to specify the interrupt number. By default, the interrupt sender connected to receiver0 interface of the IRQ mapper is highest priority with sequential receivers being successively lower priority. You can use the **IRQ Map** parameter in the parameter editor to remap the priority. For example, to reverse the priority of the four interrupt senders connected to the IRQ mapper in [Figure 7-27](#), you can type the following string for the **IRQ Map** parameter, 0:3, 1:2, 2:1, 0:3.

Merlin IRQ Clock Crosser

The Merlin IRQ clock crosser synchronizes interrupt senders and receivers that are in different clock domains. To use this component, connect the clocks for both the interrupt sender and receiver in addition to the interrupt sender and receiver interfaces. Qsys automatically inserts this component when it is required.

Clock Interfaces

You can use the **Clock Settings** tab to define external clock sources, for example an oscillator on your board. You can define separate reset sources for each clock domain, a single reset source for all clocks, or any combination in between.

Reset Interfaces

You can choose to have a single global reset domain generated by Qsys or, if your design requires more than one reset domain, you can implement your own reset logic and connectivity.

Single Global Reset Signal Implemented by Qsys

If you turn on **Global Reset** on the **Project Settings** tab in Qsys, the Qsys interconnect distributes a global reset bus. All of the reset requests are ORed together, synchronized to each clock domain, and fed to the reset inputs. The duration of the reset signal is at least one clock period.

The Qsys interconnect inserts the system-wide reset under the following conditions:

- The global reset input to the Qsys system is asserted.
- Any component asserts its `resetrequest` signal.

Multiple Reset Signals

The Qsys component library includes a reset controller and a reset bridge to implement the reset functionality. You can also design your own reset logic.



If you design your own reset circuitry you must carefully consider situations which might result in system lockup. For example, if an Avalon-MM slave is reset in the middle of a transaction, the Avalon-MM master might wait forever.

Merlin Reset Controller

If you design a system with multiple reset inputs, the Merlin reset controller, ORs all reset inputs and generates a single reset output. The reset controller has the following three parameters which you can specify to customize its behavior.

- **Number of reset inputs**—indicates the number of individual reset interfaces the controller ORs to create a signal reset output.

- **Output reset synchronous edges**—specifies the level of synchronization. You can select one the following options:
 - **None**—The reset is asserted and deasserted asynchronously. You can use this setting if you have designed internal synchronization circuitry.
 - **Both**—The reset is asserted and deasserted synchronously.
 - **Deassert**—The reset is deasserted synchronously and asserted asynchronously.
- **Synchronization depth**—specifies the number of register stages the synchronizer uses to eliminate the propagation of metastable events.

Qsys automatically inserts reset synchronizers under the following conditions:

- More than one reset source is connected to a reset sink
- There is a mismatch between the reset source's synchronous edges and the reset sinks' synchronous edges

Reset Bridge

The reset bridge allows you to use a reset signal in two or more subsystems of your Qsys system. You can connect one reset source to local components and export one or more to other subsystems as required. You to specify the number of reset outputs using the parameter editor.

Conduits

You can use the conduit interface type for interfaces that do not fit any of the interface types defined in the *Avalon Interface Specifications*. You can use conduit interfaces to group any arbitrary collection of signals. Like other interface types, you can export or connect conduit interfaces. The PCI Express link of the PCI Express IP core shown in [Figure 5-11 on page 5-15](#) is an example of the use of the conduit interface for export.

To connect two conduit interfaces inside Qsys, the following conditions must be met:

- The interfaces must match exactly with the same signal roles and widths.
- The interfaces must be the opposite directions.



Conduits connections are always point-to-point connections.

Summary: Qsys Interconnect Components

[Table 7-4](#) lists all of the Qsys components that implement the Qsys interconnect.

Table 7-4. Summary of Qsys Interconnect Components (Part 1 of 2)

Component Name	Typical Applications	
	Internal Qsys interconnect <i>(Note 1)</i>	User Designs
Avalon-MM Master and Slave Network Transformation		
Merlin Master Translator	✓	—
Merlin Master Agent	✓	—

Table 7-4. Summary of Qsys Interconnect Components (Part 2 of 2)

Component Name	Typical Applications	
	Internal Qsys interconnect <i>(Note 1)</i>	User Designs
Merlin Router	✓	—
Merlin Traffic Limiter	✓	—
Merlin Slave Translator	✓	—
Merlin Slave Agent	✓	—
Avalon-ST Components		
Avalon-ST Handshake Clock Crosser	✓	✓
Avalon-ST Pipeline Stage	✓	✓
Merlin Multiplexer	✓	✓
Merlin Demultiplexer	✓	✓
Bridges		
Clock Bridge	—	✓
Avalon-MM Clock Crossing Bridge (Qsys)	—	✓
Avalon-MM Pipeline Bridge (Qsys)	—	✓
Arbitration and Adapters		
Merlin Arbiter	✓	—
Merlin Width Adapter	✓	✓
Merlin Burst Adapter	✓	✓
Tristate Conduits		
Generic Tristate Controller	—	✓
Tristate Conduit Pin Sharer	—	✓
Tristate Conduit Bridge	—	✓
Interrupts		
IRQ Bridge	—	✓
Merlin IRQ Mapper	✓	—
Merlin IRQ Clock Crosser	✓	✓
Reset		
Merlin Reset Controller	✓	✓
Reset Bridge	—	✓

Note to Table 7-4:

- (1) These components are described to enhance your understanding of the Qsys interconnect. You probably will not need to use them in your own designs.
- (2) In this table, a ✓ means that the component is typically used for the purpose specified by the column header, a — means that the component is not typically used for the purpose specified by the column header.

Document Revision History

Table 7-5 shows the revision history for this document.

Table 7-5. Document Revision History

Date	Version	Changes
December 2010	10.1.0	Initial release.

- For previous versions of the Quartus II Handbook, refer to the [Quartus II Handbook Archive](#).
- Take an [online survey](#) to provide feedback about this handbook chapter.

Qsys Beta



Altera's Qsys system integration tool is now available as beta for evaluation in the Quartus® II software subscription edition version 10.1. Altera does not recommend using the beta release of Qsys in the Quartus II software version 10.1 for designs that are close to completion and are meeting design requirements. Before using Qsys, review the *Quartus II Software Version 10.1 Release Notes* and *AN 632: SOPC Builder to Qsys Migration Guidelines* for known issues and limitations. To submit general feedback or technical support on the beta release of Qsys, submit a service request through mysupport.altera.com. Alternatively, to submit general feedback, click **Feedback** on the Quartus II software Help menu.

You define Qsys components in the component editor by declaring their properties and behaviors or directly in a Hardware Component Description File (**_hw.tcl**). Each **_hw.tcl** file represents one component which you can add to an Qsys system. You can also share components with other designers. For your component to have maximum flexibility, you should consider what aspects of its behavior can be parameterized so that other users can change the default parameterization to address different design requirements.

An Qsys component is usually composed of the following four types of files:

- **_hw.tcl** file—describes the Qsys related characteristics, such as interface behaviors. This file is required.
- HDL files—define the component's functionality as hardware, simulation, and constraint files. These files are optional.
- **_sw.tcl**—used by the software build tools to compile the component driver code. This file is optional.
- Component driver files—defines the component register map and driver software to allow software to control the component. These files are optional.

This chapter discusses the following topics:

- "Information in a Hardware Component Description File"
- "Component Phases" on page 8–2
- "Writing a Hardware Component Description File" on page 8–3
- "Overriding Default Behaviors" on page 8–8
- "Hardware Tcl Command Reference" on page 8–14



An excellent source of information about Tcl syntax is the [Tcl Developer Xchange](#) website.

Information in a Hardware Component Description File

A typical **_hw.tcl** file contains the following information:

© 2010 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, HARDCOPY, MAX, MEGACORE, NIOS, QUARTUS and STRATIX are Reg. U.S. Pat. & Tm. Off. and/or trademarks of Altera Corporation in the U.S. and other countries. All other trademarks and service marks are the property of their respective holders as described at www.altera.com/common/legal.html. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

- Basic component information—includes the component’s name, version, and description, a link to its documentation, and pointers to HDL implementation files for synthesis and simulation.
- Parameter Declarations—Parameters are values that the user of your component can set that affect how the component is implemented, such as the size of a memory. Properties of each parameter include the parameter’s name, whether or not it is visible, and, if visible, the text to display when describing it. When the Qsys system is generated, the parameters can be applied to the component as Verilog HDL parameters or VHDL generics.
- Interface Properties—The interfaces of a component define how to connect it to the rest of the system and determine how other components in the system interact with it. When you add interfaces to a component, you declare which signals make up each interface. You also define interface properties, such as wait states for an Avalon® Memory-Mapped (Avalon-MM) interface.

Component Phases

The following section describes the distinct phases in the development of an Qsys component.

- Main Program—Qsys first discovers a component and adds it to the component library. The `_hw.tcl` file is executed and the Tcl statements provide non-instance-specific information to Qsys. During this phase, some component interfaces may be incompletely described and ports may have a width of 0 or -1 to indicate that they are variable.
- Validation—Validation allows the component to generate error, warning, or informational messages. Validation occurs when an instance of a component is created, when its parameters are changed, or when some other property of the system is changed.
- Elaboration—Elaboration occurs as Qsys queries a component for its interface information. Elaboration typically occurs immediately after validation and before generation. Interfaces defined in the main program can be enabled or disabled during elaboration. Depending on the validation callback code, elaboration and validation may alternate a few times. Elaboration and validation always occur before generation. Once elaboration is complete, the component must be completely described. For example, all port widths must have positive values.
- Generation—Generation creates all the information that the Quartus® II software and HDL simulator require. The required files typically include VHDL or Verilog HDL files, simulation models, and timing constraints.
- Editor—After an instance of your component has been added to an Qsys system, allows the user of your component to edit the GUI that displays the parameter editor. You can change the appearance of the default editor to make it easier to use. You also use the editor to instantiate a component.

- Compose—compose is not really a component phase in and of itself, but it overrides the default behavior for the validation, elaboration, and generation phases. Compose allows you to create hierarchical components, constructing new components from combinations of other components. You can include the Tcl commands to compose hierarchical components in the main program or in a compose callback.

Writing a Hardware Component Description File

This section provides detailed information about `_hw.tcl` files and describes the default behavior of a component in all phases. The following example uses a simple UART with some simple parameterization.

Providing Basic Information

A typical `_hw.tcl` file first declares basic information such as the name, location, and the files it includes. The first command in a `_hw.tcl` file should specify the version of the `_hw.tcl` API to use, with the following Tcl command:

```
package require -exact socp <version>
```

The version number is a Quartus II release version, such as 10.1. Qsys guarantees that a valid `_hw.tcl` file that requests a particular `socp` package behaves identically in future versions of the tool. Because of differences between versions of the Quartus II software, you cannot assume that an HDL file that functions correctly one `socp` package automatically functions correctly with other versions of the package.



This chapter describes the behavior of components that request the `socp 10.1` package.



An excellent source of information about Tcl syntax is the [Tcl Developer Xchange](#) website.

Example 8-1. Basic Information for `_hw.tcl` File

```
# The package command must be the first command in the file
package require -exact socp 10.1

# The name and VERSION of the component
set_module_property NAME example_uart
set_module_property VERSION 1.0

# The name of the component to display in the library
set_module_property DISPLAY_NAME "Example Component"

# The component's description.
set_module_property DESCRIPTION "An Example Component"

# The component library group that component belongs to
set_module_property GROUP Examples
```

Declaring Parameters

By including configuration parameters in your `_hw.tcl` file, you allow users of your component to parameterize it in different ways. Each parameter has a number of properties such as its name, type, display name, and default value that can be used to control how the parameter is displayed and used. [Example 8-2](#) illustrates the use of parameters that can be configured by users of your component.

Example 8-2. Declaring Parameters

```
# Declare Baud Rate parameter as an integer with a default value of 9600.
add_parameter BAUD_RATE int 9600

# Display this parameter as "Baud Rate" in the Parameter Editor.
set_parameter_property BAUD_RATE DISPLAY_NAME "Baud Rate (bps)"

# We only support three baud rates
set_parameter_property BAUD_RATE ALLOWED_RANGES {9600 19200 38400}
```

Parameters can be divided into three types: user parameters, system information parameters, and derived parameters. The following sections describe these parameter types.

User Parameters

User parameters are parameters that users have control over and that are exposed in the component parameter editor.

Derived Parameters

Derived parameters are parameters that are inferred by the component itself from user parameters or other derived parameters. For example, a clock period parameter can be derived from a data rate parameter. You can use derived parameters to perform operations that cannot be performed in HDL. For example, determining the number of address bits that a component requires using logarithmic functions is easy in Tcl and impossible in HDL.

SYSTEM_INFO Parameters

You can use `SYSTEM_INFO` parameter to request that certain parameter values are populated with information about the system. For example, you might want to know the frequency of the clock that ends up being connected to your clock input. When you declare `SYSTEM_INFO` properties, you provide an `<info-type>` and further arguments. The `<info-type>` is the type of information you want, such as `clock_rate`, and you use the additional arguments to specify things, such as which clock input interface you require. [Example 8-3](#) illustrates the use of the `SYSTEM_INFO` parameter. For more information about the `SYSTEM_INFO` parameter properties refer to [Table 8-5 on page 8-29](#).

Example 8-3. Syntax of Tcl Command using the SYSTEM_INFO Parameter

```
set_parameter_property my_parameter SYSTEM_INFO {<info-type> [<arg>]}
```

Declaring Interfaces

To declare an interface, use the `add_interface` command. Then use the `set_interface_property` and `add_interface_port` commands to set its properties and indicate which signals belong to it. The interface declaration statement includes the name of the interface, the interface direction, and the clock interface with which it is associated. For interfaces that are not associated with clocks (such as clock interfaces themselves), omit the associated clock interface, or use the word *asynchronous*.

[Example 8-4](#) illustrates interface declaration.

Example 8-4. Declare Interfaces

```
# Declare the clock sink interface, "clock_sink", type=clock, direction=sink
add_interface clock_sink clock sink

# The clock interface has two signals, named "clk" and "reset_n" of types "clk" "reset_n"
add_interface_port clock_sink clk clk input 1
add_interface_port clock_sink reset_n reset_n input 1

# Declare the Avalon slave interface, name=avalon_slave_0, type=avalon,
# direction=slave, associated with the clock_sink clock interface.
add_interface avalon_slave_0 avalon slave clock_sink

# Set a number of properties about the Avalon Slave interface
set_interface_property avalon_slave_0 writeWaitTime 0
set_interface_property avalon_slave_0 addressAlignment DYNAMIC
set_interface_property avalon_slave_0 readWaitTime 1
set_interface_property avalon_slave_0 readLatency 0

# Declare all the signals that belong to my Avalon Slave interface
add_interface_port avalon_slave_0 my_readdata readdata output 8
add_interface_port avalon_slave_0 my_read read input 1
add_interface_port avalon_slave_0 my_write write input 1
add_interface_port avalon_slave_0 my_waitrequest waitrequest output 1
add_interface_port avalon_slave_0 my_address address input 24
add_interface_port avalon_slave_0 my_writedata writedata input 8
```

Adding Files and Guiding Generation

Component description files typically provide all of the information required for generation and downstream tools, identifying the files used by the component such as HDL files. You also identify which of the added files is the top-level HDL file and specify which Verilog module or VHDL entity within that file is the top-level module for the component. [Example 8-5](#) illustrates the files that are typically required for generation and downstream tools.

Example 8-5. Add Files

```
# Add the HDL file to the component, to be used for synthesis and simulation.
add_file simple_uart.v {SYNTHESIS SIMULATION}

# Add the Timequest file with Quartus timing constraints.
add_file simple_uart.sdc SYNTHESIS

# Indicate which of the added HDL files holds the top-level module/entity
# that describes the component, name of the top-level module/entity
set_module_property TOP_LEVEL_HDL_FILE simple_uart.v
set_module_property TOP_LEVEL_HDL_MODULE simple_uart
```

Default Behaviors

The `_hw.tcl` file described in the previous section has default behaviors during the editor, validation, elaboration, and generation phases. These default behaviors apply to instances of a component. This section describes the default Qsys behaviors for each of these phases. To override these default behaviors, refer to “[Overriding Default Behaviors](#)” on page 8-8.

Validation Phase Behavior

The default Qsys validation checks each parameter value against its `ALLOWED_RANGES` property. If the values specified are outside the allowed ranges, an error message is displayed.

The `ALLOWED_RANGES` property of each parameter is a list of ranges that the parameter can take on, where each range is a single value, or a range of values defined by a start and end value separated by a colon. [Table 8-1](#) shows some examples of values the `ALLOWED_RANGES` property can take.

Table 8-1. ALLOWED_RANGES Property

ALLOWED_RANGES	Meaning
{a b c}	a or b or c
{1 2 4 8 16}	1, 2, 4, 8, or 16.
1:3	1 through 3, inclusive
{1 2 3 7:10}	1, 2, 3, or 7 through 10 inclusive

Elaboration Phase Behavior

If the main program does not explicitly define the widths of all ports to constant values or to an expression, then default Qsys elaboration process calls `quartus_map` to determine the correct port widths. If you define all port widths in the main program, `quartus_map` is not called.

Automatic Port Widths

When port widths are not specified, or have a value of '-1', `quartus_map` determines port widths as a function of the parameter set. While this process makes authoring a component easier, it slows component generation. When using automatic port widths, you can indicate that a certain parameter does not affect any port widths or interfaces by setting that parameter's `affects_elaboration` property to `false`, meaning that `quartus_map` is not called when the parameter's value is changed by your user. However, indicating that a parameter does not affect elaboration when it actually does can lead to problems that are difficult to debug.

As an alternative to the automatic port widths, you can set port widths to simple HDL expressions using the `width_expr` property. `width_expr` is a string that holds an expression describing the port width. By using the `width_expr` property, you can define port widths as an expression that is evaluated without needing to analyze the HDL file or set them in an elaboration callback. The syntax for width expressions is the same as the HDL language that you use; however, only the addition, subtraction,

multiplication, and division operators are allowed. For more complex port widths, the width of the port can be set as an arbitrary function of the component's parameters in an elaboration callback. The width expression is the last argument to the `add_interface_port` command. [Example 8-6](#) illustrates the use of mathematical operators and the `width_expr` property.

Example 8-6. Defining Port Widths Using Simple Mathematical Operators

```
add_interface_port din din_data data input {WIDTH * SYMBOLS}  
set_port_property din_data width_expr WIDTH
```

Parameterized Parameter Widths

For VHDL users, Qsys allows a `std_logic_vector` parameter to have a width that is defined by another parameter. When adding a parameter of type `std_logic_vector` you can also specify its width as a parameter property. The width can be a constant or the name of another parameter. The commands [Example 8-7](#) add a `std_logic_vector` parameter called `myParameter` whose width is set by another parameter, called `dataWidth`.

Example 8-7. Adding Parameters

```
add_parameter myParameter STD_LOGIC_VECTOR  
set_parameter_property myParameter WIDTH dataWidth
```

Generation Phase Behavior

The default Qsys generation does one of the following:

- If the component defines the `TOP_LEVEL_HDL_MODULE` property, Qsys creates a Verilog HDL or VHDL wrapper module to instantiate the top-level module and applies the parameters as selected by the user of your component. Qsys does not apply parameters in the wrapper if they are not declared in the underlying HDL file.
- or*
- If the component does not define the `TOP_LEVEL_HDL_MODULE` property, but instead sets the `INSTANTIATE_IN_SYSTEM_MODULE` module property to `false`, the module is not instantiated inside the Qsys system and a wrapper file is not created. Rather, the interface to the module is exported to the top-level of the Qsys system, and the module must be connected outside the system.

Edit Phase Behavior

The default Qsys editor phase behavior is to use all of the parameter definitions to display the parameter editor. The properties of the parameters guide Qsys when it builds the default parameter editor. [Table 8-4 on page 8-26](#) lists the properties of parameters.

You can place parameters in logical groups and provide images and text to create a custom parameter editor for your component. [Example 8-8](#) defines four parameters and illustrates the use of the `add_display_item` command and the `DISPLAY_HINT` and `ALLOWED_RANGES` parameters.

Example 8-8. Defining and Customizing the parameter editor

```
# provide an icon for the sound group
add_display_item icon Speaker speaker-image speaker.png
add_parameter sound string 0 0
add_parameter volume_control boolean 0 0
add_parameter separate_control string 0 0

# Setup display_names for the parameters
set_parameter_property sound DISPLAY_NAME Audio
set_parameter_property volume_control DISPLAY_NAME "Include Volume Control Interface"
set_parameter_property separate_control DISPLAY_NAME "Treble/Bass Controls"

# Display all parameters in the Speaker group
add_display_item Speaker sound parameter
add_display_item Speaker volume_control parameter
add_display_item Speaker separate_control parameter

# There are 4 choices for the sound parameter.
# Strings with internal spaces require double quotes
set_parameter_property sound ALLOWED_RANGES {"0:No Audio" 1:Monophonic 2:Stereo
4:Quadraphonic}
set_parameter_property separate_control ALLOWED_RANGES {"No Control" "Single Control" "Dual
Controls"

#Specify how parameters should be displayed
set_parameter_property volume_control DISPLAY_HINT boolean
set_parameter_property separate_control DISPLAY_HINT radio
```

[Figure 8-1](#) shows the parameter editor that the Tcl commands in [Example 8-8](#) produces.

Figure 8-1. parameter editor for Audio Component



Overriding Default Behaviors

You can override each of the default behaviors by using callbacks. This section explains how to write callback procedures for each phase of component development.

Validation Callback

You can use the validation callback to provide validation that extends beyond the default range checking. A validation callback is defined by setting the `VALIDATION_CALLBACK` module property to be the name of the validation callback procedure, as shown in [Example 8-9](#). This validation procedure displays an error if you select a baud rate of 38400 and odd parity.

You can also use the validation callback to set the value of derived parameters. Derived parameters are parameters that are derived from other parameters; their values are not editable and are not saved in the Qsys System File (`.qsys`). You indicate that a parameter is derived by setting the parameter's `DERIVED` property to true. In [Example 8-9](#) `BAUDRATE_PRESCALE` is a derived parameter whose value is 1/16 of the value of the `BAUDRATE` parameter.

Example 8-9. Custom Validation Callback Function

```
# Declare the validation callback.
set_module_property VALIDATION_CALLBACK my_validation_callback

# Add the BAUDRATE_PRESCALE parameter, and indicate that it's derived
add_parameter BAUDRATE_PRESCALE int 600
set_parameter_property BAUDRATE_PRESCALE DERIVED true

# Add the PARITY parameter
add_parameter PARITY string ODD
set_parameter_property PARITY ALLOWED_RANGES {EVEN ODD}

# The validation callback
proc my_validation_callback {} {
    # Get the current value of parameters we care about
    set br [get_parameter_value BAUD_RATE]
    set p [get_parameter_value PARITY]
    # Display an error for invalid combinations.
    if {($br==38400) && ($p=="ODD")} {
        send_message warning "Odd parity at 38400 bps is not supported."
    }
    # Set the value of our DERIVED parameter
    set bp [expr $br / 16]
    set_parameter_value BAUDRATE_PRESCALE $bp
}
```

Elaboration Callback

You can use an elaboration callback to change interface properties or add new interfaces as a function of parameter values. You define an elaboration callback by setting the `ELABORATION_CALLBACK` module property to the name of the elaboration callback function, as shown in [Example 8-10](#). You can enable and disable interfaces from the elaboration callback if they are only needed for some parameterizations of the component. [Example 8-10](#) shows how an Avalon-MM slave interface can be included in an instance of the component, based on the `USE_STATUS_INTERFACE` parameter. All of the functionality available in the validation callback can also be used in the elaboration callback; separate callbacks for validation and elaboration are not required.



The elaboration callback is not be called when parameters with `AFFECTS_ELABORATION=false` are changed by the user of the component.

Example 8-10. Elaboration Callback

```
# Declare the callback.
set_module_property ELABORATION_CALLBACK my_elaboration_callback

# Add the USE_STATUS_INTERFACE parameter
add_parameter USE_STATUS_INTERFACE boolean

# Declare the status slave interface
add_interface_status_slave avalon_slave clock_sink
set_interface_property status_slave ENABLED false

# The elaboration callback
# Declare signals
add_interface_port status_slave st_readdata readdata output 16
add_interface_port status_slave st_read read input 1
add_interface_port status_slave st_write write input 1
add_interface_port status_slave st_waitrequest waitrequest output 1
add_interface_port status_slave st_address address input 24
add_interface_port status_slave st_writedata writedata input 16

# The elaboration callback
proc my_elaboration_callback {} {

    # Get the current value of parameters we care about
    set use_status [get_parameter_value USE_STATUS_INTERFACE]

    # Optionally add the status interface
    if { $use_status } {
        set_interface_property status_slave ENABLED true
    }
}
```

Generation Callback

If you define a generation callback, Qsys does not generate an HDL wrapper file to apply parameter values to your component. Instead, it calls the generation callback you defined during the generation phase, allowing the component to programmatically generate its HDL. A generation callback is defined by setting the `GENERATION_CALLBACK` module property to be the name of the generation callback function, as [Example 8-11](#) illustrates.

Generation callbacks typically retrieve the current value of the component's parameters and the generation properties that guide the generation process, and then generate the HDL files and supporting files in Tcl or by calling an external program. The callback procedure also reports the required files to Qsys with the `add_file` command. Any files added in the generation callback are in addition to the files added in the main body of the `_hw.tcl` file.

The generation callback must write `<output_name>.v` or `.sv` for Verilog or `<output_name.vhd>` for VHDL to the specified `<output_directory>`. This file is a parameterized instance of the component. Other supporting files, such as `.hex` files to initialize memory, may be written to `<output_directory>`. These file names must begin with `<output_name>`. If the supporting files are the same for all parameterizations of the component, you add them from the main program rather than the generation

callback. If your system includes multiple instantiations of a component with different parameterizations, you must add the supporting files from the main program to prevent failures. If a static supporting file is only needed in some parameterizations of the component, you should add it from the main program and turn it on or off by setting its `SYNTHESIS` and `SIMULATION` properties appropriately from the elaboration callback.

Example 8-11. Generation Callback Example

```
set_module_property GENERATION_CALLBACK my_generate
# My generation method
proc my_generate {} {
    send_message info "Starting Generation"

    # get generation settings

    set language [get_generation_property HDL_LANGUAGE]
    set outdir [get_generation_property OUTPUT_DIRECTORY ]
    set outputname [get_generation_property OUTPUT_NAME ]

    # get parameter values

    set p1 [get_parameter_value PARAMETER_ONE]
    set csr [get_parameter_value CSR_ENABLED]

# Your callback needs to write $outdir$outputname.v here,
# perhaps by using exec to call an external program.

    # add_file creates files relative to the _hw.tcl directory; therefore specify $outdir
    # for synthesis and simulation files

    exec perl my_generate.pl lang=$language dir=$outdir name=$outputname p1=$p1 csr=$csr
    add_file ${outdir}${outputname}.v SYNTHESIS
    add_file ${outdir}${outputname}_sim.v SIMULATION
}
```

Compose Callback

You can use a compose callback to define components that are constructed from combinations of other components. Compose can be used in one of two ways:

- You can use compose commands such as `add_instance`, `set_instance_parameter_value`, and `add_connection` in the main program to create and parameterize subcomponent instances.
- Or, after you have set up the basic component template in the main program, you can then use a compose callback to instantiate and parameterize subcomponents as a function of the component's parameter values. You define a compose callback by setting the `COMPOSE_CALLBACK` module property to the name of the compose callback function.

When used, compose replaces elaboration and generation. Your component's interface information is collected by analyzing the interfaces on exported subcomponents. HDL is generated by generating all of your subcomponents and a top-level that stitches them all together.

Exporting an interface means that you are making the interface visible from the outside of your component, instead of connecting it internally. Set the `EXPORT_OF` property of the externally visible interface to indicate that it is an exported view of the submodule's interface. Refer to “[get_interface_properties](#)” on page 8-38 for the format of the `EXPORT_OF` property. You can set this from the main program or the compose callback.

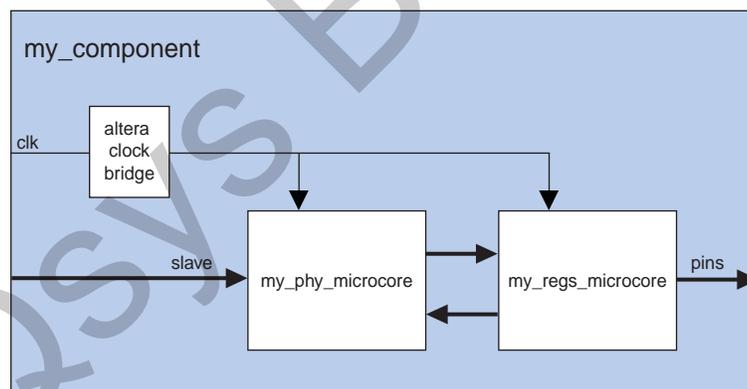
Exporting an interface is different than connecting two interfaces together—the exported interface is a copy of the subcomponent's interface. For example, if the inner interface is a 32-bit Avalon-MM master without bursting then the exported interface will be as well.

 Because the exported interface is a copy of the inner interface, no adaptation is possible between the two interfaces.

When you create an exported interface, the properties of the exported interface are copied from the subcomponent's interface without modification. Ports are copied from the subcomponents interface with only one modification—the names of the exported ports on the composed component are chosen to ensure they are unique.

[Figure 8-1](#) is a block diagram for the composed component that is shown [Example 8-1](#).

Figure 8-1. Top-Level of a Composed Component



[Example 8-1](#) provides an example of a composed `_hw.tcl` file which instantiates two subcomponents. It connects them together, also connecting the clocks and resets. Note that a clock bridge component is required to allow both subcomponents to see a common clock input.

Example 8-1. Composed Component

```
package require -exact socp 10.1
set_module_property name my_component
...
add_interface clk clock end
set_interface_property clk EXPORT_OF clk.in_clk

add_interface reset reset end
set_interface_property reset EXPORT_OF reset.in_reset

add_interface pins conduit end
set_interface_property pins EXPORT_OF phy.pins

add_interface slave avalon slave
set_interface_property slave EXPORT_OF regs.slave

add_instance clk altera_clock_bridge
add_instance reset altera_reset_bridge
set_instance_property_value reset synchronous_edges deassert
add_connection clk.out_clk reset.clk

add_instance phy my_phy_microcore
add_connection clk.out_clk phy.clk
add_connection reset.out_reset phy.clk_reset

add_instance regs my_regs_microcore
add_connection clk.out_clk regs.clk
add_connection reset.out_reset regs.reset
add_connection phy.output regs.input
add_connection regs.output phy.input
```

Editor Callback

You can use the editor callback procedure to replace the parameter editor to make it easier to use. An editor callback is defined by setting the `EDITOR_CALLBACK` module property to the name of your editor callback procedure, as shown in the [Example 8-12](#). If the editor callback is defined, Qsys calls the editor callback instead of displaying the parameter editor, typically when the component is added to a system or updated after it is in the system.

To display your custom parameter editor, the editor callback must call another program. Typically, an editor callback provides the current parameter values to your program via the command line and collects the new parameter values via `stdout`. The editor callback then uses the `set_parameter_value` command to update Qsys with the new parameter values.

The editor callback returns one of the following three values:

- `OK`—indicates that the results of the edit should be applied.
- `CANCEL`—indicates that the system should revert to the state it was in before the editor callback was called.
- `ERROR`—indicates that the parameter editor was unable to launch. An appropriate error message should be displayed.

If no value is returned, OK is assumed.

Example 8-12. Editor Callback

```
set_module_property EDITOR_CALLBACK my_editor

# Define Module parameters.
add_parameter PARAMETER_ONE integer 32 "A parameter"
add_parameter CSR_ENABLED boolean true "Enable CSR interface"

# My editor method
proc my_editor {} {
# get parameter values
  set p1 [ get_parameter_value PARAMETER_ONE ]
  set csr [ get_parameter_value CSR_ENABLED ]

# Display UI, populated with current parameter values.
# The stdout returned by the UI program includes the new paramter values.
  set result [exec my_component_ui.exe p1=$p1 csr=$csr]

# Use the fictional "parse_for_new_value" procedure to parse the returned text for the
# new parameter values.
  set p1 [parse_for_new_value $result p1]
  set csr [parse_for_new_value $result csr]

# Return the new parameter values to Qsys
  set_parameter_value PARAMETER_ONE $p1
  set_parameter_value CSR_ENABLED $csr
  return OK
}
```

Hardware Tcl Command Reference

This section provides a reference for all hardware Tcl commands, as follows:

- [“Module Definition” on page 8-17](#)
- [“Parameters” on page 8-24](#)
- [“Display Items” on page 8-33](#)
- [“Interfaces and Ports” on page 8-36](#)
- [“Compose” on page 8-43](#)
- [“Generation” on page 8-49](#)

The description of each command indicates during which phases it is available: in the main body of the program (main), or during the validation, elaboration, compose, generation, and editor callback phases, or any combination. [Table 8-2](#) summarizes the commands and provides a reference to the full description.

 All Tcl commands that you can use in the validation callback are also available in the elaboration callback. You may be able to omit the custom validation callback by including some validation commands in your elaboration callback.

Table 8-2. Command Summary (Note 1) (Part 1 of 2)

Command	Full Description
Module Definition	
package <require> -exact socp <version>	page 8-17
get_module_properties	page 8-17
get_module_property <propertyName>	page 8-19
set_module_property <propertyName> <propertyValue>	page 8-19
get_module_ports	page 8-20
get_module_assignments	page 8-20
get_module_assignment <moduleName>	page 8-21
set_module_assignment <moduleName> [value]	page 8-21
get_files	page 8-21
add_file filename [<fileProperties> . . .]	page 8-21
add_documentation_link <docType> <title> <fileOrUrl>	page 8-22
get_file_properties	page 8-22
get_file_property <filename> <propertyName>	page 8-22
set_file_property <filename> <propertyName> <propertyValue>	page 8-23
send_message <messageLevel> <messageText>	page 8-23
Parameters	
add_parameter <parameterName> <parameterType> [<defaultValue> <description>]	page 8-24
get_parameters	page 8-25
get_parameter_properties	page 8-25
get_parameter_property <parameterName> <propertyName>	page 8-30
set_parameter_property <parameterName> <propertyName> <value>	page 8-30
get_parameter_value <parameterName>	page 8-31
set_parameter_value <parameterName> <value>	page 8-31
decode_address_map <address_map_XML_string>	page 8-32
Display Items	
add_display_item <groupName> <id> <type> [<additionalInfo>]	page 8-33
get_display_items	page 8-34
get_display_item_properties	page 8-35
get_display_item_property <itemName> <propertyName>	page 8-35
set_display_item_property <itemName> <propertyName> <value>	page 8-35
Interfaces and Ports	
add_interface <interfaceName> <interfaceType> <direction> [<associatedClock>]	page 8-37
get_interfaces <interfaceName>	page 8-37

Table 8-2. Command Summary (Note 1) (Part 2 of 2)

Command	Full Description
<code>get_interface_property <interfaceName> <propertyName></code>	page 8-38
<code>set_interface_property <interfaceName> <propertyName> <value></code>	page 8-39
<code>add_interface_port <interfaceName> <portName> <portRole> [<direction> <width_expr>]</code>	page 8-39
<code>get_interface_ports [<interfaceName>]</code>	page 8-40
<code>get_port_properties</code>	page 8-40
<code>get_port_property <portName> <propertyName></code>	page 8-41
<code>set_port_property <portName> <propertyName> [<value>]</code>	page 8-42
<code>get_interface_assignments</code>	page 8-42
<code>get_interface_assignment <interfaceName> <name></code>	page 8-42
<code>set_interface_assignmet <interfaceName> <name> [<value>]</code>	page 8-43
Compose	
<code>add_instance <instanceName> <instanceType> <version></code>	page 8-43
<code>get_instances</code>	page 8-44
<code>get_instance_parameters <instanceName></code>	page 8-44
<code>set_instance_parameter <instanceName> <parameterName> <parameterValue></code>	page 8-44
<code>get_instance_parameter_value <instanceName> <parameterName></code>	page 8-44
<code>get_instance_parameter_properties <instanceName> <parameterName></code>	page 8-45
<code>get_instance_parameter_property <instanceName> <parameterName> <propertyName></code>	page 8-45
<code>get_instance_interfaces <instanceName></code>	page 8-46
<code>get_instance_interface_properties <instanceName> <interfaceName></code>	page 8-46
<code>get_instance_interface_property <instanceName> <interfaceName> <propertyName></code>	page 8-46
<code>get_instance_interface_ports <instanceName> <portName></code>	page 8-47
<code>get_instance_port_property <instanceName> <interfaceName> <propertyName></code>	page 8-47
<code>add_connection [<instanceName>] <startInterface> <endInterface></code>	page 8-47
<code>get_connections</code>	page 8-48
<code>get_connection_parameters <instanceName></code>	page 8-48
<code>get_connection_parameter <connectionName> <parameterName></code>	page 8-49
<code>set_connection_parameter_value <connectionName> <parameterName> <parameterValue></code>	page 8-49
Generation	
<code>get_generation_properties</code>	page 8-49
<code>get_generation_property <propertyName></code>	page 8-50

Note to Table 8-2:

(1) Arguments enclosed in []'s are optional

Module Definition

This section provides information about the commands that you use to define and query a module.

package

The `package` command allows you to specify a particular version of the Qsys software to avoid software compatibility issues. You should use the `package` command at the beginning of your `_hw.tcl` file. When used, the component files behave as if they are interpreted by the version of the Qsys software that you specify. When the `package` command is not used, installed version of the Qsys software is assumed. For components designed before 9.0, you can set the required package to 9.0. This document describes the behavior of component which start with `package require -exact socp 10.1`. For earlier releases, refer to the documentation for that release.

 `package` is a standard Tcl command. For more information on this command refer to the following [Package](#) page of the Altera website.

package	
Callback availability	Main (before any other commands in the file)
Usage	<code>package require -exact socp <version></code>
Returns	None
Arguments	<code>version</code> The version of Qsys that you require, specified as decimal number
Example	<code>package require -exact socp 10.0</code>

get_module_properties

This command returns the names of all the available module properties as a list of strings. You can use the `get_module_property` and `set_module_property` commands to get and set values of individual properties. The value returned by this command is always the same for a particular version of Qsys.

get_module_properties	
Callback availability	Main, validation, elaboration, generation, compose, and editor
Usage	<code>get_module_properties</code>
Returns	List of strings
Arguments	None
Example	<code>get_module_properties</code>

Table 8-3 lists the available module properties, their use, and the phases in which they can be set.

Table 8-3. Module Properties (Part 1 of 2)

Property Name	Property Type	Can Be Set	Description
ANALYZE_HDL	Boolean	Main program	When set to <code>false</code> , prevents a call to the Quartus II mapper to verify port widths and directions, speeding up generation time at the expense of fewer validation checks. If this property is set to <code>false</code> , invalid port widths and directions are discovered during Quartus II compilation.
AUTHOR	String	Main program	The module's author.
DESCRIPTION	String	Main program	The description of the module, such as "Example Qsys Module."
DISPLAY_NAME	String	Main program	The name to display when referencing the module, such as "My SOPC Component."
EDITABLE	Boolean	Main program	Indicates if the component is editable in the component editor.
EDITOR_CALLBACK	String	Main program	The name of the editor callback. The default parameterization UI is displayed if this property is not set.
ELABORATION_CALLBACK	String	Main program	The name of the elaboration callback. For static and generated components, the default elaborations used if this property is not set.
GENERATION_CALLBACK	String	Main program	The name of the generation callback.
GROUP	String	Main program	The component group that the module belongs to, such as "Example Components."
HIDE_FROM_QSYS	Boolean	Main program	When set to <code>true</code> , the component is not visible in the component library.
ICON_PATH	String	Main program	A path to an icon to display in the module's parameter editor.
INSTANTIATE_IN_SYSTEM_MODULE	Boolean	Main program	When <code>false</code> the instances of the module are not included in the generated system interconnect fabric. Instead, interfaces to the module are exported out of the top-level of the Qsys system.
INTERNAL	Boolean	Main program	A component which is marked as internal does not appear in the Qsys component library. This feature allows you to hide the submodules of a larger composed component.
MODULE_DIRECTORY	String	Can only be read, not set	The directory containing the <code>_hw.tcl</code> file. All relative file names within the Tcl file are resolved relative to this directory. This directory is set as the current directory when running the main program or a callback.
MODULE_TCL_FILE	String	Can only be read, not set	The path to the <code>_hw.tcl</code> file.

Table 8-3. Module Properties (Part 2 of 2)

Property Name	Property Type	Can Be Set	Description
NAME	String	Main program	The name of the module, such as my_soc_component.
TOP_LEVEL_HDL_FILE	String	Main program	Indicates which of the files added by the add_file command contains the module's top-level HDL.
TOP_LEVEL_HDL_MODULE	String	Main program	Indicates the name of the top-level module which must be defined in the module's top-level HDL file.
VALIDATION_CALLBACK	String	Main program	The name of the validation callback. This callback is run in addition to the default validation.
VERSION	String	Main program	The module's version, such as 10.0
COMPOSE_CALLBACK	String	Main Program	The name of the compose callback. If you define a compose callback then you must not define the generation or elaboration callbacks.

 The INSTANTIATE_IN_SYSTEM_MODULE, TOP_LEVEL_HDL_MODULE and GENERATION_CALLBACK commands are used to select the type of generation used by the component. You must set only one of these in the main program of your file.

get_module_property

This command returns the value of a single module property.

get_module_property	
Callback availability	Main, validation, elaboration, generation, compose, and editor
Usage	get_module_property <propertyName>
Returns	String, boolean, or file
Arguments	propertyName One of the properties listed in Table 8-3 on page 8-18
Example	set my_name [get_module_property NAME]

set_module_property

This command allows you to set the values for module properties.

set_module_property	
Callback availability	Main program
Usage	set_module_property <propertyName> <propertyValue>
Returns	None
Arguments	propertyName One of the properties listed in Table 8-3 on page 8-18
	propertyValue The new value of the property
Example	set_module_property VERSION 10.0

get_module_ports

This command returns a list of the names of all the ports which are currently defined.

get_module_ports	
Callback availability	Main, validation, elaboration, generation, and editor
Usage	get_module_ports
Returns	String
Arguments	None
Example	get_module_ports

get_module_assignments

This command returns names of the module assignment variables.

get_module_assignments	
Callback availability	Main, validation, elaboration, and compose
Usage	get_module_assignments
Returns	String
Arguments	None
Example	get_module_assignments

get_module_assignment

This command returns the value of the specified argument. You can use the `get_module_assignment` and `set_module_assignment` and the `get_interface_assignment` and `set_interface_assignment` commands to transfer information about hardware components to embedded software tools and applications.

get_module_assignment	
Callback availability	Main, validation, elaboration, and compose
Usage	get_module_assignment <name>
Returns	String
Arguments	name The name whose value is being retrieved
Example	get_module_assignment embedded.sw.CMacro.colorSpace

 For more information about specifying information for software tools, refer to *Publishing Component Information to Embedded Software* in the *Nios II Software Developer's Handbook*.

set_module_assignment

This command sets the value of the specified argument.

set_module_assignment		
Callback availability	Main, validation, elaboration, and compose	
Usage	set_module_assignment <name> [<value>]	
Returns	None	
Arguments	name	The name whose value is being set
	value	The value of the <name> argument
Example	set_module_assignment embedded.sw.CMacro.colorSpace CMYK	

get_files

This command returns a list of all the files that have been added to the module.

get_files	
Callback availability	Main, validation, elaboration, generation, and editor
Usage	get_files
Returns	List of strings
Arguments	None
Example	set list_of_files [get_files]

add_file

This command adds a synthesis, simulation, or TimeQuest constraints file to the module. Files added in the main program cannot be removed. Adding files in the generation callback allows the included files to be a function of the parameter set or to be a result of generation. Files added in callbacks are in addition to any files added in the main program.

add_file		
Callback availability	Main, elaborate, and generation	
Usage	add_file filename [<fileProperties> . . .]	
Returns	String	
Arguments	filename	The file name to be added, relative to the directory containing the <code>_hw.tcl</code> file
	fileProperties	Files support the following 3 properties: <ul style="list-style-type: none"> ■ SIMULATION—File for simulation ■ SYNTHESIS—File for synthesis ■ SDC—TimeQuest constraints (SDC behaves like a synthesis file)
Example	add_file my_component.v {SIMULATION SYNTHESIS}	

add_documentation_link

This command allows you to add multiple documentation links for a single component.

add_documentation_link		
Callback availability	Main	
Usage	<code>add_documentation_link filename <docType> <title> <fileOrUrl></code>	
Returns	None	
Arguments	docType	One of the following document types: USER_GUIDE, RELEASE_NOTES, WEBLINK, ERRATA, DATASHEET, REFERENCE_MANUAL, WAVEFORM, SCHEMATICS, TUTORIAL, OTHER
	title	The title of the document for use on menus and buttons.
	fileOrUrl	A path to the component documentation, using a syntax that provides the entire URL, not a relative path. For example: http://www.mydomain.com/my_memory_controller.html or <code>file:///datasheet.txt</code> .
Example	<code>add_documentation_link USER_GUIDE "Avalon Verification IP Suite User Guide" http://www.altera.com/literature/ug/ug_avalon_verification_ip.pdf</code>	

get_file_properties

This command returns the list of all properties that have been defined for a file.

get_file_properties	
Callback availability	Main, validation, elaboration, generation, compose, and editor
Usage	<code>get_file_properties</code>
Returns	List of strings
Arguments	None
Example	<code>get_file_properties</code>

get_file_property

This command returns the value of a single file property. The file name passed as an argument may be a partial as long as it is unique. For example, if the full file name is `/components/my_file.v`, `my_file.v` is sufficient.

get_file_property		
Callback availability	Main, validation, elaboration, generation, and editor	
Usage	<code>get_file_property <filename> <propertyName></code>	
Returns	Boolean	
Arguments	filename	The file name whose properties are being retrieved
	propertyName	The file name property whose value is being retrieved
Example	<code>set forSynthesis [get_file_property my_file.v SYNTHESIS]</code>	

set_file_property

This command sets the value of a single file property. The file name passed to the function can be a partial file name as long as it is unique. For example, if the full file name is `/components/my_file.v`, `my_file.v` is sufficient. The available properties are described in the `add_files` command.

set_file_property		
Callback availability	Main, elaboration, and generation	
Usage	<code>set_file_property <filename> <propertyName> <propertyValue></code>	
Returns	Boolean	
Arguments	<code>filename</code>	The file name whose properties are being retrieved
	<code>propertyName</code>	Name of the file property whose value is being retrieved
	<code>propertyValue</code>	Value to set for the file property
Example	<code>set_file_property my_file.v SYNTHESIS true</code>	

send_message

This command sends a message to the user of the component. The message text is normally interpreted as HTML. The `` element can be used to provide emphasis. If you do not want the message text to be interpreted as HTML then pass a list like `{ info text }` as the message level.

send_message		
Callback availability	Main, validation, elaboration, generation, compose, and editor	
Usage	<code>send_message <messageLevel> <messageText></code>	
Returns	None	
Arguments	<code>messageLevel</code>	The following 4 message levels are supported: <ul style="list-style-type: none"> ■ Error—provides an error message. The Qsys system cannot be generated while there are error messages. ■ Warning—provides a warning message. ■ Info—provides an informational message. ■ Debug—provides messages when debug mode is enabled.
	<code>messageText</code>	The text of the message
Example	<code>send_message Error "param1 must be greater than param2."</code>	

Parameters

Parameters allow users of your component to affect its operation in the same manner as Verilog HDL parameters or VHDL generics.

add_parameter

This command adds a parameter to your component. Most of the parameter types are self-explanatory because they are used in the C programming language or HDL. However, the `string_list` and `integer_list` parameters that are used to create tables in GUIs require some explanation.

- When you use the `add_parameter` command with a `string_list` or `integer_list` parameter type, the parameter you define is displayed in a variable-sized table that includes **add** and **remove** buttons.
- If you define multiple parameters of type `string_list` or `integer_list`, you can also use the `add_display_item` command to specify that parameters should each be displayed as a column in a table, each parameter of type `string_list` or `integer_list` becomes a column in the table. [Example 8-13](#) illustrates the use of the `integer_list` parameter types to create a multi-column table.

Example 8-13. Creating Tables Using the `string_list` and `integer_list` Parameter Types

```
add_parameter bitsWide INTEGER
add_parameter divider INTEGER
add_parameter coefficients INTEGER_LIST
add_parameter positions INTEGER_LIST
add_display_item myTable coefficients TABLE
add_display_item myTable positions TABLE
```

add_parameter		
Callback availability	Main program	
Usage	<code>add_parameter <parameterName> <parameterType> [<defaultValue> <description>]</code>	
Returns	String	
Arguments	<code>parameterName</code>	A name that you, the component author, choose for your parameter
	<code>parameterType</code>	The following types are supported: <code>Integer</code> , <code>Natural</code> , <code>Positive</code> , <code>Boolean</code> , <code>Std_logic</code> , <code>Std_logic_vector</code> , <code>String</code> , <code>String_list</code> , and <code>Integer_list</code> .
	<code>defaultValue</code>	The default length of the parameter is derived from its range.
	<code>description</code>	Explains the use of the parameter
Example	<code>add_parameter seed integer 17 "The seed to use for data generation."</code>	

get_parameters

This command returns the names of all parameters that have been previously defined by `add_parameter` as a space separated list.

get_parameters	
Callback availability	Main, validation, elaboration, generation, compose, and editor
Usage	<code>get_parameters</code>
Returns	List of strings
Arguments	None
Example	<code>set parameter_summary [get_parameters]</code>

get_parameter_properties

This command returns a list of all the available parameter properties as a list of strings. The `get_parameter_property` and `set_parameter_property` commands are used to get and set the values of these properties, respectively.

get_parameter_properties	
Callback availability	Main, validation, elaboration, generation, compose, and editor
Usage	<code>get_parameter_properties</code>
Returns	List of strings
Arguments	None
Example	<code>set property_summary [get_parameter_properties]</code>

Table 8-4 describes the properties available to describe the behaviors of each of the parameters you can specify, their use, and when they can be set.

Table 8-4. Parameter Properties (Part 1 of 3)

Property Name	Type/Default	Can Be Set	Description
AFFECTS_ELABORATION	Boolean, true	Main program	Set AFFECTS_ELABORATION to false for parameters that do not affect the external interface of the module. An example of a parameter that does not affect the external interface is isNonVolatileStorage. An example of a parameter that does affect the external interface is width. When the value of a parameter changes, if that parameter has set AFFECTS_ELABORATION=false, the elaboration phase (calling the callback or hardware analysis) is not repeated, improving performance. Because the default value of AFFECTS_ELABORATION is true, the provided HDL file is normally re-analyzed to determine the new port widths and configuration every time a parameter changes.
AFFECTS_GENERATION	Boolean, refer to description	Main program	The default value of AFFECTS_GENERATION is false if you provide a top-level HDL module, it is true if you provide a custom generation callback. Set AFFECTS_GENERATION to false if the value of a parameter does not change the results of system generation.
ALLOWED_RANGES	String, ""	Main program	Indicates the range or ranges that the parameter value can have. For integers, The ALLOWED_RANGES property is a list of ranges that the parameter can take on, where each range is a single value, or a range of values defined by a start and end value separated by a colon, such as 11:15. This property can also specify legal values and display strings for integers, such as {0:None 1:Monophonic 2:Stereo 4:Quadrophonic} meaning 0,1,2,4 are the legal values. You can also assign longer strings to be displayed in the parameter editor to string variables. For example, ALLOWED_RANGES {"dev1:Cyclone IV GX" "dev2:Stratix V GT"} Refer to Example 8-8 on page 8-8 and Figure 8-1 on page 8-8 for additional examples illustrating the use of this property.
DEFAULT_VALUE	String or Boolean	Main program	The default value.
DERIVED	Boolean, false	Validation or elaboration callback	When true, indicates that the parameter value does not need to be stored, typically because it is set from the validation callback. The default value is false.
DESCRIPTION	String, ""	Main program	A user-visible description of the parameter.

Table 8-4. Parameter Properties (Part 2 of 3)

Property Name	Type/Default	Can Be Set	Description
DISPLAY_HINT	String, ""	Main program	<p>Provides a hint about how to display a property. The following values are possible:</p> <ul style="list-style-type: none"> ■ <code>boolean</code>—for integer parameters whose value can be 0 or 1. The parameter displays as an option that you can turn on or off. ■ <code>radio</code>—displays a parameter with a list of values as radio buttons instead of a drop-down list. ■ <code>hexadecimal</code>—for integer parameters, display and interpret the value as a hexadecimal number, for example: <code>0x00000010</code> instead of 16. ■ <code>fixed_size</code>—for <code>string_list</code> and <code>integer_list</code> parameters, the <code>fixed_size</code> <code>DISPLAY_HINT</code> eliminates the add and remove buttons from tables. <p>Refer to Example 8-8 on page 8-8 and Figure 8-1 on page 8-8 for examples illustrating the use of this property.</p>
DISPLAY_NAME	String, ""	Main program	This is the GUI label that appears to the left of the parameter.
DISPLAY_UNITS	String, ""	Main program	This is the GUI label that appears to the right of the parameter.
ENABLED	Boolean, true	Main program, validation, and elaboration, callbacks	When <code>false</code> , the parameter is disabled, meaning that it is displayed, but greyed out, indicating that it is not editable on the parameter editor.
GROUP	String, ""	Main	Controls the layout of parameters in GUI. Refer to Example 8-8 for an illustration of its use.
HDL_PARAMETER	Boolean, false	Main program	When <code>true</code> , the parameter must be passed to the HDL component description. The default value is <code>false</code> .
NEW_INSTANCE_VALUE	String, ""	Main program	This property allows you to change the default value of a parameter without affecting older components that have assigned a default value to this parameter using the <code>defaultValue</code> argument. The practical result is that older components will continue to use <code>defaultValue</code> for the parameter and newer components can use the value assigned by <code>NEW_INSTANCE_VALUE</code> .

Table 8-4. Parameter Properties (Part 3 of 3)

Property Name	Type/Default	Can Be Set	Description
SYSTEM_INFO	String, ""	Main program	<p>Allows you to assign information about the instantiating system to a parameter that you define. <code>SYSTEM_INFO</code> requires a keyword argument specifying the type of information requested, <code><info-type></code>. <code><info-type></code> may also take an argument. The syntax of the Tcl command is:</p> <pre>set_parameter_property my_parameter SYSTEM_INFO <info-type> [<arg>]</pre> <p>The following values for <code><info-type></code> are predefined: <code>ADDRESS_MAP</code>, <code>ADDRESS_WIDTH</code>, <code>CLOCK_DOMAIN</code>, <code>CLOCK_RATE</code>, <code>CLOCK_RESET_INFO</code>, <code>CUSTOM_INSTRUCTION_SLAVES</code>, <code>DEVICE</code>, <code>DEVICE_FAMILY</code>, <code>DEVICE_FEATURES</code>, <code>INTERRUPTS_USED</code>, <code>GENERATION_ID</code>, <code>MAX_SLAVE_DATA_WIDTH</code>, <code>RESET_DOMAIN</code>, and <code>TRISTATE_ONDUIT_MASTERS</code></p> <p>Refer to Table 8-5 for descriptions of the <code><info-type></code> argument.</p>
TYPE	String, ""	Main program	<p>Specifies one of the following types: <code>INTEGER</code>, <code>NATURAL</code>, <code>POSITIVE</code>, <code>BOOLEAN</code>, <code>STD_LOGIC</code>, <code>STD_LOGIC_VECTOR</code>, <code>STRING</code>, <code>STRING_LIST</code>, <code>INTEGER_LIST</code>, <code>LONG</code>, or <code>FLOAT</code>.</p>
UNITS	String, ""	Main program	<p>Sets the units of the parameter. The following values are possible: <code>None</code>, <code>Picoseconds</code>, <code>Nanoseconds</code>, <code>Microseconds</code>, <code>Milliseconds</code>, <code>Seconds</code>, <code>Hertz</code>, <code>Kilohertz</code>, <code>Megahertz</code>, <code>Gigahertz</code>, <code>Address</code>, <code>Bits</code>, <code>Bytes</code>, <code>Kilobytes</code>, <code>Megabytes</code>, <code>Gigabytes</code>, <code>BitsPerSecond</code>, <code>KiloBitsPerSecond</code>, <code>MegaBitsPerSecond</code>, <code>BigaBitsPerSecond</code>, <code>Percent</code>, and <code>Cycles</code>. For example,</p> <pre>set_parameter_property frequency UNITS gigahertz</pre>
VISIBLE	Boolean, true	Main program, validation, and elaboration, callbacks	<p>Indicates whether or not to display the parameter in the parameterization GUI.</p>

- Table 8-5 describes the properties that you can use with the `system_info` parameter property. For more information about how to use the `system_info` parameter property, refer to “SYSTEM_INFO Parameters” on page 8-4.

Table 8-5. SYSTEM_INFO Properties (Part 1 of 2)

Property	Type	Description
ADDRESS_MAP	String	Assigns an XML formatted string describing the address map to the parameter you specify. <code>set_parameter_property <my_parameter> SYSTEM_INFO {ADDRESS_MAP <my_avalon-mm_master>}</code>
ADDRESS_WIDTH	Integer	Assigns an integer to the parameter that you specify that is the number of bits an Avalon-MM master must drive to address all of its slaves, using byte addresses. <code>set_parameter_property <my_parameter> SYSTEM_INFO {ADDRESS_WIDTH <my_avalon-mm_master>}</code>
CLOCK_DOMAIN	Integer	Assigns an integer representing the clock domain to the parameter you specify. You can use this command to determine whether multiple interfaces in your module are on the same clock domain. The absolute value of the integer value is arbitrary, but if two interfaces are on the same clock domain, the <code>CLOCK_DOMAIN</code> value is guaranteed to be the same and greater than zero. <code>set_parameter_property <my_parameter> SYSTEM_INFO {CLOCK_DOMAIN <my_clk>}</code>
CLOCK_RATE	Integer or String	Assigns a positive number which is the clock frequency in Hz to the clock input interface you specify. Assigns 0 if the clock rate is not known. <code>set_parameter_property <my_parameter> SYSTEM_INFO {CLOCK_RATE <my_clk>}</code>
CLOCK_RESET_INFO	String	Specifies the name of the module's clock or reset sink interface. (Specifies the clock sink interface for designs that use a global reset.)
CUSTOM_INSTRUCTION_SLAVES	String	Provides custom instruction slave information, including the name, base address, address span, and clock cycle type.
DEVICE	String	Specifies the Altera part number, for example EP2S15F484C3.
DEVICE_FAMILY	String	Assigns the family name (not the specific device part number) of the currently selected device to the parameter you specify. <code>set_parameter_property <my_parameter> SYSTEM_INFO {DEVICE_FAMILY}</code>
DEVICE_FEATURES	String	Creates a list of key/value pairs delineated by spaces indicating whether a particular device feature is available in the currently selected device family. The format of the list is suitable for passing to the Tcl array <code>set</code> command. This list is assigned to the parameter you specify. The following features are supported: <code>M512_MEMORY</code> , <code>M4K_MEMORY</code> , <code>M9K_MEMORY</code> , <code>M144K_MEMORY</code> , <code>MRAM_MEMORY</code> , <code>MLAB_MEMORY</code> , <code>ESB</code> , <code>DSP</code> , and <code>EMUL</code> . <code>set_parameter_property <my_parameter> SYSTEM_INFO {DEVICE_FEATURES}</code>

Table 8-5. SYSTEM_INFO Properties (Part 2 of 2)

Property	Type	Description
INTERRUPTS_USED	Integer or string	Creates a mask indicating which bits of the interrupt receiver vector are connected to an interrupt sender. This mask is assigned to the parameter you specify. You can use this interrupt mask to optimize logic that handles interrupts. set_parameter_property <my_parameter> SYSTEM_INFO (INTERRUPTS_USED <my_interrupt_receiver>}
GENERATION_ID	Integer	Records a unique ID for a particular generation run.
MAX_SLAVE_DATA_WIDTH	Integer	Assigns an integer to the parameter you specify that is the data width of the widest slave connected to the specified Avalon-MM master. set_parameter_property <my_parameter> SYSTEM_INFO {MAX_SLAVE_DATA_WIDTH <my_avalon_mm_master>}
RESET_DOMAIN	Integer	Assigns an integer representing the reset domain to the parameter you specify. You can use this command to determine whether multiple interfaces in your module are on the same reset domain. The absolute value of the integer value is arbitrary, but if two interfaces are on the same reset domain, the RESET_DOMAIN value is guaranteed to be the same and greater than zero. set_parameter_property <my_parameter> SYSTEM_INFO {RESET_DOMAIN <my_reset>}

get_parameter_property

This command returns a single parameter property.

get_parameter_property		
Callback availability	Main, validation, elaboration, generation, compose, and editor	
Usage	get_parameter_property <parameterName> <propertyName>	
Returns	string, boolean, or units, depending on property. Refer to Table 8-4 on page 8-26 .	
Arguments	parameterName	The name of the parameter whose property value is being retrieved
	propertyName	One of the properties listed in Table 8-4 on page 8-26
Example	get_parameter_property parameter1 GROUP	

set_parameter_property

This command sets a single parameter property.

set_parameter_property	
Callback availability	Main, validation, compose, and elaboration
Usage	set_parameter_property <parameterName> <propertyName> <value>
Returns	string, boolean, or units depending on property

set_parameter_property		
Arguments	parameterName	Specifies the parameter that is being set
	propertyName	Specifies the property of parameterName that is being set, refer to Table 8-4 on page 8-26 for a list of properties
	value	Provides the values
Example	set_parameter_property BAUD_RATE ALLOWED_RANGES {9600 19200 38400}	

get_parameter_value

This command returns the current value of a parameter defined previously with the add_parameter command.

get_parameter_value		
Callback availability	Validation, elaboration (1), compose, generation, and editor	
Usage	get_parameter_value <parameterName>	
Returns	String	
Arguments	parameterName	Specifies the parameter that is being retrieved
Example	set fifo_width [get_parameter_value fifo_width]	

Note:

- (1) If AFFECTS_ELABORATION=false for a given parameter, get_parameter_value is not available for that parameter from the elaboration callback. If affects_generation=false then it is not available from the generation callback.

set_parameter_value

This command sets a parameter value. The values of derived parameters can be set from the validation and elaboration callbacks. The values of parameters which are not marked as derived or system_info can be set from the editor callback.

set_parameter_value		
Callback availability	Validation, elaboration, compose, and editor	
Usage	set_parameter_value <parameterName> <value>	
Returns	None	
Arguments	parameterName	Specifies the parameter that is being set
	value	Specifies the value of parameterName
Example	set_parameter_value BAUD_RATE 19200	

decode_address_map

This is a utility function to convert an XML-formatted address map into a list of Tcl lists. Each inner list is in the correct format for conversion to an array. The XML code describing each slave includes: its name, start address, and end address + 1. [Figure 8-2](#) shows a portion of an Qsys system with three Avalon-MM slave devices.

Figure 8-2. Qsys System with Three Avalon-MM Slaves



[Example 8-14](#) shows the XML that describes the address map for the Avalon-MM master that accesses these slaves. The format of the XML string provided may differ from that described here, it may have different white space between the elements and could include additional attributes or elements. Using `decode_address_map` command to decode the XML representing an Avalon-MM master's address map is easier and ensures that your code will work with future versions of the XML address map.



Altera recommends that you use the code provided in the description of [Example 8-14](#) to enumerate over the components within an address map, rather than writing your own parser.

Example 8-14. Address Map for an Avalon-MM Master

```
<address-map>
  <slave name='ext_ssram' start='0x01000000' end='0x01200000' />
  <slave name='sys_clk_timer' start='0x02120800' end='0x02120820' />
  <slave name='sysid' start='0x021208B8' end='0x021208C0' />
</address-map>
```

decode_address_map			
Callback availability	Validation, compose, elaboration, and generation		
Usage	<code>decode_address_map <address_map_XML_string></code>		
Returns	List of Tcl lists, each one suitable for passing to array set		
Arguments	<table border="1"> <tr> <td><code>address_map_XML_string</code></td> <td>An XML string describing the address map of an Avalon-MM master.</td> </tr> </table>	<code>address_map_XML_string</code>	An XML string describing the address map of an Avalon-MM master.
<code>address_map_XML_string</code>	An XML string describing the address map of an Avalon-MM master.		
Example	<pre>set address_map_xml [get_parameter_value my_map_param] set address_map_dec [decode_address_map \$address_map_xml] foreach i \$address_map_dec { array set info \$i send_message info "Connected to slave \$info(name)" }</pre>		

Display Items

You specify your component GUI using the display commands.

add_display_item

You can use this command to specify the following aspects of component display:

- You can create logical groups for a component's parameters. For example, you might want to create separate groups for the component's timing, size, and simulation parameters. A component displays the groups and parameters in the order that you specify the display items for them in the `_hw.tcl` file.
- You can create multicolumn tables to present a component's parameters. Refer to [Example 8-13 on page 8-24](#) for an example that illustrates multicolumn tables.
- You can specify an image to provide a pictorial representation of a parameter or parameter group.
- You can create a button by adding a display item of type `action`. The display item includes the name of the callback to run when the action is performed.

Qsys Beta

You create a display group by adding display items to it.

add_display_item		
Callback availability	Main program	
Usage	<code>add_display_item <groupName> <id> <type> [<additionalInfo>]</code>	
Returns	String	
Arguments	<code>groupName</code>	Specifies the group to which a display item belongs.
	<code>id</code>	Specifies the parameter or icon to be displayed in a group. Each display item associated with a component must have a different ID.
	<code>type</code>	Specifies the category of the display item. The following types are defined: <ul style="list-style-type: none"> ■ <code>icon</code>—a <code>.gif</code>, <code>.jpg</code>, or <code>.png</code> file ■ <code>parameter</code>—a parameter in the instance ■ <code>text</code>—a block of text ■ <code>group</code>—a group. If the <code>groupName</code> is also defined, the new group is a child of the <code>groupName</code> group. If <code>groupName</code> is an empty string, the group is top-level. ■ <code>action</code>—an action defined by a callback procedure when you click the button labeled by <code>actionName</code>.
	<code>additionalInfo</code>	Provides extra information required for display items. The following examples illustrate how you use the <code>additionalInfo</code> argument for the various types: <ul style="list-style-type: none"> ■ <code>add_display_item groupName id icon path-to-image-file</code> ■ <code>add_display_item groupName parameterName parameter</code> (<code>additionalInfo</code> not required) ■ <code>add_display_item groupName id text "your-text"</code> The <code>your-text</code> argument is a block of text that is displayed in the GUI. Some simple HTML formatting is allowed, such as <code></code> and <code><i></code>, if the text starts with <code>"html"</code>. ■ <code>add_display_item parentGroupName childGroupName group [tab]</code> The <code>tab</code> is an optional parameter. If present, the group appears in separate tab in the GUI for the instance. ■ <code>add_display_item parentGroupName actionName action</code> <code>buttonClickCallbackProc</code>
Examples	<pre>add_display_item timing read_latency parameter add_display_item sound speaker icon speaker.jpg</pre>	

get_display_items

This command returns a list of all items to be displayed as part of the parameterization GUI.

get_display_items	
Callback availability	Main, elaboration, validation, generation, compose, and editor
Usage	<code>get_display_items</code>

get_display_items	
Returns	List of strings
Arguments	None
Example	get_display_items

get_display_item_properties

This command returns a list of names of the properties of display items that are part of the parameterization GUI.

get_display_item_properties	
Callback availability	Main
Usage	get_display_item_properties
Returns	List of strings
Arguments	None
Example	get_display_item_properties

get_display_item_property

This command returns the value of specific property of a display item that is part of the parameterization GUI.

get_display_item_property		
Callback availability	Main	
Usage	get_display_item_property <itemName> <propertyName>	
Returns	String	
Arguments	itemName	The item whose property value is being retrieved
	propertyName	The property whose value is being retrieved
Example	set my_label [get_display_item_property my_action DISPLAY_NAME]	

set_display_item_property

This command sets the value of specific property of a display item that is part of the parameterization GUI.

set_display_item_property	
Callback availability	Main
Usage	set_display_item_property <itemName> <propertyName> <value>
Returns	String

set_display_item_property		
Arguments	itemName	The item whose property value is being set
	propertyName	The property whose value is being set
	value	The value to set
Example	<pre>set_display_item_property my_action DISPLAY_NAME "Click Me" set_display_item_property my_action DESCRIPTION "clicking this button runs the click_me_callback proc in the hw.tcl file"</pre>	

Interfaces and Ports

You can use the interface and port commands to define interfaces and ports and retrieve their properties.

add_interface

This command adds an interface to your module. As the component author, you choose the name of the interface. By default, interfaces are enabled. You can set the interface property `ENABLED` to `false`, to disable a component interface. If an interface is disabled, it is hidden and its ports are automatically terminated to their default values. Signals that you designate as active low by appending a `_n` are terminated to 1. All other signals are terminated to 0.

 The properties available for each interface type are different. The common properties, `ENABLED` and `ASSOCIATED_CLOCK` apply to all interface types. Refer to the [Avalon Interface Specifications](#) for a description of other properties.

add_interface (Part 1 of 2)																				
Callback availability	Main program, elaboration, and compose																			
Usage	<code>add_interface <interfaceName> <interfaceType> <direction> [<associatedClock>]</code> (1)																			
Returns	String																			
Arguments	interfaceName	A name that you choose to identify an interface.																		
	interfaceType and direction	<p>There are 7 interfaceTypes. The following directions are possible for these interfaceTypes</p> <table border="1"> <thead> <tr> <th>Interface Type</th> <th>Direction</th> </tr> </thead> <tbody> <tr> <td>avalon</td> <td>master, slave (2)</td> </tr> <tr> <td>avalon_conduit_tristate</td> <td>master</td> </tr> <tr> <td>avalon_streaming</td> <td>source, sink</td> </tr> <tr> <td>interrupt</td> <td>sender, receiver</td> </tr> <tr> <td>conduit</td> <td>end</td> </tr> <tr> <td>clock</td> <td>source, sink</td> </tr> <tr> <td>reset</td> <td>source, sink</td> </tr> <tr> <td>nios_custom_instruction</td> <td>slave</td> </tr> </tbody> </table>	Interface Type	Direction	avalon	master, slave (2)	avalon_conduit_tristate	master	avalon_streaming	source, sink	interrupt	sender, receiver	conduit	end	clock	source, sink	reset	source, sink	nios_custom_instruction	slave
	Interface Type	Direction																		
avalon	master, slave (2)																			
avalon_conduit_tristate	master																			
avalon_streaming	source, sink																			
interrupt	sender, receiver																			
conduit	end																			
clock	source, sink																			
reset	source, sink																			
nios_custom_instruction	slave																			
associatedClock	This defines the clock associated with the interface. It is required for all interfaces except clock interfaces.																			

add_interface (Part 2 of 2)	
Example	<code>add_interface mm_slave avalon_slave clock0</code>

Notes:

- (1) For interfaces that are not associated with clocks, such as clock interfaces themselves, the `associatedClock` is omitted. Another option is to specify the `associatedClock` argument as `asynchronous`.
- (2) The terms *master*, *source*, and *start* are interchangeable. The terms *slave*, *sink*, and *end* are interchangeable.

get_interfaces

This command returns the names of all interfaces that have been previously defined by `add_interface` as a space separated list.

get_interfaces	
Callback availability	Main, validation, elaboration, generation, compose, and editor
Usage	<code>get_interfaces</code>
Returns	List of strings
Arguments	None
Example	<code>set all_interfaces [get_interfaces]</code>

get_interface_properties

This command returns the names of all the available interface properties for the specified interface as a space separated list.

get_interface_properties	
Callback availability	Main program, validation, elaborations, compose, and editor
Usage	<code>get_interface_properties <interfaceName></code>
Returns	List of strings
Arguments	<code>interfaceName</code> The name of an interface that you defined
Example	<code>get_interface_properties mm_slave</code>

 The properties available for each interface type are different. Refer to the [Avalon Interface Specifications](#) for more information about interface properties.

The interface properties that are common to all interface types are listed below in [Table 8-6](#).

Table 8-6. Interface Properties Common to All Interface Types

Property	Type	Description
EXPORT_OF	String	For composed <code>_hwl.tcl</code> files, the <code>EXPORT_OF</code> property indicates which interface of a child instance is to be exported through this interface. Before using this command, you must have created the border interface using <code>add_interface</code> . The interface to be exported is of the form <code><instanceName.interfaceName></code> . Example: <code>set_interface_property CSC_input EXPORT_OF my_colorSpaceConverter.input_port</code>
ASSOCIATED_CLOCK	String	The name of the clock interface that this interface is synchronous to.
ENABLED	Boolean	Specifies whether or not interface is enabled.

get_interface_property

This command returns the value of a single interface property from the specified interface.

get_interface_property		
Callback availability	Main program, compose, and elaboration	
Usage	<code>get_interface_property <interfaceName> <propertyName></code>	
Returns	string, boolean, or units, depending on property. Refer to the Avalon Interface Specifications for more information about interface properties	
Arguments	<code>interfaceName</code>	The name of an interface from which you want to retrieve information
	<code>propertyName</code>	The name of the property whose value you want to retrieve. This property is either <code>ENABLED</code> or <code>ASSOCIATED_CLOCK</code> or a property name defined by the interface.
Example	<code>get_interface_property mm_slave readWaitTime</code>	

set_interface_property

This command sets a single interface property for an interface.

set_interface_property		
Callback availability	Main, compose, and elaboration	
Usage	<code>set_interface_property <interfaceName> <propertyName> <value></code>	
Returns	String	
Arguments	<code>interfaceName</code>	The name of an interface that includes this property
	<code>propertyName</code>	The name of the property whose value you want to set, which is <code>ENABLED</code> or <code>ASSOCIATED_CLK</code> or a name from the Avalon Interface Specifications .
	<code>value</code>	The value to set for the specified property
Example	<code>set_interface_property mm_slave linewidthBursts false</code>	

add_interface_port

This command adds a port to an interface on your module. As the component author, you determine the name of the port. The port width and direction must be set by the end of the elaboration phase. The port width can be set with one of the following mechanisms:

- A constant width or a width expression can be set in the main program
- A constant width can be set in the elaboration callback



Without an elaboration callback, for static components `quartus_map` determines the port width from the HDL

add_interface_port		
Callback availability	Main program and elaboration	
Usage	<code>add_interface_port <interfaceName> <portName> <portRole> [<direction> <width_expr>]</code>	
Returns	String	
Arguments	<code>interfaceName</code>	The name of the interface to which the port belongs.
	<code>portName</code>	The name of the port that you, the component author, have chosen.
	<code>portRole</code>	The role of this port within the interfaces. Port roles are referred to as signal types in the <i>Avalon Interface Specification</i> . Refer to the <i>Avalon Interface Specifications</i> for the signal types available for each interface type.
	<code>direction</code>	The direction can be input, output, or bidir
	<code>width_expr</code>	The port's width expression. In simple cases, this is just the width of the port in bits.
Example	<code>add_interface_port mm_slave s0_rdata readdata output 32</code>	

get_interface_ports

This command returns the names of all of the ports that have been added to a given interface. If the interface name is omitted, all ports for all interfaces are returned.

get_interface_ports	
Callback availability	Main, validation, elaboration, generation, and editor
Usage	<code>get_interface_ports [<interfaceName>]</code>
Returns	String
Arguments	<code>interfaceName</code> The name of the interface whose ports you want to list. (Optional)
Example	<code>get_interface_ports mm_slave</code>

get_port_properties

This command returns a list of all available port properties.

get_port_properties	
Callback availability	Main, validation, elaboration, generation, compose, and editor
Usage	<code>get_port_properties <portName></code>
Returns	String, boolean, or units, depending on property. Refer to Table 8-4 on page 8-26
Arguments	<p>The name of the port whose properties are required. The following 7 port properties are supported:</p> <ul style="list-style-type: none"> ■ DIRECTION ■ TERMINATION ■ TERMINATION_VALUE ■ VHDL_TYPE ■ WIDTH ■ WIDTH_EXPR ■ DRIVEN_BY ■ ROLE <p>Refer to Table 8-7 for a description of these properties.</p>
Example	<code>get_port_properties mm_slave</code>

[Table 8-7](#) describes the available port properties

Table 8-7. Port Properties (Part 1 of 2)

Name	Type	Description
DIRECTION	input, output, bidir	The direction of the port from the component's perspective.
TERMINATION	boolean	When true, instead of connecting the port to the Qsys system, it is left unconnected for output and bidir or set to a fixed value for input. Has no effect for components that implement a generation callback instead of using the default wrapper generation.
TERMINATION_VALUE	integer	The constant value to drive an input port.
VHDL_TYPE	std_logic std_logic_vector auto	indicates the type of a VHDL port. The default value, auto, selects std_logic if the width is fixed at 1, and std_logic_vector otherwise.
WIDTH	integer	The width of the port in bits.
WIDTH_EXPR	string	<p>The width expression of a port. Setting the width and width_expr properties have the same effect; they both update the effective width expression. The width/width_expr properties can be set to an integer at any time. They can only be set to arithmetic expressions in the main program.</p> <p>The values of the width and width_expr properties behave differently when get_port_property is used. width always returns the current integer width of the port. width_expr always returns the unevaluated width expression.</p>

Table 8-7. Port Properties (Part 2 of 2)

Name	Type	Description
DRIVEN_BY	integer, input	Indicates that this output port is always driven to a constant value or by an input port. If all outputs on a component have their <code>driven_by</code> property set to a valid value then the component's HDL is generated automatically.
ROLE	string	Specifies an Avalon signal type such as <code>waitrequest</code> , <code>readdata</code> , or <code>read</code> . For a complete list of signal types, refer to the <i>Avalon Interface Specifications</i> .

get_port_property

This command returns the value of single port property for the specified port.

get_port_property		
Callback availability	Main, validation, elaboration, generation, and editor	
Usage	<code>get_port_property <portName> <propertyName></code>	
Returns	Depends on the type of the property	
Arguments	<code>portName</code>	The name of the port
	<code>propertyName</code>	One of the supported properties described in Table 8-7 .
Example	<code>get_port_property rdata WIDTH</code>	

set_port_property

This command sets a single port property.

set_port_property		
Callback availability	Main program, elaboration, and generation	
Usage	<code>set_port_property <portName> <propertyName> [<value>]</code>	
Returns	String, boolean, or units, depending on property. Refer to Table 8-4 on page 8-26 .	
Arguments	<code>portName</code>	The name of the port
	<code>propertyName</code>	One of the supported properties described in Table 8-7 .
	<code>value</code>	The value to set
Example	<code>set_port_property rdata WIDTH 32</code>	

get_interface_assignments

This command returns the value of all interface assignments for the specified interface.

get_interface_assignments	
Callback availability	Main, validation, compose, and elaboration
Usage	<code>get_interface_assignments <interfaceName></code>

get_interface_assignments		
Returns	String	
Arguments	interfaceName	The name of the Avalon interface whose assignment is being retrieved
Example	get_interface_assignments s1	

get_interface_assignment

This command returns the value of the specified name for the specified interface.

get_interface_assignment		
Callback availability	Main, validation, compose, and elaboration	
Usage	get_interface_assignments <interfaceName> <name>	
Returns	String	
Arguments	interfaceName	The name of the Avalon interface whose assignment is being retrieved
	name	The assignment whose value is being retrieved
Example	get_interface_assignment s1 embeddedsw.configuration.isFlash	

set_interface_assignment

This command sets the value of the specified assignment for the specified interface.

set_interface_assignment		
Callback availability	Main, validation, compose, and elaboration	
Usage	set_interface_assignment <interfaceName> <name> [<value>]	
Returns	None	
Arguments	interfaceName	The name of the Avalon interface whose assignment is being set
	name	The assignment whose value is being set
	value	The value to assign
Example	set_interface_assignment s1 embeddedsw.configuration.isFlash 1	

 For more information about the use of the set_interface_assignment command, refer to the “Publishing Component Information to Embedded Software” chapter in the *Nios II Software Developer’s Handbook*.

Compose

This section covers the commands that allow you to build new components by combining other components. It also includes commands to query the module instances in the system.

add_instance

The `add_instance` command adds an instance of a predefined module, referred to as a *child* or *child module*, to a new component. You can use this command to create components that are composed of other components.

add_instance		
Callback availability	Main and compose	
Usage	<code>add_instance <instanceName> <type> [<version>]</code>	
Returns	String	
Arguments	<code>instanceName</code>	Specifies a unique local name that you can use to manipulate the module. This name is used in the generated HDL to identify the module.
	<code>type</code>	The <code>type</code> refers to a module available in a library, for example <code>altera_avalon_uart</code> .
	<code>version</code>	The required version of the specified module. If no version is specified, the latest version is used.
Example	<code>add_instance my_uart altera_avalon_uart</code>	

get_instances

This command lists the instance names of all modules in the system.

get_instances	
Callback availability	Main, validation, and compose
Usage	<code>get_instances</code>
Returns	List of strings
Arguments	None
Example	<code>get_instances</code>

get_instance_parameters

This command returns the names of all parameters on a child instance that can be manipulated by the parent. It omits parameters that are derived and those that have the `SYSTEM_INFO` parameter property set.

get_instance_parameters	
Callback availability	Main, validation, and compose
Usage	<code>get_instance_parameters <instanceName></code>
Returns	List of strings
Arguments	<code>instanceName</code> Specifies the name of the instance whose parameters are being retrieved.
Example	<code>get_instance_parameters pixel_converter</code>

set_instance_parameter_value

This command sets a parameter on a child module. Derived parameters and `SYSTEM_INFO` parameters for the child module may not be set using this command.

set_instance_parameter_value	
Callback availability	Main and compose
Usage	<code>set_instance_parameter_value <instanceName> <parameterName> <parameterValue></code>
Returns	None
Arguments	<code>instanceName</code> Specifies the name of the child module
	<code>parameterName</code> Specifies the parameter that is being set
	<code>parameterValue</code> Specifies the value of the parameter that is being set
Example	<code>set_instance_parameter_value pixel_converter input_DPI 1200</code>

get_instance_parameter_value

This command returns the value of the named parameter. You cannot use this command to get the value of parameters whose values are derived or those that are defined using the `SYSTEM_INFO` parameter property.

get_instance_parameter_value	
Callback availability	Main and compose
Usage	<code>get_instance_parameter_value <instanceName> <parameterName></code>
Returns	String, boolean, or units, depending on property. Refer to Table 8-4 on page 8-26
Arguments	<code>instanceName</code> Specifies the name of the instance whose parameter is being retrieved
	<code>parameterName</code> Specifies the parameter whose value is being retrieved
Example	<code>get_instance_parameter_value pixel_converter input_DPI</code>

get_instance_parameter_properties

This command returns the names of all properties for the specified parameter. The values returned are a subset of those returned from `get_parameter_properties`. Refer to [Table 8-4](#) for a list of parameter properties.

get_instance_parameter_properties	
Callback availability	Main and compose
Usage	<code>get_instance_parameter_properties <instanceName> <parameterName></code>
Returns	List of strings
Arguments	<code>instanceName</code> Specifies the instance name of the module
	<code>parameterName</code> Specifies the parameter that is being set
Example	<code>get_instance_parameter_properties my_colorSpaceConverter colorSpace</code>

get_instance_parameter_property

This command returns the names of the specified instance parameter property. The following parameter properties on a child instance that are visible from the parent: TYPE, WIDTH, DERIVED, VISIBLE, ENABLED, UNITS, DISPLAY_NAME, ALLOWED_RANGES, and SYSTEM_INFO.

get_instance_parameter_property	
Callback availability	Main and compose
Usage	<code>get_instance_parameter_property <instanceName> <parameterName> <propertyName></code>
Returns	String, boolean, or units, depending on property. Refer to Table 8-4 on page 8-26 .
Arguments	<code>instanceName</code> Specifies the instance name of the module
	<code>parameterName</code> Specifies the parameter for which a property is being retrieved
	<code>propertyName</code> Specifies the property whose value is being retrieved
Example	<code>get_instance_parameter_property my_stereo separate_control DISPLAY_NAME</code>

get_instance_interfaces

This command returns the names of all of the interfaces of a child module as a list. The interfaces can change if the parameterization of the module changes.

get_instance_interfaces	
Callback availability	Main and compose
Usage	<code>get_instance_interfaces <instanceName></code>
Returns	String
Arguments	<code>instanceName</code> Specifies the instance name of the module
Example	<code>get_instance_interfaces my_ColorSpaceConverter</code>

get_instance_interface_properties

This command returns the names of all of the properties of the specified interface.

get_instance_interface_properties	
Callback availability	Main and compose
Usage	<code>get_instance_interface_properties <instanceName> <interfaceName></code>
Returns	String
Arguments	<code>instanceName</code> Specifies the instance name of the module
	<code>interfaceName</code> Specifies an interface of instance
Example	<code>get_instance_interface_properties my_ColorSpaceConverter inputInterface</code>

get_instance_interface_property

This command returns the value of a property associated with the specified module interface.

get_instance_interface_property	
Callback availability	Main and compose
Usage	<code>get_instance_interface_property <instanceName> <interfaceName> <propertyName></code>
Returns	String
Arguments	<code>instanceName</code> Specifies the instance name of the module
	<code>interfaceName</code> Specifies an interface of instance
	<code>propertyName</code> Specifies the property whose value is being retrieved.
Example	<code>get_instance_interface_property my_component s1 setupTime</code>

get_instance_interface_ports

This command returns a list of the names of the ports on the specified interface.

get_instance_interface_ports	
Callback availability	Main and compose
Usage	<code>get_instance_interface_ports <instanceName> <interfaceName></code>
Returns	List of Strings
Arguments	<code>instanceName</code> Specifies the instance name of the module
	<code>interfaceName</code> Specifies an interface of instance
Example	<code>get_instance_interface_ports my_ColorSpaceConverter outputInterface</code>

get_instance_port_property

This command returns a information about the port property specified.

get_instance_port_property	
Callback availability	Main and compose
Usage	<code>get_instance_port_property <instanceName> <portName> <propertyName></code>
Returns	String
Arguments	<code>instanceName</code> Specifies the instance name of the module
	<code>portName</code> Specifies a port
	<code>property</code> Specifies the property for which information is being retrieved. Not all port properties are visible from the parent. Those which are visible are ROLE, DIRECTION, WIDTH, WIDTH_EXPR and VHDL_TYPE.
Example	<code>get_instance_port_property my_uart width</code>

add_connection

This command connects the named interfaces together using an appropriate connection type. Both interface names consist of a child instance name, followed by the name of an interface provided by that module. For example, `mux0.out` is the interface named `out` on the instance named `mux0`. The command returns the name of the newly added connection in `start.point/end.point` format. Be careful to connect the start to the end, and not the other way around.

add_connection		
Callback availability	Main program and compose	
Usage	<code>add_connection <start.Interface> [<end.Interface>] [kind] [name]</code>	
Returns	String	
Arguments	<code>start.interface</code>	The start interface to be connected, of the form, <code><instance_name>.<interface_name></code>
	<code>end.interface</code>	The end interface to be connected, <code><instance_name>.<interface_name></code>
	<code>kind</code>	Indicates the interface type. For a list of interface types refer to “ add_interface ” on page 8-37.
	<code>name</code>	Specifies the name of the connection. If omitted, the name is of the form <code>start-module.start-interface/end-module.end-interface</code> .
Example	<code>add_connection dma.read_master sdram.sl</code>	

get_connections

This command lists the connectivity for all modules in the system.

get_connections	
Callback availability	Main and compose
Usage	<code>get_connections</code>
Returns	List of strings
Arguments	None
Example	<code>get_connections</code>

get_connection_parameters

This command gets the names of all parameters for the connection specified.

get_connection_parameters	
Callback availability	Main and compose
Usage	<code>get_connection_parameters <connectionName></code>
Returns	List of strings
Arguments	<code>connectionName</code> Specifies the connection whose connection parameters are required.
Example	<code>get_connection_parameters cpu0.data_master/dma0.csr</code>

get_connection_parameter_value

This command gets the value of a parameter on the connection.

set_connection_parameter_value

This command sets a property of the connection. The start and end are each interface names of the format *<instance>.<interface>*. Connection parameters depend on the type of connection, for Avalon-MM they include base addresses and arbitration priorities.

set_connection_parameter_value	
Callback availability	Main program and compose
Usage	<code>set_connection_parameter_value <connName> <parameterName> <parameterValue></code>
Returns	None
Arguments	<code>connName</code> Specifies the name of the connection as returned by the <code>add_connection</code> command. It is of the form <code>start.point/end.point</code>
	<code>parameterName</code> Specifies the parameter that is being set
	<code>parameterValue</code> Specifies the value of the parameter
Example	<code>set_connection_parameter_value cpu0.data_master/dma0.csr baseAddress 0x1000</code>

Generation

This section covers the commands that get generation properties.

get_generation_properties

This command returns the names of all the available generation properties as a space separated list. These properties cannot be changed by the module. Generation properties are provided to the generation callback to support per-instance HDL generation.

get_generation_properties	
Callback availability	Main, validation, elaboration, compose generation, and editor
Usage	<code>get_generation_properties</code>
Returns	String. The following generation properties are supported: <code>n hdl_language</code> <code>n output_directory</code> <code>n output_name</code> Refer to Table 8-8 for a description of the generation properties.
Arguments	None
Example	<code>get_generation_properties</code>

Table 8-8 describes the generation properties.

Table 8-8. Generation Properties

Name	Type	Description
HDL_LANGUAGE	enum	The HDL language to generate. Is either <code>verilog</code> or <code>vhdl</code> (lowercase). If the module cannot generate the specified language, generating in the other language is acceptable.
OUTPUT_DIRECTORY	file	The location in which files must be generated. The filename components in the directory name are separated with forward slashes.
OUTPUT_NAME	string	OUTPUT_NAME is <code>module_0</code> and the HDL_LANGUAGE is <code>verilog</code> , the file module_0.v or module_0.sv must be generated and must contain the module, <code>module_0</code> .

get_generation_property

This command returns the value of a single generation property.

get_generation_property	
Callback availability	Generation
Usage	<code>get_generation_property <propertyName></code>
Returns	String, boolean, or units, depending on property. Refer to Table 8-4 on page 8-26 .
Arguments	propertyName One of the 3 generation properties: <ul style="list-style-type: none"> ■ HDL_LANGUAGE ■ OUTPUT_DIRECTORY ■ OUTPUT_NAME
Example	<code>get_generation_property OUTPUT_DIRECTORY</code>

Document Revision History

Table 8-9 shows the revision history for this document.

Table 8-9. Document Revision History

Date	Version	Changes
December 2010	10.1.0	Initial release.

 For previous versions of the Quartus II Handbook, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

When designing for large and complex FPGAs, your design and coding styles can impact your quality of results significantly. Designs reflecting synchronous design practices behave predictably reliably, even when re-targeted to different device families or speed grades. Using recommended HDL coding styles ensures that synthesis tools can infer the optimal device hardware to implement your design. Following best practices when creating your design hierarchy and logic provides the most flexibility when partitioning the design for incremental compilation, and leads to the best results. If you create floorplan location assignments to control the placement of different design blocks (useful in team-based designs so each designer can target a different area of the device floorplan), following best practices is important to maintaining good design performance.

This section presents design and coding style recommendations in the following chapters:

- [Chapter 9, Recommended Design Practices](#)

This chapter describes synchronous design practices, and provides guidelines for combinational logic structures and clocking schemes. It also explains how to check design rules using the Quartus® II Design Assistant. Finally, it discusses use of clock and register-control features in device architecture.

- [Chapter 10, Recommended HDL Coding Styles](#)

This chapter discusses Altera megafunctions and provides specific Verilog HDL and VHDL coding examples to insure the Quartus II software infers Altera dedicated logic such as memory and DSP blocks. It also provides device-specific coding recommendations for registers and certain logic functions such as tri-state signals, multiplexers, and cyclic redundancy check (CRC) functions, and includes references to other Altera documentation for low-level logic design information.

- [Chapter 11, Managing Metastability with the Quartus II Software](#)

This chapter describes ways you can use the Quartus II software to analyze the average mean time between failures (MTBF) due to metastability caused by synchronization of asynchronous signals, and optimize the design to improve the metastability MTBF.

- [Chapter 12, Best Practices for Incremental Compilation Partitions and Floorplan Assignments](#)

This chapter provides a set of guidelines to help you set up and partition your design to take advantage of the compilation time savings, performance preservation, and hierarchical design features offered by Quartus II incremental compilation, and to help you create a design floorplan (using LogicLock™ regions) to support the flow when required.

Use this chapter when setting up your design hierarchy and determining the interfaces between logic blocks in your design, as well as if/when you create a design floorplan. You can also use this chapter to make changes to a design that was not originally set up to take advantage of incremental compilation, because it provides tips on changing a design to work better with an incremental design flow.

This chapter provides design recommendations for Altera® devices and describes the Quartus® II Design Assistant, which helps you check your design for violations of Altera's design recommendations. Current FPGA applications have reached the complexity and performance requirements of ASICs. In the development of complex system designs, good design practices have an enormous impact on the timing performance, logic utilization, and system reliability of a device. Well-coded designs behave in a predictable and reliable manner even when retargeted to different families or speed grades. Good design practices also aid in successful design migration between FPGA and HardCopy® or ASIC implementations for prototyping and production.

For optimal performance, reliability, and faster time-to-market when designing with Altera devices, you should adhere to the following guidelines:

- Understand the impact of synchronous design practices
- Follow recommended design techniques, including hierarchical design partitioning
- Take advantage of the architectural features in the targeted device

This chapter contains the following sections:

- “Synchronous FPGA Design Practices” on page 9–2
- “Design Guidelines” on page 9–4
- “Checking Design Violations With the Design Assistant” on page 9–13
- “Targeting Clock and Register-Control Architectural Features” on page 9–19
- “Targeting Embedded RAM Architectural Features” on page 9–24

 For specific HDL coding examples and recommendations, including coding guidelines for targeting dedicated device hardware, such as memory and digital signal processing (DSP) blocks, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*. For information about partitioning a hierarchical design for incremental compilation, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.

 For information about migrating designs to HardCopy devices, refer to the *Design Guidelines for HardCopy Series Devices* chapter in volume 1 of the *HardCopy Series Handbook*.

Synchronous FPGA Design Practices

The first step in good design methodology is to understand the implications of your design practices and techniques. This section outlines the benefits of optimal synchronous design practices and the hazards involved in other techniques. Good synchronous design practices can help you meet your design goals consistently. Problems with other design techniques can include reliance on propagation delays in a device, incomplete timing analysis, and possible glitches.

In a synchronous design, a clock signal triggers all events. As long as you ensure that all the timing requirements of the registers are met, a synchronous design behaves in a predictable and reliable manner for all process, voltage, and temperature (PVT) conditions. You can easily target synchronous designs to different device families or speed grades. In addition, synchronous design practices help ensure successful migration if you plan to migrate your design to a high-volume solution such as a HardCopy device or if you are prototyping an ASIC design.

Fundamentals of Synchronous Design

In a synchronous design, the clock signal controls the activities of all inputs and outputs. On every active edge of the clock (usually the rising edge), the data inputs of registers are sampled and transferred to outputs. Following an active clock edge, the outputs of combinational logic feeding the data inputs of registers change values. This change triggers a period of instability due to propagation delays through the logic as the signals go through several transitions and finally settle to new values. Changes that occur on data inputs of registers do not affect the values of their outputs until the next active clock edge.

Because the internal circuitry of registers isolates data outputs from inputs, instability in the combinational logic does not affect the operation of the design as long as you meet the following timing requirements:

- Before an active clock edge, you must ensure that the data input has been stable for at least the setup time of the register.
- After an active clock edge, you must ensure that the data input remains stable for at least the hold time of the register.

When you specify all of your clock frequencies and other timing requirements, the Quartus II TimeQuest Timing Analyzer reports actual hardware requirements for the setup times (t_{SU}) and hold times (t_H) for every pin in your design. By meeting these external pin requirements and following synchronous design techniques, you ensure that you satisfy the setup and hold times for all registers in your device.



To meet setup and hold time requirements on all input pins, any inputs to combinational logic that feed a register should have a synchronous relationship with the clock of the register. If signals are asynchronous, you can register the signals at the input of the device to help prevent a violation of the required setup and hold times.

When you violate the setup or hold time of a register, the output can be set to an intermediate voltage level between the high and low levels, called a metastable state. In this unstable state, small perturbations such as noise in power rails can cause the register to assume either the high or low voltage level, resulting in an unpredictable valid state. Various undesirable effects can occur, including increased propagation delays and incorrect output states. In some cases, the output can even oscillate between the two valid states for a relatively long period of time.

- ② For information about timing requirements and analysis in the Quartus II software, refer *About TimeQuest Timing Analysis* in Quartus II Help.

Hazards of Asynchronous Design

In the past, designers have often used asynchronous techniques such as ripple counters or pulse generators in programmable logic device (PLD) designs, enabling them to take “short cuts” to save device resources. Asynchronous design techniques have inherent problems such as relying on propagation delays in a device, which can result in incomplete timing constraints and possible glitches and spikes.

Some asynchronous design structures rely on the relative propagation delays of signals to function correctly. In these cases, race conditions can arise where the order of signal changes can affect the output of the logic. PLD designs can have varying timing delays, depending on how the design is placed and routed in the device with each compilation. Therefore, it is almost impossible to determine the timing delay associated with a particular block of logic ahead of time. As devices become faster due to device process improvements, the delays in an asynchronous design may decrease, resulting in a design that does not function as expected. Specific examples are provided in “*Design Guidelines*” on page 9-4. Relying on a particular delay also makes asynchronous designs difficult to migrate to different architectures, devices, or speed grades.

The timing of asynchronous design structures is often difficult or impossible to model with timing assignments and constraints. If you do not have complete or accurate timing constraints, the timing-driven algorithms used by your synthesis and place-and-route tools may not be able to perform the best optimizations and the reported results may not be complete.

Some asynchronous design structures can generate harmful glitches, which are pulses that are very short compared with clock periods. Most glitches are generated by combinational logic. When the inputs of combinational logic change, the outputs exhibit several glitches before they settle to their new values. These glitches can propagate through the combinational logic, leading to incorrect values on the outputs in asynchronous designs. In a synchronous design, glitches on the data inputs of registers are normal events that have no negative consequences because the data is not processed until the clock edge.

Design Guidelines

When designing with HDL code, you should understand how a synthesis tool interprets different HDL design techniques and what results to expect. Your design techniques can affect logic utilization and timing performance, as well as the design's reliability. This section describes basic design techniques that ensure optimal synthesis results for designs targeted to Altera devices while avoiding several common causes of unreliability and instability. Design your combinational logic carefully to avoid potential problems and pay attention to your clocking schemes so you can maintain synchronous functionality and avoid timing problems.

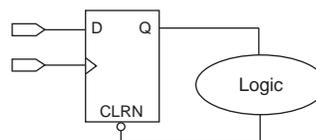
Combinational Logic Structures

Combinational logic structures consist of logic functions that depend only on the current state of the inputs. In Altera FPGAs, these functions are implemented in the look-up tables (LUTs) of the device's architecture, with either logic elements (LEs) or adaptive logic modules (ALMs). For some cases in which combinational logic feeds registers, the register control signals can implement part of the logic function to save LUT resources. By following the recommendations in this section, you can improve the reliability of your combinational design.

Combinational Loops

Combinational loops are among the most common causes of instability and unreliability in digital designs. They should be avoided whenever possible. In a synchronous design, feedback loops should include registers. Combinational loops generally violate synchronous design principles by establishing a direct feedback loop that contains no registers. For example, a combinational loop occurs when the left-hand side of an arithmetic expression also appears on the right-hand side in HDL code. A combinational loop also occurs when you feed back the output of a register to an asynchronous pin of the same register through combinational logic, as shown in [Figure 9-1](#).

Figure 9-1. Combinational Loop Through Asynchronous Control Pin



-  Use recovery and removal analysis to perform timing analysis on asynchronous ports, such as `clear` or `reset` in the Quartus II software.
-  If you are using the TimeQuest Timing Analyzer, refer to *Specifying Timing Constraints and Exceptions (TimeQuest Timing Analyzer)* in Quartus II Help for details about how the TimeQuest analyzer performs recovery and removal analysis.

Combinational loops are inherently high-risk design structures for the following reasons:

- Combinational loop behavior generally depends on relative propagation delays through the logic involved in the loop. As discussed, propagation delays can change, which means the behavior of the loop is unpredictable.
- Combinational loops can cause endless computation loops in many design tools. Most tools break open combinational loops to process the design. The various tools used in the design flow may open a given loop in a different manner, processing it in a way that is inconsistent with the original design intent.

Latches

A latch is a small circuit with combinational feedback that holds a value until a new value is assigned. You can implement latches with the Quartus II Text Editor or Block Editor. It is common for mistakes in HDL code to cause unintended latch inference; Quartus II Synthesis issues a warning message if this occurs.

Unlike other technologies, a latch in an FPGA architecture is not significantly smaller than a register. The architecture is not optimized for latch implementation and latches generally have slower timing performance compared to equivalent registered circuitry.

Latches have a transparent mode in which data flows continuously from input to output. A positive latch is in transparent mode when the enable signal is high (low for negative latch). In transparent mode, glitches on the input can pass through the output because of the direct path created. This presents significant complexity for timing analysis. Typical latch schemes use multiple enable phases to prevent long transparent paths from occurring. However, timing analysis cannot identify these safe applications.

The TimeQuest analyzer analyzes latches as synchronous elements by default, and allows you to treat latches as having nontransparent start and end points. Be aware that even an instantaneous transition through transparent mode can lead to glitch propagation. The TimeQuest analyzer cannot perform cycle-borrowing analysis; this is performed instead by the Synopsys PrimeTime third-party timing analysis tool.

Due to various timing complexities, latches have limited support in formal verification tools. Therefore, you should not rely on formal verification for a design that includes latches.



Avoid using latches to ensure that you can completely analyze the timing performance and reliability of your design.

Delay Chains

Delay chains occur when you use two or more consecutive nodes with a single fan-in and a single fan-out to cause delay. Inverters are often chained together to add delay. Delay chains are sometimes used to resolve race conditions created by other asynchronous design practices.

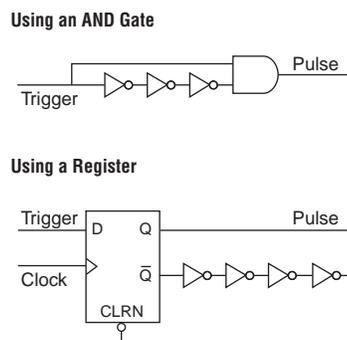
Delays in PLD designs can change with each placement and routing cycle. Effects such as rise and fall time differences and on-chip variation mean that delay chains, especially those placed on clock paths, can cause significant problems in your design. Refer to “[Hazards of Asynchronous Design](#)” on [page 9-3](#) for examples of the kinds of problems that delay chains can cause. Avoid using delay chains to prevent these kinds of problems.

In some ASIC designs, delays are used for buffering signals as they are routed around the device. This functionality is not required in FPGA devices because the routing structure provides buffers throughout the device.

Pulse Generators and Multivibrators

You can use delay chains to generate either one pulse (pulse generators) or a series of pulses (multivibrators). There are two common methods for pulse generation, as shown in [Figure 9-2](#). These techniques are purely asynchronous and must be avoided.

Figure 9-2. Asynchronous Pulse Generators



In [Figure 9-2](#), a trigger signal feeds both inputs of a 2-input AND gate, but the design inverts or adds a delay chain to one of the inputs. The width of the pulse depends on the relative delays of the path that feed the gate directly and the path that goes through the delay. This is the same mechanism responsible for the generation of glitches in combinational logic following a change of input values. This technique artificially increases the width of the glitch by using a delay chain.

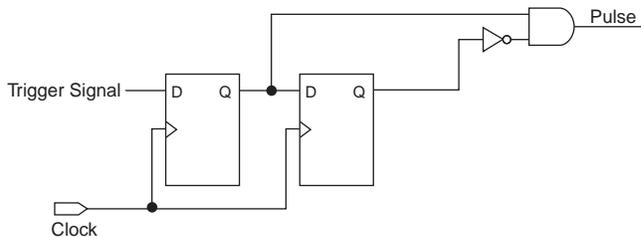
As also shown in [Figure 9-2](#), a register’s output drives the same register’s asynchronous reset signal through a delay chain. The register resets itself asynchronously after a certain delay.

The width of pulses generated in this way are difficult for synthesis and place-and-route software to determine, set, or verify. The actual pulse width can only be determined after placement and routing, when routing and propagation delays are known. You cannot reliably determine the width of the pulse when creating HDL code, and it cannot be set by EDA tools. The pulse may not be wide enough for the application under all PVT conditions. Also, the pulse width changes if you change to a different device. Additionally, static timing analysis cannot be used to verify the pulse width, so verification is difficult.

Multivibrators use a glitch generator to create pulses, together with a combinational loop that turns the circuit into an oscillator. This creates additional problems because of the number of pulses involved. Additionally, when the structures generate multiple pulses, they also create a new artificial clock in the design that has to be analyzed by the design tools.

When you must use a pulse generator, use synchronous techniques, as shown in [Figure 9-3](#).

Figure 9-3. Recommended Pulse-Generation Technique



In this design, the pulse width is always equal to the clock period. This pulse generator is predictable, can be verified with timing analysis, and is easily moved to other architectures, devices, or speed grades.

Clocking Schemes

Like combinational logic, clocking schemes have a large effect on the performance and reliability of a design. Avoid using internally generated clocks wherever possible because they can cause functional and timing problems in the design. Clocks generated with combinational logic can introduce glitches that create functional problems and the delay inherent in combinational logic can lead to timing problems.



Specify all clock relationships in the Quartus II software to allow for the best timing-driven optimizations during fitting and to allow correct timing analysis. Use clock setting assignments on any derived or internal clocks to specify their relationship to the base clock.

You use global device-wide, low-skew dedicated routing for all internally-generated clocks, instead of routing clocks on regular routing lines. For more information, refer to [“Clock Network Resources” on page 9-20](#).

Avoid data transfers between different clocks wherever possible. If you require a data transfer between different clocks, use FIFO circuitry. You can use the clock uncertainty features in the Quartus II software to compensate for the variable delays between clock domains. Consider setting a **Clock Setup Uncertainty** and **Clock Hold Uncertainty** value of 10% to 15% of the clock delay.

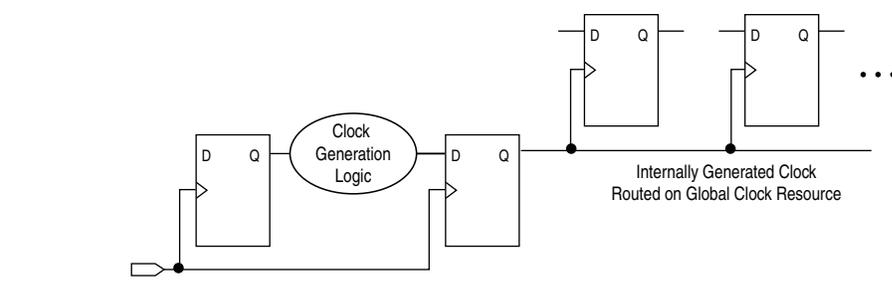
The following sections provide some specific examples and recommendations for avoiding clocking scheme problems.

Internally Generated Clocks

If you use the output from combinational logic as a clock signal or as an asynchronous reset signal, expect to see glitches in your design. In a synchronous design, glitches on data inputs of registers are normal events that have no consequences. However, a glitch or a spike on the clock input (or an asynchronous input) to a register can have significant consequences. Narrow glitches can violate the register's minimum pulse width requirements. Setup and hold times might also be violated if the data input of the register is changing when a glitch reaches the clock input. Even if the design does not violate timing requirements, the register output can change value unexpectedly and cause functional hazards elsewhere in the design.

To avoid these problems, you should always register the output of combinational logic before you use it as a clock signal (Figure 9-4).

Figure 9-4. Recommended Clock-Generation Technique



Registering the output of combinational logic ensures that the glitches generated by the combinational logic are blocked at the data input of the register.

Divided Clocks

Designs often require clocks created by dividing a master clock. Most Altera FPGAs provide dedicated phase-locked loop (PLL) circuitry for clock division. Using dedicated PLL circuitry can help you to avoid many of the problems that can be introduced by asynchronous clock division logic.

When you must use logic to divide a master clock, always use synchronous counters or state machines. Additionally, create your design so that registers always directly generate divided clock signals, as described in “Internally Generated Clocks”, and route the clock on global clock resources. To avoid glitches, do not decode the outputs of a counter or a state machine to generate clock signals.

Ripple Counters

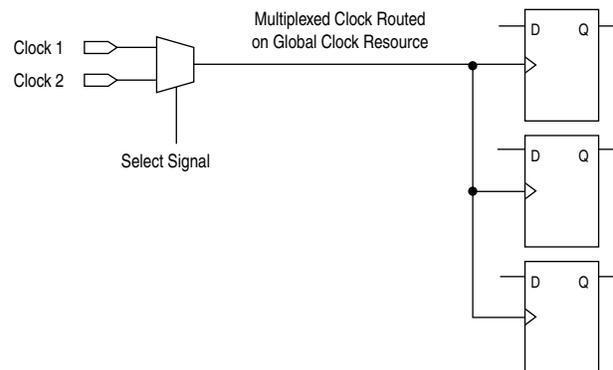
To simplify verification, avoid ripple counters in your design. In the past, FPGA designers implemented ripple counters to divide clocks by a power of two because the counters are easy to design and may use fewer gates than their synchronous counterparts. Ripple counters use cascaded registers, in which the output pin of each register feeds the clock pin of the register in the next stage. This cascading can cause problems because the counter creates a ripple clock at each stage. These ripple clocks must be handled properly during timing analysis, which can be difficult and may require you to make complicated timing assignments in your synthesis and placement and routing tools.

Ripple clock structures are often used to make ripple counters out of the smallest amount of logic possible. However, in all Altera devices supported by the Quartus II software, using a ripple clock structure to reduce the amount of logic used for a counter is unnecessary because the device allows you to construct a counter using one logic element per counter bit. You should avoid using ripple counters completely.

Multiplexed Clocks

Use clock multiplexing to operate the same logic function with different clock sources. In these designs, multiplexing selects a clock source, as shown in Figure 9-5. For example, telecommunications applications that deal with multiple frequency standards often use multiplexed clocks.

Figure 9-5. Multiplexing Logic and Clock Sources



Adding multiplexing logic to the clock signal can create the problems addressed in the previous sections, but requirements for multiplexed clocks vary widely, depending on the application. Clock multiplexing is acceptable when the clock signal uses global clock routing resources and if the following criteria are met:

- The clock multiplexing logic does not change after initial configuration
- The design uses multiplexing logic to select a clock for testing purposes
- Registers are always reset when the clock switches
- A temporarily incorrect response following clock switching has no negative consequences

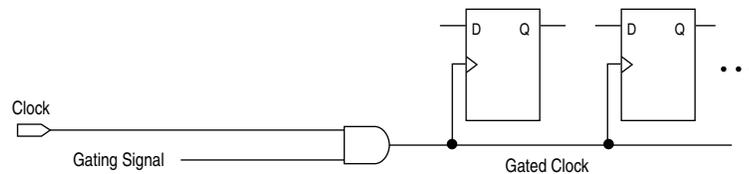
If the design switches clocks in real time with no reset signal, and your design cannot tolerate a temporarily incorrect response, you must use a synchronous design so that there are no timing violations on the registers, no glitches on clock signals, and no race conditions or other logical problems. By default, the Quartus II software optimizes and analyzes all possible paths through the multiplexer and between both internal clocks that may come from the multiplexer. This may lead to more restrictive analysis than required if the multiplexer is always selecting one particular clock. If you do not require the more complete analysis, you can assign the output of the multiplexer as a base clock in the Quartus II software, so that all register-to-register paths are analyzed using that clock.

-  Use dedicated hardware to perform clock multiplexing when it is available, instead of using multiplexing logic. For example, you can use the clock-switchover feature or clock control block available in certain Altera devices. These dedicated hardware blocks ensure that you use global low-skew routing lines and avoid any possible hold time problems on the device due to logic delay on the clock line.
-  For device-specific information about clocking structures, refer to the appropriate device data sheet or handbook on the [Literature](#) page of the Altera website.

Gated Clocks

Gated clocks turn a clock signal on and off using an enable signal that controls gating circuitry, as shown in [Figure 9-6](#). When a clock is turned off, the corresponding clock domain is shut down and becomes functionally inactive.

Figure 9-6. Gated Clock



You can use gated clocks to reduce power consumption in some device architectures by effectively shutting down portions of a digital circuit when they are not in use. When a clock is gated, both the clock network and the registers driven by it stop toggling, thereby eliminating their contributions to power consumption. However, gated clocks are not part of a synchronous scheme and therefore can significantly increase the effort required for design implementation and verification. Gated clocks contribute to clock skew and make device migration difficult. These clocks are also sensitive to glitches, which can cause design failure.

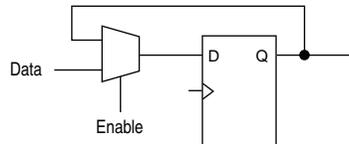
Use dedicated hardware to perform clock gating rather than an AND or OR gate. For example, you can use the clock control block in newer Altera devices to shut down an entire clock network. Dedicated hardware blocks ensure that you use global routing with low skew and avoid any possible hold time problems on the device due to logic delay on the clock line.

From a functional point of view, you can shut down a clock domain in a purely synchronous manner using a synchronous clock enable signal. However, when using a synchronous clock enable scheme, the clock network continues toggling. This practice does not reduce power consumption as much as gating the clock at the source does. In most cases, use a synchronous scheme such as those described in [“Synchronous Clock Enables”](#). For improved power reduction when gating clocks with logic, refer to [“Recommended Clock-Gating Methods”](#) on page 9-11.

Synchronous Clock Enables

To turn off a clock domain in a synchronous manner, use a synchronous clock enable signal. FPGAs efficiently support clock enable signals because there is a dedicated clock enable signal available on all device registers. This scheme does not reduce clock enable signal consumption as much as gating the clock at the source because the clock network keeps toggling, but it performs the same function as a gated clock by disabling a set of registers. Insert a multiplexer in front of the data input of every register to either load new data or copy the output of the register (Figure 9-7).

Figure 9-7. Synchronous Clock Enable

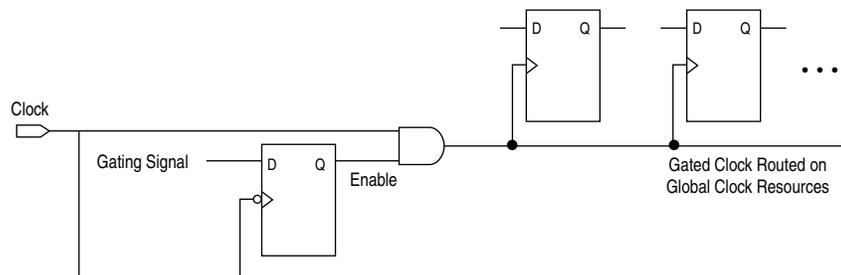


Recommended Clock-Gating Methods

Use gated clocks only when your target application requires power reduction and when gated clocks are able to provide the required reduction in your device architecture. If you must use clocks gated by logic, implement these clocks using the robust clock-gating technique shown in Figure 9-8 and ensure that the gated clock signal uses dedicated global clock routing.

You can gate a clock signal at the source of the clock network, at each register, or somewhere in between. Because the clock network contributes to switching power consumption, gate the clock at the source whenever possible, so you can shut down the entire clock network instead of gating it further along the clock network at the registers.

Figure 9-8. Recommended Clock-Gating Technique



In the technique shown in Figure 9-8, a register generates the enable signal to ensure that the signal is free of glitches and spikes. The register that generates the enable signal is triggered on the inactive edge of the clock to be gated. Use the falling edge when gating a clock that is active on the rising edge, as shown in Figure 9-8. Using this technique, only one input of the gate that turns the clock on and off changes at a time. This prevents any glitches or spikes on the output. Use an AND gate to gate a clock that is active on the rising edge. For a clock that is active on the falling edge, use an OR gate to gate the clock and register the enable command with a positive edge-triggered register.

When using this technique, pay attention to the duty cycle of the clock and the delay through the logic that generates the enable signal because the enable command must be generated in one-half the clock cycle. This situation might cause problems if the logic that generates the enable command is particularly complex, or if the duty cycle of the clock is severely unbalanced. However, careful management of the duty cycle and logic delay may be an acceptable solution when compared with problems created by other methods of gating clocks.

Ensure that you apply a clock setting to the gated clock in the Quartus II software. As shown in [Figure 9-8 on page 9-11](#), apply a clock setting to the output of the AND gate. Otherwise, the timing analyzer might analyze the circuit using the clock path through the register as the longest clock path and the path that skips the register as the shortest clock path, resulting in artificial clock skew.

In certain cases, converting the gated clocks to clock enables may help to reduce glitch and clock skew, and eventually produce a more accurate timing analysis. You can set the Quartus II software to automatically convert gated clocks to clock enables by turning on the **Auto Gated Clock Conversion** option. The conversion applies to two types of gated clocking schemes: single-gated clock and cascaded-gated clock. This option is available for all devices that are supported by the TimeQuest analyzer (Arria® II, Arria II GX, Cyclone® II, Cyclone III, Cyclone IV, HardCopy series, Stratix® II, Stratix II GX, Stratix III, Stratix IV, and Stratix V devices).

-  For information about the settings and limitations of this option, refer to the “Auto Gated Clock Conversion” section of the [Quartus II Integrated Synthesis](#) chapter in volume 1 of the *Quartus II Handbook*.

Power Optimization

The total FPGA power consumption is comprised of I/O power, core static power, and core dynamic power. Knowledge of the relationship between these components is fundamental in calculating the overall total power consumption. You can use various optimization techniques and tools to minimize power consumption when applied during FPGA design implementation. The Quartus II software offers power-driven compilation features to fully optimize device power consumption. Power-driven compilation focuses on reducing your design’s total power consumption using power-driven synthesis and power-driven placement and routing.

-  For information about power-driven compilation flow and low-power design guidelines, refer to the [Power Optimization](#) chapter in volume 2 of the *Quartus II Handbook*.
-  For information about power optimization techniques available for Stratix III devices, refer to [AN 437: Power Optimization in Stratix III FPGAs](#). For information about power optimization techniques available for Stratix IV devices, refer to [AN 514: Power Optimization in Stratix IV FPGAs](#).
-  Additionally, you can use the Quartus II PowerPlay suite of power analysis and optimization tools to help you during the design process by delivering fast and accurate estimations of power consumption. For information about the Quartus II PowerPlay suite of power analysis and optimization tools, refer to [About Power Estimation and Analysis](#) in Quartus II Help.

Metastability

Metastability in Altera designs can be caused by the synchronization of asynchronous signals. You can use the Quartus II software to analyze the mean time between failures (MTBF) due to metastability, thus optimizing the design to improve the metastability MTBF. A high metastability MTBF indicates a more robust design.

-  For more information about how to ensure complete and accurate metastability analysis, refer to the *Managing Metastability With the Quartus II Software* chapter in volume 1 of the *Quartus II Handbook*.
-  For more information about viewing metastability reports, refer to *Viewing Metastability Reports* in Quartus II Help.

Incremental Compilation

The incremental compilation feature in the Quartus II software allows you to partition your design, separately compile partitions, and reuse the results for unchanged partitions. Incremental compilation flows require more up-front planning than flat compilations, and generally require you to be more rigorous about following good design practices than flat compilations.

-  For more information about incremental compilation and floorplan assignments, refer to the *Best Practices for Incremental Compilation Partitions and Floorplan Assignments* chapter in volume 1 of the *Quartus II Handbook*.
-  For more information about incremental compilation, refer to *About Incremental Compilation* in Quartus II Help.

Checking Design Violations With the Design Assistant

To improve the reliability, timing performance, and logic utilization of your design, practicing good design methodology and understanding how to avoid design rule violations are important. The Quartus II software provides the Design Assistant tool that automatically checks for design rule violations and reports their location.

The Design Assistant is a design rule checking tool that allows you to check for design issues early in the design flow. The Design Assistant checks your design for adherence to Altera-recommended design guidelines. You can specify which rules you want the Design Assistant to apply to your design. This is useful if you know that your design violates particular rules that are not critical, so you can allow these rule violations. The Design Assistant generates design violation reports with clear details about each violation, based on the settings that you specified.

This section provides an introduction to the Quartus II design flow with the Design Assistant, message severity levels, and an explanation about how to set up the Design Assistant. The last parts of the section describe the design rules and the reports generated by the Design Assistant. The Design Assistant supports all Altera devices supported by the Quartus II software.

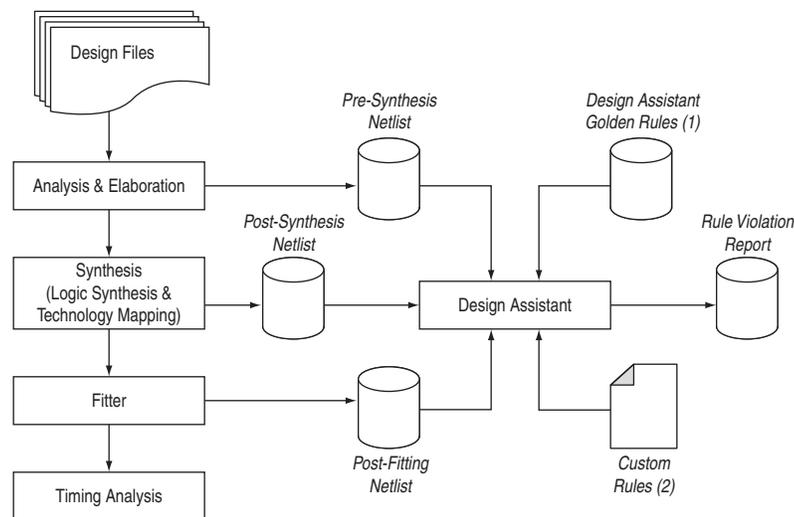
Quartus II Design Flow with the Design Assistant

You can run the Design Assistant after Analysis and Elaboration, Analysis and Synthesis, fitting, or a full compilation. If you set the Design Assistant to run automatically during compilation, the Design Assistant performs a post-fitting netlist analysis of your design. The default is to apply all of the rules to your project. If there are some rules that are unimportant to your design, you can turn off the rules that you do not want the Design Assistant to use.

- ④ For more information about running the Design Assistant, refer to *About the Design Assistant* in Quartus II Help.

Figure 9-9 shows the Quartus II software design flow with the Design Assistant.

Figure 9-9. Quartus II Design Flow with the Design Assistant



Notes to Figure 9-9:

- (1) Database of the default rules for the Design Assistant.
- (2) A file that contains the **.xml** codes of the custom rules for the Design Assistant. For more details about how to create this file, refer to “Custom Rules” on page 9-15.

The Design Assistant analyzes your design netlist at different stages of the compilation flow and may yield different warnings or errors, even though the netlists are functionally the same. Your pre-synthesis, post-synthesis, and post-fitting netlists might be different due to optimizations performed by the Quartus II software. For example, a warning message in a pre-synthesis netlist may be removed after the netlist has been synthesized into a post-synthesis or post-fitting netlist.

The exact operation of the Design Assistant depends on when you run it:

- When you run the Design Assistant after running a full compilation or fitting, the Design Assistant performs a post-fitting analysis on the design.
- When you start the Design Assistant after performing Analysis and Synthesis, the Design Assistant performs post-synthesis analysis on the design.

- When you start the Design Assistant after performing Analysis and Elaboration, the Design Assistant performs a pre-synthesis analysis on the design. You can also perform pre-synthesis analysis with the Design Assistant using the command-line. You can use the `-rtl` option with the `quartus_drc` executable, as shown in the following example:

```
quartus_drc <project_name> --rtl=on ←
```

- ④ For more information about Design Assistant settings, refer to *About the Design Assistant* and *Design Assistant Page (Settings Dialog Box)* in Quartus II Help.

Enabling and Disabling Design Assistant Rules

- ④ For more information about enabling or disabling Design Assistant rules on individual nodes by making an assignment in the Assignment Editor, in the Quartus II Settings File (`.qsf`), with the `altera_attribute` synthesis attribute in Verilog HDL or VHDL, or with a Tcl command, refer to *Enabling Design Assistant Rules on Nodes, Entities, or Instances*, or *Disabling Design Assistant Rules on Nodes, Entities, or Instances* in Quartus II Help.

Viewing Design Assistant Results

If your design violates a design rule, the Design Assistant generates warning messages and information messages about the violated design rule. The Design Assistant displays these messages in the Messages window, in the Design Assistant Messages report, and in the Design Assistant report files. You can find the Design Assistant report files called `<project_name>.drc.rpt` in the `<project_name>` subdirectory of the project directory.

- ④ For information about the contents of the reports generated by the Design Assistant, refer to *Design Assistant Reports* in Quartus II Help.

Custom Rules

In addition to the existing design rules that the Design Assistant offers, you can also create your own rules and specify your own reporting format in a text file (with any file extension) with the XML format. You then specify the path to that file in the Design Assistant settings page and run the Design Assistant for violation checking.

The file that contains the default rules for the Design Assistant is located at `<Quartus II install path>\quartus\libraries\design-assistant\da_golden_rule.xml`.

- ④ For more information about how to set the file path to your custom rules, refer to *Custom Rules Settings Dialog Box* in Quartus II Help. For more information about the basics of writing custom rules, the Design Assistant settings, and coding examples on how to check for clock relationship and node relationship in a design, refer to *Creating Custom Design Assistant Rules* in Quartus II Help. To specify the rules that you want the Design Assistant to use when checking for violations, refer to *Design Assistant Page (Settings Dialog Box)* in Quartus II Help.

Custom Rules Coding Examples

The following examples of custom rules show how to check node relationships and clock relationships in a design.

Checking SR Latch Structures In a Design

[Example 9-1](#) shows the XML codes for checking SR latch structures in a design.

Example 9-1. Detecting SR Latches in a Design

```
<DA_RULE ID="EX01" SEVERITY="CRITICAL" NAME="Checking Design for SR Latch"
DEFAULT_RUN="YES">
<RULE_DEFINITION>
  <FORBID>
    <OR>
      <NODE NAME="NODE_1" TYPE="SRLATCH" />
      <HAS_NODE NODE_LIST="NODE_1" />
      <NODE NAME="NODE_1" TOTAL_FANIN="EQ2" />
      <NODE NAME="NODE_2" TOTAL_FANIN="EQ2" />
    <AND>
      <NODE_RELATIONSHIP FROM_NAME="NODE_1" FROM_TYPE="NAND" TO_NAME="NODE_2"
TO_TYPE="NAND" />
      <NODE_RELATIONSHIP FROM_NAME="NODE_2" FROM_TYPE="NAND" TO_NAME="NODE_1"
TO_TYPE="NAND" />
    </AND>
    <AND>
      <NODE_RELATIONSHIP FROM_NAME="NODE_1" FROM_TYPE="NOR" TO_NAME="NODE_2"
TO_TYPE="NOR" />
      <NODE_RELATIONSHIP FROM_NAME="NODE_2" FROM_TYPE="NOR" TO_NAME="NODE_1"
TO_TYPE="NOR" />
    </AND>
  </OR>
</FORBID>
</RULE_DEFINITION>

<REPORTING_ROOT>
  <MESSAGE NAME="Rule %ARG1%: Found %ARG2% node(s) related to this rule.">
    <MESSAGE_ARGUMENT NAME="ARG1" TYPE="ATTRIBUTE" VALUE="ID" />
    <MESSAGE_ARGUMENT NAME="ARG2" TYPE="TOTAL_NODE" VALUE="NODE_1" />
  </MESSAGE>
</REPORTING_ROOT>
</DA_RULE>
```

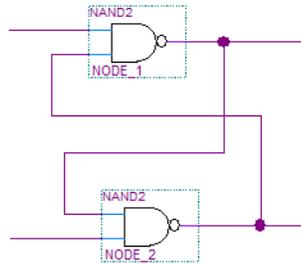
In [Example 9-1](#), the possible SR latch structures are specified in the rule definition section. Codes defined in the `<AND></AND>` block are tied together, meaning that each statement in the block must be true for the block to be fulfilled (AND gate similarity). In the `<OR></OR>` block, as long as one statement in the block is true, the block is fulfilled (OR gate similarity). If no `<AND></AND>` or `<OR></OR>` block are specified, the default is `<AND></AND>`.

The `<FORBID></FORBID>` section contains the undesirable condition for the design, which in this case is the SR latch structures. If the condition is fulfilled, the Design Assistant highlights a rule violation.

The following examples are the undesired conditions from [Example 9-1](#) with their equivalent block diagrams ([Figure 9-10](#) and [Figure 9-11](#)):

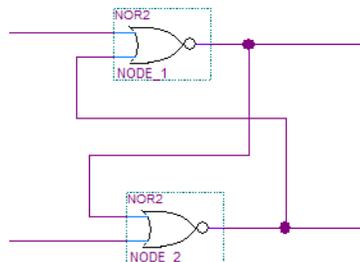
```
<AND>
  <NODE_RELATIONSHIP FROM_NAME="NODE_1" FROM_TYPE="NAND" TO_NAME="NODE_2"
TO_TYPE="NAND" />
  <NODE_RELATIONSHIP FROM_NAME="NODE_2" FROM_TYPE="NAND" TO_NAME="NODE_1"
TO_TYPE="NAND" />
</AND>
```

Figure 9-10. Undesired Condition 1



```
<AND>
  <NODE_RELATIONSHIP FROM_NAME="NODE_1" FROM_TYPE="NOR" TO_NAME="NODE_2" TO_TYPE="NOR" />
  <NODE_RELATIONSHIP FROM_NAME="NODE_2" FROM_TYPE="NOR" TO_NAME="NODE_1" TO_TYPE="NOR" />
</AND>
```

Figure 9-11. Undesired Condition 2



Relating Nodes to a Clock Domain

[Example 9-2](#) shows how to use the `CLOCK_RELATIONSHIP` attribute to relate nodes to clock domains. This example checks for correct synchronization in data transfer between asynchronous clock domains. Synchronization is done with cascaded registers, also called synchronizers, at the receiving clock domain. The code in [Example 9-2](#) checks for the synchronizer configuration based on the following guidelines:

- The cascading registers need to be triggered on the same clock edge

- There is no logic between the register output of the transmitting clock domain and the cascaded registers in the receiving asynchronous clock domain

Example 9-2. Detecting Incorrect Synchronizer Configuration

```
<DA_RULE ID="EX02" SEVERITY="HIGH" NAME="Data Transfer Not Synch Correctly"
DEFAULT_RUN="YES">

<RULE_DEFINITION>
<DECLARE>
  <NODE NAME="NODE_1" TYPE="REG" />
  <NODE NAME="NODE_2" TYPE="REG" />
  <NODE NAME="NODE_3" TYPE="REG" />
</DECLARE>
<FORBID>
  <NODE_RELATIONSHIP FROM_NAME="NODE_1" TO_NAME="NODE_2" TO_PORT="D_PORT"
CLOCK_RELATIONSHIP="ASYN" />
  <NODE_RELATIONSHIP FROM_NAME="NODE_2" TO_NAME="NODE_3" TO_PORT="D_PORT"
CLOCK_RELATIONSHIP="!ASYN" />
  <OR>
    <NODE_RELATIONSHIP FROM_NAME="NODE_1" TO_NAME="NODE_2" TO_PORT="D_PORT"
REQUIRED_THROUGH="YES" THROUGH_TYPE="COMB" CLOCK_RELATIONSHIP="ASYN" />
    <CLOCK_RELATIONSHIP NAME="SEQ_EDGE|ASYN" NODE_LIST="NODE_2, NODE_3" />
  </OR>
</FORBID>
</RULE_DEFINITION>

<REPORTING_ROOT>
<MESSAGE NAME="Rule %ARG1%: Found %ARG2% node(s) related to this rule.">
  <MESSAGE_ARGUMENT NAME="ARG1" TYPE="ATTRIBUTE" VALUE="ID" />
  <MESSAGE_ARGUMENT NAME="ARG2" TYPE="TOTAL_NODE" VALUE="NODE_1" />
  <MESSAGE NAME="Source node(s): %ARG3%, Destination node(s): %ARG4%">
    <MESSAGE_ARGUMENT NAME="ARG3" TYPE="NODE" VALUE="NODE_1" />
    <MESSAGE_ARGUMENT NAME="ARG4" TYPE="NODE" VALUE="NODE_2" />
  </MESSAGE>
</MESSAGE>
</REPORTING_ROOT>
</DA_RULE>
```

The codes differentiate the clock domains. ASYN means asynchronous, and !ASYN means non-asynchronous. This notation is useful for describing nodes that are in different clock domains. The following lines from [Example 9-2](#) state that NODE_2 and NODE_3 are in the same clock domain, but NODE_1 is not.

```
<NODE_RELATIONSHIP FROM_NAME="NODE_1" TO_NAME="NODE_2" TO_PORT="D_PORT"
CLOCK_RELATIONSHIP="ASYN" />

<NODE_RELATIONSHIP FROM_NAME="NODE_2" TO_NAME="NODE_3" TO_PORT="D_PORT"
CLOCK_RELATIONSHIP="!ASYN" />
```

The next line of code states that NODE_2 and NODE_3 have a clock relationship of either sequential edge or asynchronous.

```
<CLOCK_RELATIONSHIP NAME="SEQ_EDGE|ASYN" NODE_LIST="NODE_2, NODE_3" />
```

The <FORBID></FORBID> section contains the undesirable condition for the design, which in this case is the undesired configuration of the synchronizer. If the condition is fulfilled, the Design Assistant highlights a rule violation.

The following examples are the undesired conditions from [Example 9-2](#) with their equivalent block diagrams ([Figure 9-12](#) and [Figure 9-13](#)):

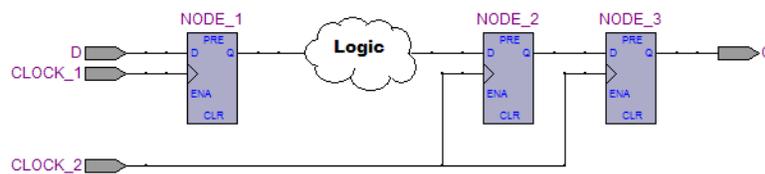
Example 9-3.

```
<NODE_RELATIONSHIP FROM_NAME="NODE_1" TO_NAME="NODE_2" TO_PORT="D_PORT"
CLOCK_RELATIONSHIP="ASYN" />

<NODE_RELATIONSHIP FROM_NAME="NODE_2" TO_NAME="NODE_3" TO_PORT="D_PORT"
CLOCK_RELATIONSHIP="!ASYN" />

<NODE_RELATIONSHIP FROM_NAME="NODE_1" TO_NAME="NODE_2" TO_PORT="D_PORT"
REQUIRED_THROUGH="YES" THROUGH_TYPE="COMB" CLOCK_RELATIONSHIP="ASYN" />
```

Figure 9-12. Undesired Condition 3



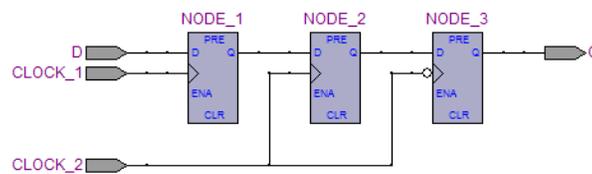
Example 9-4.

```
<NODE_RELATIONSHIP FROM_NAME="NODE_1" TO_NAME="NODE_2" TO_PORT="D_PORT"
CLOCK_RELATIONSHIP="ASYN" />

<NODE_RELATIONSHIP FROM_NAME="NODE_2" TO_NAME="NODE_3" TO_PORT="D_PORT"
CLOCK_RELATIONSHIP="!ASYN" />

<CLOCK_RELATIONSHIP NAME="SEQ_EDGE|ASYN" NODE_LIST="NODE_2, NODE_3" />
```

Figure 9-13. Undesired Condition 4



Targeting Clock and Register-Control Architectural Features

In addition to following general design guidelines, you must code your design with the device architecture in mind. FPGAs provide device-wide clocks and register control signals that can improve performance.

Clock Network Resources

Altera FPGAs provide device-wide global clock routing resources and dedicated inputs. Use the FPGA's low-skew, high fan-out dedicated routing where available. By assigning a clock input to one of these dedicated clock pins or with a Quartus II logic option to assign global routing, you can take advantage of the dedicated routing available for clock signals.

In an ASIC design, you should balance the clock delay as it is distributed across the device. Because Altera FPGAs provide device-wide global clock routing resources and dedicated inputs, there is no need to manually balance delays on the clock network.

You should limit the number of clocks in your design to the number of dedicated global clock resources available in your FPGA. Clocks feeding multiple locations that do not use global routing may exhibit clock skew across the device that could lead to timing problems. In addition, when you use combinational logic to generate an internal clock, it adds delays on the clock line. In some cases, delay on a clock line can result in a clock skew greater than the data path length between two registers. If the clock skew is greater than the data delay, the timing parameters of the register (such as hold time requirements) are violated and the design does not function correctly.

FPGAs offer increasing numbers of global clocks to address large designs with many clock domains. Many large FPGA devices provide dedicated global clock networks, regional clock networks, and dedicated fast regional clock networks. These clocks are organized into a hierarchical clock structure that allows many clocks in each device region with low skew and delay. There are typically several dedicated clock pins to drive either global or regional clock networks, and both PLL outputs and internal clocks can drive various clock networks.

To reduce clock skew in a given clock domain and ensure that hold times are met in that clock domain, assign each clock signal to one of the global high fan-out, low-skew clock networks in the FPGA device. The Quartus II software automatically uses global routing for high fan-out control signals, PLL outputs, and signals feeding the global clock pins on the device. You can make explicit Global Signal logic option settings by turning on the **Global Signal** option settings. Use this option when it is necessary to force the software to use the global routing for particular signals.

To take full advantage of these routing resources, the sources of clock signals in a design (input clock pins or internally-generated clocks) need to drive only the clock input ports of registers. In older Altera device families (such as FLEX[®] 10K and ACEX[®] 1K), if a clock signal feeds the data ports of a register, the signal may not be able to use dedicated routing, which can lead to decreased performance and clock skew problems. In general, allowing clock signals to drive the data ports of registers is not considered synchronous design and can complicate timing analysis.

Reset Resources

ASIC designs may use local resets to avoid long routing delays. Take advantage of the device-wide asynchronous reset pin available on most FPGAs to eliminate these problems. This reset signal provides low-skew routing across the device.

Three types of resets are used in synchronous circuits:

- Synchronous Reset
- Asynchronous Reset
- Synchronized Asynchronous Reset—this type is preferred when designing an FPGA circuit

Synchronous Reset

The synchronous reset ensures that the circuit is fully synchronous. You can easily time it with static timing analyzer tools, such as the Quartus II TimeQuest analyzer. The synchronous reset is easier to use with cycle-based simulators.

However, the synchronous reset might require pulse stretchers to guarantee a reset pulse width wide enough to ensure that reset is present during an active edge of the clock. The synchronous reset requires a clock to reset a circuit. If the clock fails to launch, the resulting circuit is not reset.

Asynchronous Reset

The asynchronous reset is the most common form of reset used in circuit designs. Typically, you can insert the asynchronous reset into the device, turn on the global buffer, and connect to the asynchronous reset pin of every register in the device. This method is only advantageous under certain circumstances—you do not need to always reset the register. Unlike the synchronous reset, the asynchronous reset is not inserted in the data path, and does not negatively impact the data arrival times between registers. Reset takes effect immediately, and as soon as the registers receive the reset pulse, the registers are reset. The asynchronous reset is not dependent on the clock.

However, when the reset is deasserted and does not pass the recovery (μt_{SU}) or removal (μt_{H}) time check (both times are checked by the TimeQuest analyzer Recovery and Removal Analysis), the edge is said to have fallen in the metastability zone. Additional time is required to determine the correct state, and the delay can cause the setup time to fail to register downstream, leading to system failure.

The asynchronous reset is susceptible to noise, and a noisy asynchronous reset can cause a spurious reset. You must ensure that the asynchronous reset is debounced and filtered. You can easily enter into a reset asynchronously, but releasing a reset asynchronously can lead to potential problems (also referred to as “reset removal”) with metastability, including the hazards of unwanted situations with synchronous circuits involving feedback.

Example 9–5 shows the equivalent Verilog code. The active edge of the reset is used in the sensitivity list for the blocks in Figure 9–14.

Example 9–5. Verilog Code for Synchronized Asynchronous Reset

```
module sync_async_reset (
    input    clock,
    input    reset_n,
    input    data_a,
    input    data_b,
    output   out_a,
    output   out_b
);

reg    reg1, reg2;
reg    reg3, reg4;

assign out_a    = reg1;
assign out_b    = reg2;
assign rst_n    = reg4;

always @ (posedge clock, negedge reset_n)
begin
    if (!reset_n)
    begin
        reg3    <= 1'b0;
        reg4    <= 1'b0;
    end
    else
    begin
        reg3    <= 1'b1;
        reg4    <= reg3;
    end
end

always @ (posedge clock, negedge rst_n)
begin
    if (!rst_n)
    begin
        reg1    <= 1'b0;
        reg2    <= 1'b0;
    end
    else
    begin
        reg1    <= data_a;
        reg2    <= data_b;
    end
end

endmodule // sync_async_reset
```

To minimize the metastability effect between the two synchronization registers, and to increase the MTBF, the registers should be located as close as possible in the device to minimize routing delay. If possible, locate the registers in the same logic array block (LAB). The input reset signal (`reset_n`) must be cut with a `set_false_path` command, so the reset that comes from the synchronization register (`rst_n`) can be timed in the TimeQuest analyzer with Recovery and Removal Analysis.



For more information about specifying the minimum routing delay, refer to the *Best Practices for the Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

The circuit in [Figure 9-14 on page 9-22](#) ensures that the synchronized asynchronous reset is at least one full clock period in length. To extend this time to n clock periods, you must increase the number of synchronizer registers to $n + 1$. You must connect the asynchronous input reset (`reset_n`) to the `CLRN` pin of all the synchronizer registers to maintain the asynchronous assertion of the synchronized asynchronous reset.

Register Control Signals

Avoid using an asynchronous load signal if the design target device architecture does not include registers with dedicated circuitry for asynchronous loads. Also, avoid using both asynchronous clear and preset if the architecture provides only one of these control signals. Stratix III devices, for example, directly support an asynchronous clear function, but not a preset or load function. When the target device does not directly support the signals, the synthesis or placement and routing software must use combinational logic to implement the same functionality. In addition, if you use signals in a priority other than the inherent priority in the device architecture, combinational logic may be required to implement the necessary control signals. Combinational logic is less efficient and can cause glitches and other problems; it is best to avoid these implementations.



For Verilog HDL and VHDL examples of registers with various control signals, and information about the inherent priority order of register control signals in Altera device architecture, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*.

Targeting Embedded RAM Architectural Features

Altera's dedicated memory architecture offers many advanced features that you can target easily with the MegaWizard™ Plug-In Manager or with the recommended HDL coding styles that infer the appropriate RAM megafunction (ALTSYNCRAM or ALTDPRAM). Use synchronous memory blocks for your design, so the blocks can be mapped directly into the device dedicated memory blocks. You can use single-port, dual-port, or three-port RAM with a single- or dual-clocking method. Asynchronous memory logic is not inferred as a memory block or placed in the dedicated memory block, but is implemented in regular logic cells.

Altera memory blocks have different read-during-write behaviors, depending on the targeted device family, memory mode, and block type. Read-during-write behavior refers to read and write from the same memory address in the same clock cycle; for example, you read from the same address to which you write in the same clock cycle.

It is important to check how you specify the memory in your HDL code when you use read-during-write behavior. The HDL code that describes the read returns either the old data stored at the memory location, or the new data being written to the memory location.

In some cases, when the device architecture cannot implement the memory behavior described in your HDL code, the memory block is not mapped to the dedicated RAM blocks, or the memory block is implemented using extra logic in addition to the dedicated RAM block. Implement the read-during-write behavior using single-port RAM in Arria GX devices and the Cyclone and Stratix series of devices to avoid this extra logic implementation.

- For Verilog HDL and VHDL examples and guidelines for inferring RAM functions that match the dedicated memory architecture in Altera devices, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*.

In many synthesis tools, you can specify that the read-during-write behavior is not important to your design; if, for example, you never read and write from the same address in the same clock cycle. For Quartus II integrated synthesis, add the synthesis attribute `ramstyle="no_rw_check"` to allow the software to choose the read-during-write behavior of a RAM, rather than using the read-during-write behavior specified in your HDL code. Using this type of attribute prevents the synthesis tool from using extra logic to implement the memory block and, in some cases, can allow memory inference when it would otherwise be impossible.

- For details about using the `ramstyle` attribute, refer to the *Quartus II Integrated Synthesis* chapter in volume 1 of the *Quartus II Handbook*. For information about the synthesis attributes in other synthesis tools, refer to your synthesis tool documentation, or to the appropriate chapter in the *Synthesis* section in volume 1 of the *Quartus II Handbook*.

Conclusion

Following the design practices described in this chapter can help you to consistently meet your design goals. Asynchronous design techniques may result in incomplete timing analysis, may cause glitches on data signals, and may rely on propagation delays in a device leading to race conditions and unpredictable results. Taking advantage of the architectural features in your FPGA device can also improve the quality of your results.

Document Revision History

Table 9-1 shows the revision history for this chapter.

Table 9-1. Document Revision History

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Title changed from Design Recommendations for Altera Devices and the Quartus II Design Assistant. ■ Updated to new template. ■ Added references to Quartus II Help for “Metastability” on page 9-13 and “Incremental Compilation” on page 9-13. ■ Removed duplicated content and added references to Quartus II Help for “Custom Rules” on page 9-15.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Removed duplicated content and added references to Quartus II Help for Design Assistant settings, Design Assistant rules, Enabling and Disabling Design Assistant Rules, and Viewing Design Assistant reports. ■ Removed information from “Combinational Logic Structures” on page 5-4 ■ Changed heading from “Design Techniques to Save Power” to “Power Optimization” on page 5-12 ■ Added new “Metastability” section ■ Added new “Incremental Compilation” section ■ Added information to “Reset Resources” on page 5-23 ■ Removed “Referenced Documents” section
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Removed documentation of obsolete rules.
March 2009	9.0.0	<ul style="list-style-type: none"> ■ No change to content.
November 2008	8.1.0	<ul style="list-style-type: none"> ■ Changed to 8-1/2 x 11 page size ■ Added new section “Custom Rules Coding Examples” on page 5-18 ■ Added paragraph to “Recommended Clock-Gating Methods” on page 5-11 ■ Added new section: “Design Techniques to Save Power” on page 5-12
May 2008	8.0.0	<ul style="list-style-type: none"> ■ Updated Figure 5-9 on page 5-13; added custom rules file to the flow ■ Added notes to Figure 5-9 on page 5-13 ■ Added new section: “Custom Rules Report” on page 5-34 ■ Added new section: “Custom Rules” on page 5-34 ■ Added new section: “Targeting Embedded RAM Architectural Features” on page 5-38 ■ Minor editorial updates throughout the chapter ■ Added hyperlinks to referenced documents throughout the chapter

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter provides Hardware Description Language (HDL) coding style recommendations to ensure optimal synthesis results when targeting Altera® devices.

HDL coding styles can have a significant effect on the quality of results that you achieve for programmable logic designs. Synthesis tools optimize HDL code for both logic utilization and performance, however, synthesis tools have no information about the purpose or intent of the design. The best optimizations often require conscious interaction by you, the designer.

This chapter includes the following sections:

- “Quartus II Language Templates”
- “Using Altera Megafunctions” on page 10–2
- “Instantiating Altera Megafunctions in HDL Code” on page 10–3
- “Inferring Multiplier and DSP Functions from HDL Code” on page 10–5
- “Inferring Memory Functions from HDL Code” on page 10–13
- “Coding Guidelines for Registers and Latches” on page 10–43
- “General Coding Guidelines” on page 10–53
- “Designing with Low-Level Primitives” on page 10–73

 For additional guidelines about structuring your design, refer to the *Design Recommendations for Altera Devices and the Quartus II Design Assistant* chapter in volume 1 of the *Quartus II Handbook*. For additional handcrafted techniques you can use to optimize design blocks for the adaptive logic modules (ALMs) in many Altera devices, including a collection of circuit building blocks and related discussions, refer to the *Advanced Synthesis Cookbook: A Design Guide for Stratix II, Stratix III, and Stratix IV Devices*.

 The Altera website also provides design examples for other types of functions and to target specific applications. For more information about design examples, refer to the [Design Examples](#) page and the [Reference Designs](#) page on the Altera website.

For style recommendations, options, or HDL attributes specific to your synthesis tool (including Quartus® II integrated synthesis and other EDA tools), refer to the tool vendor’s documentation or the appropriate chapter in the *Synthesis* section in volume 1 of the *Quartus II Handbook*.

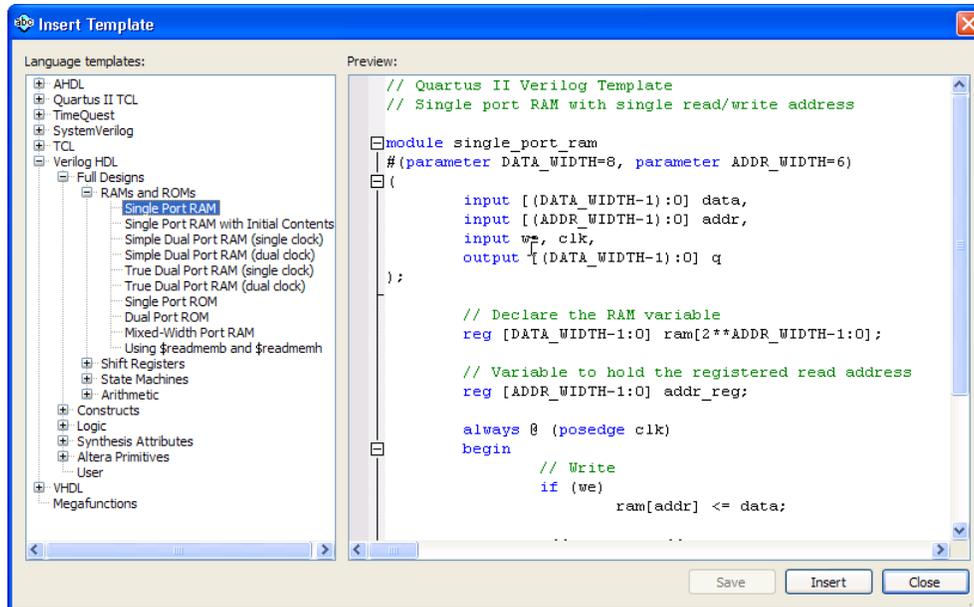
Quartus II Language Templates

Many of the Verilog HDL and VHDL examples in this document correspond with examples in the Full Designs section of the Quartus II Templates. You can easily insert examples into your HDL source code using the **Insert Template** dialog box in the Quartus II software user interface, shown in [Figure 10–1](#).

© 2010 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, HARDCOPY, MAX, MEGACORE, NIOS, QUARTUS and STRATIX are Reg. U.S. Pat. & Tm. Off. and/or trademarks of Altera Corporation in the U.S. and other countries. All other trademarks and service marks are the property of their respective holders as described at www.altera.com/common/legal.html. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera’s standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

To open the **Insert Template** dialog box when you have a file open in the Text Editor of the Quartus II software, on the Edit menu, click **Insert Template**. Alternatively, you can right-click in the Text Editor window and click **Insert Template**.

Figure 10-1. Insert Template Dialog Box



Using Altera Megafunctions

Altera provides parameterizable megafunctions that are optimized for Altera device architectures. Using megafunctions instead of coding your own logic saves valuable design time. Additionally, the Altera-provided megafunctions may offer more efficient logic synthesis and device implementation. You can scale the megafunction's size and specify various options by setting parameters. Megafunctions include the library of parameterized modules (LPM) and Altera device-specific megafunctions.

To use megafunctions in your HDL code, you can instantiate them as described in [“Instantiating Altera Megafunctions in HDL Code”](#) on page 10-3.

Sometimes it is preferable to make your code independent of device family or vendor. In this case, you might not want to instantiate megafunctions directly. For some types of logic functions, such as memories and DSP functions, you can infer device-specific dedicated architecture blocks instead of instantiating a megafunction. Synthesis tools, including Quartus II integrated synthesis, recognize certain types of HDL code and automatically infer the appropriate megafunction or map directly to device atoms. Synthesis tools infer megafunctions to take advantage of logic that is optimized for Altera devices or to target dedicated architectural blocks.

In cases where you prefer to use generic HDL code instead of instantiating a specific function, follow the guidelines and coding examples in [“Inferring Multiplier and DSP Functions from HDL Code”](#) on page 10-5 and [“Inferring Memory Functions from HDL Code”](#) on page 10-13 to ensure your HDL code infers the appropriate function.

-  You can infer or instantiate megafunctions to target some Altera device-specific architecture features such as memory and DSP blocks. You must instantiate megafunctions to target certain other device and high-speed features, such as LVDS drivers, phase-locked loops (PLLs), transceivers, and double-data rate input/output (DDIO) circuitry.

Instantiating Altera Megafunctions in HDL Code

The following sections describe how to use megafunctions by instantiating them in your HDL code with the following methods:

- “Instantiating Megafunctions Using the MegaWizard Plug-In Manager”—You can use the MegaWizard™ Plug-In Manager to parameterize the function and create a wrapper file.
- “Creating a Netlist File for Other Synthesis Tools”—You can optionally create a netlist file instead of a wrapper file.
- “Instantiating Megafunctions Using the Port and Parameter Definition”—You can instantiate the function directly in your HDL code.

Instantiating Megafunctions Using the MegaWizard Plug-In Manager

Use the MegaWizard Plug-In Manager as described in this section to create megafunctions in the Quartus II software that you can instantiate in your HDL code. The MegaWizard Plug-In Manager provides a GUI to customize and parameterize megafunctions, and ensures that you set all megafunction parameters properly. When you finish setting parameters, you can specify which files you want generated. Depending on which language you choose, the MegaWizard Plug-In Manager instantiates the megafunction with the correct parameters and generates a megafunction variation file (wrapper file) in Verilog HDL (.v), VHDL (.vhd), or AHDL (.tdf), along with other supporting files.

The MegaWizard Plug-In Manager provides options to create the files listed in Table 10-1.

Table 10-1. MegaWizard Plug-In Manager Generated Files (Part 1 of 2)

File	Description
<code><output file>.v .vhd .tdf (1)</code>	Verilog HDL Variation Wrapper File—Megafunction wrapper file for instantiation in a Verilog HDL, VHDL, or AHDL design respectively.
<code><output file>.inc</code>	AHDL Include File—Used in AHDL Text Design Files (.tdf).
<code><output file>.cmp</code>	Component Declaration File—Used in VHDL design files.
<code><output file>.bsf</code>	Block Symbol File—Used in Quartus II schematic Block Design Files (.bdf).
<code><output file>_inst.v .vhd .tdf</code>	HDL Instantiation Template for the language of the variation file—Sample instantiation of the Verilog HDL module, VHDL entity, or AHDL subdesign.
<code><output file>_bb.v</code>	Black box Verilog HDL Module Declaration—Hollow-body module declaration that can be used in Verilog HDL designs to specify port directions when instantiating the megafunction as a black box in third-party synthesis tools.

Table 10-1. MegaWizard Plug-In Manager Generated Files (Part 2 of 2)

File	Description
<code><output file>_syn.v</code>	Synthesis timing and resource estimation netlist—Additional synthesis netlist file created if you enable the option to generate a synthesis timing and resource estimation netlist. Refer to “ Creating a Netlist File for Other Synthesis Tools ” for details.

Note to Table 10-1:

- (1) The MegaWizard Plug-In Manager generates a `.v`, `.vhd`, or `.tdf` file, depending on the language you select for the output file on the megafunction-selection page of the wizard.

Creating a Netlist File for Other Synthesis Tools

When you use certain megafunctions with other EDA synthesis tools (that is, tools other than Quartus II integrated synthesis), you can optionally create a netlist for timing and resource estimation instead of a wrapper file.

The netlist file is a representation of the customized logic used in the Quartus II software. The file provides the connectivity of architectural elements in the megafunction but may not represent true functionality. This information enables certain other EDA synthesis tools to better report timing and resource estimates. In addition, synthesis tools can use the timing information to focus timing-driven optimizations and improve the quality of results.

To generate the netlist, turn on **Generate netlist** under **Timing and resource estimation** on the **EDA** page of the MegaWizard Plug-In Manager. The netlist file is called `<output file>_syn.v`. If you use this netlist for synthesis, you must include the megafunction wrapper file, either `<output file>.v` or `<output file>.vhd`, for placement and routing in the project created with the Quartus II software.

Because your synthesis tool may call the Quartus II software in the background to generate this netlist, turning on this option might not be required.

-  For information about support for timing and resource estimation netlists in your synthesis tool, refer to the tool vendor’s documentation or the appropriate chapter in the *Synthesis* section in volume 1 of the *Quartus II Handbook*.

Instantiating Megafunctions Using the Port and Parameter Definition

You can instantiate the megafunction directly in your Verilog HDL, VHDL, or AHDL code by calling the megafunction and setting its parameters as you would any other module, component, or subdesign.

-  For a list of the megafunction ports and parameters, refer to the specific megafunction in the Quartus II Help. You can also refer to the [IP and Megafunction](#) page on the Altera website.

-  Altera strongly recommends that you use the MegaWizard Plug-In Manager for complex megafunctions such as PLLs, transceivers, and LVDS drivers. For details about using the MegaWizard Plug-In Manager, refer to “[Instantiating Megafunctions Using the MegaWizard Plug-In Manager](#)” on page 10-3.

Inferring Multiplier and DSP Functions from HDL Code

The following sections describe how to infer multiplier and DSP functions from generic HDL code, and, if applicable, how to target the dedicated DSP block architecture in Altera devices:

- “Inferring Multipliers from HDL Code”
- “Inferring Multiply-Accumulators and Multiply-Adders from HDL Code” on page 10-8

-  For synthesis tool features and options, refer to your synthesis tool documentation or the appropriate chapter in the *Synthesis* section in volume 1 of the *Quartus II Handbook*.
-  For more design examples involving advanced multiply functions and complex DSP functions, refer to the [DSP Design Examples](#) page on the Altera website.

Inferring Multipliers from HDL Code

To infer multiplier functions, synthesis tools look for multipliers and convert them to LPM_MULT or ALTMULT_ADD megafunctions, or may map them directly to device atoms. For devices with DSP blocks, the software can implement the function in a DSP block instead of logic, depending on device utilization. The Quartus II Fitter can also place input and output registers in DSP blocks (that is, perform register packing) to improve performance and area utilization.

-  For additional information about the DSP block and supported functions, refer to the appropriate Altera device family handbook and the Altera [DSP Solutions Center](#) website.

[Example 10-1](#) and [Example 10-2](#) show Verilog HDL code examples, and [Example 10-3](#) and [Example 10-4](#) show VHDL code examples, for unsigned and signed multipliers that synthesis tools can infer as a megafunction or DSP block atoms. Each example fits into one DSP block element. In addition, when register packing occurs, no extra logic cells for registers are required.

-  The signed declaration in Verilog HDL is a feature of the Verilog 2001 Standard.

Example 10-1. Verilog HDL Unsigned Multiplier

```
module unsigned_mult (out, a, b);  
    output [15:0] out;  
    input  [7:0] a;  
    input  [7:0] b;  
    assign out = a * b;  
endmodule
```

Example 10-2. Verilog HDL Signed Multiplier with Input and Output Registers (Pipelining = 2)

```
module signed_mult (out, clk, a, b);
    output [15:0] out;
    input clk;
    input signed [7:0] a;
    input signed [7:0] b;

    reg signed [7:0] a_reg;
    reg signed [7:0] b_reg;
    reg signed [15:0] out;
    wire signed [15:0] mult_out;

    assign mult_out = a_reg * b_reg;

    always @ (posedge clk)
    begin
        a_reg <= a;
        b_reg <= b;
        out <= mult_out;
    end
endmodule
```

Example 10-3. VHDL Unsigned Multiplier with Input and Output Registers (Pipelining = 2)

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.numeric_std.all;

ENTITY unsigned_mult IS
  PORT (
    a: IN UNSIGNED (7 DOWNTO 0);
    b: IN UNSIGNED (7 DOWNTO 0);
    clk: IN STD_LOGIC;
    aclr: IN STD_LOGIC;
    result: OUT UNSIGNED (15 DOWNTO 0)
  );
END unsigned_mult;

ARCHITECTURE rtl OF unsigned_mult IS
  SIGNAL a_reg, b_reg: UNSIGNED (7 DOWNTO 0);
BEGIN
  PROCESS (clk, aclr)
  BEGIN
    IF (aclr = '1') THEN
      a_reg <= (OTHERS => '0');
      b_reg <= (OTHERS => '0');
      result <= (OTHERS => '0');
    ELSIF (clk'event AND clk = '1') THEN
      a_reg <= a;
      b_reg <= b;
      result <= a_reg * b_reg;
    END IF;
  END PROCESS;
END rtl;
```

Example 10-4. VHDL Signed Multiplier

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.numeric_std.all;

ENTITY signed_mult IS
  PORT (
    a: IN SIGNED (7 DOWNTO 0);
    b: IN SIGNED (7 DOWNTO 0);
    result: OUT SIGNED (15 DOWNTO 0)
  );
END signed_mult;

ARCHITECTURE rtl OF signed_mult IS
BEGIN
  result <= a * b;
END rtl;
```

Inferring Multiply-Accumulators and Multiply-Adders from HDL Code

Synthesis tools detect multiply-accumulators or multiply-adders and convert them to `ALTMULT_ACCUM` or `ALTMULT_ADD` megafunctions, respectively, or may map them directly to device atoms. The Quartus II software then places these functions in DSP blocks during placement and routing.



Synthesis tools infer multiply-accumulator and multiply-adder functions only if the Altera device family has dedicated DSP blocks that support these functions.

A simple multiply-accumulator consists of a multiplier feeding an addition operator. The addition operator feeds a set of registers that then feeds the second input to the addition operator. A simple multiply-adder consists of two to four multipliers feeding one or two levels of addition, subtraction, or addition/subtraction operators. Addition is always the second-level operator, if it is used. In addition to the multiply-accumulator and multiply-adder, the Quartus II Fitter also places input and output registers into the DSP blocks to pack registers and improve performance and area utilization.

Some device families offer additional advanced multiply-add and accumulate functions, such as complex multiplication, input shift register, or larger multiplications.



For details about advanced DSP block features, refer to the appropriate device handbook. For more design examples of DSP functions and inferring advanced features in the multiply-add and multiply-accumulate circuitry, refer to the [DSP Design Examples](#) page on Altera's website.

The Verilog HDL and VHDL code samples in [Example 10-5](#) through [Example 10-8](#) on [pages 10-9](#) through [10-12](#) infer multiply-accumulators and multiply-adders with input, output, and pipeline registers, as well as an optional asynchronous clear signal. Using the three sets of registers provides the best performance through the function, with a latency of three. You can remove the registers in your design to reduce the latency.

Example 10-5. Verilog HDL Unsigned Multiply-Accumulator

```
module unsig_altmult_accum (dataout, dataa, datab, clk, aclr, clken);
    input [7:0] dataa, datab;
    input clk, aclr, clken;
    output reg[16:0] dataout;

    reg [7:0] dataa_reg, datab_reg;
    reg [15:0] multa_reg;
    wire [15:0] multa;
    wire [16:0] adder_out;
    assign multa = dataa_reg * datab_reg;
    assign adder_out = multa_reg + dataout;

    always @ (posedge clk or posedge aclr)
    begin
        if (aclr)
            begin
                dataa_reg <= 8'b0;
                datab_reg <= 8'b0;
                multa_reg <= 16'b0;
                dataout <= 17'b0;
            end
        else if (clken)
            begin
                dataa_reg <= dataa;
                datab_reg <= datab;
                multa_reg <= multa;
                dataout <= adder_out;
            end
        end
    end
endmodule
```

Example 10-6. Verilog HDL Signed Multiply-Adder

```
module sig_altmult_add (dataa, datab, datac, datad, clock, aclr,
result);
    input signed [15:0] dataa, datab, datac, datad;
    input clock, aclr;
    output reg signed [32:0] result;

    reg signed [15:0] dataa_reg, datab_reg, datac_reg, datad_reg;
    reg signed [31:0] mult0_result, mult1_result;

    always @ (posedge clock or posedge aclr) begin
        if (aclr) begin
            dataa_reg <= 16'b0;
            datab_reg <= 16'b0;
            datac_reg <= 16'b0;
            datad_reg <= 16'b0;
            mult0_result <= 32'b0;
            mult1_result <= 32'b0;
            result <= 33'b0;
        end
        else begin
            dataa_reg <= dataa;
            datab_reg <= datab;
            datac_reg <= datac;
            datad_reg <= datad;
            mult0_result <= dataa_reg * datab_reg;
            mult1_result <= datac_reg * datad_reg;
            result <= mult0_result + mult1_result;
        end
    end
end
endmodule
```

Example 10-7. VHDL Signed Multiply-Accumulator

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.numeric_std.all;

ENTITY sig_altmult_accum IS
  PORT (
    a: IN SIGNED(7 DOWNTO 0);
    b: IN SIGNED (7 DOWNTO 0);
    clk: IN STD_LOGIC;
    aclr: IN STD_LOGIC;
    accum_out: OUT SIGNED (15 DOWNTO 0)
  ) ;
END sig_altmult_accum;

ARCHITECTURE rtl OF sig_altmult_accum IS
  SIGNAL a_reg, b_reg: SIGNED (7 DOWNTO 0);
  SIGNAL pdt_reg: SIGNED (15 DOWNTO 0);
  SIGNAL adder_out: SIGNED (15 DOWNTO 0);
BEGIN
  PROCESS (clk, aclr)
  BEGIN
    IF (aclr = '1') then
      a_reg <= (others => '0');
      b_reg <= (others => '0');
      pdt_reg <= (others => '0');
      adder_out <= (others => '0');
    ELSIF (clk'event and clk = '1') THEN
      a_reg <= (a);
      b_reg <= (b);
      pdt_reg <= a_reg * b_reg;
      adder_out <= adder_out + pdt_reg;
    END IF;
  END process;
  accum_out <= adder_out;
END rtl;
```

Example 10-8. VHDL Unsigned Multiply-Adder

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.numeric_std.all;

ENTITY unsignedmult_add IS
  PORT (
    a: IN UNSIGNED (7 DOWNTO 0);
    b: IN UNSIGNED (7 DOWNTO 0);
    c: IN UNSIGNED (7 DOWNTO 0);
    d: IN UNSIGNED (7 DOWNTO 0);
    clk: IN STD_LOGIC;
    aclr: IN STD_LOGIC;
    result: OUT UNSIGNED (15 DOWNTO 0)
  );
END unsignedmult_add;

ARCHITECTURE rtl OF unsignedmult_add IS
  SIGNAL a_reg, b_reg, c_reg, d_reg: UNSIGNED (7 DOWNTO 0);
  SIGNAL pdt_reg, pdt2_reg: UNSIGNED (15 DOWNTO 0);
  SIGNAL result_reg: UNSIGNED (15 DOWNTO 0);
BEGIN
  PROCESS (clk, aclr)
  BEGIN
    IF (aclr = '1') THEN
      a_reg <= (OTHERS => '0');
      b_reg <= (OTHERS => '0');
      c_reg <= (OTHERS => '0');
      d_reg <= (OTHERS => '0');
      pdt_reg <= (OTHERS => '0');
      pdt2_reg <= (OTHERS => '0');

    ELSIF (clk'event AND clk = '1') THEN
      a_reg <= a;
      b_reg <= b;
      c_reg <= c;
      d_reg <= d;
      pdt_reg <= a_reg * b_reg;
      pdt2_reg <= c_reg * d_reg;
      result_reg <= pdt_reg + pdt2_reg;
    END IF;
  END PROCESS;
  result <= result_reg;
END rtl;

```

Inferring Memory Functions from HDL Code

The following sections describe how to infer memory functions from generic HDL code and, if applicable, to target the dedicated memory architecture in Altera devices:

- “Inferring RAM functions from HDL Code” on page 10-13
- “Inferring ROM Functions from HDL Code” on page 10-36
- “Shift Registers—Inferring the ALTSHIFT_TAPS Megafunction from HDL Code” on page 10-40



For synthesis tool features and options, refer to your synthesis tool documentation or the appropriate chapter in the *Synthesis* section in volume 1 of the *Quartus II Handbook*.

Altera’s dedicated memory architecture offers a number of advanced features that can be easily targeted using the MegaWizard Plug-In Manager, as described in “Instantiating Altera Megafunctions in HDL Code” on page 10-3. The coding recommendations in the following sections provide portable examples of generic HDL code that infer the appropriate megafunction. However, if you want to use some of the advanced memory features in Altera devices, consider using the megafunction directly so that you can control the ports and parameters more easily.

Inferring RAM functions from HDL Code

To infer RAM functions, synthesis tools detect sets of registers and logic that can be replaced with the ALTSYNCRAM or ALTDPRAM megafunctions for device families that have dedicated RAM blocks, or may map them directly to device memory atoms. Tools typically consider all signals and variables that have a multi-dimensional array type and then create a RAM block, if applicable, based on the way the signals, variables, or both are assigned, referenced, or both in the HDL source description.

Standard synthesis tools recognize single-port and simple dual-port (one read port and one write port) RAM blocks. Some tools (such as the Quartus II software) also recognize true dual-port (two read ports and two write ports) RAM blocks that map to the memory blocks in certain Altera devices.

Some tools (such as the Quartus II software) also infer memory blocks for array variables and signals that are referenced (read/written) by two indices, to recognize mixed-width and byte-enabled RAMs for certain coding styles.



If your design contains a RAM block that your synthesis tool does not recognize and infer, the design might require a large amount of system memory that can potentially cause compilation problems.

When you use a formal verification flow, Altera recommends that you create RAM blocks in separate entities or modules that contain only the RAM logic. In certain formal verification flows, for example, when using Quartus II integrated synthesis, the entity or module containing the inferred RAM is put into a black box automatically because formal verification tools do not support RAM blocks. The Quartus II software issues a warning message when this situation occurs. If the entity or module contains any additional logic outside the RAM block, this logic cannot be verified because it also must be treated as a black box for formal verification.

The following sections present several guidelines for inferring RAM functions that match the dedicated memory architecture in Altera devices, and then provide recommended HDL code for different types of memory logic.

Use Synchronous Memory Blocks

Altera recommends using synchronous memory blocks for Altera designs. Because memory blocks in the newest devices from Altera are synchronous, RAM designs that are targeted towards architectures that contain these dedicated memory blocks must be synchronous to be mapped directly into the device architecture. For these devices, asynchronous memory logic is implemented in regular logic cells.

Synchronous memory offers several advantages over asynchronous memory, including higher frequencies and thus higher memory bandwidth, increased reliability, and less standby power. In many designs with asynchronous memory, the memory interfaces with synchronous logic so that the conversion to synchronous memory design is straightforward. To convert asynchronous memory you can move registers from the data path into the memory block.

Synchronous memories are supported in all Altera device families. A memory block is considered synchronous if it uses one of the following read behaviors:

- Memory read occurs in a Verilog `always` block with a clock signal or a VHDL clocked process. The recommended coding style for synchronous memories is to create your design with a registered read output.
- Memory read occurs outside a clocked block, but there is a synchronous read address (that is, the address used in the read statement is registered). This type of logic is not always inferred as a memory block, or may require external bypass logic, depending on the target device architecture.



The synchronous memory structures in Altera devices can differ from the structures in other vendors' devices. For best results, match your design to the target device architecture.

Later sections provide coding recommendations for various memory types. All of these examples are synchronous to ensure that they can be directly mapped into the dedicated memory architecture available in Altera FPGAs.



For additional information about the dedicated memory blocks in your specific device, refer to the appropriate Altera device family data sheet on the Altera website at www.altera.com.

Avoid Unsupported Reset and Control Conditions

To ensure that your HDL code can be implemented in the target device architecture, avoid unsupported reset conditions or other control logic that does not exist in the device architecture.

The RAM contents of Altera memory blocks cannot be cleared with a reset signal during device operation. If your HDL code describes a RAM with a reset signal for the RAM contents, the logic is implemented in regular logic cells instead of a memory block. Altera recommends against putting RAM read or write operations in an always block or process block with a reset signal. If you want to specify memory contents, initialize the memory as described in [“Specifying Initial Memory Contents at Power-Up” on page 10-33](#) or write the data to the RAM during device operation.

[Example 10-9](#) shows an example of undesirable code where there is a reset signal that clears part of the RAM contents. Avoid this coding style because it is not supported in Altera memories.

Example 10-9. Verilog RAM with Reset Signal that Clears RAM Contents: Not Supported in Device Architecture

```
module clear_ram
(
    input clock, reset, we,
    input [7:0] data_in,
    input [4:0] address,
    output reg [7:0] data_out
);

    reg [7:0] mem [0:31];
    integer i;

    always @ (posedge clock or posedge reset)
    begin
        if (reset == 1'b1)
            mem[address] <= 0;
        else if (we == 1'b1)
            mem[address] <= data_in;

        data_out <= mem[address];
    end
endmodule
```

Example 10-10 shows an example of undesirable code where the reset signal affects the RAM, although the effect may not be intended. Avoid this coding style because it is not supported in Altera memories.

Example 10-10. Verilog RAM with Reset Signal that Affects RAM: Not Supported in Device Architecture

```

module bad_reset
(
    input clock,
    input reset,
    input we,
    input [7:0] data_in,
    input [4:0] address,
    output reg [7:0] data_out,
    input d,
    output reg q
);

    reg [7:0] mem [0:31];
    integer i;

    always @ (posedge clock or posedge reset)
    begin
        if (reset == 1'b1)
            q <= 0;
        else
            begin
                if (we == 1'b1)
                    mem[address] <= data_in;

                data_out <= mem[address];
                q <= d;
            end
        end
    end
endmodule

```

In addition to reset signals, other control logic can prevent memory logic from being inferred as a memory block. For example, you cannot use a clock enable on the read address registers in Stratix® devices because doing so affects the output latch of the RAM, and therefore the synthesized result in the device RAM architecture would not match the HDL description. You can use the address stall feature as a read address clock enable in Stratix II, Cyclone® II, Arria® GX, and other newer devices to avoid this limitation. Check the documentation for your device architecture to ensure that your code matches the hardware available in the device.

Check Read-During-Write Behavior

It is important to check the read-during-write behavior of the memory block described in your HDL design as compared to the behavior in your target device architecture. Your HDL source code specifies the memory behavior when you read and write from the same memory address in the same clock cycle. The code specifies that the read returns either the old data at the address, or the new data being written to the address. This behavior is referred to as the read-during-write behavior of the memory block. Altera memory blocks have different read-during-write behavior depending on the target device family, memory mode, and block type.

Synthesis tools map an HDL design into the target device architecture, with the goal of maintaining the functionality described in your source code. Therefore, if your source code specifies unsupported read-during-write behavior for the device RAM blocks, the software must implement the logic outside the RAM hardware in regular logic cells.

One common problem occurs when there is a continuous read in the HDL code, as in the following examples. You should avoid using these coding styles:

```
//Verilog HDL concurrent signal assignment
assign q = ram[raddr_reg];

-- VHDL concurrent signal assignment
q <= ram(raddr_reg);
```

When a write operation occurs, this type of HDL implies that the read should immediately reflect the new data at the address, independent of the read clock. However, that is not the behavior of synchronous memory blocks. In the device architecture, the new data is not available until the next edge of the read clock. Therefore, if the synthesis tool mapped the logic directly to a synchronous memory block, the device functionality and gate-level simulation results would not match the HDL description or functional simulation results. If the write clock and read clock are the same, the synthesis tool can infer memory blocks and add extra bypass logic so that the device behavior matches the HDL behavior. If the write and read clocks are different, the synthesis tool cannot reliably add bypass logic, so the logic is implemented in regular logic cells instead of dedicated RAM blocks. The examples in the following sections discuss some of these differences for read-during-write conditions.

In addition, the MLAB feature in certain device logic array blocks (LABs) does not easily support old data or new data behavior for a read-during-write in the dedicated device architecture. Implementing the extra logic to support this behavior significantly reduces timing performance through the memory.



For best performance in MLAB memories, your design should not depend on the read data during a write operation.

In many synthesis tools, you can specify that the read-during-write behavior is not important to your design; for example, if you never read from the same address to which you write in the same clock cycle. For Quartus II integrated synthesis, add the synthesis attribute `ramstyle` set to `"no_rw_check"` to allow the software to choose the read-during-write behavior of a RAM, rather than use the behavior specified by your HDL code. In some cases, this attribute prevents the synthesis tool from using extra logic to implement the memory block, or can allow memory inference when it would otherwise be impossible.

Synchronous RAM blocks require a synchronous read, so Quartus II integrated synthesis packs either data output registers or read address registers into the RAM block. When the read address registers are packed into the RAM block, the read address signals connected to the RAM block contain the next value of the read address signals indexing the HDL variable, which impacts which clock cycle the read and the write occur, and changes the read-during-write conditions. Therefore, bypass logic may still be added to the design to preserve the read-during-write behavior, even if the `"no_rw_check"` attribute is set.

 For more information about attribute syntax, the `no_rw_check` attribute value, or specific options for your synthesis tool, refer to your synthesis tool documentation or the appropriate chapter in the *Synthesis* section in volume 1 of the *Quartus II Handbook*.

The next section describes how you control the logic implementation in the Altera device, and the following sections provide coding recommendations for various memory types. Each example describes the read-during-write behavior and addresses the support for the memory type in Altera devices.

Controlling Inference and Implementation in Device RAM Blocks

Tools usually do not infer small RAM blocks because small RAM blocks typically can be implemented more efficiently using the registers in regular logic. If you are using Quartus II integrated synthesis, you can direct the software to infer RAM blocks for all sizes with the **Allow Any RAM Size for Recognition** option in the **More Analysis & Synthesis Settings** dialog box.

Some synthesis tools provide options to control the implementation of inferred RAM blocks for Altera devices with synchronous memory blocks. For example, Quartus II integrated synthesis provides the `ramstyle` synthesis attribute to specify the type of memory block or to specify the use of regular logic instead of a dedicated memory block. Quartus II integrated synthesis does not map inferred memory into MLABs unless the HDL code specifies the appropriate `ramstyle` attribute, although the Fitter may map some memories to MLABs.

 For details about using the `ramstyle` attribute, refer to the *Quartus II Integrated Synthesis* chapter in volume 1 of the *Quartus II Handbook*. For information about synthesis attributes in other synthesis tools, refer to the appropriate chapter in the *Synthesis* section in volume 1 of the *Quartus II Handbook*.

If you want to control the implementation after the RAM function is inferred during synthesis, you can set the `ram_block_type` parameter of the `ALTSYNCRAM` megafunction. In the Assignment Editor, select **Parameters** in the **Categories** list. You can use the **Node Finder** or drag the appropriate instance from the Project Navigator window to enter the RAM hierarchical instance name. Type `ram_block_type` as the **Parameter Name** and type one of the following memory types supported by your target device family in the **Value** field: "M-RAM", "M4K", "M20K", "M512", "M9K", "M144K", or "MLAB".

You can also specify the maximum depth of memory blocks used to infer RAM or ROM in your design. Apply the `max_depth` synthesis attribute to the declaration of a variable that represents a RAM or ROM in your design file. For example:

```
// Limit the depth of the memory blocks implement "ram" to 512
// This forces the software to use two M512 blocks instead of one M4K block to
implement this RAM
(* max_depth = 512 *) reg [7:0] ram[0:1023];
```

Single-Clock Synchronous RAM with Old Data Read-During-Write Behavior

The code examples in this section show Verilog HDL and VHDL code that infers simple dual-port, single-clock synchronous RAM. Single-port RAM blocks use a similar coding style.

The read-during-write behavior in these examples is to read the old data at the memory address. Refer to “[Check Read-During-Write Behavior](#)” on page 10-16 for details. Altera recommends that you use the Old Data Read-During-Write coding style for most RAM blocks as long as your design does not require the RAM location’s new value when you perform a simultaneous read and write to that RAM location. For best performance in MLAB memories, use the appropriate attribute so that your design does not depend on the read data during a write operation.

If you require that the read-during-write results in new data, refer to “[Single-Clock Synchronous RAM with New Data Read-During-Write Behavior](#)” on page 10-20.

The simple dual-port RAM code samples in [Example 10-11](#) and [Example 10-12](#) map directly into Altera synchronous memory.

Single-port versions of memory blocks (that is, using the same read address and write address signals) can allow better RAM utilization than dual-port memory blocks, depending on the device family.

Example 10-11. Verilog HDL Single-Clock Simple Dual-Port Synchronous RAM with Old Data Read-During-Write Behavior

```
module single_clk_ram(  
    output reg [7:0] q,  
    input [7:0] d,  
    input [6:0] write_address, read_address,  
    input we, clk  
);  
    reg [7:0] mem [127:0];  
  
    always @ (posedge clk) begin  
        if (we)  
            mem[write_address] <= d;  
        q <= mem[read_address]; // q doesn't get d in this clock cycle  
    end  
endmodule
```

Example 10-12. VHDL Single-Clock Simple Dual-Port Synchronous RAM with Old Data Read-During-Write Behavior

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;

ENTITY single_clock_ram IS
    PORT (
        clock: IN STD_LOGIC;
        data: IN STD_LOGIC_VECTOR (2 DOWNTO 0);
        write_address: IN INTEGER RANGE 0 to 31;
        read_address: IN INTEGER RANGE 0 to 31;
        we: IN STD_LOGIC;
        q: OUT STD_LOGIC_VECTOR (2 DOWNTO 0)
    );
END single_clock_ram;

ARCHITECTURE rtl OF single_clock_ram IS
    TYPE MEM IS ARRAY(0 TO 31) OF STD_LOGIC_VECTOR(2 DOWNTO 0);
    SIGNAL ram_block: MEM;
BEGIN
    PROCESS (clock)
    BEGIN
        IF (clock'event AND clock = '1') THEN
            IF (we = '1') THEN
                ram_block(write_address) <= data;
            END IF;
            q <= ram_block(read_address);
            -- VHDL semantics imply that q doesn't get data
            -- in this clock cycle
        END IF;
    END PROCESS;
END rtl;

```

Single-Clock Synchronous RAM with New Data Read-During-Write Behavior

The examples in this section describe RAM blocks in which a simultaneous read and write to the same location reads the new value that is currently being written to that RAM location.

To implement this behavior in the target device, synthesis software adds bypass logic around the RAM block. This bypass logic increases the area utilization of the design and decreases the performance if the RAM block is part of the design's critical path. Refer to [“Check Read-During-Write Behavior” on page 10-16](#) for details. If this behavior is not required for your design, use the examples from [“Single-Clock Synchronous RAM with Old Data Read-During-Write Behavior” on page 10-18](#).

The simple dual-port RAM in [Example 10-13](#) and [Example 10-14](#) require the software to create bypass logic around the RAM block.

Single-port versions of the Verilog memory block (that is, using the same read address and write address signals) do not require any logic cells to create bypass logic in the Arria, Stratix, and Cyclone series of devices, because the device memory supports new data read-during-write behavior when in single-port mode (same clock, same read address, and same write address).

Example 10–13. Verilog HDL Single-Clock Simple Dual-Port Synchronous RAM with New Data Read-During-Write Behavior

```
module single_clock_wr_ram(  
    output reg [7:0] q,  
    input [7:0] d,  
    input [6:0] write_address, read_address,  
    input we, clk  
);  
    reg [7:0] mem [127:0];  
  
    always @ (posedge clk) begin  
        if (we)  
            mem[write_address] = d;  
        q = mem[read_address]; // q does get d in this clock cycle if  
                               // we is high  
    end  
endmodule
```



Example 10–13 is similar to Example 10–11, but Example 10–13 uses a blocking assignment for the write so that the data is assigned immediately.

An alternative way to create a single-clock RAM is to use an assign statement to read the address of mem to create the output q, as shown in the following coding style example. By itself, the code describes new data read-during-write behavior. However, if the RAM output feeds a register in another hierarchy, a read-during-write results in the old data. Synthesis tools may not infer a RAM block if the tool cannot determine which behavior is described, such as when the memory feeds a hard hierarchical partition boundary. For this reason, avoid using this alternate type of coding style:

```
reg [7:0] mem [127:0];  
reg [6:0] read_address_reg;  
  
always @ (posedge clk) begin  
    if (we)  
        mem[write_address] <= d;  
  
    read_address_reg <= read_address;  
end  
  
assign q = mem[read_address_reg];
```

The VHDL sample in [Example 10-14](#) uses a concurrent signal assignment to read from the RAM. By itself, this example describes new data read-during-write behavior. However, if the RAM output feeds a register in another hierarchy, a read-during-write results in the old data. Synthesis tools may not infer a RAM block if the tool cannot determine which behavior is described, such as when the memory feeds a hard hierarchical partition boundary.

Example 10-14. VHDL Single-Clock Simple Dual-Port Synchronous RAM with New Data Read-During-Write Behavior

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;

ENTITY single_clock_rw_ram IS
  PORT (
    clock: IN STD_LOGIC;
    data: IN STD_LOGIC_VECTOR (2 DOWNTO 0);
    write_address: IN INTEGER RANGE 0 to 31;
    read_address: IN INTEGER RANGE 0 to 31;
    we: IN STD_LOGIC;
    q: OUT STD_LOGIC_VECTOR (2 DOWNTO 0)
  );
END single_clock_rw_ram;

ARCHITECTURE rtl OF single_clock_rw_ram IS
  TYPE MEM IS ARRAY(0 TO 31) OF STD_LOGIC_VECTOR(2 DOWNTO 0);
  SIGNAL ram_block: MEM;
  SIGNAL read_address_reg: INTEGER RANGE 0 to 31;
BEGIN
  PROCESS (clock)
  BEGIN
    IF (clock'event AND clock = '1') THEN
      IF (we = '1') THEN
        ram_block(write_address) <= data;
      END IF;
      read_address_reg <= read_address;
    END IF;
  END PROCESS;
  q <= ram_block(read_address_reg);
END rtl;
```

For Quartus II integrated synthesis, if you do not require the read-through-write capability, add the synthesis attribute `ramstyle="no_rw_check"` to allow the software to choose the read-during-write behavior of a RAM, rather than using the behavior specified by your HDL code. As discussed in [“Check Read-During-Write Behavior” on page 10-16](#), this attribute may prevent generation of extra bypass logic but it is not always possible to eliminate the requirement for bypass logic.

Simple Dual-Port, Dual-Clock Synchronous RAM

In dual clock designs, synthesis tools cannot accurately infer the read-during-write behavior because it depends on the timing of the two clocks within the target device. Therefore, the read-during-write behavior of the synthesized design is undefined and may differ from your original HDL code. Refer to [“Check Read-During-Write Behavior” on page 10-16](#) for details.

When Quartus II integrated synthesis infers this type of RAM, it issues a warning because of the undefined read-during-write behavior. If this functionality is acceptable in your design, you can avoid the warning by adding the synthesis attribute `ramstyle="no_rw_check"` to allow the software to choose the read-during-write behavior of a RAM.

The code samples in [Example 10-15](#) and [Example 10-16](#) show Verilog HDL and VHDL code that infers dual-clock synchronous RAM. The exact behavior depends on the relationship between the clocks.

Example 10-15. Verilog HDL Simple Dual-Port, Dual-Clock Synchronous RAM

```
module dual_clock_ram(  
    output reg [7:0] q,  
    input [7:0] d,  
    input [6:0] write_address, read_address,  
    input we, clk1, clk2  
);  
    reg [6:0] read_address_reg;  
    reg [7:0] mem [127:0];  
  
    always @ (posedge clk1)  
    begin  
        if (we)  
            mem[write_address] <= d;  
    end  
  
    always @ (posedge clk2) begin  
        q <= mem[read_address_reg];  
        read_address_reg <= read_address;  
    end  
endmodule
```

Example 10-16. VHDL Simple Dual-Port, Dual-Clock Synchronous RAM

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;
ENTITY dual_clock_ram IS
    PORT (
        clock1, clock2: IN STD_LOGIC;
        data: IN STD_LOGIC_VECTOR (3 DOWNTO 0);
        write_address: IN INTEGER RANGE 0 to 31;
        read_address: IN INTEGER RANGE 0 to 31;
        we: IN STD_LOGIC;
        q: OUT STD_LOGIC_VECTOR (3 DOWNTO 0)
    );
END dual_clock_ram;
ARCHITECTURE rtl OF dual_clock_ram IS
    TYPE MEM IS ARRAY(0 TO 31) OF STD_LOGIC_VECTOR(3 DOWNTO 0);
    SIGNAL ram_block: MEM;
    SIGNAL read_address_reg : INTEGER RANGE 0 to 31;
BEGIN
    PROCESS (clock1)
    BEGIN
        IF (clock1'event AND clock1 = '1') THEN
            IF (we = '1') THEN
                ram_block(write_address) <= data;
            END IF;
        END IF;
    END PROCESS;
    PROCESS (clock2)
    BEGIN
        IF (clock2'event AND clock2 = '1') THEN
            q <= ram_block(read_address_reg);
            read_address_reg <= read_address;
        END IF;
    END PROCESS;
END rtl;

```

True Dual-Port Synchronous RAM

The code examples in this section show Verilog HDL and VHDL code that infers true dual-port synchronous RAM. Different synthesis tools may differ in their support for these types of memories. This section describes the inference rules for Quartus II integrated synthesis. This type of RAM inference is supported for the Arria GX, Stratix, and Cyclone series of devices.

Altera synchronous memory blocks have two independent address ports, allowing for operations on two unique addresses simultaneously. A read operation and a write operation can share the same port if they share the same address. The Quartus II software infers true dual-port RAMs in Verilog HDL and VHDL with any combination of independent read or write operations in the same clock cycle, with at most two unique port addresses, performing two reads and one write, two writes and one read, or two writes and two reads in one clock cycle with one or two unique addresses.

In the synchronous RAM block architecture, there is no priority between the two ports. Therefore, if you write to the same location on both ports at the same time, the result is indeterminate in the device architecture. You must ensure your HDL code does not imply priority for writes to the memory block, if you want the design to be implemented in a dedicated hardware memory block. For example, if both ports are defined in the same process block, the code is synthesized and simulated sequentially so that there is a priority between the two ports. If your code does imply a priority, the logic cannot be implemented in the device RAM blocks and is implemented in regular logic cells.

You must also consider the read-during-write behavior of the RAM block to ensure that it can be mapped directly to the device RAM architecture. Refer to “[Check Read-During-Write Behavior](#)” on page 10-16 for details.

When a read and write operation occurs on the same port for the same address, the read operation may behave as follows:

- **Read new data**—This mode matches the behavior of synchronous memory blocks.
- **Read old data**—This mode is supported only by synchronous memory blocks in Arria II GX, Cyclone III, Stratix III, and newer device families. This behavior is not possible in memory blocks of older families.

When a read and write operation occurs on different ports for the same address (also known as mixed port), the read operation may behave as follows:

- **Read new data**—Quartus II integrated synthesis supports this mode by creating bypass logic around the synchronous memory block.
- **Read old data**—This behavior is supported by synchronous memory blocks.

The Verilog HDL single-clock code sample in [Example 10-17](#) maps directly into Altera synchronous memory. When a read and write operation occurs on the same port for the same address, the new data being written to the memory is read. When a read and write operation occurs on different ports for the same address, the old data in the memory is read. Simultaneous writes to the same location on both ports results in indeterminate behavior.

A dual-clock version of this design describes the same behavior, but the memory in the target device will have undefined mixed port read-during-write behavior because it depends on the relationship between the clocks.

Example 10-17. Verilog HDL True Dual-Port RAM with Single Clock

```

module true_dual_port_ram_single_clock
(
  input [(DATA_WIDTH-1):0] data_a, data_b,
  input [(ADDR_WIDTH-1):0] addr_a, addr_b,
  input we_a, we_b, clk,
  output reg [(DATA_WIDTH-1):0] q_a, q_b
);

  parameter DATA_WIDTH = 8;
  parameter ADDR_WIDTH = 6;

  // Declare the RAM variable
  reg [DATA_WIDTH-1:0] ram[2**ADDR_WIDTH-1:0];

  always @ (posedge clk)
  begin // Port A
    if (we_a)
    begin
      ram[addr_a] <= data_a;
      q_a <= data_a;
    end
    else
      q_a <= ram[addr_a];
    end
  always @ (posedge clk)
  begin // Port b
    if (we_b)
    begin
      ram[addr_b] <= data_b;
      q_b <= data_b;
    end
    else
      q_b <= ram[addr_b];
    end
  end
endmodule

```

If you use the following Verilog HDL read statements instead of the if-else statements in [Example 10-17](#), the HDL code specifies that the read results in old data when a read operation and write operation occurs at the same time for the same address on the same port or mixed ports. This behavior is supported only in the memory blocks of Arria II GX, Cyclone III, Stratix III, and newer device families, and is not inferred as memory for older device families.

```

always @ (posedge clk)
begin // Port A
  if (we_a)
    ram[addr_a] <= data_a;

  q_a <= ram[addr_a];
end

always @ (posedge clk)
begin // Port B
  if (we_b)
    ram[addr_b] <= data_b;

  q_b <= ram[addr_b];
end

```

The VHDL single-clock code sample in [Example 10-18](#) maps directly into Altera synchronous memory. When a read and write operation occurs on the same port for the same address, the new data being written to the memory is read. When a read and write operation occurs on different ports for the same address, the old data in the memory is read. Simultaneous write operations to the same location on both ports results in indeterminate behavior.

A dual-clock version of this design describes the same behavior, but the memory in the target device will have undefined mixed port read-during-write behavior because it depends on the relationship between the clocks.

Example 10-18. VHDL True Dual-Port RAM with Single Clock (Part 1 of 2)

```
library ieee;
use ieee.std_logic_1164.all;

entity true_dual_port_ram_single_clock is
  generic (
    DATA_WIDTH : natural := 8;
    ADDR_WIDTH  : natural := 6;
  );
  port (
    clk: in std_logic;
    addr_a: in natural range 0 to 2**ADDR_WIDTH - 1;
    addr_b: in natural range 0 to 2**ADDR_WIDTH - 1;
    data_a: in std_logic_vector((DATA_WIDTH-1) downto 0);
    data_b: in std_logic_vector((DATA_WIDTH-1) downto 0);
    we_a: in std_logic := '1';
    we_b: in std_logic := '1';
    q_a: out std_logic_vector((DATA_WIDTH - 1) downto 0);
    q_b: out std_logic_vector((DATA_WIDTH - 1) downto 0)
  );
end true_dual_port_ram_single_clock;

architecture rtl of true_dual_port_ram_single_clock is
  -- Build a 2-D array type for the RAM
  subtype word_t is std_logic_vector((DATA_WIDTH-1) downto 0);
  type memory_t is array((2**ADDR_WIDTH - 1) downto 0) of word_t;
  -- Declare the RAM signal.
  shared variable ram : memory_t;
```

Example 10-19. VHDL True Dual-Port RAM with Single Clock (Part 2 of 2)

```

begin
  process(clk)
  begin
    if(rising_edge(clk)) then -- Port A
      if(we_a = '1') then
        ram(addr_a) <= data_a;

        -- Read-during-write on the same port returns NEW data
        q_a <= data_a;
      else
        -- Read-during-write on the mixed port returns OLD data
        q_a <= ram(addr_a);
      end if;
    end if;
  end process;

  process(clk)
  begin
    if(rising_edge(clk)) then -- Port B
      if(we_b = '1') then
        ram(addr_b) <= data_b;
        -- Read-during-write on the same port returns NEW data
        q_b <= data_b;
      else
        -- Read-during-write on the mixed port returns OLD data
        q_b <= ram(addr_b);
      end if;
    end if;
  end process;
end rtl;

```

Mixed-Width Dual-Port RAM

The RAM code examples in [Example 10-20](#) through [Example 10-23](#) show SystemVerilog and VHDL code that infers RAM with data ports with different widths. This type of logic is not supported in Verilog-1995 or Verilog-2001 because of the requirement for a multi-dimensional array to model the different read width, write width, or both. Different synthesis tools may differ in their support for these memories. This section describes the inference rules for Quartus II integrated synthesis.

The first dimension of the multi-dimensional packed array represents the ratio of the wider port to the narrower port, and the second dimension represents the narrower port width. The read and write port widths must specify a read or write ratio supported by the memory blocks in the target device, or the synthesis tool does not infer a RAM.

Refer to the Quartus II Templates for parameterized examples that you can use for supported combinations of read and write widths, and true dual port RAM examples with two read ports and two write ports for mixed-width writes and reads.

Example 10–20. SystemVerilog Mixed-Width RAM with Read Width Smaller than Write Width

```
module mixed_width_ram    // 256x32 write and 1024x8 read
(
    input [7:0] waddr,
    input [31:0] wdata,
    input we, clk,
    input [9:0] raddr,
    output [7:0] q
);
    logic [3:0][7:0] ram[0:255];
    always_ff@(posedge clk)
        begin
            if(we) ram[waddr] <= wdata;
            q <= ram[raddr / 4][raddr % 4];
        end
endmodule : mixed_width_ram
```

Example 10–21. SystemVerilog Mixed-Width RAM with Read Width Larger than Write Width

```
module mixed_width_ram    // 1024x8 write and 256x32 read
(
    input [9:0] waddr,
    input [31:0] wdata,
    input we, clk,
    input [7:0] raddr,
    output [9:0] q
);
    logic [3:0][7:0] ram[0:255];
    always_ff@(posedge clk)
        begin
            if(we) ram[waddr / 4][waddr % 4] <= wdata;
            q <= ram[raddr];
        end
endmodule : mixed_width_ram
```

Example 10-22. VHDL Mixed-Width RAM with Read Width Smaller than Write Width

```
library ieee;
use ieee.std_logic_1164.all;

package ram_types is
    type word_t is array (0 to 3) of std_logic_vector(7 downto 0);
    type ram_t is array (0 to 255) of word_t;
end ram_types;

library ieee;
use ieee.std_logic_1164.all;
library work;
use work.ram_types.all;

entity mixed_width_ram is
    port (
        we, clk : in std_logic;
        waddr   : in integer range 0 to 255;
        wdata   : in word_t;
        raddr   : in integer range 0 to 1023;
        q       : out std_logic_vector(7 downto 0));
end mixed_width_ram;

architecture rtl of mixed_width_ram is
    signal ram : ram_t;
begin -- rtl
    process(clk, we)
    begin
        if(rising_edge(clk)) then
            if(we = '1') then
                ram(waddr) <= wdata;
            end if;
            q <= ram(raddr / 4 )(raddr mod 4);
        end if;
    end process;
end rtl;
```

Example 10–23. VHDL Mixed-Width RAM with Read Width Larger than Write Width

```
library ieee;
use ieee.std_logic_1164.all;

package ram_types is
    type word_t is array (0 to 3) of std_logic_vector(7 downto 0);
    type ram_t is array (0 to 255) of word_t;
end ram_types;

library ieee;
use ieee.std_logic_1164.all;
library work;
use work.ram_types.all;

entity mixed_width_ram is
    port (
        we, clk : in  std_logic;
        waddr   : in  integer range 0 to 1023;
        wdata   : in  std_logic_vector(7 downto 0);
        raddr   : in  integer range 0 to 255;
        q       : out word_t);
end mixed_width_ram;

architecture rtl of mixed_width_ram is
    signal ram : ram_t;
begin -- rtl
    process(clk, we)
    begin
        if(rising_edge(clk)) then
            if(we = '1') then
                ram(waddr / 4)(waddr mod 4) <= wdata;
            end if;
            q <= ram(raddr);
        end if;
    end process;
end rtl;
```

RAM with Byte-Enable Signals

The RAM code examples in [Example 10–24](#) and [Example 10–25](#) show SystemVerilog and VHDL code that infers RAM with controls for writing single bytes into the memory word, or byte-enable signals. Byte enables are modeled by creating write expressions with two indices and writing part of a RAM "word." With these implementations, you can also write more than one byte at once by enabling the appropriate byte enables.

This type of logic is not supported in Verilog-1995 or Verilog-2001 because of the requirement for a multidimensional array. Different synthesis tools may differ in their support for these memories. This section describes the inference rules for Quartus II integrated synthesis.

Refer to the Quartus II Templates for parameterized examples that you can use for different address widths, and true dual port RAM examples with two read ports and two write ports.

Example 10-24. SystemVerilog Simple Dual-Port Synchronous RAM with Byte Enable

```
module byte_enabled_simple_dual_port_ram
(
    input we, clk,
    input [5:0] waddr, raddr, // address width = 6
    input [3:0] be, // 4 bytes per word
    input [31:0] wdata, // byte width = 8, 4 bytes per word
    output reg [31:0] q // byte width = 8, 4 bytes per word
);
    // use a multi-dimensional packed array
    //to model individual bytes within the word
    logic [3:0][7:0] ram[0:63]; // # words = 1 << address width

    always_ff@(posedge clk)
    begin
        if(we) begin
            if(be[0]) ram[waddr][0] <= wdata[7:0];
            if(be[1]) ram[waddr][1] <= wdata[15:8];
            if(be[2]) ram[waddr][2] <= wdata[23:16];
            if(be[3]) ram[waddr][3] <= wdata[31:24];
        end
        q <= ram[raddr];
    end
endmodule
```

Example 10–25. VHDL Simple Dual-Port Synchronous RAM with Byte Enable

```
library ieee;
use ieee.std_logic_1164.all;
library work;

entity byte_enabled_simple_dual_port_ram is
port (
    we, clk : in  std_logic;
    waddr, raddr : in  integer range 0 to 63 ;      -- address width = 6
    be       : in  std_logic_vector(3 downto 0);  -- 4 bytes per word
    wdata    : in  std_logic_vector(31 downto 0);  -- byte width = 8
    q        : out std_logic_vector(31 downto 0) ); -- byte width = 8
end byte_enabled_simple_dual_port_ram;

architecture rtl of byte_enabled_simple_dual_port_ram is
    -- build up 2D array to hold the memory
    type word_t is array (0 to 3) of std_logic_vector(7 downto 0);
    type ram_t is array (0 to 63) of word_t;

    signal ram : ram_t;
    signal q_local : word_t;

begin -- Re-organize the read data from the RAM to match the output
    unpack: for i in 0 to 3 generate
        q(8*(i+1) - 1 downto 8*i) <= q_local(i);
    end generate unpack;

    process(clk)
    begin
        if(rising_edge(clk)) then
            if(we = '1') then
                if(be(0) = '1') then
                    ram(waddr)(0) <= wdata(7 downto 0);
                end if;
                if be(1) = '1' then
                    ram(waddr)(1) <= wdata(15 downto 8);
                end if;
                if be(2) = '1' then
                    ram(waddr)(2) <= wdata(23 downto 16);
                end if;
                if be(3) = '1' then
                    ram(waddr)(3) <= wdata(31 downto 24);
                end if;
            end if;
            q_local <= ram(raddr);
        end if;
    end process;
end rtl;
```

Specifying Initial Memory Contents at Power-Up

Your synthesis tool may offer various ways to specify the initial contents of an inferred memory.



Certain device memory types do not support initialized memory, such as the M-RAM blocks in Stratix and Stratix II devices.

There are slight power-up and initialization differences between dedicated RAM blocks and the MLAB memory due to the continuous read of the MLAB. Altera dedicated RAM block outputs always power-up to zero and are set to the initial value on the first read. For example, if address 0 is pre-initialized to FF, the RAM block powers up with the output at 0. A subsequent read after power-up from address 0 outputs the pre-initialized value of FF. Therefore, if a RAM is powered up and an enable (read enable or clock enable) is held low, the power-up output of 0 is maintained until the first valid read cycle. The MLAB is implemented using registers that power-up to 0, but are initialized to their initial value immediately at power-up or reset. Therefore, the initial value is seen, regardless of the enable status. The Quartus II software maps inferred memory to MLABs when the HDL code specifies an appropriate `ramstyle` attribute.

Quartus II integrated synthesis supports the `ram_init_file` synthesis attribute that allows you to specify a Memory Initialization File (`.mif`) for an inferred RAM block.

 For information about the `ram_init_file` attribute, refer to the *Quartus II Integrated Synthesis* chapter in volume 1 of the *Quartus II Handbook*. For information about synthesis attributes in other synthesis tools, refer to the tool vendor's documentation.

In Verilog HDL, you can use an initial block to initialize the contents of an inferred memory. Quartus II integrated synthesis automatically converts the initial block into a `.mif` file for the inferred RAM. [Example 10-26](#) shows Verilog HDL code that infers a simple dual-port RAM block and corresponding `.mif` file.

Example 10-26. Verilog HDL RAM with Initialized Contents

```
module ram_with_init(
    output reg [7:0] q,
    input [7:0] d,
    input [4:0] write_address, read_address,
    input we, clk
);
    reg [7:0] mem [0:31];
    integer i;

    initial begin
        for (i = 0; i < 32; i = i + 1)
            mem[i] = i[7:0];
        end

    always @ (posedge clk) begin
        if (we)
            mem[write_address] <= d;
        q <= mem[read_address];
        end
endmodule
```

Quartus II integrated synthesis and other synthesis tools also support the `$readmemb` and `$readmemh` commands so that RAM initialization and ROM initialization work identically in synthesis and simulation. [Example 10-27](#) shows an initial block that initializes an inferred RAM block using the `$readmemb` command.

 Refer to the *Verilog Language Reference Manual (LRM) 1364-2001* Section 17.2.8 or the example in the Templates for the Quartus II software for details about the format of the `ram.txt` file.

Example 10-27. Verilog HDL RAM Initialized with the readmemb Command

```
reg [7:0] ram[0:15];
initial
begin
    $readmemb("ram.txt", ram);
end
```

In VHDL, you can initialize the contents of an inferred memory by specifying a default value for the corresponding signal. Quartus II integrated synthesis automatically converts the default value into a `.mif` file for the inferred RAM. [Example 10-28](#) shows VHDL code that infers a simple dual-port RAM block and corresponding `.mif` file.

Example 10-28. VHDL RAM with Initialized Contents

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;
use ieee.numeric_std.all;

ENTITY ram_with_init IS
    PORT(
        clock: IN STD_LOGIC;
        data: IN UNSIGNED (7 DOWNTO 0);
        write_address: IN integer RANGE 0 to 31;
        read_address: IN integer RANGE 0 to 31;
        we: IN std_logic;
        q: OUT UNSIGNED (7 DOWNTO 0));
END;

ARCHITECTURE rtl OF ram_with_init IS

    TYPE MEM IS ARRAY(31 DOWNTO 0) OF unsigned(7 DOWNTO 0);
    FUNCTION initialize_ram
        return MEM is
        variable result : MEM;
    BEGIN
        FOR i IN 31 DOWNTO 0 LOOP
            result(i) := to_unsigned(natural(i), natural'(8));
        END LOOP;
        RETURN result;
    END initialize_ram;

    SIGNAL ram_block : MEM := initialize_ram;
BEGIN
    PROCESS (clock)
    BEGIN
        IF (clock'event AND clock = '1') THEN
            IF (we = '1') THEN
                ram_block(write_address) <= data;
            END IF;
            q <= ram_block(read_address);
        END IF;
    END PROCESS;
END rtl;
```

Inferring ROM Functions from HDL Code

To infer ROM functions, synthesis tools detect sets of registers and logic that can be replaced with the ALTSYNCRAM or LPM_ROM megafunctions, depending on the target device family, and only for device families that have dedicated memory blocks.

ROMs are inferred when a CASE statement exists in which a value is set to a constant for every choice in the case statement. Because small ROMs typically achieve the best performance when they are implemented using the registers in regular logic, each ROM function must meet a minimum size requirement to be inferred and placed into memory.



If you use Quartus II integrated synthesis, you can direct the software to infer ROM blocks for all sizes with the **Allow Any ROM Size for Recognition** option in the **More Analysis & Synthesis Settings** dialog box.

Some synthesis tools provide options to control the implementation of inferred ROM blocks for Altera devices with synchronous memory blocks. For example, Quartus II integrated synthesis provides the `romstyle` synthesis attribute to specify the type of memory block or to specify the use of regular logic instead of a dedicated memory block.



For details about using the `romstyle` attribute, refer to the *Quartus II Integrated Synthesis* chapter in volume 1 of the *Quartus II Handbook*. For information about synthesis attributes in other synthesis tools, refer to the appropriate chapter in the *Synthesis* section in volume 1 of the *Quartus II Handbook*.



Because formal verification tools do not support ROM megafunctions, Quartus II integrated synthesis does not infer ROM megafunctions when a formal verification tool is selected. When you are using a formal verification flow, Altera recommends that you instantiate ROM megafunction blocks in separate entities or modules that contain only the ROM logic, because you may need to treat the entity or module as a black box during formal verification.

The Verilog HDL and VHDL code samples in [Example 10-29](#) through [Example 10-32](#) on [pages 10-37](#) through [10-39](#) infer synchronous ROM blocks. Depending on the device family's dedicated RAM architecture, the ROM logic may have to be synchronous; refer to the device family handbook for details.

For device architectures with synchronous RAM blocks, such as the Arria series, Cyclone series, or Stratix series devices and newer device families, either the address or the output must be registered for synthesis software to infer a ROM block. When your design uses output registers, the synthesis software implements registers from the input registers of the RAM block without affecting the functionality of the ROM. If you register the address, the power-up state of the inferred ROM can be different from the HDL design. In this scenario, the synthesis software issues a warning. The Quartus II Help explains the condition under which the functionality changes when you use Quartus II integrated synthesis.

The ROM code examples in [Example 10–29](#) through [Example 10–32](#) on pages 10–37 through 10–39 map directly to the Altera memory architecture.

Example 10–29. Verilog HDL Synchronous ROM

```

module sync_rom (clock, address, data_out);
    input clock;
    input [7:0] address;
    output [5:0] data_out;

    reg [5:0] data_out;

    always @ (posedge clock)
    begin
        case (address)
            8'b00000000: data_out = 6'b1011111;
            8'b00000001: data_out = 6'b1101110;
            ...
            8'b11111110: data_out = 6'b0000001;
            8'b11111111: data_out = 6'b101010;
        endcase
    end
endmodule
    
```

Example 10–30. VHDL Synchronous ROM

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;

ENTITY sync_rom IS
    PORT (
        clock: IN STD_LOGIC;
        address: IN STD_LOGIC_VECTOR(7 downto 0);
        data_out: OUT STD_LOGIC_VECTOR(5 downto 0)
    );
END sync_rom;

ARCHITECTURE rtl OF sync_rom IS
BEGIN
    PROCESS (clock)
    BEGIN
        IF rising_edge (clock) THEN
            CASE address IS
                WHEN "00000000" => data_out <= "1011111";
                WHEN "00000001" => data_out <= "1101110";
                ...
                WHEN "11111110" => data_out <= "0000001";
                WHEN "11111111" => data_out <= "101010";
                WHEN OTHERS => data_out <= "1011111";
            END CASE;
        END IF;
    END PROCESS;
END rtl;
    
```

Example 10-31. Verilog HDL Dual-Port Synchronous ROM Using readmemb

```
module dual_port_rom (
    input [(addr_width-1):0] addr_a, addr_b,
    input clk,
    output reg [(data_width-1):0] q_a, q_b
);
    parameter data_width = 8;
    parameter addr_width = 8;

    reg [data_width-1:0] rom[2**addr_width-1:0];

    initial // Read the memory contents in the file
        //dual_port_rom_init.txt.
    begin
        $readmemb("dual_port_rom_init.txt", rom);
    end

    always @ (posedge clk)
    begin
        q_a <= rom[addr_a];
        q_b <= rom[addr_b];
    end
endmodule
```

Example 10-32. VHDL Dual-Port Synchronous ROM Using Initialization Function

```

library ieee;
use ieee.std_logic_1164.all;
use ieee.numeric_std.all;

entity dual_port_rom is
    generic (
        DATA_WIDTH : natural := 8;
        ADDR_WIDTH  : natural := 8
    );
    port (
        clk      : in std_logic;
        addr_a   : in natural range 0 to 2**ADDR_WIDTH - 1;
        addr_b   : in natural range 0 to 2**ADDR_WIDTH - 1;
        q_a      : out std_logic_vector((DATA_WIDTH - 1) downto 0);
        q_b      : out std_logic_vector((DATA_WIDTH - 1) downto 0)
    );
end entity;

architecture rtl of dual_port_rom is
    -- Build a 2-D array type for the ROM
    subtype word_t is std_logic_vector((DATA_WIDTH-1) downto 0);
    type memory_t is array(addr_a'high downto 0) of word_t;

    function init_rom
        return memory_t is
        variable tmp : memory_t := (others => (others => '0'));
    begin
        for addr_pos in 0 to 2**ADDR_WIDTH - 1 loop
            -- Initialize each address with the address itself
            tmp(addr_pos) := std_logic_vector(to_unsigned(addr_pos,
DATA_WIDTH));
        end loop;
        return tmp;
    end init_rom;

    -- Declare the ROM signal and specify a default initialization value.
    signal rom : memory_t := init_rom;
begin
    process(clk)
    begin
        if (rising_edge(clk)) then
            q_a <= rom(addr_a);
            q_b <= rom(addr_b);
        end if;
    end process;
end rtl;

```

Shift Registers—Inferring the ALTSHIFT_TAPS Megafunction from HDL Code

To infer shift registers, synthesis tools detect a group of shift registers of the same length and convert them to an ALTSHIFT_TAPS megafunction. To be detected, all the shift registers must have the following characteristics:

- Use the same clock and clock enable
- Do not have any other secondary signals
- Have equally spaced taps that are at least three registers apart

When you use a formal verification flow, Altera recommends that you create shift register blocks in separate entities or modules containing only the shift register logic, because you might have to treat the entity or module as a black box during formal verification.



Because formal verification tools do not support shift register megafunctions, Quartus II integrated synthesis does not infer the ALTSHIFT_TAPS megafunction when a formal verification tool is selected. You can select EDA tools for use with your design on the **EDA Tool Settings** page of the **Settings** dialog box in the Quartus II software.



For more information about the ALTSHIFT_TAPS megafunction, refer to the [ALTSHIFT_TAPS Megafunction User Guide](#).

Synthesis software recognizes shift registers only for device families that have dedicated RAM blocks, and the software uses certain guidelines to determine the best implementation.

Quartus II integrated synthesis uses the following guidelines which are common in other EDA tools. The Quartus II software determines whether to infer the ALTSHIFT_TAPS megafunction based on the width of the registered bus (W), the length between each tap (L), and the number of taps (N). If the **Auto Shift Register Recognition** setting is set to **Auto**, Quartus II integrated synthesis uses the **Optimization Technique** setting, logic and RAM utilization information about the design, and timing information from **Timing-Driven Synthesis** to determine which shift registers are implemented in RAM blocks for logic.

- If the registered bus width is one ($W = 1$), the software infers ALTSHIFT_TAPS if the number of taps times the length between each tap is greater than or equal to 64 ($N \times L \geq 64$).
- If the registered bus width is greater than one ($W > 1$), the software infers ALTSHIFT_TAPS if the registered bus width times the number of taps times the length between each tap is greater than or equal to 32 ($W \times N \times L \geq 32$).

If the length between each tap (L) is not a power of two, the software uses more logic to decode the read and write counters. This situation occurs because for different sizes of shift registers, external decode logic that uses logic elements (LEs) or ALMs is required to implement the function. This decode logic eliminates the performance and utilization advantages of implementing shift registers in memory.

The registers that the software maps to the ALTSHIFT_TAPS megafunction and places in RAM are not available in a Verilog HDL or VHDL output file for simulation tools because their node names do not exist after synthesis.

Simple Shift Register

The code samples in [Example 10-33](#) and [Example 10-34](#) show a simple, single-bit wide, 64-bit long shift register. The synthesis software implements the register ($N = 1$ and $M = 64$) in an ALTSHIFT_TAPS megafunction for supported devices and maps it to RAM in supported devices, which may be placed in dedicated RAM blocks or MLAB memory. If the length of the register is less than 64 bits, the software implements the shift register in logic.

Example 10-33. Verilog HDL Single-Bit Wide, 64-Bit Long Shift Register

```
module shift_1x64 (clk, shift, sr_in, sr_out);
    input clk, shift;
    input sr_in;
    output sr_out;

    reg [63:0] sr;

    always @ (posedge clk)
    begin
        if (shift == 1'b1)
        begin
            sr[63:1] <= sr[62:0];
            sr[0] <= sr_in;
        end
    end
    assign sr_out = sr[63];
endmodule
```

Example 10-34. VHDL Single-Bit Wide, 64-Bit Long Shift Register

```
LIBRARY IEEE;
USE IEEE.STD_LOGIC_1164.all;
ENTITY shift_1x64 IS
    PORT (
        clk: IN STD_LOGIC;
        shift: IN STD_LOGIC;
        sr_in: IN STD_LOGIC;
        sr_out: OUT STD_LOGIC
    );
END shift_1x64;

ARCHITECTURE arch OF shift_1x64 IS
    TYPE sr_length IS ARRAY (63 DOWNTO 0) OF STD_LOGIC;
    SIGNAL sr: sr_length;
BEGIN
    PROCESS (clk)
    BEGIN
        IF (clk'EVENT and clk = '1') THEN
            IF (shift = '1') THEN
                sr(63 DOWNTO 1) <= sr(62 DOWNTO 0);
                sr(0) <= sr_in;
            END IF;
        END IF;
    END PROCESS;
    sr_out <= sr(63);
END arch;
```

Shift Register with Evenly Spaced Taps

The code samples in [Example 10-35](#) and [Example 10-36](#) show a Verilog HDL and VHDL 8-bit wide, 64-bit long shift register ($W > 1$ and $M = 64$) with evenly spaced taps at 15, 31, and 47. The synthesis software implements this function in a single ALTSHIFT_TAPS megafunction and maps it to RAM in supported devices, which is allowed placement in dedicated RAM blocks or MLAB memory.

Example 10-35. Verilog HDL 8-Bit Wide, 64-Bit Long Shift Register with Evenly Spaced Taps

```
module shift_8x64_taps (clk, shift, sr_in, sr_out, sr_tap_one,
sr_tap_two, sr_tap_three );
    input clk, shift;
    input [7:0] sr_in;
    output [7:0] sr_tap_one, sr_tap_two, sr_tap_three, sr_out;

    reg [7:0] sr [63:0];
    integer n;

    always @ (posedge clk)
    begin
        if (shift == 1'b1)
        begin
            for (n = 63; n>0; n = n-1)
            begin
                sr[n] <= sr[n-1];
            end
            sr[0] <= sr_in;
        end

        end

        assign sr_tap_one = sr[15];
        assign sr_tap_two = sr[31];
        assign sr_tap_three = sr[47];
        assign sr_out = sr[63];
    endmodule
```

Example 10-36. VHDL 8-Bit Wide, 64-Bit Long Shift Register with Evenly Spaced Taps

```
LIBRARY IEEE;
USE IEEE.STD_LOGIC_1164.all;
ENTITY shift_8x64_taps IS
    PORT (
        clk: IN STD_LOGIC;
        shift: IN STD_LOGIC;
        sr_in: IN STD_LOGIC_VECTOR(7 DOWNTO 0);
        sr_tap_one: OUT STD_LOGIC_VECTOR(7 DOWNTO 0);
        sr_tap_two : OUT STD_LOGIC_VECTOR(7 DOWNTO 0);
        sr_tap_three: OUT STD_LOGIC_VECTOR(7 DOWNTO 0);
        sr_out: OUT STD_LOGIC_VECTOR(7 DOWNTO 0)
    );
END shift_8x64_taps;

ARCHITECTURE arch OF shift_8x64_taps IS
    SUBTYPE sr_width IS STD_LOGIC_VECTOR(7 DOWNTO 0);
    TYPE sr_length IS ARRAY (63 DOWNTO 0) OF sr_width;
    SIGNAL sr: sr_length;
BEGIN
    PROCESS (clk)
    BEGIN
        IF (clk'EVENT and clk = '1') THEN
            IF (shift = '1') THEN
                sr(63 DOWNTO 1) <= sr(62 DOWNTO 0);
                sr(0) <= sr_in;
            END IF;
        END IF;
        sr_tap_one <= sr(15);
        sr_tap_two <= sr(31);
        sr_tap_three <= sr(47);
        sr_out <= sr(63);
    END PROCESS;
END arch;
```

Coding Guidelines for Registers and Latches

This section provides device-specific coding recommendations for Altera registers and latches. Understanding the architecture of the target Altera device helps ensure that your code produces the expected results and achieves the optimal quality of results.

This section provides guidelines in the following areas:

- [“Register Power-Up Values in Altera Devices”](#)
- [“Secondary Register Control Signals Such as Clear and Clock Enable” on page 10-45](#)
- [“Latches” on page 10-49](#)

Register Power-Up Values in Altera Devices

Registers in the device core always power up to a low (0) logic level on all Altera devices. However, there are ways to implement logic such that registers behave as if they were powering up to a high (1) logic level.

If you use a preset signal on a device that does not support presets in the register architecture, your synthesis tool may convert the preset signal to a clear signal, which requires synthesis to perform an optimization referred to as NOT gate push-back. NOT gate push-back adds an inverter to the input and the output of the register so that the reset and power-up conditions will appear to be high, and the device operates as expected. In this case, your synthesis tool may issue a message informing you about the power-up condition. The register itself powers up low, but the register output is inverted, so the signal that arrives at all destinations is high.

Due to these effects, if you specify a non-zero reset value, you may cause your synthesis tool to use the asynchronous clear (`ac1r`) signals available on the registers to implement the high bits with NOT gate push-back. In that case, the registers look as though they power up to the specified reset value.

When an asynchronous load (`aload`) signal is available in the device registers, your synthesis tools can implement a reset of 1 or 0 value by using an asynchronous load of 1 or 0. When the synthesis tool uses a load signal, it is not performing NOT gate push-back, so the registers power up to a 0 logic level.

 For additional details, refer to the appropriate device family handbook or the appropriate handbook on the Altera website.

Designers typically use an explicit reset signal for the design, which forces all registers into their appropriate values after reset. Altera recommends this practice to reset the device after power-up to restore the proper state if there is any doubt about the power-up conditions of the device.

You can make your design more stable and avoid potential glitches by synchronizing external or combinational logic of the device architecture before you drive the asynchronous control ports of registers.

 For additional information about good synchronous design practices, refer to the *Design Recommendations for Altera Devices and the Quartus II Design Assistant* chapter in volume 1 of the *Quartus II Handbook*.

If you want to force a particular power-up condition for your design, you can use the synthesis options available in your synthesis tool. With Quartus II integrated synthesis, you can apply the **Power-Up Level** logic option. You can also apply the option with an `altera_attribute` assignment in your source code. Using this option forces synthesis to perform NOT gate push-back because synthesis tools cannot actually change the power-up states of core registers.

You can apply the Quartus II integrated synthesis **Power-Up Level** logic option to a specific register or to a design entity, module, or subdesign. If you do so, every register in that block receives the value. Registers power up to 0 by default; therefore, you can use this assignment to force all registers to power up to 1 using NOT gate push-back.

 Using NOT gate push-back as a global assignment could slightly degrade the quality of results due to the number of inverters that are required. In some situations, issues are caused by enable signal inference or secondary control logic inference. It may also be more difficult to migrate such a design to an ASIC or a HardCopy® device. You can simulate the power-up behavior in a functional simulation if you use initialization.

 The **Power-Up Level** option and the `altera_attribute` assignment are described in the *Quartus II Integrated Synthesis* chapter in volume 1 of the *Quartus II Handbook*.

Some synthesis tools can also read the default or initial values for registered signals and implement this behavior in the device. For example, Quartus II integrated synthesis converts default values for registered signals into **Power-Up Level** settings. When the Quartus II software reads the default values, the synthesized behavior matches the power-up state of the HDL code during a functional simulation.

For example, the code samples in [Example 10-37](#) and [Example 10-38](#) both infer a register for `q` and set its power-up level to high.

Example 10-37. Verilog Register with High Power-Up Value

```
reg q = 1'b1; //q has a default value of '1'

always @ (posedge clk)
begin
    q <= d;
end
```

Example 10-38. VHDL Register with High Power-Up Level

```
SIGNAL q : STD_LOGIC := '1'; -- q has a default value of '1'

PROCESS (clk, reset)
BEGIN
    IF (rising_edge(clk)) THEN
        q <= d;
    END IF;
END PROCESS;
```

 If the target device architecture does not support two asynchronous control signals, such as `aclr` and `aload`, you cannot set a different power-up state and reset state. If the NOT gate push-back algorithm creates logic to set a register to 1, that register will power-up high. If you set a different power-up condition through a synthesis assignment or initial value, the power-up level is ignored during synthesis.

Secondary Register Control Signals Such as Clear and Clock Enable

The registers in Altera FPGAs provide a number of secondary control signals (such as clear and enable signals) that you can use to implement control logic for each register without using extra logic cells. Device families vary in their support for secondary signals, so consult the device family data sheet to verify which signals are available in your target device.

To make the most efficient use of the signals in the device, your HDL code should match the device architecture as closely as possible. The control signals have a certain priority due to the nature of the architecture, so your HDL code should follow that priority where possible.

Your synthesis tool can emulate any control signals using regular logic, so achieving functionally correct results is always possible. However, if your design requirements are flexible in terms of which control signals are used and in what priority, match your design to the target device architecture to achieve the most efficient results. If the priority of the signals in your design is not the same as that of the target architecture, extra logic may be required to implement the control signals. This extra logic uses additional device resources and can cause additional delays for the control signals.

In addition, there are certain cases where using logic other than the dedicated control logic in the device architecture can have a larger impact. For example, the clock enable signal has priority over the synchronous reset or clear signal in the device architecture. The clock enable turns off the clock line in the LAB, and the clear signal is synchronous. Therefore, in the device architecture, the synchronous clear takes effect only when a clock edge occurs.

If you code a register with a synchronous clear signal that has priority over the clock enable signal, the software must emulate the clock enable functionality using data inputs to the registers. Because the signal does not use the clock enable port of a register, you cannot apply a Clock Enable Multicycle constraint. In this case, following the priority of signals available in the device is clearly the best choice for the priority of these control signals, and using a different priority causes unexpected results with an assignment to the clock enable signal.



The priority order for secondary control signals in Altera devices differs from the order for other vendors' devices. If your design requirements are flexible regarding priority, verify that the secondary control signals meet design performance requirements when migrating designs between FPGA vendors and try to match your target device architecture to achieve the best results.

The signal order is the same for all Altera device families, although, as noted previously, not all device families provide every signal. The following priority order is observed:

1. Asynchronous Clear, `aclr`—highest priority
2. Preset, `pre`
3. Asynchronous Load, `aload`
4. Enable, `ena`
5. Synchronous Clear, `sclr`
6. Synchronous Load, `sload`
7. Data In, `data`—lowest priority

The following examples provide Verilog HDL and VHDL code that creates a register with the `aclr`, `aload`, and `ena` control signals.



The Verilog HDL example (Example 10-39) does not have `adata` on the sensitivity list, but the VHDL example (Example 10-40) does. This is a limitation of the Verilog HDL language—there is no way to describe an asynchronous load signal (in which `q` toggles if `adata` toggles while `aload` is high). All synthesis tools should infer an `aload` signal from this construct despite this limitation. When they perform such inference, you may see information or warning messages from the synthesis tool.

Example 10-39. Verilog HDL D-Type Flipflop (Register) with `ena`, `aclr`, and `aload` Control Signals

```
module dff_control(clk, aclr, aload, ena, data, adata, q);
    input clk, aclr, aload, ena, data, adata;
    output q;

    reg q;

    always @ (posedge clk or posedge aclr or posedge aload)
    begin
        if (aclr)
            q <= 1'b0;
        else if (aload)
            q <= adata;
        else if (ena)
            q <= data;
    end
endmodule
```

Example 10-40. VHDL D-Type Flipflop (Register) with `ena`, `aclr`, and `aload` Control Signals

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;

ENTITY dff_control IS
    PORT (
        clk: IN STD_LOGIC;
        aclr: IN STD_LOGIC;
        aload: IN STD_LOGIC;
        adata: IN STD_LOGIC;
        ena: IN STD_LOGIC;
        data: IN STD_LOGIC;
        q: OUT STD_LOGIC
    );
END dff_control;

ARCHITECTURE rtl OF dff_control IS
BEGIN
    PROCESS (clk, aclr, aload, adata)
    BEGIN
        IF (aclr = '1') THEN
            q <= '0';
        ELSIF (aload = '1') THEN
            q <= adata;
        ELSE
            IF (clk = '1' AND clk'event) THEN
                IF (ena = '1') THEN
                    q <= data;
                END IF;
            END IF;
        END IF;
    END PROCESS;
END rtl;
```

The dedicated preset signal is available only in MAX 3000 and MAX 7000 devices; therefore, it is not included in the examples.

Creating many registers with different `sload` and `sclr` signals can make packing the registers into LABs difficult for the Quartus II Fitter because the `sclr` and `sload` signals are LAB-wide signals. In addition, using the LAB-wide `sload` signal prevents the Fitter from packing registers using the quick feedback path in the device architecture, which means that some registers cannot be packed with other logic.

Synthesis tools typically restrict use of `sload` and `sclr` signals to cases in which there are enough registers with common signals to allow good LAB packing. Using the look-up table (LUT) to implement the signals is always more flexible if it is available. Because different device families offer different numbers of control signals, inference of these signals is also device-specific. For example, because Stratix II devices have more flexibility than Stratix devices with respect to secondary control signals, synthesis tools might infer more `sload` and `sclr` signals for Stratix II devices.

If you use these additional control signals, use them in the priority order that matches the device architecture. To achieve the most efficient results, ensure the `sclr` signal has a higher priority than the `sload` signal in the same way that `ac1r` has higher priority than `aload` in the previous examples. Remember that the register signals are not inferred unless the design meets the conditions described previously. However, if your HDL described the desired behavior, the software always implements logic with the correct functionality.

In Verilog HDL, the following code for `sload` and `sclr` could replace the `if (ena) q <= data;` statements in the Verilog HDL in [Example 10-39](#) (after adding the control signals to the module declaration).

Example 10-41. Verilog HDL `sload` and `sclr` Control Signals

```
if (ena) begin
    if (sclr)
        q <= 1'b0;
    else if (sload)
        q <= sdata;
    else
        q <= data;
end
```

In VHDL, the following code for `sload` and `sclr` could replace the `IF (ena = '1') THEN q <= data; END IF;` statements in the VHDL in [Example 10-40](#) on [page 10-47](#) (after adding the control signals to the entity declaration).

Example 10-42. VHDL `sload` and `sclr` Control Signals

```
IF (ena = '1') THEN
    IF (sclr = '1') THEN
        q <= '0';
    ELSIF (sload = '1') THEN
        q <= sdata;
    ELSE
        q <= data;
    END IF;
END IF;
```

Latches

A latch is a small combinational loop that holds the value of a signal until a new value is assigned.

 Altera recommends that you design without the use of latches whenever possible.

 For additional information about the issues involved in designing with latches and combinational loops, refer to the *Design Recommendations for Altera Devices and the Quartus II Design Assistant* chapter in volume 1 of the *Quartus II Handbook*.

Latches can be inferred from HDL code when you did not intend to use a latch, as described in “[Unintentional Latch Generation](#)”. If you do intend to infer a latch, it is important to infer it correctly to guarantee correct device operation as detailed in “[Inferring Latches Correctly](#)” on page 10-50.

Unintentional Latch Generation

When you are designing combinational logic, certain coding styles can create an unintentional latch. For example, when `CASE` or `IF` statements do not cover all possible input conditions, latches may be required to hold the output if a new output value is not assigned. Check your synthesis tool messages for references to inferred latches. If your code unintentionally creates a latch, make code changes to remove the latch.

A latch is required if a signal is assigned a value outside of a clock edge (for example, with an asynchronous reset), but is not assigned a value in an edge-triggered design block. An unintentional latch may be generated if your HDL code assigns a value to a signal in an edge-triggered design block, but that logic is removed during synthesis. For example, when a `CASE` or `IF` statement tests the value of a condition with a parameter or generic that evaluates to `FALSE`, any logic or signal assignment in that statement is not required and is optimized away during synthesis. This optimization may result in a latch being generated for the signal.

 Latches have limited support in formal verification tools. Therefore, ensure that you do not infer latches unintentionally.

The `full_case` attribute can be used in Verilog HDL designs to treat unspecified cases as don't care values (X). However, using the `full_case` attribute can cause simulation mismatches because this attribute is a synthesis-only attribute, so simulation tools still treat the unspecified cases as latches.

 Refer to the appropriate chapter in the *Synthesis* section in volume 1 of the *Quartus II Handbook* for more information about using attributes in your synthesis tool. The *Quartus II Integrated Synthesis* chapter in volume 1 of the *Quartus II Handbook* provides an example explaining possible simulation mismatches.

Omitting the final `else` or `when others` clause in an `if` or `case` statement can also generate a latch. Don't care (X) assignments on the default conditions are useful in preventing latch generation. For the best logic optimization, assign the default case or final `else` value to don't care (X) instead of a logic value.

The VHDL code sample in [Example 10-43](#) prevents unintentional latches. Without the final `else` clause, this code creates unintentional latches to cover the remaining combinations of the `sel` inputs. When you are targeting a Stratix device with this code, omitting the final `else` condition can cause the synthesis software to use up to six LEs, instead of the three it uses with the `else` statement. Additionally, assigning the final `else` clause to 1 instead of X can result in slightly more LEs, because the synthesis software cannot perform as much optimization when you specify a constant value compared to a don't care value.

Example 10-43. VHDL Code Preventing Unintentional Latch Creation

```
LIBRARY ieee;
USE IEEE.std_logic_1164.all;

ENTITY nolatch IS
  PORT (a,b,c: IN STD_LOGIC;
        sel: IN STD_LOGIC_VECTOR (4 DOWNTO 0);
        oput: OUT STD_LOGIC);
END nolatch;

ARCHITECTURE rtl OF nolatch IS
BEGIN
  PROCESS (a,b,c,sel) BEGIN
    if sel = "00000" THEN
      oput <= a;
    ELSIF sel = "00001" THEN
      oput <= b;
    ELSIF sel = "00010" THEN
      oput <= c;
    ELSE
      --- Prevents latch inference
      oput <= 'X'; --/
    END if;
  END PROCESS;
END rtl;
```

Inferring Latches Correctly

Synthesis tools can infer a latch that does not exhibit the glitch and timing hazard problems typically associated with combinational loops.



Any use of latches generates warnings and is flagged if the design is migrated to a HardCopy ASIC. In addition, timing analysis does not completely model latch timing in some cases. Do not use latches unless required by your design, and you fully understand the impact of using the latches.

When using Quartus II integrated synthesis, latches that are inferred by the software are reported in the **User-Specified and Inferred Latches** section of the Compilation Report. This report indicates whether the latch is considered safe and free of timing hazards.

If a latch or combinational loop in your design is not listed in the **User-Specified and Inferred Latches** section, it means that it was not inferred as a safe latch by the software and is not considered glitch-free.

All combinational loops listed in the **Analysis & Synthesis Logic Cells Representing Combinational Loops** table in the Compilation Report are at risk of timing hazards. These entries indicate possible problems with your design that you should investigate. However, it is possible to have a correct design that includes combinational loops. For example, it is possible that the combinational loop cannot be sensitized. This can occur in cases where there is an electrical path in the hardware, but either the designer knows that the circuit never encounters data that causes that path to be activated, or the surrounding logic is set up in a mutually exclusive manner that prevents that path from ever being sensitized, independent of the data input.

For macrocell-based devices, such as MAX[®] 7000AE and MAX 3000A, all data (D-type) latches and set-reset (S-R) latches listed in the **Analysis & Synthesis User-Specified and Inferred Latches** table have an implementation free of timing hazards, such as glitches. The implementation includes both a cover term to ensure there is no glitching and a single macrocell in the feedback loop.

For 4-input LUT-based devices, such as Stratix devices, the Cyclone series, and MAX II devices, all latches in the **User-Specified and Inferred Latches** table with a single LUT in the feedback loop are free of timing hazards when a single input changes. Because of the hardware behavior of the LUT, the output does not glitch when a single input toggles between two values that are supposed to produce the same output value, such as a D-type input toggling when the enable input is inactive or a set input toggling when a reset input with higher priority is active. This hardware behavior of the LUT means that no cover term is required for a loop around a single LUT. The Quartus II software uses a single LUT in the feedback loop whenever possible. A latch that has data, enable, set, and reset inputs in addition to the output fed back to the input cannot be implemented in a single 4-input LUT. If the Quartus II software cannot implement the latch with a single-LUT loop because there are too many inputs, the **User-Specified and Inferred Latches** table indicates that the latch is not free of timing hazards.

For 6-input LUT-based devices, the software can implement all latch inputs with a single adaptive look-up table (ALUT) in the combinational loop. Therefore, all latches in the **User-Specified and Inferred Latches** table are free of timing hazards when a single input changes.

If a latch is listed as a safe latch, other optimizations performed by the Quartus II software, such as physical synthesis netlist optimizations in the Fitter, maintain the hazard-free performance.

To ensure hazard-free behavior, only one control input can change at a time. Changing two inputs simultaneously, such as deasserting set and reset at the same time, or changing data and enable at the same time, can produce incorrect behavior in any latch.

Quartus II integrated synthesis infers latches from `always` blocks in Verilog HDL and `process` statements in VHDL, but not from continuous assignments in Verilog HDL or concurrent signal assignments in VHDL. These rules are the same as for register inference. The software infers registers or flipflops only from `always` blocks and `process` statements.

The Verilog HDL code sample in [Example 10-44](#) infers a S-R latch correctly in the Quartus II software.

Example 10-44. Verilog HDL Set-Reset Latch

```
module simple_latch (
    input SetTerm,
    input ResetTerm,
    output reg LatchOut
);

    always @ (SetTerm or ResetTerm) begin
        if (SetTerm)
            LatchOut = 1'b1
        else if (ResetTerm)
            LatchOut = 1'b0
        end
    endmodule
```

The VHDL code sample in [Example 10-45](#) infers a D-type latch correctly in the Quartus II software.

Example 10-45. VHDL Data Type Latch

```
LIBRARY IEEE;
USE IEEE.std_logic_1164.all;

ENTITY simple_latch IS
    PORT (
        enable, data      : IN STD_LOGIC;
        q                 : OUT STD_LOGIC
    );
END simple_latch;

ARCHITECTURE rtl OF simple_latch IS
BEGIN

    latch : PROCESS (enable, data)
    BEGIN
        IF (enable = '1') THEN
            q <= data;
        END IF;
    END PROCESS latch;
END rtl;
```

The following example shows a Verilog HDL continuous assignment that does not infer a latch in the Quartus II software:

```
assign latch_out = (~en & latch_out) | (en & data);
```

The behavior of the assignment is similar to a latch, but it may not function correctly as a latch, and its timing is not analyzed as a latch.

Quartus II integrated synthesis also creates safe latches when possible for instantiations of the LPM_LATCH megafunction. You can use this megafunction to create a latch with any combination of data, enable, set, and reset inputs. The same limitations apply for creating safe latches as for inferring latches from HDL code.

Inferring the Altera LPM_LATCH function in another synthesis tool ensures that the implementation is also recognized as a latch in the Quartus II software. If a third-party synthesis tool implements a latch using the LPM_LATCH megafunction, the Quartus II integrated synthesis lists the latch in the **User-Specified and Inferred Latches** table in the same way as it lists latches created in HDL source code. The coding style necessary to produce an LPM_LATCH implementation may depend on your synthesis tool. Some third-party synthesis tools list the number of LPM_LATCH functions that are inferred.

For LUT-based families, the Fitter uses global routing for control signals, including signals that Analysis and Synthesis identifies as latch enables. In some cases the global insertion delay may decrease the timing performance. If necessary, you can turn off the **Quartus II Global Signal** logic option to manually prevent the use of global signals. Global latch enables are listed in the **Global & Other Fast Signals** table in the Compilation Report.

General Coding Guidelines

This section helps you understand how synthesis tools map various types of HDL code into the target Altera device. Following Altera recommended coding styles, and in some cases designing logic structures to match the appropriate device architecture, can provide significant improvements in the design's quality of results.

This section provides coding guidelines for the following logic structures:

- **“Tri-State Signals”**. This section explains how to create tri-state signals for bidirectional I/O pins.
- **“Clock Multiplexing” on page 10-54**. This section provides recommendations for multiplexing clock signals.
- **“Adder Trees” on page 10-58**. This section explains the different coding styles that lead to optimal results for devices with 4-input LUTs and 6-input ALUTs.
- **“State Machines” on page 10-60**. This section helps ensure the best results when you use state machines.
- **“Multiplexers” on page 10-67**. This section explains how multiplexers can be synthesized, addresses common problems, and provides guidelines to achieve optimal resource utilization.
- **“Cyclic Redundancy Check Functions” on page 10-70**. This section provides guidelines for getting good results when designing Cyclic Redundancy Check (CRC) functions.
- **“Comparators” on page 10-72**. This section explains different comparator implementations and provides suggestions for controlling the implementation.
- **“Counters” on page 10-73**. This section provides guidelines to ensure your counter design targets the device architecture optimally.

Tri-State Signals

When you target Altera devices, you should use tri-state signals only when they are attached to top-level bidirectional or output pins. Avoid lower-level bidirectional pins, and avoid using the Z logic value unless it is driving an output or bidirectional pin.

Synthesis tools implement designs with internal tri-state signals correctly in Altera devices using multiplexer logic, but Altera does not recommend this coding practice.



In hierarchical block-based or incremental design flows, a hierarchical boundary cannot contain any bidirectional ports, unless the lower-level bidirectional port is connected directly through the hierarchy to a top-level output pin without connecting to any other design logic. If you use boundary tri-states in a lower-level block, synthesis software must push the tri-states through the hierarchy to the top level to make use of the tri-state drivers on output pins of Altera devices. Because pushing tri-states requires optimizing through hierarchies, lower-level tri-states are restricted with block-based design methodologies.

The code in [Example 10-46](#) and [Example 10-47](#) show Verilog HDL and VHDL code that creates tri-state bidirectional signals.

Example 10-46. Verilog HDL Tri-State Signal

```
module tristate (myinput, myenable, mybidir);
    input myinput, myenable;
    inout mybidir;
    assign mybidir = (myenable ? myinput : 1'bZ);
endmodule
```

Example 10-47. VHDL Tri-State Signal

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.std_logic_arith.all;

ENTITY tristate IS
PORT (
    mybidir : INOUT STD_LOGIC;
    myinput : IN STD_LOGIC;
    myenable : IN STD_LOGIC
);
END tristate;

ARCHITECTURE rtl OF tristate IS
BEGIN
    mybidir <= 'Z' WHEN (myenable = '0') ELSE myinput;
END rtl;
```

Clock Multiplexing

Clock multiplexing is sometimes used to operate the same logic function with different clock sources. This type of logic can introduce glitches that create functional problems, and the delay inherent in the combinational logic can lead to timing problems. Clock multiplexers trigger warnings from a wide range of design rule check and timing analysis tools.

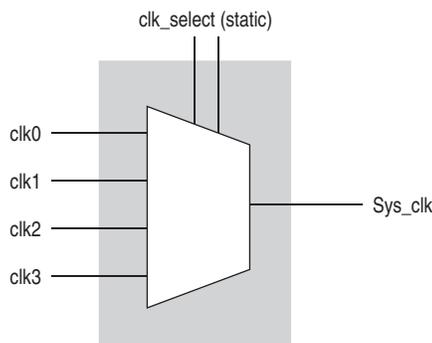
Altera recommends using dedicated hardware to perform clock multiplexing when it is available, instead of using multiplexing logic. For example, you can use the Clock Switchover feature or the Clock Control Block available in certain Altera devices. These dedicated hardware blocks avoid glitches, ensure that you use global low-skew routing lines, and avoid any possible hold time problems on the device due to logic delay on the clock line. Many Altera devices also support dynamic PLL reconfiguration, which is the safest and most robust method of changing clock rates during device operation.

Refer to the appropriate device data sheet or handbook for device-specific information about clocking structures. Also refer to the *ALTCLKCTRL Megafunction User Guide*, the *ALTPLL Megafunction User Guide*, and the *Phase-Locked Loops Reconfiguration (ALTPLL_RECONFIG) Megafunction User Guide*.

If you implement a clock multiplexer in logic cells because the design has too many clocks to use the clock control block, or if dynamic reconfiguration is too complex for your design, it is important to consider simultaneous toggling inputs and ensure glitch-free transitions.

Figure 10-2 shows a simple representation of a clock multiplexer (mux) in a device with 6-input LUTs.

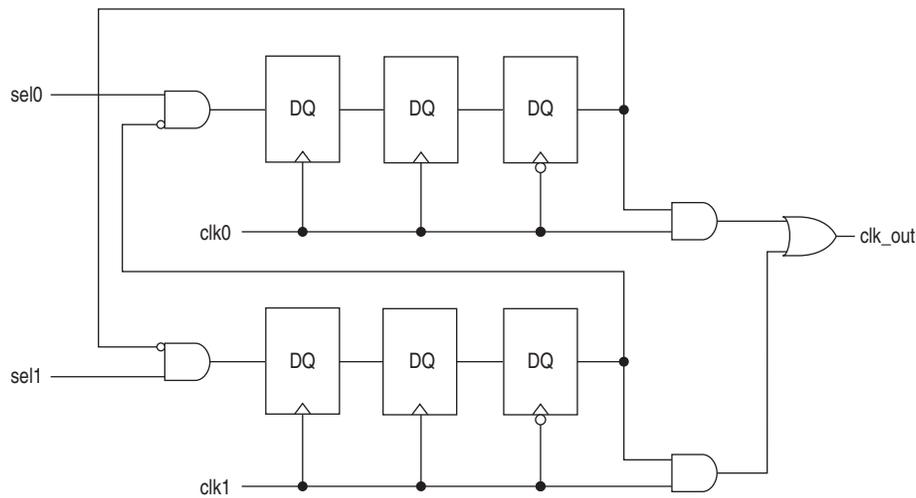
Figure 10-2. Simple Clock Multiplexer in a 6-Input LUT



The data sheet for your target device describes how LUT outputs may glitch during a simultaneous toggle of input signals, independent of the LUT function. Although, in practice, the 4:1 MUX function does not generate detectable glitches during simultaneous data input toggles, it is possible to construct cell implementations that do exhibit significant glitches, so this simple clock mux structure is not recommended. An additional problem with this implementation is that the output behaves erratically during a change in the `clk_select` signals. This behavior could create timing violations on all registers fed by the system clock and result in possible metastability.

A more sophisticated clock select structure can eliminate the simultaneous toggle and switching problems, as in [Figure 10-3](#).

Figure 10-3. Glitch-Free Clock Multiplexer Structure



This structure can be generalized for any number of clock channels. [Example 10-48](#) contains a parameterized version in Verilog HDL. The design enforces that no clock activates until all others have been inactive for at least a few cycles, and that activation occurs while the clock is low. The design applies a `synthesis_keep` directive to the AND gates on the right side of the figure, which ensures there are no simultaneous toggles on the input of the `clk_out` OR gate.



Switching from clock A to clock B requires that clock A continue to operate for at least a few cycles. If the old clock stops immediately, the design sticks. The select signals are implemented as a “one-hot” control in this example, but you can use other encoding if you prefer. The input side logic is asynchronous and is not critical. This design can tolerate extreme glitching during the switch process.

Example 10-48. Verilog HDL Clock Multiplexing Design to Avoid Glitches

```
module clock_mux (clk,clk_select,clk_out);

    parameter num_clocks = 4;

    input [num_clocks-1:0] clk;
    input [num_clocks-1:0] clk_select; // one hot
    output clk_out;

    genvar i;

    reg [num_clocks-1:0] ena_r0;
    reg [num_clocks-1:0] ena_r1;
    reg [num_clocks-1:0] ena_r2;
    wire [num_clocks-1:0] qualified_sel;

    // A look-up-table (LUT) can glitch when multiple inputs
    // change simultaneously. Use the keep attribute to
    // insert a hard logic cell buffer and prevent
    // the unrelated clocks from appearing on the same LUT.

    wire [num_clocks-1:0] gated_clks /* synthesis keep */;

    initial begin
        ena_r0 = 0;
        ena_r1 = 0;
        ena_r2 = 0;
    end

    generate
        for (i=0; i<num_clocks; i=i+1)
            begin : lp0
                wire [num_clocks-1:0] tmp_mask;
                assign tmp_mask = {num_clocks{1'b1}} ^ (1 << i);

                assign qualified_sel[i] = clk_select[i] & (~|(ena_r2 & tmp_mask));

                always @(posedge clk[i]) begin
                    ena_r0[i] <= qualified_sel[i];
                    ena_r1[i] <= ena_r0[i];
                end

                always @(negedge clk[i]) begin
                    ena_r2[i] <= ena_r1[i];
                end

                assign gated_clks[i] = clk[i] & ena_r2[i];
            end
    endgenerate

    // These will not exhibit simultaneous toggle by construction
    assign clk_out = |gated_clks;

endmodule
```

Adder Trees

Structuring adder trees appropriately to match your targeted Altera device architecture can result in significant performance and density improvements. A good example of an application using a large adder tree is a finite impulse response (FIR) correlator. Using a pipelined binary or ternary adder tree appropriately can greatly improve the quality of your results.

This section explains why coding recommendations are different for Altera 4-input LUT devices and 6-input LUT devices.

Architectures with 4-Input LUTs in Logic Elements

Architectures such as Stratix devices and the Cyclone series of devices contain 4-input LUTs as the standard combinational structure in the LE.

If your design can tolerate pipelining, the fastest way to add three numbers A, B, and C in devices that use 4-input lookup tables is to add A + B, register the output, and then add the registered output to C. Adding A + B takes one level of logic (one bit is added in one LE), so this runs at full clock speed. This can be extended to as many numbers as desired.

Example 10-49 shows five numbers A, B, C, D, and E are added. Adding five numbers in devices that use 4-input lookup tables requires four adders and three levels of registers for a total of 64 LEs (for 16-bit numbers).

Example 10-49. Verilog HDL Pipelined Binary Tree

```

module binary_adder_tree (a, b, c, d, e, clk, out);
    parameter width = 16;
    input [width-1:0] a, b, c, d, e;
    input clk;
    output [width-1:0] out;

    wire [width-1:0] sum1, sum2, sum3, sum4;
    reg [width-1:0] sumreg1, sumreg2, sumreg3, sumreg4;
    // Registers

    always @ (posedge CLK)
        begin
            sumreg1 <= sum1;
            sumreg2 <= sum2;
            sumreg3 <= sum3;
            sumreg4 <= sum4;
        end

    // 2-bit additions
    assign sum1 = A + B;
    assign sum2 = C + D;
    assign sum3 = sumreg1 + sumreg2;
    assign sum4 = sumreg3 + E;
    assign out = sumreg4;
endmodule

```

Architectures with 6-Input LUTs in Adaptive Logic Modules

High-performance Altera device families use a 6-input LUT in their basic logic structure, so these devices benefit from a different coding style from the previous example presented for 4-input LUTs. Specifically, in these devices, ALMs can simultaneously add three bits. Therefore, the tree in [Example 10-49](#) must be two levels deep and contain just two add-by-three inputs instead of four add-by-two inputs.

Although the code in the previous example compiles successfully for 6-input LUT devices, the code is inefficient and does not take advantage of the 6-input adaptive ALUT. By restructuring the tree as a ternary tree, the design becomes much more efficient, significantly improving density utilization. Therefore, when you are targeting with ALUTs and ALMs, large pipelined binary adder trees designed for 4-input LUT architectures should be rewritten to take advantage of the advanced device architecture.

[Example 10-50](#) uses just 32 ALUTs in a Stratix II device—more than a 4:1 advantage over the number of LUTs in the prior example implemented in a Stratix device.



You cannot pack a LAB full when using this type of coding style because of the number of LAB inputs. However, in a typical design, the Quartus II Fitter can pack other logic into each LAB to take advantage of the unused ALMs.

Example 10-50. Verilog HDL Pipelined Ternary Tree

```
module ternary_adder_tree (a, b, c, d, e, clk, out);
    parameter width = 16;
    input [width-1:0] a, b, c, d, e;
    input clk;
    output [width-1:0] out;

    wire [width-1:0] sum1, sum2;
    reg [width-1:0] sumreg1, sumreg2;
    // registers

    always @ (posedge clk)
        begin
            sumreg1 <= sum1;
            sumreg2 <= sum2;
        end

    // 3-bit additions
    assign sum1 = a + b + c;
    assign sum2 = sumreg1 + d + e;
    assign out = sumreg2;
endmodule
```

These examples show pipelined adders, but partitioning your addition operations can help you achieve better results in nonpipelined adders as well. If your design is not pipelined, a ternary tree provides much better performance than a binary tree. For example, depending on your synthesis tool, the HDL code $sum = (A + B + C) + (D + E)$ is more likely to create the optimal implementation of a 3-input adder for $A + B + C$ followed by a 3-input adder for $sum1 + D + E$ than the code without the parentheses. If you do not add the parentheses, the synthesis tool may partition the addition in a way that is not optimal for the architecture.

State Machines

Synthesis tools can recognize and encode Verilog HDL and VHDL state machines during synthesis. This section presents guidelines to ensure the best results when you use state machines. Ensuring that your synthesis tool recognizes a piece of code as a state machine allows the tool to recode the state variables to improve the quality of results, and allows the tool to use the known properties of state machines to optimize other parts of the design. When synthesis recognizes a state machine, it is often able to improve the design area and performance.

To achieve the best results on average, synthesis tools often use one-hot encoding for FPGA devices and minimal-bit encoding for CPLD devices, although the choice of implementation can vary for different state machines and different devices. Refer to your synthesis tool documentation for specific ways to control the manner in which state machines are encoded.



For information about state machine encoding in Quartus II integrated synthesis, refer to the *State Machine Processing* section in the *Quartus II Integrated Synthesis* chapter in volume 1 of the *Quartus II Handbook*.

To ensure proper recognition and inference of state machines and to improve the quality of results, Altera recommends that you observe the following guidelines, which apply to both Verilog HDL and VHDL:

- Assign default values to outputs derived from the state machine so that synthesis does not generate unwanted latches.
- Separate the state machine logic from all arithmetic functions and data paths, including assigning output values.
- If your design contains an operation that is used by more than one state, define the operation outside the state machine and cause the output logic of the state machine to use this value.
- Use a simple asynchronous or synchronous reset to ensure a defined power-up state. If your state machine design contains more elaborate reset logic, such as both an asynchronous reset and an asynchronous load, the Quartus II software generates regular logic rather than inferring a state machine.

If a state machine enters an illegal state due to a problem with the device, the design likely ceases to function correctly until the next reset of the state machine. Synthesis tools do not provide for this situation by default. The same issue applies to any other registers if there is some kind of fault in the system. A default or when others clause does not affect this operation, assuming that your design never deliberately enters this state. Synthesis tools remove any logic generated by a default state if it is not reachable by normal state machine operation.

Many synthesis tools (including Quartus II integrated synthesis) have an option to implement a safe state machine. The software inserts extra logic to detect an illegal state and force the state machine's transition to the reset state. It is commonly used when the state machine can enter an illegal state. The most common cause of this situation is a state machine that has control inputs that come from another clock domain, such as the control logic for a dual-clock FIFO.

This option protects only state machines by forcing them into the reset state. All other registers in the design are not protected this way. If the design has asynchronous inputs, Altera recommends using a synchronization register chain instead of relying on the safe state machine option.

-  For additional information about tool-specific options for implementing state machines, refer to the tool vendor's documentation or the appropriate chapter in the *Synthesis* section in volume 1 of the *Quartus II Handbook*.

The following two sections, “Verilog HDL State Machines” and “VHDL State Machines” on page 10-65, describe additional language-specific guidelines and coding examples.

Verilog HDL State Machines

To ensure proper recognition and inference of Verilog HDL state machines, observe the following additional Verilog HDL guidelines. Some of these guidelines may be specific to Quartus II integrated synthesis. Refer to your synthesis tool documentation for specific coding recommendations.

If the state machine is not recognized and inferred by the synthesis software (such as Quartus II integrated synthesis), the state machine is implemented as regular logic gates and registers, and the state machine is not listed as a state machine in the **Analysis & Synthesis** section of the Quartus II Compilation Report. In this case, the software does not perform any of the optimizations that are specific to state machines.

- If you are using the SystemVerilog standard, use enumerated types to describe state machines. For more information, refer to “[SystemVerilog State Machine Coding Example](#)” on page 10-64.
- Represent the states in a state machine with the parameter data types in Verilog-1995 and Verilog-2001, and use the parameters to make state assignments. For more information, refer to “[Verilog-2001 State Machine Coding Example](#)” on page 10-62. This parameter implementation makes the state machine easier to read and reduces the risk of errors during coding.

 Altera recommends against the direct use of integer values for state variables, such as `next_state <= 0`. However, using an integer does not prevent inference in the Quartus II software.

- No state machine is inferred in the Quartus II software if the state transition logic uses arithmetic similar to that in the following example:

```
case (state)
  0: begin
    if (ena) next_state <= state + 2;
    else next_state <= state + 1;
  end
  1: begin
    ...
  endcase
```

- No state machine is inferred in the Quartus II software if the state variable is an output.
- No state machine is inferred in the Quartus II software for signed variables.

Verilog-2001 State Machine Coding Example

The following module `verilog_fsm` is an example of a typical Verilog HDL state machine implementation (Example 10-51).

This state machine has five states. The asynchronous reset sets the variable `state` to `state_0`. The sum of `in_1` and `in_2` is an output of the state machine in `state_1` and `state_2`. The difference (`in_1 - in_2`) is also used in `state_1` and `state_2`. The temporary variables `tmp_out_0` and `tmp_out_1` store the sum and the difference of `in_1` and `in_2`. Using these temporary variables in the various states of the state machine ensures proper resource sharing between the mutually exclusive states.

Example 10-51. Verilog-2001 State Machine

```
module verilog_fsm (clk, reset, in_1, in_2, out);
    input clk, reset;
    input [3:0] in_1, in_2;
    output [4:0] out;
    parameter state_0 = 3'b000;
    parameter state_1 = 3'b001;
    parameter state_2 = 3'b010;
    parameter state_3 = 3'b011;
    parameter state_4 = 3'b100;

    reg [4:0] tmp_out_0, tmp_out_1, tmp_out_2;
    reg [2:0] state, next_state;

    always @ (posedge clk or posedge reset)
    begin
        if (reset)
            state <= state_0;
        else
            state <= next_state;
    end
    always @ (state or in_1 or in_2)
    begin
        tmp_out_0 = in_1 + in_2;
        tmp_out_1 = in_1 - in_2;
        case (state)
            state_0: begin
                tmp_out_2 <= in_1 + 5'b00001;
                next_state <= state_1;
            end
            state_1: begin
                if (in_1 < in_2) begin
                    next_state <= state_2;
                    tmp_out_2 <= tmp_out_0;
                end
                else begin
                    next_state <= state_3;
                    tmp_out_2 <= tmp_out_1;
                end
            end
            state_2: begin
                tmp_out_2 <= tmp_out_0 - 5'b00001;
                next_state <= state_3;
            end
            state_3: begin
                tmp_out_2 <= tmp_out_1 + 5'b00001;
                next_state <= state_0;
            end
            state_4: begin
                tmp_out_2 <= in_2 + 5'b00001;
                next_state <= state_0;
            end
            default: begin
                tmp_out_2 <= 5'b00000;
                next_state <= state_0;
            end
        endcase
        assign out = tmp_out_2;
    end
endmodule
```

An equivalent implementation of this state machine can be achieved by using `'define` instead of the parameter data type, as follows:

```
'define state_0 3'b000
'define state_1 3'b001
'define state_2 3'b010
'define state_3 3'b011
'define state_4 3'b100
```

In this case, the state and next_state assignments are assigned a `'state_x` instead of a `state_x`, for example:

```
next_state <= 'state_3;
```



Although the `'define` construct is supported, Altera strongly recommends the use of the parameter data type because doing so preserves the state names throughout synthesis.

SystemVerilog State Machine Coding Example

The module `enum_fsm` in [Example 10-52](#) is an example of a SystemVerilog state machine implementation that uses enumerated types. Altera recommends using this coding style to describe state machines in SystemVerilog.



In Quartus II integrated synthesis, the enumerated type that defines the states for the state machine must be of an unsigned integer type as in [Example 10-52](#). If you do not specify the enumerated type as `int unsigned`, a signed `int` type is used by default. In this case, the Quartus II integrated synthesis synthesizes the design, but does not infer or optimize the logic as a state machine.

Example 10-52. SystemVerilog State Machine Using Enumerated Types

```
module enum_fsm (input clk, reset, input int data[3:0], output int o);  
  
enum int unsigned { S0 = 0, S1 = 2, S2 = 4, S3 = 8 } state, next_state;  
  
always_comb begin : next_state_logic  
    next_state = S0;  
    case(state)  
        S0: next_state = S1;  
        S1: next_state = S2;  
        S2: next_state = S3;  
        S3: next_state = S3;  
    endcase  
end  
  
always_comb begin  
    case(state)  
        S0: o = data[3];  
        S1: o = data[2];  
        S2: o = data[1];  
        S3: o = data[0];  
    endcase  
end  
  
always_ff@(posedge clk or negedge reset) begin  
    if(~reset)  
        state <= S0;  
    else  
        state <= next_state;  
    end  
end  
endmodule
```

VHDL State Machines

To ensure proper recognition and inference of VHDL state machines, represent the states in a state machine with enumerated types and use the corresponding types to make state assignments. This implementation makes the state machine easier to read and reduces the risk of errors during coding. If the state is not represented by an enumerated type, synthesis software (such as Quartus II integrated synthesis) does not recognize the state machine. Instead, the state machine is implemented as regular logic gates and registers and the state machine is not listed as a state machine in the **Analysis & Synthesis** section of the Quartus II Compilation Report. In this case, the software does not perform any of the optimizations that are specific to state machines.

VHDL State Machine Coding Example

The following entity, `vhd1_fsm`, is an example of a typical VHDL state machine implementation (Example 10-53).

This state machine has five states. The asynchronous reset sets the variable `state` to `state_0`. The sum of `in1` and `in2` is an output of the state machine in `state_1` and `state_2`. The difference (`in1 - in2`) is also used in `state_1` and `state_2`. The temporary variables `tmp_out_0` and `tmp_out_1` store the sum and the difference of `in1` and `in2`. Using these temporary variables in the various states of the state machine ensures proper resource sharing between the mutually exclusive states.

Example 10-53. VHDL State Machine

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.numeric_std.all;
ENTITY vhdl_fsm IS
  PORT(
    clk: IN STD_LOGIC;
    reset: IN STD_LOGIC;
    in1: IN UNSIGNED(4 downto 0);
    in2: IN UNSIGNED(4 downto 0);
    out_1: OUT UNSIGNED(4 downto 0)
  );
END vhdl_fsm;
ARCHITECTURE rtl OF vhdl_fsm IS
  TYPE Tstate IS (state_0, state_1, state_2, state_3, state_4);
  SIGNAL state: Tstate;
  SIGNAL next_state: Tstate;
BEGIN
  PROCESS(clk, reset)
  BEGIN
    IF reset = '1' THEN
      state <= state_0;
    ELSIF rising_edge(clk) THEN
      state <= next_state;
    END IF;
  END PROCESS;
  PROCESS (state, in1, in2)
    VARIABLE tmp_out_0: UNSIGNED (4 downto 0);
    VARIABLE tmp_out_1: UNSIGNED (4 downto 0);
  BEGIN
    tmp_out_0 := in1 + in2;
    tmp_out_1 := in1 - in2;
    CASE state IS
      WHEN state_0 =>
        out_1 <= in1;
        next_state <= state_1;
      WHEN state_1 =>
        IF (in1 < in2) then
          next_state <= state_2;
          out_1 <= tmp_out_0;
        ELSE
          next_state <= state_3;
          out_1 <= tmp_out_1;
        END IF;
      WHEN state_2 =>
        IF (in1 < "0100") then
          out_1 <= tmp_out_0;
        ELSE
          out_1 <= tmp_out_1;
        END IF;
        next_state <= state_3;
      WHEN state_3 =>
        out_1 <= "1111";
        next_state <= state_4;
      WHEN state_4 =>
        out_1 <= in2;
        next_state <= state_0;
      WHEN OTHERS =>
        out_1 <= "0000";
        next_state <= state_0;
    END CASE;
  END PROCESS;
END rtl;

```

Multiplexers

Multiplexers form a large portion of the logic utilization in many FPGA designs. By optimizing your multiplexer logic, you ensure the most efficient implementation in your Altera device. This section addresses common problems and provides design guidelines to achieve optimal resource utilization for multiplexer designs. The section also describes various types of multiplexers, and how they are implemented.

For more information, refer to the *Advanced Synthesis Cookbook: A Design Guide for Stratix II, Stratix III, and Stratix IV Devices*.

Quartus II Software Option for Multiplexer Restructuring

Quartus II integrated synthesis provides the **Restructure Multiplexers** logic option that extracts and optimizes buses of multiplexers during synthesis. The default setting **Auto** for this option uses the optimization when it is most likely to benefit the optimization targets for your design. You can turn the option on or off specifically to have more control over its use.

 For details, refer to the *Restructure Multiplexers* section in the *Quartus II Integrated Synthesis* chapter in volume 1 of the *Quartus II Handbook*.

Even with this Quartus II-specific option turned on, it is beneficial to understand how your coding style can be interpreted by your synthesis tool, and avoid the situations that can cause problems in your design.

Multiplexer Types

This section addresses how multiplexers are created from various types of HDL code. CASE statements, IF statements, and state machines are all common sources of multiplexer logic in designs. These HDL structures create different types of multiplexers, including binary multiplexers, selector multiplexers, and priority multiplexers. Understanding how multiplexers are created from HDL code, and how they might be implemented during synthesis, is the first step toward optimizing multiplexer structures for best results.

Binary Multiplexers

Binary multiplexers select inputs based on binary-encoded selection bits.

Example 10-54 shows Verilog HDL code for two ways to describe a simple 4:1 binary multiplexer.

Example 10-54. Verilog HDL Binary-Encoded Multiplexers

```
case (sel)
  2'b00: z = a;
  2'b01: z = b;
  2'b10: z = c;
  2'b11: z = d;
endcase
```

Stratix series devices starting with the Stratix II device family feature 6-input look up tables (LUTs) which are perfectly suited for 4:1 multiplexer building blocks (4 data and 2 select inputs). The extended input mode facilitates implementing 8:1 blocks, and the fractured mode handles residual 2:1 multiplexer pairs. For device families using 4-input LUTs, such as the Cyclone series and Stratix devices, the 4:1 binary multiplexer is efficiently implemented by using two 4-input LUTs. Larger binary multiplexers are decomposed by the synthesis tool into 4:1 multiplexer blocks, possibly with a residual 2:1 multiplexer at the head.

Selector Multiplexers

Selector multiplexers have a separate select line for each data input. The select lines for the multiplexer are one-hot encoded. [Example 10-55](#) shows a simple Verilog HDL code example describing a one-hot selector multiplexer.

Example 10-55. Verilog HDL One-Hot-Encoded Case Statement

```
case (sel)
  4'b0001: z = a;
  4'b0010: z = b;
  4'b0100: z = c;
  4'b1000: z = d;
  default: z = 1'bx;
endcase
```

Selector multiplexers are commonly built as a tree of AND and OR gates. An N-input selector multiplexer of this structure is slightly less efficient in implementation than a binary multiplexer. However, in many cases the select signal is the output of a decoder, in which case Quartus II Synthesis will try to combine the selector and decoder into a binary multiplexer.

Priority Multiplexers

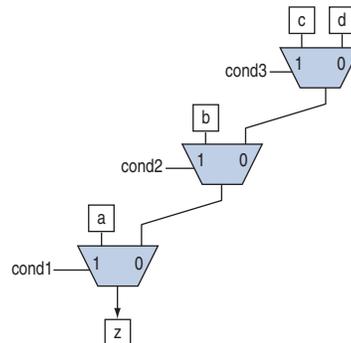
In priority multiplexers, the select logic implies a priority. The options to select the correct item must be checked in a specific order based on signal priority. These structures commonly are created from IF, ELSE, WHEN, SELECT, and ?: statements in VHDL or Verilog HDL. The example VHDL code in [Example 10-56](#) probably results in the schematic implementation illustrated in [Figure 10-4](#).

Example 10-56. VHDL IF Statement Implying Priority

```
IF cond1 THEN z <= a;
ELSIF cond2 THEN z <= b;
ELSIF cond3 THEN z <= c;
ELSE z <= d;
END IF;
```

The multiplexers in Figure 10-4 form a chain, evaluating each condition or select bit sequentially.

Figure 10-4. Priority Multiplexer Implementation of an IF Statement



Depending on the number of multiplexers in the chain, the timing delay through this chain can become large, especially for device families with 4-input LUTs.

To improve the timing delay through the multiplexer, avoid priority multiplexers if priority is not required. If the order of the choices is not important to the design, use a CASE statement to implement a binary or selector multiplexer instead of a priority multiplexer. If delay through the structure is important in a multiplexed design requiring priority, consider recoding the design to reduce the number of logic levels to minimize delay, especially along your critical paths.

Implicit Defaults in If Statements

The IF statements in Verilog HDL and VHDL can be a convenient way to specify conditions that do not easily lend themselves to a CASE-type approach. However, using IF statements can result in complicated multiplexer trees that are not easy for synthesis tools to optimize. In particular, every IF statement has an implicit ELSE condition, even when it is not specified. These implicit defaults can cause additional complexity in a multiplexed design.

There are several ways you can simplify multiplexed logic and remove unneeded defaults. The optimal method may be to recode the design so the logic takes the structure of a 4:1 CASE statement. Alternatively, if priority is important, you can restructure the code to reduce default cases and flatten the multiplexer. Examine whether the default "ELSE IF" conditions are don't care cases. You may be able to create a default ELSE statement to make the behavior explicit. Avoid unnecessary default conditions in your multiplexer logic to reduce the complexity and logic utilization required to implement your design.

Default or Others Case Assignment

To fully specify the cases in a CASE statement, include a default (Verilog HDL) or OTHERS (VHDL) assignment. This assignment is especially important in one-hot encoding schemes where many combinations of the select lines are unused. Specifying a case for the unused select line combinations gives the synthesis tool information about how to synthesize these cases, and is required by the Verilog HDL and VHDL language specifications.

Some designs do not require that the outcome in the unused cases be considered, often because designers assume these cases will not occur. For these types of designs, you can specify any value for the default or OTHERS assignment. However, be aware that the assignment value you choose can have a large effect on the logic utilization required to implement the design due to the different ways synthesis tools treat different values for the assignment, and how the synthesis tools use different speed and area optimizations.

To obtain best results, explicitly define invalid CASE selections with a separate default or OTHERS statement instead of combining the invalid cases with one of the defined cases.

If the value in the invalid cases is not important, specify those cases explicitly by assigning the X (don't care) logic value instead of choosing another value. This assignment allows your synthesis tool to perform the best area optimizations.

Cyclic Redundancy Check Functions

CRC computations are used heavily by communications protocols and storage devices to detect any corruption of data. These functions are highly effective; there is a very low probability that corrupted data can pass a 32-bit CRC check.

CRC functions typically use wide XOR gates to compare the data. The way synthesis tools flatten and factor these XOR gates to implement the logic in FPGA LUTs can greatly impact the area and performance results for the design. XOR gates have a cancellation property that creates an exceptionally large number of reasonable factoring combinations, so synthesis tools cannot always choose the best result by default.

The 6-input ALUT has a significant advantage over 4-input LUTs for these designs. When properly synthesized, CRC processing designs can run at high speeds in devices with 6-input ALUTs.

The following guidelines help you improve the quality of results for CRC designs in Altera devices.

If Performance is Important, Optimize for Speed

Synthesis tools flatten XOR gates to minimize area and depth of levels of logic. Synthesis tools such as Quartus II integrated synthesis target area optimization by default for these logic structures. Therefore, for more focus on depth reduction, set the synthesis optimization technique to speed.

Flattening for depth sometimes causes a significant increase in area.

Use Separate CRC Blocks Instead of Cascaded Stages

Some designers optimize their CRC designs to use cascaded stages (for example, four stages of 8 bits). In such designs, intermediate calculations are used as required (such as the calculations after 8, 24, or 32 bits) depending on the data width. This design is not optimal in FPGA devices. The XOR cancellations that can be performed in CRC designs mean that the function does not require all the intermediate calculations to

determine the final result. Therefore, forcing the use of intermediate calculations increases the area required to implement the function, as well as increasing the logic depth because of the cascading. It is typically better to create full separate CRC blocks for each data width that you require in the design, and then multiplex them together to choose the appropriate mode at a given time

Use Separate CRC Blocks Instead of Allowing Blocks to Merge

Synthesis tools often attempt to optimize CRC designs by sharing resources and extracting duplicates in two different CRC blocks because of the factoring options in the XOR logic. As addressed previously, the CRC logic allows significant reductions, but this works best when each CRC function is optimized separately. Check for duplicate extraction behavior if you have different CRC functions that are driven by common data signals or that feed the same destination signals.

If you are having problems with the quality of results and you see that two CRC functions are sharing logic, ensure that the blocks are synthesized independently using one of the following methods:

- Define each CRC block as a separate design partition in an incremental compilation design flow.
 - For details, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.
- Synthesize each CRC block as a separate project in your third-party synthesis tool and then write a separate Verilog Quartus Mapping (.vqm) or EDIF netlist file for each.

Take Advantage of Latency if Available

If your design can use more than one cycle to implement the CRC functionality, adding registers and retiming the design can help reduce area, improve performance, and reduce power utilization. If your synthesis tool offers a retiming feature (such as the Quartus II software **Perform gate-level register retiming** option), you can insert an extra bank of registers at the input and allow the retiming feature to move the registers for better results. You can also build the CRC unit half as wide and alternate between halves of the data in each clock cycle.

Save Power by Disabling CRC Blocks When Not in Use

CRC designs are heavy consumers of dynamic power because the logic toggles whenever there is a change in the design. To save power, use clock enables to disable the CRC function for every clock cycle that the logic is not required. Some designs don't check the CRC results for a few clock cycles while other logic is performed. It is valuable to disable the CRC function even for this short amount of time.

Use the Device Synchronous Load (sload) Signal to Initialize

The data in many CRC designs must be initialized to 1's before operation. If your target device supports the use of the `sload` signal, you should use it to set all the registers in your design to 1's before operation. To enable use of the `sload` signal, follow the coding guidelines presented in “[Secondary Register Control Signals Such as Clear and Clock Enable](#)” on page 10–45. You can check the register equations in the Chip Planner to ensure that the signal was used as expected.



If you must force a register implementation using an `sload` signal, you can use low-level device primitives as described in the *Designing with Low-Level Primitives User Guide*.

Comparators

Synthesis software, including Quartus II integrated synthesis, uses device and context-specific implementation rules for comparators (`<`, `>`, or `==`) and selects the best one for your design. This section provides some information about the different types of implementations available and provides suggestions on how you can code your design to encourage a specific implementation.

The `==` comparator is implemented in general logic cells. The `<` comparison can be implemented using the carry chain or general logic cells. In devices with 6-input ALUTs, the carry chain is capable of comparing up to three bits per cell. In devices with 4-input LUTs, the capacity is one bit of comparison per cell, which is similar to an add/subtract chain. The carry chain implementation tends to be faster than the general logic on standalone benchmark test cases, but can result in lower performance when it is part of a larger design due to the increased restriction on the Fitter. The area requirement is similar for most input patterns. The synthesis software selects an appropriate implementation based on the input pattern.

If you are using Quartus II integrated synthesis, you can guide the synthesis by using specific coding styles. To select a carry chain implementation explicitly, rephrase your comparison in terms of addition. As a simple example, the following coding style allows the synthesis tool to select the implementation, which is most likely using general logic cells in modern device families:

```
wire [6:0] a,b;  
wire alb = a<b;
```

In the following coding style, the synthesis tool uses a carry chain (except for a few cases, such as when the chain is very short or the signals `a` and `b` minimize to the same signal):

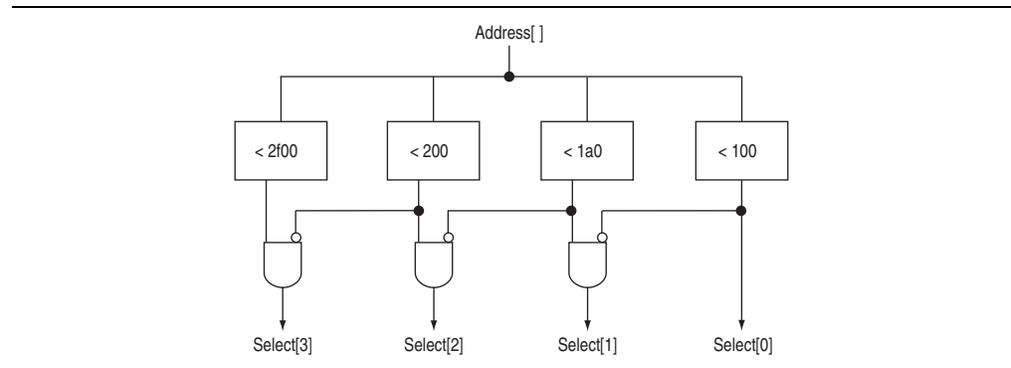
```
wire [6:0] a,b;  
wire [7:0] tmp = a - b;  
wire alb = tmp[7]
```

This second coding style uses the top bit of the `tmp` signal, which is 1 in twos complement logic if `a` is less than `b`, because the subtraction `a - b` results in a negative number.

If you have any information about the range of the input, you have “don’t care” values that you can use to optimize the design. Because this information is not available to the synthesis tool, you can often reduce the device area required to implement the comparator with specific hand implementation of the logic.

You can also check whether a bus value is within a constant range with a small amount of logic area by using the logic structure in Figure 10-5. This type of logic occurs frequently in address decoders.

Figure 10-5. Example Logic Structure for Using Comparators to Check a Bus Value Range



Counters

Implementing counters in HDL code is easy; they are implemented with an adder followed by registers. Remember that the register control signals, such as enable (*ena*), synchronous clear (*sclr*), and synchronous load (*sload*), are available. For the best area utilization, ensure that the up/down control or controls are expressed in terms of one addition instead of two separate addition operators.

If you use the following coding style, your synthesis tool may implement two separate carry chains for addition (if it doesn't detect the issue and optimize the logic):

```
out <= count_up ? out + 1 : out - 1;
```

The following coding style requires only one adder along with some other logic:

```
out <= out + (count_up ? 1 : -1);
```

In this case, the coding style better matches the device hardware because there is only one carry chain adder, and the -1 constant logic is implemented in the LUT in front of the adder without adding extra area utilization.

Designing with Low-Level Primitives

Low-level HDL design is the practice of using low-level primitives and assignments to dictate a particular hardware implementation for a piece of logic. Low-level primitives are small architectural building blocks that assist you in creating your design. With the Quartus II software, you can use low-level HDL design techniques to force a specific hardware implementation that can help you achieve better resource utilization or faster timing results.



Using low-level primitives is an advanced technique to help with specific design challenges, and is optional in the Altera design flow. For many designs, synthesizing generic HDL source code and Altera megafunctions gives you the best results.

Low-level primitives allow you to use the following types of coding techniques:

- Instantiate the logic cell or LCELL primitive to prevent Quartus II integrated synthesis from performing optimizations across a logic cell

- Create carry and cascade chains using `CARRY`, `CARRY_SUM`, and `CASCADE` primitives
- Instantiate registers with specific control signals using `DFF` primitives
- Specify the creation of LUT functions by identifying the LUT boundaries
- Use I/O buffers to specify I/O standards, current strengths, and other I/O assignments
- Use I/O buffers to specify differential pin names in your HDL code, instead of using the automatically-generated negative pin name for each pair

 For details about and examples of using these types of assignments, refer to the *Designing with Low-Level Primitives User Guide*.

Conclusion

Because coding style and megafunction implementation can have such a large effect on your design performance, it is important to match the coding style to the device architecture from the very beginning of the design process. To improve design performance and area utilization, take advantage of advanced device features, such as memory and DSP blocks, as well as the logic architecture of the targeted Altera device by following the coding recommendations presented in this chapter.

 For additional optimization recommendations, refer to the *Area and Timing Optimization* chapter in volume 2 of the *Quartus II Handbook*.

Document Revision History

Table 10-2 shows the revision history for this document.

Table 10-2. Document Revision History (Part 1 of 2)

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Changed to new document template. ■ Updated Unintentional Latch Generation content. ■ Code update for Example 10-18.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Added support for mixed-width RAM ■ Updated support for <code>no_rw_check</code> for inferring RAM blocks ■ Added support for byte-enable
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Updated support for Controlling Inference and Implementation in Device RAM Blocks ■ Updated support for Shift Registers
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Corrected and updated several examples ■ Added support for Arria II GX devices ■ Other minor changes to chapter

Table 10–2. Document Revision History (Part 2 of 2)

Date	Version	Changes
November 2008	8.1.0	Changed to 8-1/2 x 11 page size. No change to content.
May 2008	8.0.0	Updates for the Quartus II software version 8.0 release, including: <ul style="list-style-type: none"> ■ Added information to “RAM Functions—Inferring ALTSYNCRAM and ALTDPRAM Megafunctions from HDL Code” on page 6–13 ■ Added information to “Avoid Unsupported Reset and Control Conditions” on page 6–14 ■ Added information to “Check Read-During-Write Behavior” on page 6–16 ■ Added two new examples to “ROM Functions—Inferring ALTSYNCRAM and LPM_ROM Megafunctions from HDL Code” on page 6–28: Example 6–24 and Example 6–25 ■ Added new section: “Clock Multiplexing” on page 6–46 ■ Added hyperlinks to references within the chapter ■ Minor editorial updates

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter describes the industry-leading analysis, reporting, and optimization features that can help you manage metastability in Altera® devices. You can use the Quartus® II software to analyze the average mean time between failures (MTBF) due to metastability caused by synchronization of asynchronous signals, and optimize the design to improve the metastability MTBF. This chapter explains how to take advantage of these features in the Quartus II software, and provides guidelines to help you reduce the chance of metastability effects caused by signal synchronization.

Introduction

All registers in digital devices, such as FPGAs, have defined signal-timing requirements that allow each register to correctly capture data at its input ports and produce an output signal. To ensure reliable operation, the input to a register must be stable for a minimum amount of time before the clock edge (register setup time or t_{SU}) and a minimum amount of time after the clock edge (register hold time or t_H). The register output is available after a specified clock-to-output delay (t_{CO}).

If the data violates the setup or hold time requirements, the output of the register might go into a metastable state. In a metastable state, the voltage at the register output hovers at a value between the high and low states, which means the output transition to a defined high or low state is delayed beyond the specified t_{CO} . Different destination registers might capture different values for the metastable signal, which can cause the system to fail.

In synchronous systems, the input signals must always meet the register timing requirements, so that metastability does not occur. Metastability problems commonly occur when a signal is transferred between circuitry in unrelated or asynchronous clock domains, because the signal can arrive at any time relative to the destination clock.

The MTBF due to metastability is an estimate of the average time between instances when metastability could cause a design failure. A high MTBF (such as hundreds or thousands of years between metastability failures) indicates a more robust design. You should determine an acceptable target MTBF in the context of your entire system and taking in account that MTBF calculations are statistical estimates.

The metastability MTBF for a specific signal transfer, or all the transfers in a design, can be calculated using information about the design and the device characteristics. Improving the metastability MTBF for your design reduces the chance that signal transfers could cause metastability problems in your device.



For more information about metastability due to signal synchronization, its effects in FPGAs, and how MTBF is calculated, refer to the *Understanding Metastability in FPGAs* white paper on the Altera website. Your overall device MTBF is also affected by other FPGA failure mechanisms that you cannot control with your design. For information about Altera device reliability, refer to the *Reliability Report* on the Altera website.

The Quartus II software provides analysis, optimization, and reporting features to help manage metastability in Altera designs. These metastability features are supported only for designs constrained with the Quartus II Timing Analyzer. Both typical and worst-case MBTF values are generated for select device families.

- ② For information about device and version support for the metastability features in the Quartus II software, refer to the Quartus II Help.

This chapter contains the following topics:

- [“Metastability Analysis in the Quartus II Software”](#)
- [“Metastability and MTBF Reporting” on page 11-5](#)
- [“MTBF Optimization” on page 11-8](#)
- [“Reducing Metastability Effects” on page 11-9](#)
- [“Scripting Support” on page 11-11](#)

Metastability Analysis in the Quartus II Software

When a signal transfers between circuitry in unrelated or asynchronous clock domains, the first register in the new clock domain acts as a synchronization register. To minimize the failures due to metastability in asynchronous signal transfers, circuit designers typically use a sequence of registers (a synchronization register chain or synchronizer) in the destination clock domain to resynchronize the signal to the new clock domain and allow additional time for a potentially metastable signal to resolve to a known value. Designers commonly use two registers to synchronize a new signal, but a standard of three registers provides better metastability protection.

The timing analyzer can analyze and report the MTBF for each identified synchronizer that meets its timing requirements, and can generate an estimate of the overall design MTBF. The software uses this information to optimize the design MTBF, and you can use this information to determine whether your design requires longer synchronizer chains.

This section contains the following subsections:

- [“Synchronization Register Chains”](#)
- [“Identifying Synchronizers for Metastability Analysis” on page 11-4](#)
- [“How Timing Constraints Affect Synchronizer Identification and Metastability Analysis” on page 11-4](#)

For information about the reports generated by the timing analyzer, refer to [“Metastability and MTBF Reporting” on page 11-5](#). For more information about optimizing the MTBF, refer to [“MTBF Optimization” on page 11-8](#).

Synchronization Register Chains

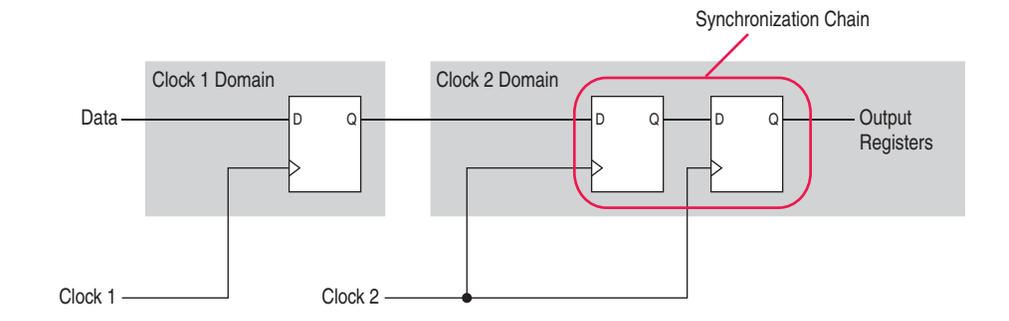
A synchronization register chain, or synchronizer, is defined as a sequence of registers that meets the following requirements:

- The registers in the chain are all clocked by the same clock or phase-related clocks.

- The first register in the chain is driven asynchronously or from an unrelated clock domain.
- Each register fans out to only one register, except the last register in the chain.

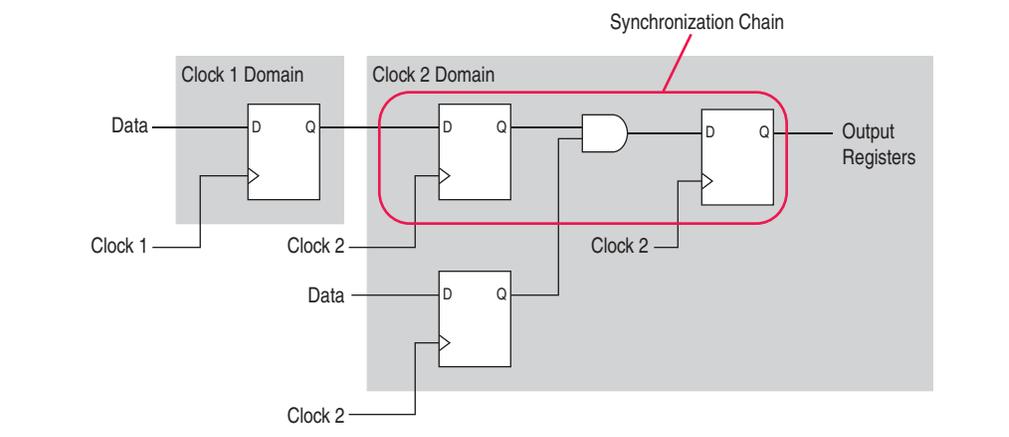
The length of the synchronization register chain is the number of registers in the synchronizing clock domain that meet the above requirements. Figure 11-1 shows a sample two-register synchronization chain.

Figure 11-1. Sample Synchronization Register Chain



The path between synchronization registers can contain combinational logic as long as all registers of the synchronization register chain are in the same clock domain. Figure 11-2 shows an example of a synchronization register chain that includes logic between the registers.

Figure 11-2. Sample Synchronization Register Chain Containing Logic



The Quartus II software uses the design timing constraints to determine which connections are asynchronous signal transfers, as described in “How Timing Constraints Affect Synchronizer Identification and Metastability Analysis” on page 11-4.

The timing slack available in the register-to-register paths of the synchronizer allows a metastable signal to settle, and is referred to as the available settling time. The available settling time in the MTBF calculation for a synchronizer is the sum of the output timing slacks for each register in the chain. Adding available settling time with additional synchronization registers improves the metastability MTBF.

Identifying Synchronizers for Metastability Analysis

The first step in enabling metastability MTBF analysis and optimization in the Quartus II software is to identify which registers are part of a synchronization register chain. You can apply synchronizer identification settings globally to automatically list possible synchronizers with the **Synchronizer identification** option on the **Timing Analyzer** page in the **Settings** dialog box.

Synchronization chains are already identified within most Altera intellectual property (IP) cores.

- ④ For more information about how to enable metastability MTBF analysis and optimization in the Quartus II software, and more detailed descriptions of the synchronizer identification settings, refer to *Identifying Synchronizers for Metastability Analysis* in Quartus II Help.

How Timing Constraints Affect Synchronizer Identification and Metastability Analysis

The timing analyzer can analyze metastability MTBF only if a synchronization chain meets its timing requirements. The metastability failure rate depends on the timing slack available in the synchronizer's register-to-register connections, because that slack is the available settling time for a potential metastable signal. Therefore, you must ensure that your design is correctly constrained with the real application frequency requirements to get an accurate MTBF report.

In addition, the **Auto** and **Forced If Asynchronous** synchronizer identification options use timing constraints to automatically detect the synchronizer chains in the design. These options check for signal transfers between circuitry in unrelated or asynchronous clock domains, so clock domains must be related correctly with timing constraints.

The timing analyzer views input ports as asynchronous signals unless they are associated correctly with a clock domain. If an input port fans out to registers that are not acting as synchronization registers, apply a `set_input_delay` constraint to the input port; otherwise, the input register might be reported as a synchronization register. Constraining a synchronous input port with a `set_max_delay` constraint for a setup (t_{SU}) requirement does not prevent synchronizer identification because the constraint does not associate the input port with a clock domain.

Instead, use the following command to specify an input setup requirement associated with a clock:

```
set_input_delay -max -clock <clock name> <latch - launch - tsu requirement> <input port name>
```

Registers that are at the end of false paths are also considered synchronization registers because false paths are not timing-analyzed. Because there are no timing requirements for these paths, the signal may change at any point, which may violate the t_{SU} and t_H of the register. Therefore, these registers are identified as synchronization registers. If these registers are not used for synchronization, you can turn off synchronizer identification and analysis. To do so, set **Synchronizer Identification** to **Off** for the first synchronization register in these register chains.

Metastability and MTBF Reporting

The Quartus II software reports the metastability analysis results in the Compilation Report and Timing Analyzer reports as described in “[Metastability Reports](#)”. The MTBF calculation uses timing and structural information about the design, silicon characteristics, and operating conditions, along with the data toggle rate described in “[Synchronizer Data Toggle Rate in MTBF Calculation](#)” on page 11-7.

If you change the **Synchronizer Identification** settings, you can generate new metastability reports by rerunning the timing analyzer. However, you should rerun the Fitter first so that the registers identified with the new setting can be optimized for metastability MTBF. For information about metastability optimization, refer to “[MTBF Optimization](#)” on page 11-8.

For more information about how metastability MTBF is calculated, refer to the [Understanding Metastability in FPGAs](#) white paper.

Metastability Reports

Metastability reports provide summaries of the metastability analysis results. In addition to the MTBF Summary and Synchronizer Summary reports, the Timing Analyzer tool reports additional statistics in a report for each synchronizer chain.

 For more information about how to access metastability reports in the Quartus II software, refer to [Viewing Metastability Reports](#) in Quartus II Help.

 If the design uses only the **Auto Synchronizer Identification** setting, the reports list likely synchronizers but do not report MTBF. To obtain an MTBF for each register chain, force identification of synchronization registers as described in “[Identifying Synchronizers for Metastability Analysis](#)” on page 11-4.

 If the synchronizer chain does not meet its timing requirements, the reports list identified synchronizers but do not report MTBF. To obtain MTBF calculations, ensure that the design is properly constrained and that the synchronizer meets its timing requirements, as described in “[How Timing Constraints Affect Synchronizer Identification and Metastability Analysis](#)” on page 11-4.

MTBF Summary Report

The MTBF Summary reports an estimate of the overall robustness of cross-clock domain and asynchronous transfers in the design. This estimate uses the MTBF results of all synchronization chains in the design to calculate an MTBF for the entire design.

The MTBF Summary Report reports the **Typical MTBF of Design** and the **Worst-Case MTBF of Design** for supported fully-characterized devices. The typical MTBF result assumes typical conditions, defined as nominal silicon characteristics for the selected device speed grade, as well as nominal operating conditions. The worst case MTBF result uses the worst case silicon characteristics for the selected device speed grade.

When you analyze multiple timing corners in the timing analyzer, the MTBF calculation may vary because of changes in the operating conditions, and the timing slack or available metastability settling time. Altera recommends running multi-corner timing analysis to ensure that you analyze the worst MTBF results, because the worst timing corner for MTBF does not necessarily match the worst corner for timing performance.

- For more information about turning on multicorner timing analysis in the Quartus II software, refer to the *Timing Analyzer page* in Quartus II Help.

The MTBF Summary report also lists the **Number of Synchronizer Chains Found** and the length of the **Shortest Synchronizer Chain**, which can help you identify whether the report is based on accurate information. If the number of synchronizer chains found is different from what you expect, or if the length of the shortest synchronizer chain is less than you expect, you might have to add or change **Synchronizer Identification** settings for the design. The report also provides the **Worst Case Available Settling Time**, defined as the available settling time for the synchronizer with the worst MTBF.

You can use the reported **Fraction of Chains for which MTBFs Could Not be Calculated** to determine whether a high proportion of chains are missing in the metastability analysis. A fraction of 1, for example, means that MTBF could not be calculated for any chains in the design. MTBF is not calculated if you have not identified the chain with the appropriate **Synchronizer identification** option, or if paths are not timing-analyzed and therefore have no valid slack for metastability analysis. You might have to correct your timing constraints to enable complete analysis of the applicable register chains.

Finally, the MTBF Summary report specifies how an increase of 100ps in available settling time increases the MTBF values. If your MTBF is not satisfactory, this metric can help you determine how much extra slack would be required in your synchronizer chain to allow you to reach the desired design MTBF.

Synchronizer Summary Report

The **Synchronizer Summary** lists the synchronization register chains detected in the design depending on the Synchronizer Identification setting. The **Source Node** is the register or input port that is the source of the asynchronous transfer. The **Synchronization Node** is the first register of the synchronization chain. The **Source Clock** is the clock domain of the source node, and the **Synchronization Clock** is the clock domain of the synchronizer chain.

This summary reports the calculated **Worst-Case MTBF**, if available, and the **Typical MTBF**, for each appropriately identified synchronization register chain that meets its timing requirement. To see more detail about each synchronizer, refer to the statistics report described in the following section.

Synchronizer Chain Statistics Report in the Timing Analyzer

The timing analyzer provides an additional report for each synchronizer chain. The **Chain Summary** tab matches the Synchronizer Summary information described in the previous section, while the **Statistics** tab adds more details, including whether the **Method of Synchronizer Identification** was **User Specified** (with the **Forced if Asynchronous** or **Forced** settings for the **Synchronizer Identification** setting), or **Automatic** (with the **Auto** setting). The **Number of Synchronization Registers in Chain** report provides information about the parameters that affect the MTBF calculation, including the **Available Settling Time** for the chain and the **Data Toggle Rate Used in MTBF Calculation**.

 For information about the toggle rate, see “[Synchronizer Data Toggle Rate in MTBF Calculation](#)” on page 11-7.

The following information is also included to help you locate the chain is in your design:

- **Source Clock** and **Asynchronous Source** node of the signal.
- **Synchronization Clock** in the destination clock domain.
- Node names of the **Synchronization Registers** in the chain.

Synchronizer Data Toggle Rate in MTBF Calculation

The MTBF calculations assume the data being synchronized is switching at a toggle rate of 12.5% of the source clock frequency. That is, the arriving data is assumed to switch once every eight source clock cycles. If multiple clocks apply, the highest frequency is used. If no source clocks can be determined, the data rate is taken as 12.5% of the synchronization clock frequency.

If you know an approximate rate at which the data changes, specify it with the **Synchronizer Toggle Rate** assignment in the Assignment Editor. You can also apply this assignment to an entity or the entire design. Set the data toggle rate, in number of transitions per second, on the first register of a synchronization chain. The timing analyzer takes the specified rate into account when computing the MTBF of that particular register chain. If a data signal never toggles and does not affect the reliability of the design, you can set the **Synchronizer Toggle Rate** to **0** for the synchronization chain so the MTBF is not reported. To apply the assignment with Tcl, use the following command:

```
set_instance_assignment -name SYNCHRONIZER_TOGGLE_RATE <toggle rate in transitions/second> -to <register name>
```

 There are two other assignments associated with toggle rates, which are not used for metastability MTBF calculations. The I/O Maximum Toggle Rate is only used for pins, and specifies the worst-case toggle rates used for signal integrity purposes. The Power Toggle Rate assignment is used to specify the expected time-averaged toggle rate, and is used by the PowerPlay Power Analyzer to estimate time-averaged power consumption.

MTBF Optimization

In addition to reporting synchronization register chains and MTBF values found in the design, the Quartus II software can also protect these registers from optimizations that might negatively impact MTBF and can optimize the register placement and routing if the MTBF is too low. Synchronization register chains must first be explicitly identified as synchronizers, as described in “[Identifying Synchronizers for Metastability Analysis](#)” on page 11-4. Altera recommends that you set **Synchronizer Identification** to **Forced If Asynchronous** for all registers that are part of a synchronizer chain.

Optimization algorithms, such as register duplication and logic retiming in physical synthesis, are not performed on identified synchronization registers. The Fitter protects the number of synchronization registers specified by the **Synchronizer Register Chain Length** option which is described in the next section.

In addition, the Fitter optimizes identified synchronizers for improved MTBF by placing and routing the registers to increase their output setup slack values. Adding slack in the synchronizer chain increases the available settling time for a potentially metastable signal, which improves the chance that the signal resolves to a known value, and exponentially increases the design MTBF. The Fitter optimizes the number of synchronization registers specified by the **Synchronizer Register Chain Length** option.

Metastability optimization is **on** by default. To view or change the option, on the Assignments menu, click **Settings**. Under **Fitter Settings**, click **More Settings**. From the **More Settings** dialog box, you can turn on or off the **Optimize Design for Metastability** option. To turn the optimization on or off with Tcl, use the following command:

```
set_global_assignment -name OPTIMIZE_FOR_METASTABILITY <ON|OFF>
```

Synchronization Register Chain Length

The **Synchronization Register Chain Length** option specifies how many registers should be protected from optimizations that might reduce MTBF for each register chain, and controls how many registers should be optimized to increase MTBF with the **Optimize Design for Metastability** option. For example, if the **Synchronization Register Chain Length** option is set to **2**, optimizations such as register duplication or logic retiming are prevented from being performed on the first two registers in all identified synchronization chains. The first two registers are also optimized to improve MTBF when the **Optimize Design for Metastability** option is turned on.

The default setting for the **Synchronization Register Chain Length** option is **2**. The first register of a synchronization chain is always protected from operations that might reduce MTBF, but you should set the protection length to protect more of the synchronizer chain. Altera recommends that you set this option to the maximum length of synchronization chains you have in your design so that all synchronization registers are preserved and optimized for MTBF.

To change the global **Synchronization Register Chain Length** option, on the Assignments menu, click **Settings**. Under **Analysis & Synthesis Settings**, click **More Settings**. From the **More Settings** dialog box, you can set the **Synchronization Register Chain Length**.

You can also set the **Synchronization Register Chain Length** on a node or an entity in the Assignment Editor. You can set this value on the first register in a synchronization chain to specify how many registers to protect and optimize in this chain. This individual setting is useful if you want to protect and optimize extra registers that you have created in a specific synchronization chain that has low MTBF, or optimize less registers for MTBF in a specific chain where the maximum frequency or timing performance is not being met. To make the global setting with Tcl, use the following command:

```
set_global_assignment -name SYNCHRONIZATION_REGISTER_CHAIN_LENGTH  
<number of registers>
```

To apply the assignment to a design instance or the first register in a specific chain with Tcl, use the following command:

```
set_instance_assignment -name SYNCHRONIZATION_REGISTER_CHAIN_LENGTH  
<number of registers> -to <register or instance name>
```

Reducing Metastability Effects

You can check your design's metastability MTBF in the Metastability Summary report described in [“Metastability Reports” on page 11-5](#), and determine an acceptable target MTBF given the context of your entire system and the fact that MTBF calculations are statistical estimates. A high metastability MTBF (such as hundreds or thousands of years between metastability failures) indicates a more robust design.

This section provides guidelines to ensure complete and accurate metastability analysis, and some suggestions to follow if the Quartus II metastability reports calculate an unacceptable MTBF value. The Timing Optimization Advisor (available from the Tools menu) gives similar suggestions in the Metastability Optimization section.

Apply Complete System-Centric Timing Constraints for the Timing Analyzer

To enable the Quartus II metastability features, make sure that the timing analyzer is used for timing analysis.

Ensure that the design is fully timing constrained and that it meets its timing requirements. If the synchronization chain does not meet its timing requirements, MTBF cannot be calculated. If the clock domain constraints are set up incorrectly, the signal transfers between circuitry in unrelated or asynchronous clock domains might be identified incorrectly.

Use industry-standard system-centric I/O timing constraints instead of using FPGA-centric timing constraints. As described in [“How Timing Constraints Affect Synchronizer Identification and Metastability Analysis” on page 11-4](#), you should use `set_input_delay` constraints in place of `set_max_delay` constraints to associate each input port with a clock domain to help eliminate false positives during synchronization register identification.

Force the Identification of Synchronization Registers

Use the guidelines in [“Identifying Synchronizers for Metastability Analysis” on page 11-4](#) to ensure the software reports and optimizes the appropriate register chains.

In summary, identify synchronization registers with the **Synchronizer Identification** set to **Forced If Asynchronous** in the Assignment Editor. If there are any registers that the software detects as synchronous but you want to be analyzed for metastability, apply the **Forced** setting to the first synchronizing register. Set **Synchronizer Identification** to **Off** for registers that are not synchronizers for asynchronous signals or unrelated clock domains.

To help you find the synchronizers in your design, you can set the global **Synchronizer Identification** setting on the **Timing Analyzer** page of the **Settings** dialog box to **Auto** to generate a list of all the possible synchronization chains in your design.

Set the Synchronizer Data Toggle Rate

The MTBF calculations assume the data being synchronized is switching at a toggle rate of 12.5% of the source clock frequency. To obtain a more accurate MTBF for a specific chain or all chains in your design, set the **Synchronizer Toggle Rate** as described in [“Synchronizer Data Toggle Rate in MTBF Calculation” on page 11-7](#).

Optimize Metastability During Fitting

Ensure that the **Optimize Design for Metastability** setting described in [“MTBF Optimization” on page 11-8](#) is turned on.

Increase the Length of Synchronizers to Protect and Optimize

Increase the Synchronizer Chain Length parameter to the maximum length of synchronization chains in your design, as described in [“Synchronization Register Chain Length” on page 11-8](#). If you have synchronization chains longer than 2 identified in your design, you can protect the entire synchronization chain from operations that might reduce MTBF and allow metastability optimizations to improve the MTBF.

Set Fitter Effort to Standard Fit instead of Auto Fit

If your design MTBF is too low after following the previous guidelines in this section, you can try increasing the Fitter effort to perform more metastability optimization. The default **Auto Fit** setting reduces the Fitter’s effort after meeting the design’s timing and routing requirements to reduce compilation time. This effort reduction can result in less metastability optimization if the timing requirements are easy to meet. If **Auto Fit** reduces the Fitter’s effort during your design compilation, setting the Fitter effort to **Standard Fit** might improve the design’s MTBF results. In the **Settings** dialog box, on the **Fitter Settings** page, set **Fitter effort** to **Standard Fit**.

Increase the Number of Stages Used in Synchronizers, If Possible

Designers commonly use two registers in a synchronization chain to minimize the occurrence of metastable events, and a standard of three registers provides better metastability protection. However, synchronization chains with two or even three registers may not be enough to produce a high enough MTBF when the design runs at high clock and data frequencies.

If a synchronization chain is reported to have a low MTBF, consider adding an additional register stage to your synchronization chain. This additional stage increases the settling time of the synchronization chain, allowing more opportunity for the signal to resolve to a known state during a metastable event. Additional settling time increases the MTBF of the chain and improves the robustness of your design. However, adding a synchronization stage introduces an additional stage of latency on the signal.

If you use the Altera FIFO megafunction with separate read and write clocks to cross clock domains, increase the metastability protection (and latency) for better MTBF. In the MegaWizard™ Plug-In Manager for the DCFIFO function, choose the **Best metastability protection, best fmax, unsynchronized clocks** option to add three or more synchronization stages. You can increase the number of stages to more than three using the **How many sync stages?** setting.

Select a Faster Speed Grade Device, if Possible

The design MTBF depends on process parameters of the device used. Faster devices are less susceptible to metastability issues. If the design MTBF falls significantly below the target MTBF, switching to a faster speed grade can improve the MTBF substantially.

Scripting Support

You can run procedures and make settings described in this chapter in a Tcl script. You can also run some procedures at a command prompt. For detailed information about scripting command options, refer to the Quartus II Command-Line and Tcl API Help browser. To run the Help browser, type the following command at the command prompt:

```
quartus_sh --qhelp ←
```



For more information about Tcl scripting, refer to the *Tcl Scripting* chapter in volume 2 of the *Quartus II Handbook*. For more information about settings and constraints in the Quartus II software, refer to the *Quartus II Settings File Reference Manual*. For more information about command-line scripting, refer to the *Command-Line Scripting* chapter in volume 2 of the *Quartus II Handbook* and *About Quartus II Scripting* in Quartus II Help.

Identifying Synchronizers for Metastability Analysis

To apply the global Synchronizer Identification assignment described in “[Identifying Synchronizers for Metastability Analysis](#)” on page 11-4, use the following command:

```
set_global_assignment -name SYNCHRONIZER_IDENTIFICATION  
<OFF|AUTO|"FORCED IF ASYNCHRONOUS">
```

To apply the **Synchronizer Identification** assignment to a specific register or instance, use the following command:

```
set_instance_assignment -name SYNCHRONIZER_IDENTIFICATION  
<AUTO|"FORCED IF ASYNCHRONOUS"|FORCED|OFF> -to <register or instance  
name>
```

Synchronizer Data Toggle Rate in MTBF Calculation

To specify a toggle rate for MTBF calculations as described on page “[Synchronizer Data Toggle Rate in MTBF Calculation](#)” on page 11-7, use the following command:

```
set_instance_assignment -name SYNCHRONIZER_TOGGLE_RATE <toggle rate in transitions/second> -to <register name>
```

report_metastability and Tcl Command

If you use a command-line or scripting flow, you can generate the metastability analysis reports described in “[Metastability Reports](#)” on page 11-5 outside of the Quartus II and user interfaces. [Table 11-1](#) describes the options for the `report_metastability` and Tcl command.

Table 11-1. report_metastability Command Options

Option	Description
<code>-append</code>	If output is sent to a file, this option appends the result to that file. Otherwise, the file is overwritten.
<code>-file <name></code>	Sends the results to an ASCII or HTML file. The extension specified in the file name determines the file type—either <code>*.txt</code> or <code>*.html</code> .
<code>-panel_name <name></code>	Sends the results to the panel and specifies the name of the new panel.
<code>-stdout</code>	Indicates the report be sent to the standard output, via messages. This option is required only if you have selected another output format, such as a file, and would also like to receive messages.

MTBF Optimization

To ensure that metastability optimization described on page “[MTBF Optimization](#)” on page 11-8 is turned on (or to turn it off), use the following command:

```
set_global_assignment -name OPTIMIZE_FOR_METASTABILITY <ON|OFF>
```

Synchronization Register Chain Length

To globally set the number of registers in a synchronization chain to be protected and optimized as described on page “[Synchronization Register Chain Length](#)” on page 11-8, use the following command:

```
set_global_assignment -name SYNCHRONIZATION_REGISTER_CHAIN_LENGTH <number of registers>
```

To apply the assignment to a design instance or the first register in a specific chain, use the following command:

```
set_instance_assignment -name SYNCHRONIZATION_REGISTER_CHAIN_LENGTH <number of registers> -to <register or instance name>
```

Conclusion

Altera's Quartus II software provides industry-leading analysis and optimization features to help you manage metastability in your FPGA designs. Set up your Quartus II project with the appropriate constraints and settings to enable the software to analyze, report, and optimize the design MTBF. Take advantage of these features in the Quartus II software and follow the guidelines in this chapter to make your design more robust with respect to metastability.

Document Revision History

Table 11-2 shows the revision history for this document.

Table 11-2. Document Revision History

Date	Version	Changes
December 2010	10.0.1	Changed to new document template.
July 2010	10.0.0	Technical edit.
November 2009	9.1.0	Clarified description of synchronizer identification settings. Minor changes to text and figures throughout document.
March 2009	9.0.0	Initial release.

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter provides a set of guidelines to help you partition your design to take advantage of Quartus II incremental compilation, and to help you create a design floorplan using LogicLock™ regions when they are recommended to support the flow.

The Quartus® II incremental compilation feature allows you to partition a design, compile partitions separately, and reuse results for unchanged partitions. It provides the following benefits:

- Reduces compilation times by an average of 75% for large design changes
- Preserves performance for unchanged design blocks
- Provides repeatable results and reduces the number of compilations
- Enables team-based design flows

 For more information about the incremental compilation feature and application examples, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*. For feature support, refer to *About Incremental Compilation* in Quartus II Help.

This document contains the following sections:

- “Overview: Incremental Compilation” on page 12–2
- “Design Flows Using Incremental Compilation” on page 12–3
- “Why to Plan Partitions and Floorplan Assignments” on page 12–5
- “General Partitioning Guidelines” on page 12–7
- “Design Partition Guidelines” on page 12–10
- “Design Partition Guidelines for Third-Party IP Delivery” on page 12–26
- “Checking Partition Quality” on page 12–31
- “Including SDC Constraints from Lower-Level Partitions for Third-Party IP Delivery” on page 12–37
- “Introduction to Design Floorplans” on page 12–41
- “Design Floorplan Placement Guidelines” on page 12–44
- “Checking Floorplan Quality” on page 12–50
- “Recommended Design Flows and Application Examples” on page 12–52

Overview: Incremental Compilation

Quartus II incremental compilation is an optional compilation flow that enhances the default Quartus II compilation. If you do not partition your design for incremental compilation, your design is compiled using the default “flat” compilation flow. This section provides an overview of the incremental compilation flow and highlights several best practices.

To prepare your design for incremental compilation, you first determine which logical hierarchy boundaries should be treated as separate partitions in your design, and ensure your design hierarchy and source code is set up to support this partitioning. Then create design partition assignments in the Quartus II software to specify which hierarchy blocks are compiled independently as partitions (including empty partitions for missing or incomplete logic blocks).

During compilation, Quartus II Analysis & Synthesis and the Fitter create separate netlists for each partition. These netlists are internal post-synthesis and post-fit database representations of your design.

In subsequent compilations, you can select which netlist type to preserve for each partition. You can either reuse the synthesis or fitting netlist, or instruct the Quartus II software to resynthesize the source files. You can also use compilation results exported from another Quartus II project.

When you make changes to your design, the Quartus II software recompiles only the required partitions and merges the new compilation results with existing netlists for other partitions, according to the degree of results preservation you set with the netlist for each design partition.

In some cases, as described in [“Introduction to Design Floorplans” on page 12-41](#), Altera recommends that you create a design floorplan with placement assignments to constrain parts of the design to specific regions of the device.

- ❓ For step-by-step information on using incremental compilation to recompile only changed parts of your design, refer to [Using the Incremental Compilation Design Flow](#) in Quartus II Help.

Recommendations for the Netlist Type

For subsequent compilations, you must specify which post-compilation netlist you want to use by specifying the netlist type for each partition.

Use the following general guidelines to set the netlist type for partitions:

- **Source File**—Use this setting to resynthesize the source code (with any new assignments and replace any previous synthesis or Fitter results)
 - If you modify the design source, the software automatically resynthesizes the partitions with the appropriate netlist type settings, which makes the **Source File** setting optional in this case.
 - Most assignments do not trigger an automatic recompilation, so setting the netlist type to **Source File** is required to compile the source files with new assignments or constraints that affect synthesis.

- **Post-Synthesis** (default)—Use this setting to re-fit the design (with any new Fitter assignments), but preserve the synthesis results when the source files have not changed. If it is difficult to meet the required timing performance, you can use this setting to allow the Fitter the most flexibility in placement and routing. This setting does not reduce compilation time as much as the **Post-Fit** setting or preserve timing performance from the previous compilation.
- **Post-Fit**—Use this setting to preserve Fitter and performance results when the source files have not changed. This setting reduces compilation time the most, and preserves timing performance from the previous compilation.
- **Post-Fit with Fitter Preservation Level set to Placement**—Use the advanced Fitter Preservation Level setting on the **Advanced** tab in the **Design Partition Properties** dialog box to allow more flexibility to find the best routing for all partitions given their placement.

The Quartus II software Rapid Recompile feature instructs the compiler to reuse the compatible compilation results if most of the design has not changed since the last compilation. This feature reduces compilation times and preserves performance when there are small and isolated design changes within a partition, and works with all netlist type settings. You do not have control over which parts of the design are recompiled using this option; the Compiler determines which parts of the design must be recompiled. You can turn on the **Rapid Recompile** option in the Quartus II software on the **Incremental Compilation** page of the **Settings** dialog box.

Design Flows Using Incremental Compilation

The Quartus II incremental compilation feature supports various design flows. Your design flow affects the impact design partitions have on design optimization, and the amount of design planning required to obtain optimal results.

In the standard incremental compilation flow, the top-level design is divided into partitions, which can be compiled and optimized together in one Quartus II project. If another team member or IP provider is developing source code for the top-level design, they can functionally verify their partition independently, and then simply provide the partition's source code to the project lead for integration into the top-level design. If the project lead wants to compile the top-level design when source code is not yet complete for a partition, they can create an empty placeholder for the partition until the code is ready and added to the top-level design.

Compiling all design partitions in a single Quartus II project ensures that all design logic is compiled with a consistent set of assignments and allows the software to perform global placement and routing optimizations. Compiling all design logic together is beneficial for FPGA design flows because all parts of the design must use the same shared set of device resources. Therefore, it is often easier to ensure good quality of results when partitions are developed within a single top-level Quartus II project.

To enable team-based development, you can design and optimize partitions by accessing the top-level project from a shared source control system or creating copies of the top-level Quartus II project framework, and later exporting each partition so that the post-synthesis netlist or post-fitting results can be integrated into the top-level design.

If required for third-party IP delivery, or in cases where designers can not access a shared or copied top-level project framework, you can create and compile a design partition logic in isolation and export a partition that is included in the top-level project. If this type of design flow is necessary, planning and rigorous design guidelines may be required to ensure that designers have a consistent view of project assignments and resource allocations. Therefore, developing partitions in completely separate Quartus II projects can be more challenging than having all source code within one project or developing design partitions within the same top-level project framework. See the following subsection “[Project Management in Team-Based Design Flows](#)” on page 12-4 for more information.

You can also combine design flows and use exported partitions only when it is necessary to support your design environment. For example, if the top-level design includes one or more design blocks that will be optimized by remote designers or IP providers, you can integrate those blocks into the reserved partitions in top-level design when the code is complete, but also have other partitions that will be developed within the top-level design.

If any partitions are developed independently, the project lead must ensure that top-level constraints (such as timing constraints, any relevant floorplan or pin assignments, and optimization settings) are consistent with those used by all designers working independently.

Project Management in Team-Based Design Flows

If possible, each team member should work within the same top-level project framework. As mentioned in the beginning of this section, using the same project framework amongst team members ensures that designers have the settings and constraints needed for their partition, and allows designers to analyze how their design block interacts with other partitions in the top-level design.

In a team-based environment where designers have access to the project through source control software, each designer can use project files as read-only and develop their partition within the source control system. As designers check in their completed designs, other team members can see how their design partitions interact. If designs do not have access to a source control system, the project lead can provide each designer with a snapshot copy of the top-level project framework to use as they develop their partitions. In both cases, each designer exports their completed design as a partition, and then the project lead integrates the partition into the top-level design. The project lead can choose to use just the post-synthesis netlist and rerun placement and routing, or to use the post-fitting results to preserve the placement and routing results from the other designer's project. Using post-synthesis partitions gives the Fitter the most flexibility and is likely to achieve a good result for all partitions, but if one partition has difficulty meeting timing, then the designer can preserve their successful fitting results.

Alternatively, designers can use their own Quartus II project for their independent design block. You might use this design flow if a designer, such as a third-party IP provider, does not have access to the entire top-level project framework. In this case, each designer must create their own project with all the relevant assignments and constraints. As mentioned at the beginning of this section, this type of design flow requires more planning and rigorous design guidelines. If the project lead plans to incorporate the post-fitting compilation results for the partition, then this design flow requires especially careful planning to avoid resource conflicts.

The project lead also has the option to generate design partition scripts to manage resource and timing budgets in the top-level design when partitions are developed outside the top-level project framework. Scripts make it easier for any designers of independent Quartus II projects to implement instructions from the project lead. The Quartus II design partition scripts feature creates Tcl scripts or Tcl Script Files (.tcl) and makefiles that an independent designer can run to set up an independent Quartus II project.

- For more information on how to generate design partition scripts, refer to *Generating Design Partition Scripts for Project Management* in Quartus II Help.

If designers create Quartus II assignments or timing constraints for their partitions, they must ensure that the constraints are integrated into the top-level design. If partition designers use the same top-level project framework (and design hierarchy), then the constraints or Synopsys Design Constraints File (.sdc) can be easily copied or included in the top-level design. If any partition designers use a separate Quartus II project with a different design hierarchy, they must ensure that constraints are applied to the appropriate level of hierarchy in the top-level design, and should design the .sdc for easy delivery to the project lead, as described in *“Including SDC Constraints from Lower-Level Partitions for Third-Party IP Delivery”* on page 12–37.

You cannot use an exported partition if you want to migrate to a HardCopy ASIC. The Revision Compare feature requires that the HardCopy and FPGA netlists are the same, and all operations performed on one revision must also occur on the other revision. Unfortunately, exporting partitions does not support this requirement.

- For more information about the various types of incremental design flows and example applications, as well as documented restrictions and limitations, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.

Why to Plan Partitions and Floorplan Assignments

Incremental design flows typically require more up-front planning than flat compilations, and require you to be more rigorous about following good design practices. For example, you might have to structure your source code or design hierarchy to ensure that logic is grouped correctly for optimization. It is easier to implement the correct logic grouping early in the design cycle than to restructure the code later.

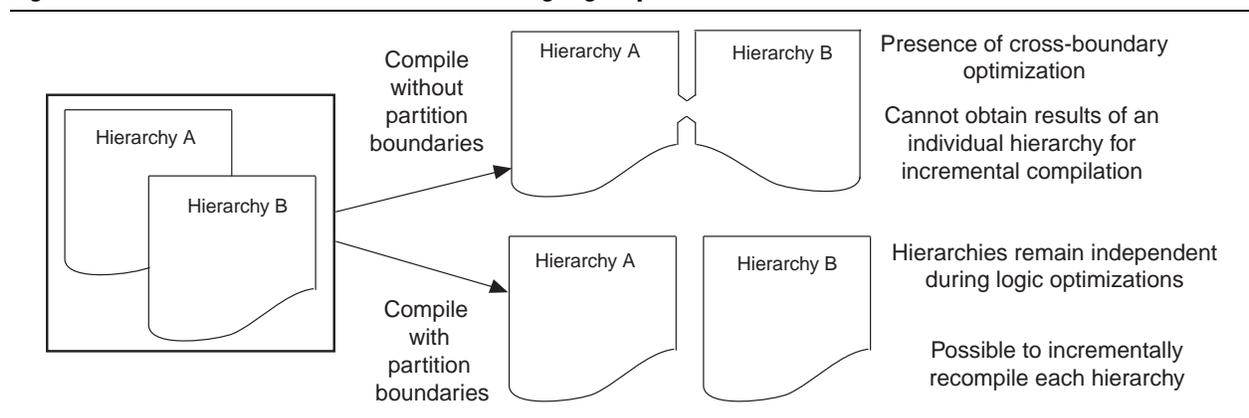
Planning involves setting up the design logic for partitioning, and may also involve planning placement assignments to create a floorplan. Not all design flows require floorplan assignments. For more information, refer to *“Introduction to Design Floorplans”* on page 12–41. If you decide to add floorplan assignments later, when the design is close to completion, well-planned partitions make floorplan creation much easier. Poor partition or floorplan assignments can worsen design area utilization and performance, making timing closure more difficult.

As FPGA devices get larger and more complex, following good design practices becomes more important for all design flows. Adhering to recommended synchronous design practices makes designs more robust and easier to debug. Using an incremental compilation flow adds additional steps and requirements to your project, but can provide significant benefits in design productivity by preserving the performance of critical blocks and reducing compilation times.

Partition Boundaries and Optimization

If there are cross-boundary optimizations between partitions, the Quartus II software cannot obtain separate results for each individual partition. The logical hierarchical boundaries between partitions are treated as hard boundaries for logic optimization to allow the software to size and place each partition independently. [Figure 12-1](#) shows the effects of partition boundaries during logic optimization. It is important to understand this effect so that you can effectively plan your design partitions.

Figure 12-1. Effects of Partition Boundaries During Logic Optimization



To prevent any cross-boundary optimizations, the Quartus II software synthesizes each partition without using any information about logic in other partitions. In a flat compilation, the software uses unconnected signals, constants, inversions, and other design information to perform optimizations. When you partition a design, these types of optimizations do not take place on partition I/O ports. Good design partitions do not rely on these types of logic optimizations.

You can use the **Merge** command in Design Partitions window to combine hierarchical partitions into a single partition, as long as they share the same immediate parent partition. Merging partitions allows additional optimizations for partition I/O ports that connect between or feed more than one of the merged hierarchical design blocks.

When all partitions are placed together, the Fitter can perform placement optimizations on the design as a whole to optimize the placement of cross-boundary paths. However, the Fitter can never perform any logic optimizations such as physical synthesis across the partition boundary. If partitions are fit separately in different projects, or if some partitions use previous post-fitting results, the Fitter does not place and route the entire cross-boundary path at the same time and cannot fully optimize placement across the partition boundaries. Good design partitions can be placed independently because cross-partition paths are not the critical timing paths in the design.

There are possible timing performance utilization effects because of partitioning and creating a floorplan. Not all designs encounter these issues, but you should consider these effects if a flat version of your design is very close to meeting its timing requirements, or is close to using all the device resources, before adding partition or floorplan assignments:

- Partitions can increase resource utilization due to cross-boundary optimization limitations if the design does not follow partitioning guidelines, for example when partitions include unconnected ports or ports driven by constants. For more information, refer to [“Design Partition Guidelines” on page 12–10](#). Floorplan assignments can also increase resource utilization because regions can lead to unused logic. If your device is full with the flat version of your design, you can focus on creating partitions and floorplan assignments for timing-critical or often-changing blocks to benefit most from the Quartus II features. For more information, refer to [“Checking Floorplan Quality” on page 12–50](#).
- Partitions and floorplan assignments might increase routing utilization compared to a flat design. If long compilation times are due to routing congestion, you might not be able to use the incremental flow to reduce compilation time. Review the Fitter messages to check how much time is spent during routing optimizations and to determine the percentage of routing utilization. When routing is difficult, you can use incremental compilation to lock the routing for routing-critical blocks only (with other partitions empty), and then compile the rest of the design after the critical block meets its requirements.
- Partitions can reduce timing performance in some cases because of the optimization and resource effects described above, causing longer logic delays. Floorplan assignments restrict logic placement, which can make it more difficult for the Fitter to meet timing requirements. Use the guidelines in this chapter to reduce any effect on your design performance.

Because cross-boundary logic and placement optimizations cannot occur, the quality of results may decrease as the number of partitions increases. Although more partitions allow for greater reduction in compilation time, consider limiting the number of partitions to prevent degradation in the quality of results.

Creating good design partitions and good floorplan location assignments helps to improve the design resource utilization and timing performance results for cross-partition paths. Guidelines for creating these assignments are discussed in the following sections.

General Partitioning Guidelines

The first stage in planning your design partitions is to organize your source code so that it supports good partition assignments. Although you can assign any hierarchical block of your design as a design partition or merge hierarchical blocks into the same partition, following the design guidelines presented in this section ensures better results.

Plan Design Hierarchy and Source Design Files

You begin the partitioning process by planning the design hierarchy. When you assign a hierarchical instance as a design partition, the partition includes the assigned instance and any entities instantiated below it that are not defined as separate partitions. You can also use the **Merge** command in the Design Partitions window to combine hierarchical partitions into a single partition, as long as they have the same immediate parent partition.

Take advantage of the design hierarchy to provide flexibility for partitioning and to support different design flows. Keep logic in the “leaves” of the hierarchy tree instead of having a lot of logic at the top-level of the design so that you can isolate partitions if required.

Create entities that can form partitions of approximately equal size. For example, do not instantiate a lot of small entities at the same hierarchy level because it is more difficult to group them to form reasonably-sized partitions.

Create each entity in an independent file. The software uses a file checksum to detect changes, and automatically recompiles a partition if its source file changes and its netlist type is set to either post-synthesis or post-fit. If the design entities for two partitions are defined in the same file, changes to the logic in one partition initiates recompilation for both partitions.

Design dependencies also affect which partitions are compiled when a source file changes. If two partitions rely on the same lower-level entity definition, changes in that lower-level entity affect both partitions. Commands such as VHDL `use` and Verilog HDL `include` create dependencies between files, so that changes to one file can trigger recompilations in all dependent files. Avoid these types of file dependencies if possible. The Partition Dependent Files report for each partition in the Analysis & Synthesis section of the Compilation report lists which files contribute to each partition.

-  For more information about what changes initiate an automatic recompilation of a partition, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.

Using Partitions with Third-Party Synthesis Tools

Incremental compilation works well with third-party synthesis tools in addition to Quartus II Integrated Synthesis. If you use a third-party synthesis tool, set up your tool to create a separate Verilog Quartus Mapping File (**.vqm**) or EDIF Input File (**.edf**) netlist for each hierarchical partition. In the Quartus II software, designate the top-level entity from each netlist as a design partition. The **.vqm** or **.edf** netlist file is treated as the source file for the partition in the Quartus II software.

-  For more information about incremental synthesis in third-party tools, refer to the documentation provided by your tool vendor, or the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.

Partition Design by Functionality and Block Size

Initially, you should partition your design along functional boundaries. In a top-level system block diagram, each block is often a natural design partition. Typically, each block of a system is relatively independent and has more signal interaction internally than interaction between blocks, which helps reduce optimizations between partition boundaries. Keeping functional blocks together means that synthesis and fitting can optimize related logic as a whole, which can lead to improved optimization.

Consider how many partitions you want to maintain in your design to determine the size of each partition. Your compilation time reduction goal is also a factor, because compiling small partitions is typically faster than compiling large partitions.

There is no minimum size for partitions; however, having too many partitions can reduce the quality of results by limiting optimization. Ensure that the design partitions are not too small. As a general guideline, each partition should contain more than approximately 2,000 logic elements (LEs) or adaptive logic modules (ALMs). If your design is incomplete when you partition the design, use previous designs to help you estimate the size that each block is likely to be.

Partition Design by Clock Domain and Timing Criticality

Consider which clock in your design feeds the logic in each partition. If possible, keep clock domains within one partition. When a clock signal is isolated to one partition, it reduces dependence on other partitions for timing optimization. Isolating a clock domain to one partition also allows better use of regional clock routing networks if the partition logic will be constrained to one region of the design. Additionally, limiting the number of clocks within each partition simplifies the timing requirements for each partition during optimization. Use an appropriate subsystem to implement the required logic for any clock domain transfers (such as a synchronization circuit, dual-port RAM, or FIFO). You can include this logic inside the partition at one side of the transfer.

Try to isolate timing-critical logic from logic that you expect to meet its timing requirements easily. Doing so allows you to preserve the satisfactory results for non-critical partitions and focus optimization iterations on just the timing-critical portions of the design to minimize compilation time.

Consider What Is Changing

When assigning partitions, you should consider what is changing in the design. Is there intellectual property (IP) or reused logic for which the source code will not change during future design iterations? If so, define the logic in its own partition so that you can compile one time and immediately preserve the results, so that you will not have to compile that part of the design again. Is logic being tuned or optimized, or are specifications changing for part of the design? If so, define changing logic in its own partition so that you can recompile only the changing part while the rest of the design remains unchanged.

As a general rule, create partitions to isolate logic that will change from logic that will not change. Partitioning a design in this way maximizes the preservation of unchanged logic and minimizes compilation time.

Design Partition Guidelines

Follow the partitioning guidelines presented in this section when you create or modify the HDL code for each design block that you might want to assign as a design partition. You do not have to follow all the recommendations exactly to achieve a good quality of results with the incremental compilation flow, but adhering to as many as possible maximizes your chances for success.

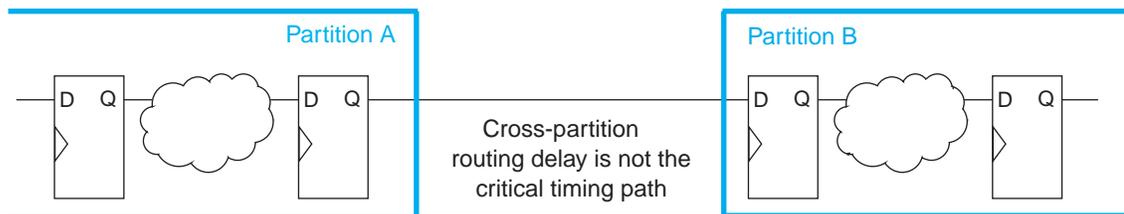
This section includes examples of the types of optimizations that are prevented by partition boundaries and describes how you can structure or modify your partitions to avoid these limitations.

Register Partition Inputs and Outputs

Use registers at partition input and output connections that are potentially timing-critical. Registers minimize the delays on inter-partition paths and prevent the need for cross-boundary logic optimizations.

If every partition boundary has a register as shown in [Figure 12-2](#), every register-to-register timing path between partitions includes only routing delay. Therefore, the timing paths between partitions are likely not timing-critical, and the Fitter can generally place each partition independently from other partitions. This advantage makes it easier to create floorplan location assignments for each separate partition, and is especially important for flows in which partitions are placed independently in separate Quartus II projects. Additionally, the partition boundary does not affect combinational logic optimization because each register-to-register logic path is contained within a single partition.

Figure 12-2. Registering Partition I/O



If a design cannot include both input and output registers for each partition due to latency or resource utilization concerns, choose to register one end of each connection. If you register every partition output, for example, the combinational logic that occurs in each cross-partition path is included in one partition so that it can be optimized together.

It is a good synchronous design practice to include registers for every output of a design block. Registered outputs ensure that the input timing performance for each design block is controlled exclusively within the destination logic block. For more information about I/O ports and registers for each partition, refer to [“Partition Statistics Report” on page 12-35](#), and [“Incremental Compilation Advisor” on page 12-50](#).

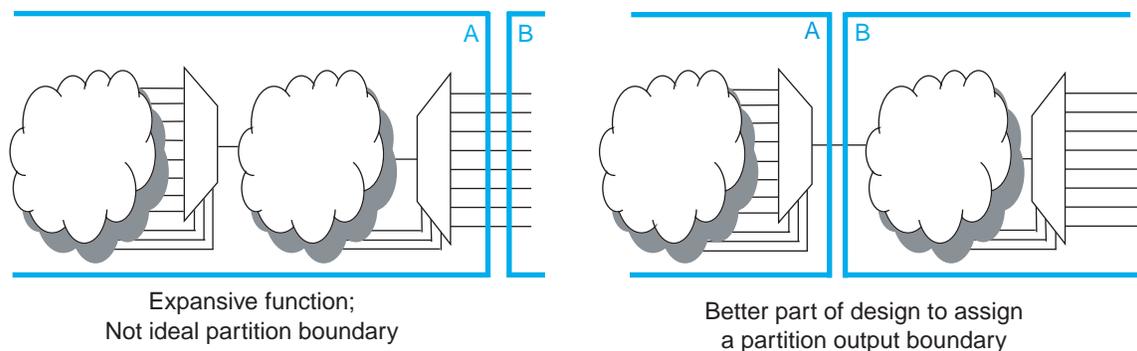
Minimize Cross-Partition-Boundary I/O

Minimize the number of I/O paths that cross between partition boundaries to keep logic paths within a single partition for optimization. Doing so makes partitions more independent for both logic and placement optimization.

This guideline is most important for timing-critical and high-speed connections between partitions, especially in cases where the input and output of each partition is not registered. Slow connections that are not timing-critical are acceptable because they should not impact the overall timing performance of the design. If there are timing-critical paths between partitions, rework or merge the partitions to avoid these inter-partition paths.

When dividing your design into partitions, consider the types of functions at the partition boundaries. [Figure 12-3](#) shows an expansive function with more outputs than inputs in the left diagram, which makes a poor partition boundary, and, on the right side, a better place to assign the partition boundary that minimizes cross-partition I/Os. Adding registers to one or both sides of the cross-partition path in this example would further improve the partition quality.

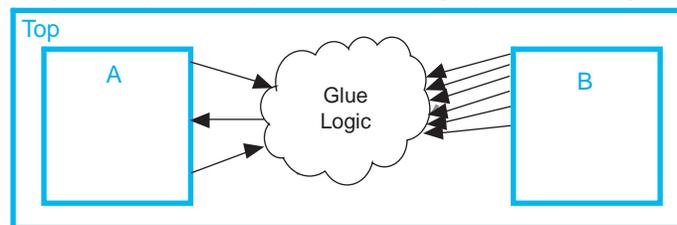
Figure 12-3. Minimizing I/O Between Partitions by Moving the Partition Boundary



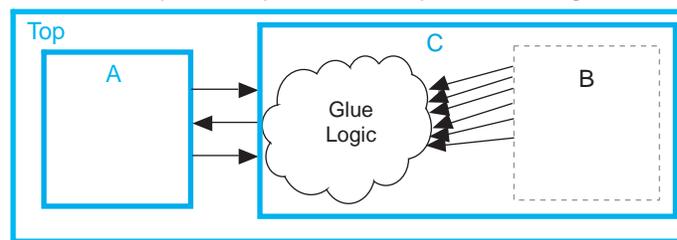
Another way to minimize connections between partitions is to avoid using combinational “glue logic” between partitions. You can often move the logic to the partition at one end of the connection to keep more logic paths within one partition. For example, the bottom diagram in [Figure 12-4](#) includes a new level of hierarchy C defined as a partition instead of block B. Clearly, there are fewer I/O connections between partitions A and C than between partitions A and B.

Figure 12-4. Minimizing I/O between Partitions by Modifying Glue Logic

Many cross-partition paths: Poor design partition assignment



Fewer cross-partition paths: Better partition assignment



For more information about the number of I/O ports, as well as the number of inter-partition connections for each partition, refer to [“Partition Statistics Report”](#) on [page 12-35](#). For more information about the number of intra-partition (within a partition) and inter-partition (between partitions) timing edges, refer to [“Incremental Compilation Advisor”](#) on [page 12-50](#).

Avoid the Need for Logic Optimization Across Partitions

As discussed in [“Partition Boundaries and Optimization”](#) on [page 12-6](#), partition boundaries prevent logic optimizations across partitions. Remember this rule: Logic cannot be optimized or merged across a partition boundary.

To ensure correct and optimal logic optimization, follow the guidelines in this section. In some cases, especially if part of the design is complete or comes from another designer, the designer might not have followed these guidelines when the source code was created. These guidelines are not mandatory to implement an incremental compilation flow, but can improve the quality of results. If assigning a partition affects resource utilization or timing performance of a design block as compared to the flat design, it might be due to one of the issues described in this section. Many of the examples suggest simple changes to your partition definitions or hierarchy to move the partition boundary and improve your results.

The following guidelines ensure that your design does not require logic optimization across partition boundaries:

- “Keep Logic in the Same Partition for Optimization and Merging” on page 12-13
- “Keep Constants in the Same Partition as Logic” on page 12-14
- “Avoid Unconnected Partition I/O” on page 12-15
- “Avoid Signals That Drive Multiple Partition I/O or Connect I/O Together” on page 12-16
- “Invert Clocks in Destination Partitions” on page 12-17
- “Connect I/O Pin Directly to I/O Register for Packing Across Partition Boundaries” on page 12-18
- “Do Not Use Internal Tri-States” on page 12-22
- “Include All Tri-State and Enable Logic in the Same Partition” on page 12-23
- “Include Bidirectional I/O Registers in the Same Partition (For Older Device Families)” on page 12-24

Keep Logic in the Same Partition for Optimization and Merging

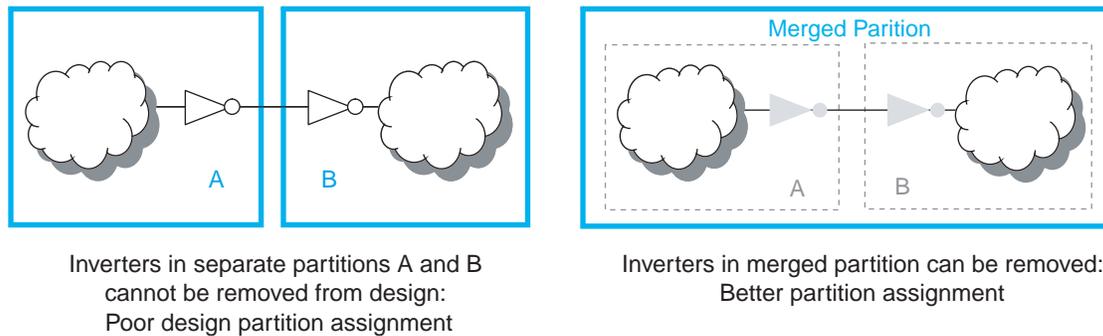
If any design logic requires logic optimization or merging to obtain optimal results, ensure that all the logic is part of the same partition.

If a combinational logic path is split across two partitions, the logic cannot be optimized or merged into one logic cell in the device. This effect can result in an extra logic cell in the path, increasing the logic delay. As a very simple example, consider two inverters on the same signal in two different partitions, A and B, as shown in the left diagram of [Figure 12-5](#). To maintain correct incremental functionality, these two inverters cannot be removed from the design during optimization because they occur in different design partitions. The Quartus II software cannot use information about other partitions when it compiles each partition, because each partition is allowed to change independently from the other.

On the right side of the figure, partitions A and B are merged to group the logic in blocks A and B into one partition. If the two blocks A and B are not under the same immediate parent partition, you can create a wrapper file to define a new level of hierarchy that contains both blocks, and set this new hierarchy block as the partition. With the logic contained in one partition, the software can optimize the logic and

remove the two inverters (shown in gray), which reduces the delay for that logic path. Removing two inverters is not a significant reduction in resource utilization because inversion logic is readily available in Altera device architecture. However, this example is a simple demonstration of the types of logic optimization that are prevented by partition boundaries.

Figure 12-5. Keeping Logic in the Same Partition for Optimization



In a flat design, the Fitter can also merge logical instantiations into the same physical device resource. With incremental compilation, logic defined in different partitions cannot be merged to use the same physical device resource.

For example, the Fitter can merge two single-port RAMs from a design into one dedicated RAM block in the device. If the two RAMs are defined in different partitions, the Fitter cannot merge them into one dedicated device RAM block.

This limitation is only a concern if merging is required to fit the design in the target device. Therefore, you are more likely to encounter this issue during troubleshooting rather than during planning, if your design uses more logic than is available in the device.

Merging PLLs and Transceivers (GXB)

Multiple instances of the ALTPLL megafunction can use the same PLL resource on the device. Similarly, GXB transceiver instances can share high-speed serial interface (HSSI) resources in the same quad as other instances. The Fitter can merge multiple instantiations of these blocks into the same device resource, even if it requires optimization across partitions. Therefore, there are no restrictions for PLLs and high-speed transceiver blocks when setting up partitions.

Keep Constants in the Same Partition as Logic

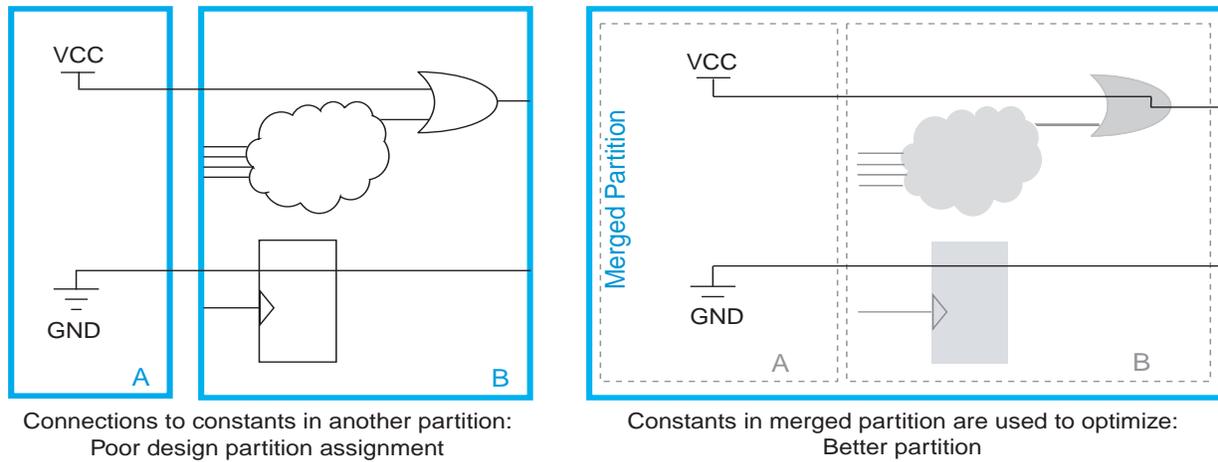
Because the Quartus II software cannot optimize across a partition boundary, constants are not propagated across partition boundaries. A signal that is constant ($1/V_{CC}$ or $0/GND$) in one partition cannot affect another partition.

For example, in the left diagram of [Figure 12-6](#) shows part of a design in which partition A defines some signals as constants (and assumes that the other input connections come from elsewhere in the design and are not shown in the figure). Constants such as these could appear due to parameter or generic settings or configurations with parameters, setting a bus to a specific set of values, or could result from optimizations that occur within a group of logic. Because the blocks are independent, the software cannot optimize the logic in block B based on the

information from block A. The right side of Figure 12-6 shows a merged partition that groups the logic in blocks A and B. If the two blocks A and B are not under the same immediate parent partition, you can create a wrapper file to define a new level of hierarchy that contains both blocks, and set this new hierarchical block as the partition.

Within the single merged partition, the Quartus II software can use the constants to optimize and remove much of the logic in block B (shown in gray), as shown in Figure 12-6 below.

Figure 12-6. Keeping Constants in the Same Partition as the Logic They Feed



For more information about how many input ports are fed by GND or V_{CC}, refer to “Partition Statistics Report” on page 12-35. For more information about port connections, refer to “Incremental Compilation Advisor” on page 12-50.

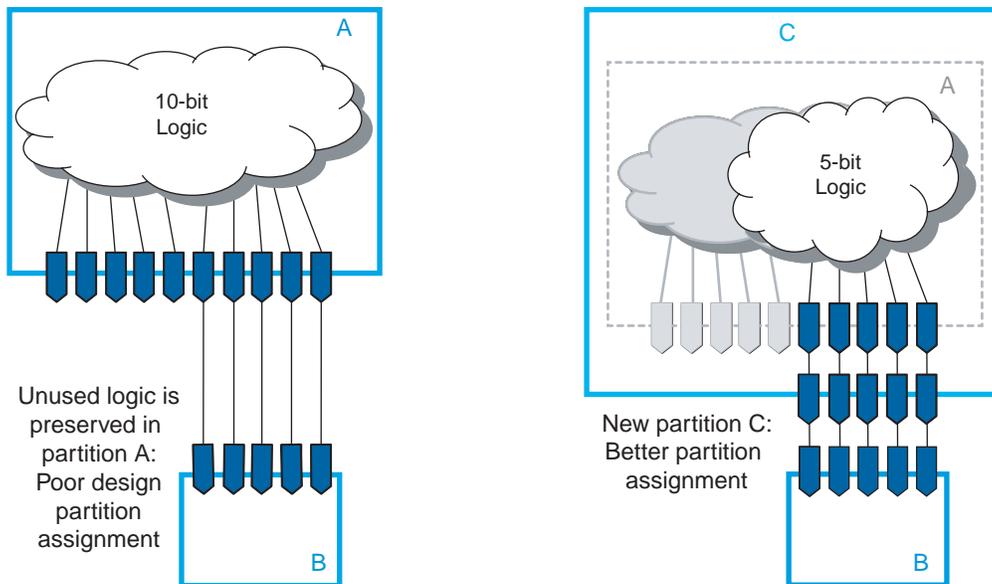
Avoid Unconnected Partition I/O

When a port is left unconnected, optimizations might be able to remove logic driving that port and improve results, similar to a constant connection. However, these optimizations are not allowed across partitions in incremental compilation, because they could create cross-partition dependence. For best results, connect ports to an appropriate node or remove them from the partition. If you know that a port will not be used, consider defining a wrapper module with a different port interface.

For example, in the left diagram of Figure 12-7 shows a design that has a 10-bit function defined in partition A, but has only 5 bits connected in partition B. In a flat design, you expect the logic for the other unused 5 bits to be removed during synthesis. With incremental compilation, synthesis does not remove the unused logic from partition A because partition B is allowed to change independently from partition A. Therefore, you can later connect all 10 bits in partition B and use all 10 bits from partition A. In this design example, if you know that you will not use the other 5 bits of partition A, you should remove the unconnected ports and replace them with ground signals inside partition A. You can create a new wrapper file in the design hierarchy to do this, as shown on the right side of the figure. A new partition C contains the logic from block A, but includes only the 5 output ports required for connection with partition B. Within this new partition C, the logic for the unused 5 bits can be removed from the design, reducing area utilization.

Figure 12-7 shows a design with a 10-bit function defined in partition A, but has only 5 bits connected in partition B.

Figure 12-7. Avoiding Unconnected Partition I/O by Creating a Wrapper File



For more information about how many I/Os are unconnected, refer to “[Partition Statistics Report](#)” on page 12-35. For more information about unconnected ports, refer to “[Incremental Compilation Advisor](#)” on page 12-50.

Avoid Signals That Drive Multiple Partition I/O or Connect I/O Together

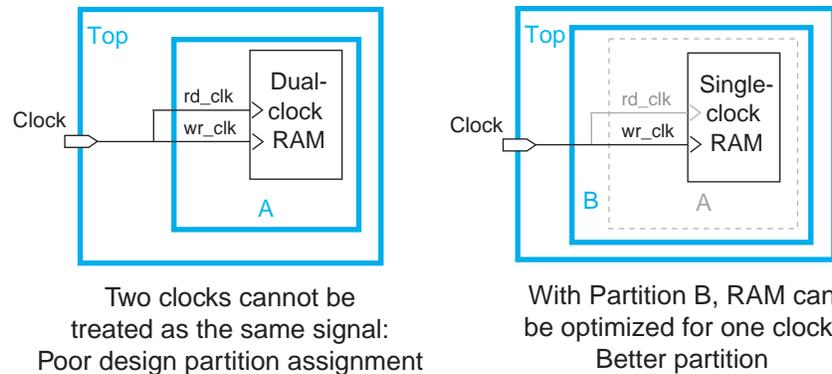
Do not use the same signal to drive multiple ports of a single partition or directly connect two ports of a partition. If the same signal drives multiple ports of a partition, or if two ports of a partition are directly connected, those ports are logically equivalent. However, because the software has no information about connections made in another partition (including the top-level partition), the compilation cannot take advantage of the equivalence. This restriction usually results in sub-optimal results.

If your design has these types of connections, redefine the partition boundaries to remove the affected ports. If one signal from a higher-level partition feeds two input ports of the same partition, feed the one signal into the partition, and then make the two connections within the partition. If an output port drives an input port of the same partition, the connection can be made internally without going through any I/O ports. If an input port drives an output port directly, the connection can likely be implemented without the ports in the lower-level partition by connecting the signals in a higher-level partition.

Figure 12-8 shows an example of one signal driving more than one port. The left diagram shows a design where a single clock signal is used to drive both the read and write clocks of a RAM block. Because the RAM block is compiled as a separate partition A, the RAM block is implemented as though there are two unique clocks. If you know that the port connectivity will not change (that is, the ports will always be driven by the same signal in the top-level partition), redefine the port interface so that

there is only a single port that can drive both connections inside the partition. You can create a wrapper file to define a partition that has fewer ports, as shown in the diagram on the right side. With the single clock fed into the partition, the RAM can be optimized into a single-clock RAM instead of a dual-clock RAM. Single-clock RAM can provide better performance in the device architecture. Additionally, partition A might use two global routing lines for the two copies of the clock signal. Partition B can use one global line that fans out to all destinations. Using just the single port connection prevents overuse of global routing resources.

Figure 12-8. Preventing One Signal from Driving Multiple Partition Inputs



For more information about partition ports that have the same driving signal and ports that are directly connected together, refer to [“Incremental Compilation Advisor” on page 12-50](#).

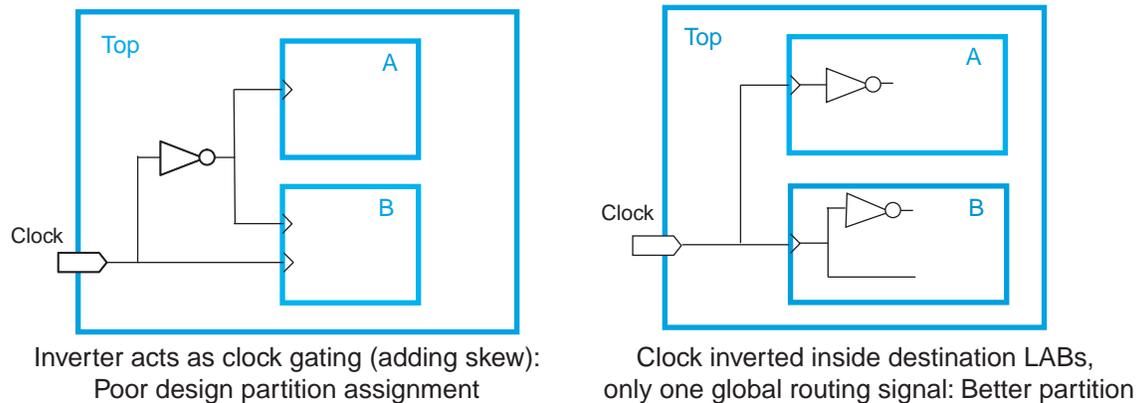
Invert Clocks in Destination Partitions

For best results, clock inversion should be done in the destination logic array block (LAB) because each LAB contains clock inversion circuitry in the device architecture. In a flat compilation, the Quartus II software can optimize a clock inversion to propagate it to the destination LABs regardless of where the inversion takes place in the design hierarchy. However, clock inversion cannot propagate through a partition boundary to take advantage of the inversion architecture in the destination LABs.

With partition boundaries as shown in the left diagram of [Figure 12-9](#), the Quartus II software uses logic to invert the signal in the partition that defines the inversion (the top-level partition in this example), and then routes the signal on a global clock resource to its destinations (in partitions A and B). The inverted clock acts as a gated clock with high skew. A better solution is to invert the clock signal in the destination partitions as shown on the right side of the figure. In this case, the correct logic and routing resources can be used, and the signal does not behave like a gated clock.

Figure 12-9 shows the clock signal inversion in the destination partitions.

Figure 12-9. Inverting Clock Signal in Destination Partitions



Notice that this diagram also shows another example of a single pin feeding two ports of a partition boundary. In the left diagram, partition B does not have the information that the clock and inverted clock come from the same source. In the right diagram, partition B has more information to help optimize the design because the clock is connected as one port of the partition.

Connect I/O Pin Directly to I/O Register for Packing Across Partition Boundaries

The Quartus II software allows cross-partition register packing of I/O registers in certain cases where your input and output pins are defined in the top-level hierarchy (and the top-level partition), but the corresponding I/O registers are defined in other partitions.

Input pin cross-partition register packing requires the following specific circumstances:

- The input pin feeds exactly one register.
- The path between the input pin and register includes only input ports of partitions that have one fan-out each.

Output pin cross-partition register packing requires the following specific circumstances:

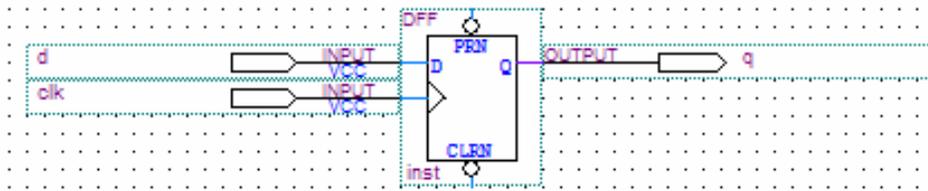
- The register feeds exactly one output pin.
- The output pin is fed by only one signal.
- The path between the register and output pin includes only output ports of partitions that have one fan-out each.

The following examples of I/O register packing illustrate this point using Block Design File (.bdf) schematics to describe the design logic.

Example 1—Output Register in Partition Feeding Multiple Output Pins

In this example, a subdesign contains a single register, as shown in Figure 12-10.

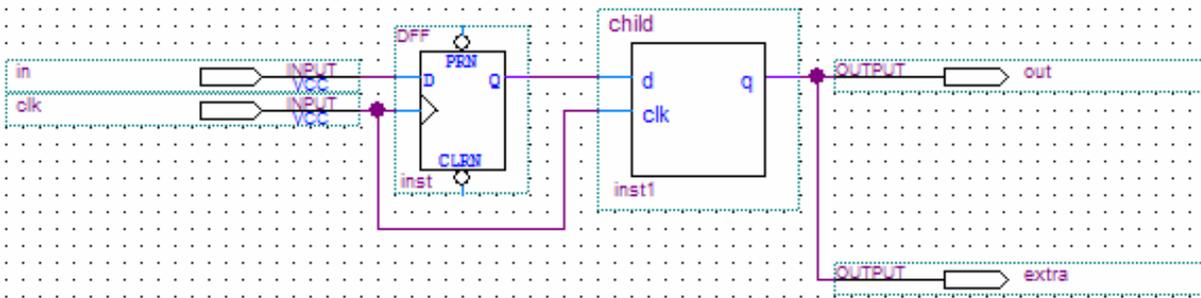
Figure 12-10. Subdesign with One Register, Designated as a Separate Partition



If the top-level design instantiates the subdesign with a single fan-out directly feeding an output pin, and designates the subdesign as a separate design partition, the Quartus II software can perform cross-partition register packing because the single partition port feeds the output pin directly.

In Example 1, the top-level design instantiates the subdesign in Figure 12-10 as an output register with more than one fan-out signal, as shown in Figure 12-11.

Figure 12-11. Top-Level Design Instantiating the Subdesign in Figure 12-10 with Two Output Pins



In this case, the Quartus II software does not perform output register packing. If there is a **Fast Output Register** assignment on pin out, the software issues a warning that the Fitter cannot pack the node to an I/O pin because the node and the I/O cell are connected across a design partition boundary.

This type of cross-partition register packing is not allowed because it requires modification to the interface of the subdesign partition. To perform incremental compilation, the Quartus II software must preserve the interface of design partitions.

To allow the Quartus II software to pack the register in the subdesign from Figure 12-10 with the output pin out in Figure 12-11, restructure your HDL code so that output registers directly connect to output pins by making one of the following changes:

- Place the register in the same partition as the output pin. The simplest method is to move the register from the subdesign partition into the partition containing the output pin. Doing so guarantees that the Fitter can optimize the

two nodes without violating partition boundaries.

- Duplicate the register in your subdesign HDL as in [Figure 12-12](#) so that each register feeds only one pin, then connect the extra output pin to the new port in the top-level design as shown in [Figure 12-13](#). Doing so converts the cross-partition register packing into the simplest case where each register has a single fan-out.

Figure 12-12. Modified Subdesign from [Figure 12-10](#) with Two Output Registers and Two Output Ports

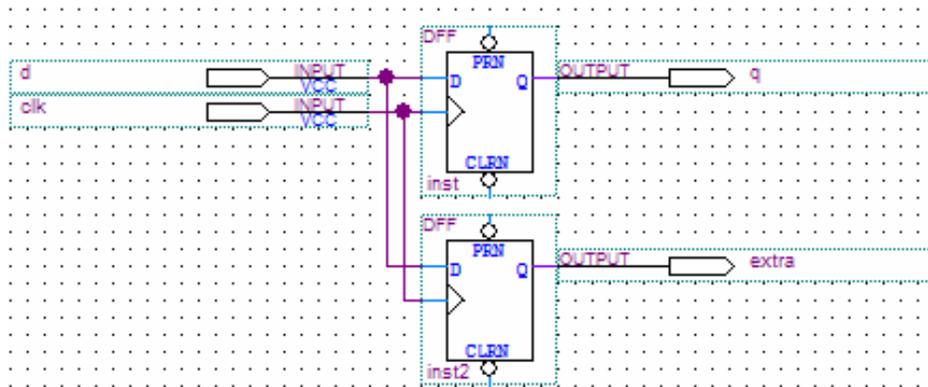
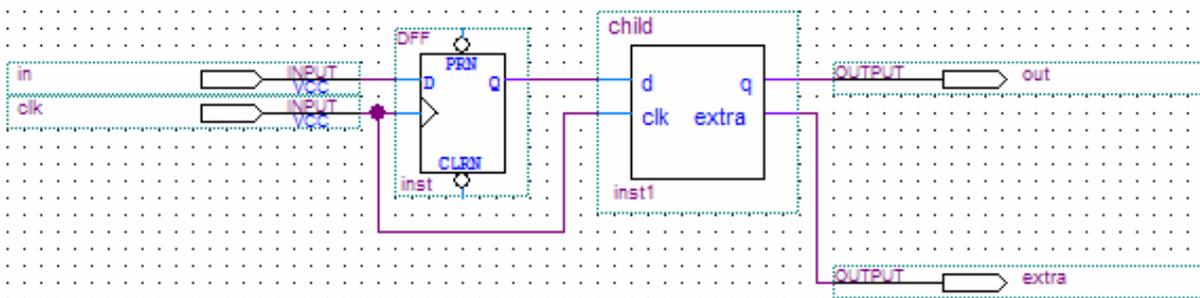


Figure 12-13. Modified Top-Level Design from [Figure 12-11](#) Connecting Two Output Ports to Output Pins



Example 2—Input Register in Partition Fed by an Inverted Input Pin or Output Register in Partition Feeding an Inverted Output Pin

In this example, a subdesign designated as a separate partition contains a register, as shown in Figure 12-10. The top-level design in Figure 12-14 instantiates the subdesign as an input register with the input pin inverted. The top-level design in Figure 12-15 instantiates the subdesign as an output register with the signal inverted before feeding an output pin.

Figure 12-14. Top-Level Design Instantiating the Subdesign in Figure 12-10 as an Input Register with an Inverted Input Pin

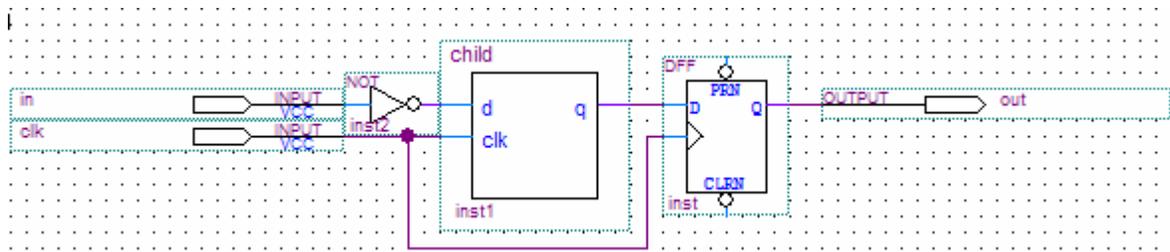
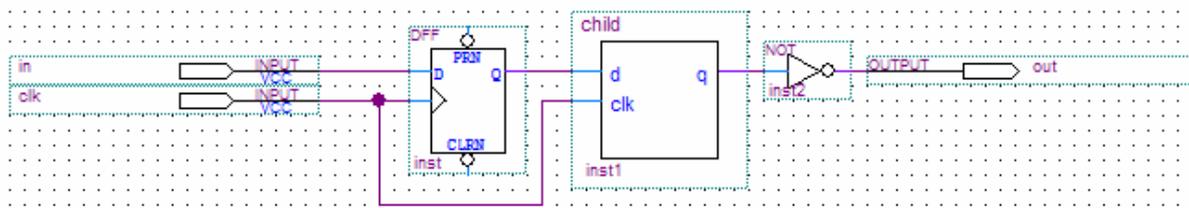


Figure 12-15. Top-Level Design Instantiating the Subdesign in Figure 12-10 as an Output Register Feeding an Inverted Output Pin



In these cases, the Quartus II software does not perform register packing. If there is a **Fast Input Register** assignment on pin *in*, as in Figure 12-14, or a **Fast Output Register** assignment on pin *out*, as in Figure 12-15, the Quartus II software issues a warning that the Fitter cannot pack the node to an I/O pin because the node and I/O cell are connected across a design partition boundary.

This type of register packing is not allowed because it requires moving logic across a design partition boundary to place into a single I/O device atom. To perform register packing, either the register must be moved out of the subdesign partition, or the inverter must be moved into the subdesign partition to be implemented in the register.

To allow the Quartus II software to pack the register in the subdesign from Figure 12-10 with the input pin *in*, as in Figure 12-14 or the output pin *out*, as in Figure 12-15, restructure your HDL code to place the register in the same partition as the inverter by making one of the following changes:

- Move the register from the subdesign partition into the top-level partition containing the pin. Doing so ensures that the Fitter can optimize the I/O register and inverter without violating partition boundaries.

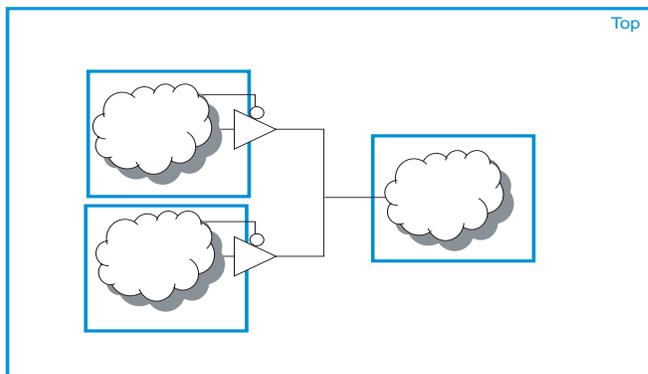
- Move the inverter from the top-level block into the subdesign, then connect the subdesign directly to a pin in the top-level design. Doing so allows the Fitter to optimize the inverter into the register implementation, so that the register is directly connected to a pin, which enables register packing.

Do Not Use Internal Tri-States

Internal tri-state signals are not recommended for FPGAs because the device architecture does not include internal tri-state logic. If designs use internal tri-states in a flat design, the tri-state logic is usually converted to OR gates or multiplexing logic. If tri-state logic occurs on a hierarchical partition boundary, the Quartus II software cannot convert the logic to combinational gates because the partition could be connected to a top-level device I/O through another partition.

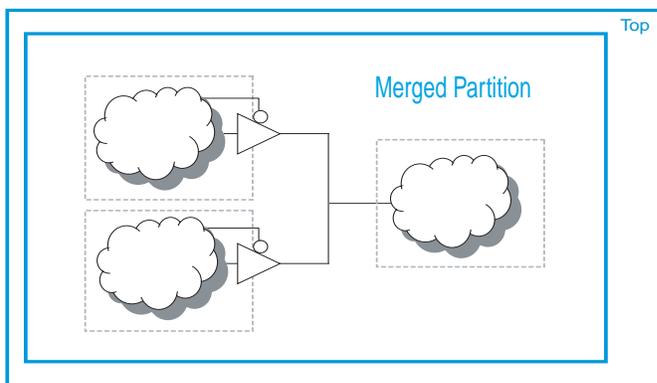
Figure 12-16 and Figure 12-17 show a design with partitions that are not supported for incremental compilation due to the internal tri-state output logic on the partition boundaries. Instead of using internal tri-state logic for partition outputs, implement the correct logic to select between the two signals. Doing so is good practice even when there are no partitions, because such logic explicitly defines the behavior for the internal signals instead of relying on the Quartus II software to convert the tri-state signals into logic.

Figure 12-16. Unsupported Internal Tri-State Signals



Design results in Quartus II error message:
The software cannot synthesize this design and maintain incremental functionality

Figure 12-17. Merged Partition Allows Synthesis to Convert Internal Tri-State Logic to Combinational Logic



Merged partition allows synthesis to convert tri-state logic into combinational logic.

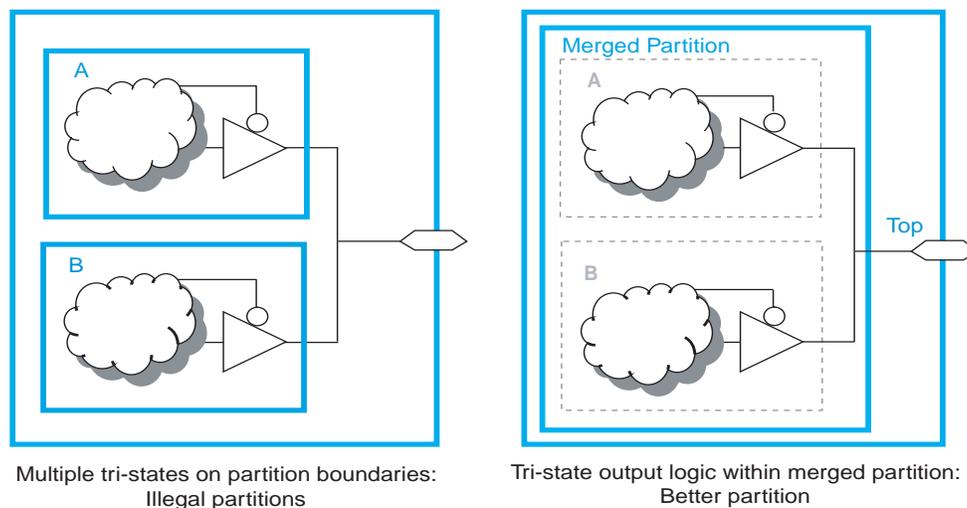
Do not use tri-state signals or bidirectional ports on hierarchical partition boundaries, unless the port is connected directly to a top-level I/O pin on the device. If you must use internal tri-state logic, ensure that all the control and destination logic is contained in the same partition, in which case the Quartus II software can convert the internal tri-state signals into combinational logic like in a flat design. In this example, you can also merge all three partitions into one partition, as shown in Figure 12-17, to allow the Quartus II software to treat the logic as internal tri-state and perform the same type of optimization as a flat design. If possible, you should avoid using internal tri-state logic in any Altera FPGA design to ensure that you get the desired implementation when the design is compiled for the target device architecture.

Include All Tri-State and Enable Logic in the Same Partition

When multiple output signals use tri-state logic to drive a device output pin, the Quartus II software merges the logic into one tri-state output pin. The Quartus II software cannot merge tri-state outputs into one output pin if any of the tri-state logic occurs on a partition boundary. Similarly, output pins with an output enable signal cannot be packed into the device I/O cell if the output enable logic is part of a different partition from the output register. To allow register packing for output pins with an output enable signal, structure your HDL code or design partition assignments so that the register and enable logic are defined in the same partition.

Figure 12-18 shows a design with tri-state output signals that feed a device bidirectional I/O pin (assuming that the input connection feeds elsewhere in the design and is not shown in the figure). In the left diagram of the figure, the tri-state output signals appear as the outputs of two separate partitions. In this case, the Quartus II software cannot implement the specified logic and maintain incremental functionality. On the right side, partitions A and B are merged to group the logic from the two blocks. With this single partition, the Quartus II software can merge the two tri-state output signals and implement them in the tri-state logic available in the device I/O element.

Figure 12-18. Including All Tri-State Output Logic in the Same Partition

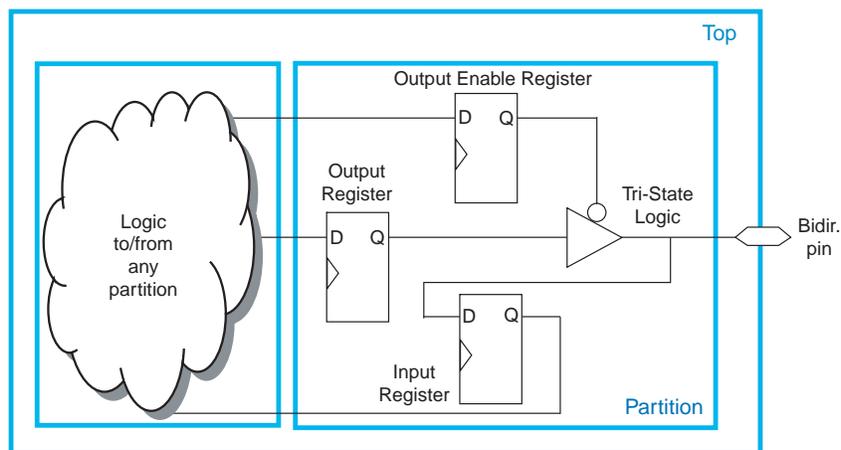


Include Bidirectional I/O Registers in the Same Partition (For Older Device Families)

For a bidirectional partition port that feeds a bidirectional I/O pin at the top level, all logic that forms the bidirectional I/O cell must reside in the same partition in the Stratix II, Stratix, Cyclone® II, and Cyclone device families (this restriction does not apply to newer devices). Additionally, as discussed in the previous two recommendations, the I/O logic must feed the I/O pin without any intervening logic.

In Figure 12-19, all the I/O logic must be defined inside the same partition for the Quartus II software to implement all three registers in the I/O element along with the tri-state logic in the affected devices. The logic connected to the registers can occur in the same partition or any other partition; only the I/O registers must be grouped with the tri-state logic definition. The bidirectional I/O port of the partition must be directly connected to the bidirectional device pin at the top level. The signal can go through several partition boundaries if necessary, as long as the connection path contains no logic.

Figure 12-19. Including All Bidirectional I/O Registers in the Same Partition (for Older Devices)



Bidirectional logic is within one partition, and I/O logic directly feeds I/O pin

Summary of Guidelines Related to Logic Optimization Across Partitions

To ensure that your design does not require logic optimization across partitions, follow the guidelines in this section:

- Include logic in the same partition for optimization and merging
- Include constants in the same partition as logic
- Avoid unconnected partition I/O
- Avoid signals that drive multiple partition I/O or connect I/O together
- Invert clocks in destination partitions
- Connect I/O directly to I/O register for packing across partition boundaries
- Do not use internal tri-states
- Include all tri-state and enable logic in the same partition
- Include bidirectional I/O registers in the same partition (in older device families)

Remember that these guidelines are not mandatory when implementing an incremental compilation flow, but can improve the quality of results. When creating source design code, follow these guidelines and organize your HDL code to support good partition boundaries. For designs that are complete, assess whether assigning a partition affects the resource utilization or timing performance of a design block as compared to the flat design. Make the appropriate changes to your design or hierarchy, or merge partitions as required, to improve your results.

Consider a Cascaded Reset Structure

Designs typically have a global asynchronous reset signal where a top-level signal feeds all partitions. To minimize skew for the high fan-out signal, the global reset signal is typically placed onto a global routing resource.

In some cases, having one global reset signal can lead to recovery and removal time problems. This issue is not specific to incremental flows; it could be applicable in any large high-speed design. In an incremental flow, the global reset signal creates a timing dependency between the top-level partition and lower-level partitions.

For incremental compilation, it is helpful to minimize the impact of global structures. To isolate each partition, consider adding reset synchronizers. Using cascaded reset structures, the intent is to reduce the inter-partition fan-out of the reset signal, thereby minimizing the effect of the global signal. Reducing the fan-out of the global reset signal also provides more flexibility in routing the cascaded signals, and might help recovery and removal times in some cases.

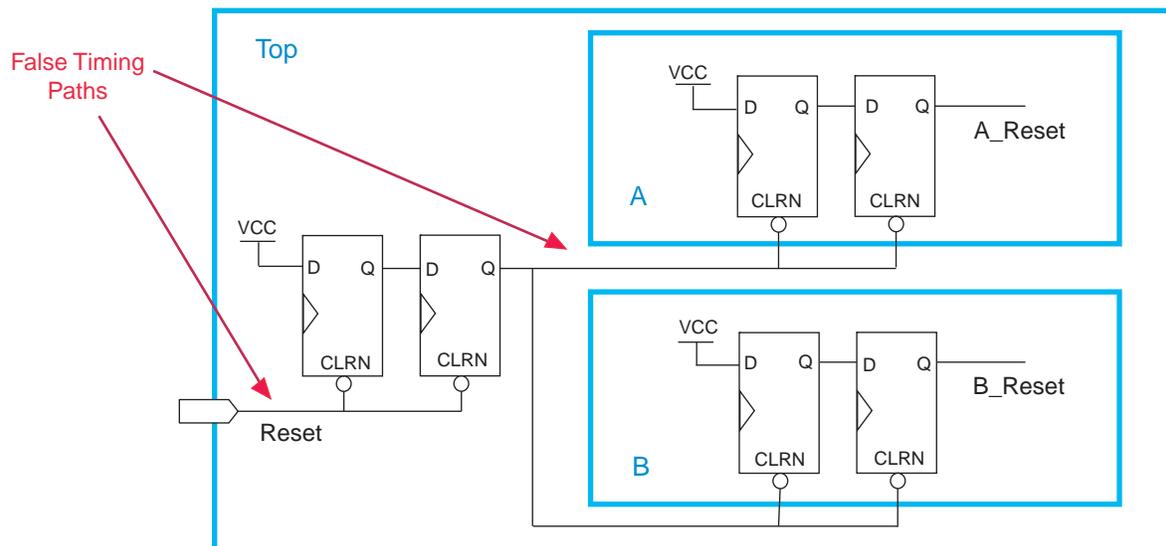
This recommendation can help in large designs, regardless of whether you are using incremental compilation. However, if one global signal can feed all the logic in its domain and meet recovery and removal times, then this recommendation may not be applicable for your design. Minimizing global structures is more relevant for high-performance designs where meeting timing on the reset logic can be challenging. Isolating each partition and allowing more flexibility in global routing structures is an additional advantage in incremental flows.

If you add additional reset synchronizers to your design, latency is also added to the reset path, so ensure that this is acceptable in your design. Additionally, parts of the design may come out of the reset state in different clock cycles. You can balance the latency or add hand-shaking logic between partitions, if necessary, to accommodate these differences.

The signal is first synchronized on the chip following good synchronous design practices, meaning that the design asynchronously resets, but synchronously releases from reset to avoid any race conditions or metastability problems. Then, to minimize the impact of global structures, the circuit employs a divide-and-conquer approach for the reset structure. By implementing a cascaded reset structure, the reset paths for each partition are independent. This structure reduces the effect of inter-partition dependency because the inter-partition reset signals can now be treated as false paths for timing analysis. In some cases, the reset signal of the partition can be placed on local lines to reduce the delay added by routing to a global routing line. In other cases, the signal can be routed on a regional or quadrant clock signal.

Figure 12-20 shows a cascaded reset structure.

Figure 12-20. Cascaded Reset Structure



This circuit design can help you achieve timing closure and partition independence for your global reset signal. Evaluate the circuit and consider how it works for your design.

Design Partition Guidelines for Third-Party IP Delivery

This section includes additional design guidelines that can improve incremental compilation flows where exported partitions are developed independently. These guidelines are not always required, but are usually recommended if the design includes partitions compiled in a separate Quartus II project, such as when delivering intellectual property (IP). A unique challenge of IP delivery for FPGAs is the fact that the partitions developed independently must share a common set of resources. To minimize issues that might arise from sharing a common set of resources, you can design partitions within a single Quartus II project, or a copy of the top-level design. A common project ensures that designers have a consistent view of the top-level design framework, as described in “[Project Management in Team-Based Design Flows](#)” on page 12-4.

Alternatively, an IP designer can export just the post-synthesis results to be integrated in the top-level design when the post-fitting results from the IP project are not required. Using a post-synthesis netlist provides more flexibility to the Quartus II Fitter so that less resource allocation is required. If a common project is not possible, especially when the project lead plans to integrate the IP's post-fitting results, it is important that the project lead and IP designer clearly communicate their requirements.

Allocate Logic Resources

In an incremental compilation design flow in which designers, such as third-party IP providers, optimize partitions and then export them to a top-level design, the Quartus II software places and routes each partition separately. In some cases, partitions can use conflicting resources when combined at the top level. Allocation of logic resources requires that you to decide on a set of logic resources (including I/O, LAB logic blocks, RAM and DSP blocks) that the IP block will “own”. This process can be interactive; the project lead and the IP designer might work together to determine what resources are required for the IP block and are available in the top-level design.

You can constrain logic utilization for the IP core using design floorplan location assignments, as described in [“Introduction to Design Floorplans” on page 12–41](#). The design should specify I/O pin locations with pin assignments.

You can also specify limits for Quartus II synthesis to allocate and balance resources. This procedure can also help if device resources are overused in the individual partitions during synthesis.

In the standard synthesis flow, the Quartus II software can perform automated resource balancing for DSP blocks or RAM blocks and convert some of the logic into regular logic cells to prevent overuse.

You can use the Quartus II synthesis options to control inference of megafunctions that use the DSP or RAM blocks. You can also use the MegaWizard™ Plug-In Manager to customize your RAM or DSP megafunctions to use regular logic instead of the dedicated hardware blocks.

You can, in certain device families, assign a number of LAB, DSP or RAM blocks for each partition. Use the following logic options to specify the maximum number of logic blocks that the Quartus II software can use in the specified partition: **Maximum Number of LABs**, **Maximum DSP Block Usage**, **Maximum Number of M4K/M9K Memory Blocks**, or **Maximum Number of M-RAM/M144K Memory Blocks**. You can set these options globally for all partitions in the **More Analysis & Synthesis Settings** dialog box.

- ❓ For information on how to set global logic options for partitions, refer to [More Analysis & Synthesis Settings Dialog Box](#) in Quartus II Help.

You can also set an option for a specific partition with the Assignment Editor by selecting the assignment name, applying it to the root entity of a partition, and then setting an integer as the value. Partition-specific assignments override global assignments, if any. However, each partition that does not have a partition-specific assignment can use the number of LAB, DSP, or RAM blocks set by the global assignment. You should note that this behavior can lead to over-allocation of logic blocks, eventually resulting in a no-fit error.

-  For more information about resource balancing DSP and RAM blocks when using Quartus II synthesis, refer to the “Megafunction Inference Control” section in the [Quartus II Integrated Synthesis](#) chapter in volume 1 of the *Quartus II Handbook*. For tips about resource balancing and reducing resource utilization, refer to the appropriate “Resource Utilization Optimization Techniques” section in the [Area and Timing Optimization](#) chapter in volume 2 of the *Quartus II Handbook*.

Allocate Global Routing Signals and Clock Networks if Required

In most cases, you do not have to allocate global routing signals because the Quartus II software finds the best solution for the global signals. However, if your design is complex and has multiple clocks, especially for a partition developed by a third-party IP designer, you may have to allocate global routing resources between various partitions.

Global routing signals can cause conflicts when independent partitions are integrated into a top-level design. The Quartus II software automatically promotes high fan-out signals to use global routing resources available in the device. Third-party partitions can use the same global routing resources, thus causing conflicts in the top-level design. Additionally, LAB placement depends on whether the inputs to the logic cells within the LAB use a global clock signal. Problems can occur if a design does not use a global signal in a lower-level partition, but does use a global signal in the top-level design.

If the exported IP core is small, you can reduce the potential for problems by using constraints to promote clock and high fan-out signals to regional routing signals that cover only part of the device, instead of global routing signals. In this case, the Quartus II software is likely to find a routing solution in the top-level design because there are many regional routing signals available on most Altera devices, and designs do not typically overuse regional resources.

To ensure that an IP block can utilize a regional clock signal, you can view the resource coverage of regional clocks in the Chip Planner, and then align LogicLock regions that constrain partition placement with available global clock routing resources. For example, if the LogicLock region for a particular partition is limited to one device quadrant, that partition's clock can use a regional clock routing type that covers only one device quadrant. When all partition logic is available, the project lead can compile the entire design at the top level with floorplan assignments to allow the use of regional clocks that span only a part of the device.

If global resources are heavily used in the overall design, or the IP designer requires global clocks for their partition, you can set up constraints to avoid signal overuse at the top-level by assigning the appropriate type of global signals or setting a maximum number of clock signals for the partition.

You can use the **Global Signal** assignment to force or prevent the use of a global routing line, making the assignment to a clock source node or signal. You can also assign certain types of global clock resources in some device families, such as regional clocks. For example, if you have an IP core, such as a memory interface that specifies the use of a dual regional clock, you can constrain the IP to part of the device covered by a regional clock and change the **Global Signal** assignment to use a regional clock. This type of assignment can reduce clocking congestion and conflicts.

Alternatively, partition designers can specify the number of clocks allowed in the project using the maximum clocks allowed options in the **More Fitter Settings** dialog box. Specify **Maximum number of clocks of any type allowed**, or use the **Maximum number of global clocks allowed**, **Maximum number of regional clocks allowed**, and **Maximum number of periphery clocks allowed** options to restrict the number of clock resources of a particular type in your design.

If you require more control when planning a design with integrated partitions, you can assign a specific signal to use a particular clock network in Stratix II and newer device families by assigning the clock control block instance called CLKCTRL. You can make a point-to-point assignment from a clock source node to a destination node, or a single-point assignment to a clock source node with the **Global Clock CLKCTRL Location** logic option. Set the assignment value to the name of the clock control block: CLKCTRL_G<global network number> for a global routing network, or CLKCTRL_R<regional network number> for a dedicated regional routing network in the device.

If you want to disable the automatic global promotion performed in the Fitter to prevent other signals from being placed on global (or regional) routing networks, turn off the **Auto Global Clock** and **Auto Global Register Control Signals** options in the **More Fitter Settings** dialog box.

- ① For information on how to disable automatic global promotion, refer to *More Fitter Settings Dialog Box* in Quartus II Help.

If you are using design partition scripts for independent partitions, the Quartus II software can automatically write the commands to pass global constraints and turn off automatic options.

- ① For more information on how to generate design partition scripts, refer to *Generating Design Partition Scripts for Project Management* in Quartus II Help.

Alternatively, to avoid problems when integrating partitions into the top-level design, you can direct the Fitter to discard the placement and routing of the partition netlist by using the post-synthesis netlist, which forces the Fitter to reassign all the global signals for the partition when compiling the top-level design.

Assign Virtual Pins

Virtual pins map lower-level design I/Os to internal cells. If you are developing an IP block in an independent Quartus II project, use virtual pins when the number of I/Os on a partition exceeds the device I/O count, and to increase the timing accuracy of cross-partition paths.

You can create a virtual pin assignment in the Assignment Editor for partition I/Os that will become internal nodes in the top-level design. Leave the clock pins mapped to I/O pins to ensure proper routing.

You can specify locations for the virtual pins that correspond to the placement of other partitions, and also make timing assignments to the virtual pins to define a timing budget, as described in the following section. Virtual pins are created automatically from the top-level design if you use design partition scripts. The scripts place the virtual pins to correspond with the placement of the other partitions in the top-level design.

- ① For more information on how to generate design partition scripts, refer to *Generating Design Partition Scripts for Project Management* in Quartus II Help.



Tri-state outputs cannot be assigned as virtual pins because internal tri-state signals are not supported in Altera devices. Connect the signal in the design with regular logic, or allow the software to implement the signal as an external device I/O pin.

Perform Timing Budgeting if Required

If you optimize partitions independently and integrate them to the top-level design, or compile with empty partitions, any unregistered paths that cross between partitions are not optimized as entire paths. In these cases, the Quartus II software has no information about the placement of the logic that connects to the I/O ports. If the logic in one partition is placed far away from logic in another partition, the routing delay between the logic can lead to problems in meeting timing requirements. You can reduce this effect by ensuring that input and output ports of the partitions are registered whenever possible. Additionally, using the same top-level project framework helps to avoid this problem by providing the software with full information about other design partitions in the top-level design.

To ensure that the software correctly optimizes the input and output logic in any independent partitions, you might be required to perform some manual timing budgeting. For each unregistered timing path that crosses between partitions, make timing assignments on the corresponding I/O path in each partition to constrain both ends of the path to the budgeted timing delay. Assigning a timing budget for each part of the connection ensures that the software optimizes the paths appropriately.

When performing manual timing budgeting in a partition for I/O ports that become internal partition connections in a top-level design, you can assign location and timing constraints to the virtual pin that represents each connection to further improve the quality of the timing budget. Refer to [“Assign Virtual Pins” on page 12-29](#) for a description of virtual pins.



If you are using design partition scripts, the Quartus II software can write I/O timing budget constraints automatically for virtual pins.



For more on how to generate design partition scripts, refer to [Generating Design Partition Scripts for Project Management](#) in Quartus II Help.

Drive Clocks Directly

When partitions are exported from another Quartus II project, you should drive partition clock inputs directly with device clock input pins.

Connecting the clock signal directly avoids any timing analysis difficulties with gated clocks. Clock gating is never recommended for FPGA designs because of potential glitches and clock skew. Clock gating can be especially problematic with exported partitions because the partitions have no information about gating that takes place at the top-level design or in another partition. If a gated clock is required in a partition, perform the gating within that partition, as described for clock inversion in [“Invert Clocks in Destination Partitions” on page 12-17](#).

Direct connections to input clock pins also allows design partition scripts to send constraints from the top-level device pin to lower-level partitions.

Recreate PLLs for Lower-Level Partitions if Required

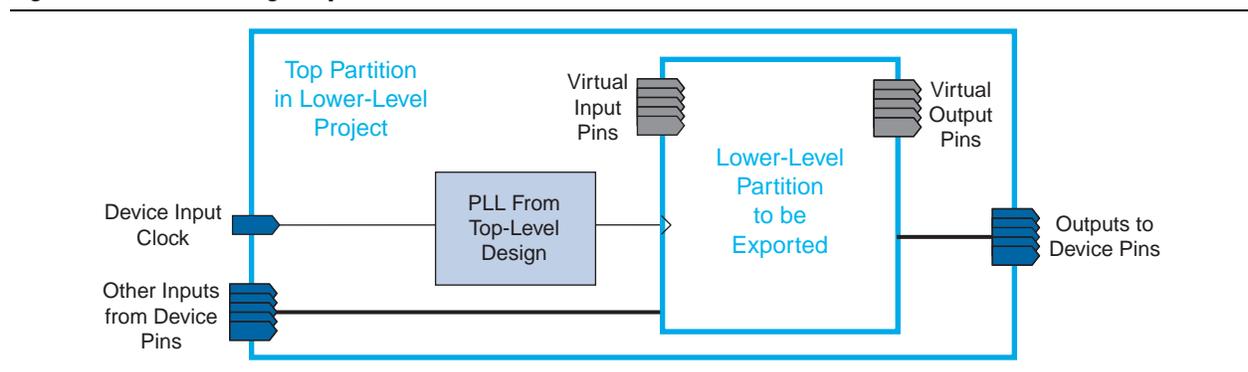
If you use a PLL in your top-level design and connect it to partitions designed in separate Quartus II projects by third-party IP designers, the IP partitions do not have information about the multiplication, phase shift, or compensation delays for the PLL in the top-level design. To accommodate the PLL timing, you can make appropriate timing assignments in the projects created by IP designers to ensure that clocks are not left unconstrained or constrained with an incorrect frequency. Alternatively, you can duplicate the top-level PLL (or other derived clock logic) in the design file for the project created by the IP designer to ensure that you have the correct PLL parameters and clock delays for a complete and accurate timing analysis.

If the project lead creates a copy of the top-level project framework that includes all the settings and constraints needed for the design, this framework should include PLLs and other interface logic if this information is important to optimize partitions.

If you use a separate Quartus II project for an independent design block (such as when a designer or third-party IP provider does not have access to the entire design framework), include a copy of the top-level PLL in the lower-level partition as shown in Figure 12-21.

In either case, the IP partition in the separate Quartus II project should contain just the partition logic that will be exported to the top-level design, while the full project includes more information about the top-level design. When the partition is complete, you can export just the partition without exporting the auxiliary PLL components to the top-level design. When you export a partition, the Quartus II software exports any hierarchy under the specified partition into the Quartus II Exported Partition File (.qxp), but does not include logic defined outside the partition (the PLL in this example).

Figure 12-21. Recreating a Top-Level PLL in a Lower-Level Partition



Checking Partition Quality

This section provides an overview of tools you can use to create and analyze partitions in the Quartus II software. Take advantage of these tools to assess your partition quality, and use the information to improve your design or assignments as required to achieve the best results.

Incremental Compilation Advisor

You can use the Incremental Compilation Advisor to check that your design follows Altera's recommendations for creating design partitions and implementing the incremental compilation design flow methodology. Each recommendation in the Incremental Compilation Advisor provides an explanation, describes the effect of the recommendation, and provides the action required to make the suggested change.

 For more information about the Incremental Compilation Advisor, refer to *Incremental Compilation Advisor Command* and *Example of Using the Incremental Compilation Advisor to Identify Non-Global Ports That Are Not Registered* in Quartus II Help, and the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.

Design Partition Planner

The Design Partition Planner allows you to view design connectivity and hierarchy, and can assist you in creating effective design partitions that follow the guidelines in this chapter. You can also use the Design Partition Planner to optimize design performance by isolating and resolving failing paths on a partition-by-partition basis.

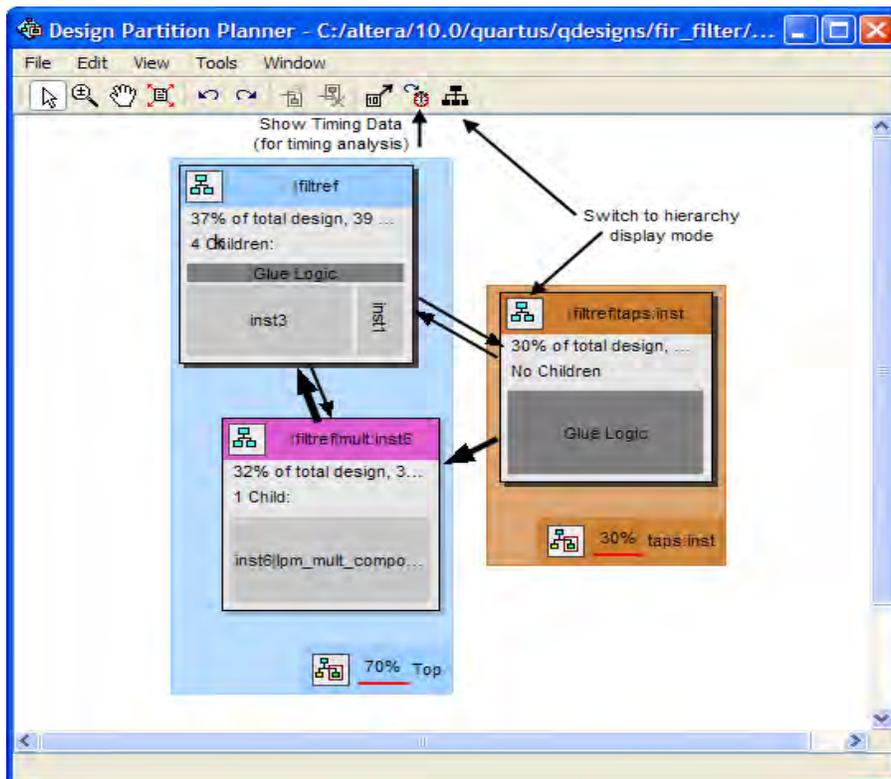
To view a design and create design partitions in the Design Partition Planner, you must first compile the design, or perform Analysis & Synthesis. In the Design Partition Planner, the design appears as a single top-level design block, with lower-level instances displayed as color-specific boxes.

In the Design Partition Planner, you can show connectivity between blocks and extract instances from the top-level design block. When you extract entities, connection bundles are drawn between entities, showing the number of connections existing between pairs of entities. When you have extracted a design block that you want to set as a design partition, right-click that design block, and then click **Create Design Partition**.

The Design Partition Planner also has an auto-partition feature that creates partitions based on the size and connectivity of the hierarchical design blocks. You can right-click the design block you want to partition (such as the top-level design hierarchy), and then click **Auto-Partition Children**. You can then analyze and adjust the partition assignments as required.

Figure 12-22 shows the Design Partition Planner after making a design partition assignment to one instance, and dragging another instance away from the top-level block within the same partition (two design blocks in the pale blue shaded box). The figure shows the connections between each partition and information about the size of each design instance.

Figure 12-22. Design Partition Planner



You can switch between connectivity display mode and hierarchical display mode, or temporarily to a view-only hierarchy display. You can also remove the connection lines between partitions and I/O banks by turning off **Display connections to I/O banks**, or use the settings on the **Connection Counting** tab in the **Bundle Configuration** dialog box to adjust how the connections are counted in the bundles.

To optimize design performance, confine failing paths within individual design partitions so that there are no failing paths passing between partitions, as discussed in earlier sections. In the top-level entity, child entities that contain failing paths are marked by a small red dot in the upper right corner of the entity box.

To view the critical timing paths from a timing analyzer report, first perform a timing analysis on your design, and then in the Design Partition Planner, click **Show Timing Data** on the View menu.

? For more information about the Design Partition Planner, refer to *About the Design Partition Planner* and *Using the Design Partition Planner* in Quartus II Help.

Viewing Design Partition Planner and Floorplan Side-by-Side

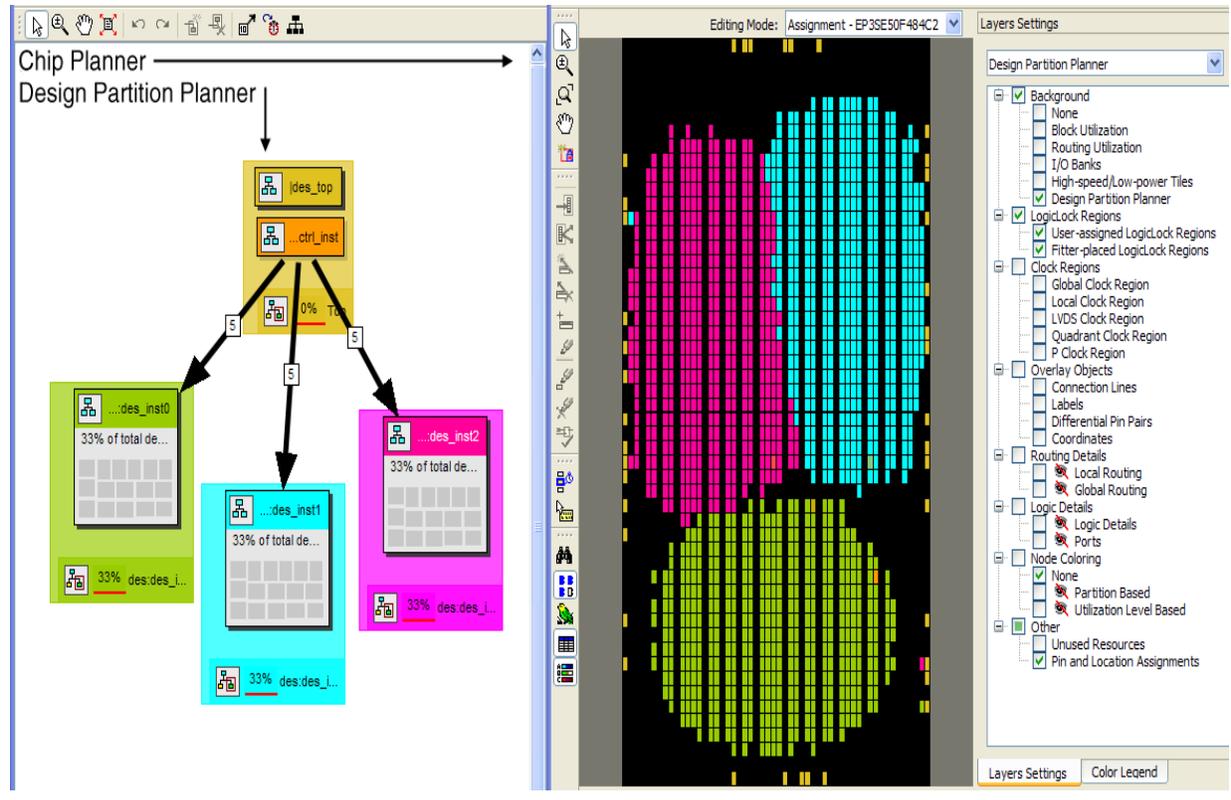
You can use the Design Partition Planner together with the Chip Planner to analyze natural placement groupings. This information can help you decide whether the design blocks should be grouped together in one partition, or whether they will make good partitions in the next compilation. It can also help determine whether the logic can easily be constrained by a LogicLock region. If logic naturally groups together when compiled without placement constraints, you can probably assign a reasonably sized LogicLock region to constrain the placement for subsequent compilations. You can experiment by extracting different design blocks in the Design Partition Planner and viewing the placement results of those design blocks from the previous compilation.

To view the Design Partition Planner and Chip Planner side-by-side, open the Design Partition Planner, and then open the Chip Planner and select the **Design Partition Planner** task. The **Design Partition Planner** task displays the physical locations of design entities with the same colors as in the Design Partition Planner.

In the Design Partition Planner, you can extract instances of interest from their parents by dragging and dropping, or with the **Extract from Parent** command. Evaluate the physical locations of instances in the Chip Planner and the connectivity between instances displayed in the Design Partition Planner. An entity is generally not suitable to be set as a separate design partition or constrained in a LogicLock region if the Chip Planner shows it physically dispersed over a noncontiguous area of the device after compilation. Use the Design Partition Planner to analyze the design connections. Child instances that are unsuitable to be set as separate design partitions or placed in LogicLock regions can be returned to their parent by dragging and dropping, or with the **Collapse to Parent** command.

Figure 12-23 shows a design displayed in the Design Partition Planner and the Chip Planner with different colors for the top-level design and the three major design instances.

Figure 12-23. Design Partition Planner and Chip Planner



For more information about the Design Partition Planner, refer to *About the Design Partition Planner* and *Using the Design Partition Planner* in Quartus II Help.

Partition Statistics Report

You can view statistics about design partitions in the Partition Merge Partition Statistics compilation report and the **Statistics** tab of the **Design Partitions Properties** dialog box. These reports are useful when optimizing your design partitions, or when compiling the completed top-level design in a team-based compilation flow to ensure that partitions meet the guidelines discussed in this chapter.

The Partition Merge Partition Statistics report in the Partition Merge section of the Compilation report lists statistics about each partition. The statistics for each partition (each row in the table) include the number of logic cells it contains, as well as the number of input and output pins and how many are registered. This report also lists how many ports are unconnected, or driven by a constant V_{CC} or GND. You can use this information to assess whether you have followed the guidelines for partition boundaries.

You can also view statistics about the resource and port connections for a particular partition on the **Statistics** tab of the **Design Partition Properties** dialog box. The **Show All Partitions** button allows you to view all the partitions in the same report. The Partition Merge Partition Statistics report also shows statistics for the **Internal Congestion: Total Connections and Registered Connections**. This information represents how many signals are connected within the partition. It then lists the inter-partition connections for each partition, which helps to see how partitions are connected to each other.

- ② For more information about the Partition Merge Reports, refer to *Partition Merge Reports* in Quartus II Help.

Report Partition Timing in the TimeQuest Timing Analyzer

The Report Partitions diagnostic report and the `report_partitions` SDC command in the TimeQuest Timing Analyzer produce a **Partition Timing Overview** and **Partition Timing Details** table, which lists the partitions and the number of failing paths and the worst case timing slack within each partition.

You can use these reports to analyze the location of the critical timing paths in the design are in relation to partitions. If a certain partition contains many failing paths, or failing inter-partition paths, you might be able to change your partitioning scheme and improve timing performance.

-  For more information about the TimeQuest `report_timing` command and reports, see the *Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

Check if Partition Assignments Impact the Quality of Results

You can ensure that you limit negative effect on the quality of results by following an iterative methodology during the partitioning process. In any incremental compilation flow where you can compile the source code for every partition during the partition planning phase, Altera recommends the following iterative flow:

1. Start with a complete design that is not partitioned and has no location or LogicLock region assignments.

After Analysis & Synthesis and Partition Merge, perform a placement and timing analysis estimate with the **Start Early Timing Estimate** command. To run a full compilation instead, use the **Start Compilation** command.
2. Record the quality of results from the Compilation report (timing slack or f_{MAX} , area, and any other relevant results).
3. Create design partitions following the guidelines described in this chapter.
4. Perform another Early Timing Estimate or a full compilation.
5. Record the quality of results from the Compilation report. If the quality of results is significantly worse than those obtained in the previous compilation, repeat step 3 through step 5 to change your partition assignments and use a different partitioning scheme.

6. Even if the quality of results is acceptable, you can repeat step 3 through step 5 by further dividing a large partition into several smaller partitions, which can improve compilation time in subsequent incremental compilations. You can repeat these steps until you achieve a good trade-off point (that is, all critical paths are localized within partitions, the quality of results is not negatively affected, and the size of each partition is reasonable).

You can also remove or disable partition assignments defined in the top-level design at any time during the design flow to compile the design as one flat compilation and get all possible design optimizations to assess the results. To disable the partitions without deleting the assignments, use the **Ignore partition assignments during compilation** option on the **Incremental Compilation** page of the **Settings** dialog box in the Quartus II software. The **Ignore partition assignments during compilation** option disables all design partition assignments in your project and runs a full compilation ignoring all partition boundaries and netlists. This option can be useful if you are using partitions to reduce compilation time as you develop various parts of the design, but can run a long compilation near the end of the design cycle to ensure the design meets its timing requirements.

Including SDC Constraints from Lower-Level Partitions for Third-Party IP Delivery

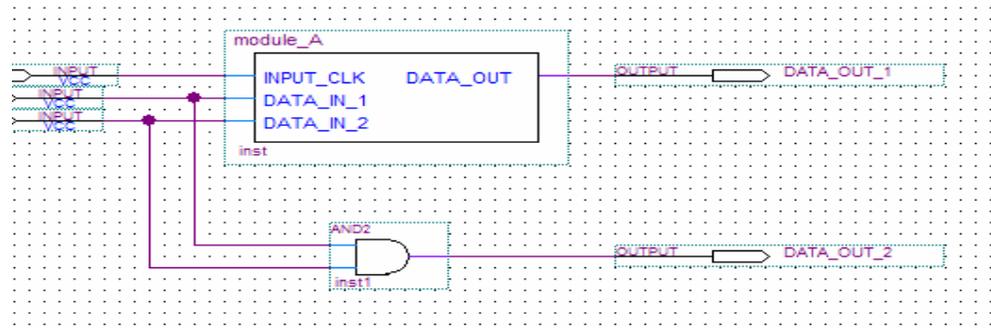
When exported partitions are compiled in a separate Quartus II project, such as when a third-party designer is delivering IP, the project lead must transfer the top-level project framework information and constraints to the partitions, so that each designer has a consistent view of the constraints that apply to the entire design. If the independent partition designers make any changes or add any constraints, they might have to transfer new constraints back to the project lead, so that these constraints are included in final timing sign-off of the entire design. Many assignments from the partition are carried with the partition into the top-level design, however, SDC format constraints for the TimeQuest Timing Analyzer are not copied into the top-level design automatically.

Passing additional timing constraints from a partition to the top-level design must be managed carefully. This section provides recommendations for managing the timing constraints in a third-party IP delivery flow. Note that you can design within a single Quartus II project or a copy of the top-level design to simplify constraint management.

To ensure that there are no conflicts between the project lead's top-level constraints and those added by the third-party IP designer, use two Synopsys Design Constraint Files (.sdc) for each separate Quartus II project: an .sdc created by the project lead that includes project-wide constraints, and an .sdc created by the IP designer that includes partition-specific constraints.

This section uses the example design shown in [Figure 12-24](#) to illustrate these recommendations. The top-level design instantiates a lower-level design block called `module_A` that is set as a design partition and developed by an IP designer in a separate Quartus II project.

Figure 12-24. Example Design to Illustrate SDC Constraints



In this top-level design, there is a single clock setting called `clk` associated with the FPGA input called `top_level_clk`. The top-level `.sdc` contains the following constraint for the clock:

```
create_clock -name {clk} -period 3.000 -waveform { 0.000 1.500 }
[get_ports {TOP_LEVEL_CLK}]
```

Creating an `.sdc` File With Project-Wide Constraints

The `.sdc` with project-wide constraints for the separate Quartus II project should contain all constraints that are not completely localized to the partition. The `.sdc` should be maintained by the project lead. The project lead must ensure that these timing constraints are delivered to the individual partition owners and that they are syntactically correct for each of the separate Quartus II projects. This communication can be challenging when the design is in flux and hierarchies change. The project lead can use design partition scripts to automatically pass some of these constraints to the separate Quartus II projects.

- ❓ For more information on how to generate design partition scripts, refer to [Generating Design Partition Scripts for Project Management](#) in Quartus II Help.

The `.sdc` with project-wide constraints is used in the partition, but is not exported back to the top-level design. The partition designer should not modify this file. If changes are necessary, they should be communicated to the project lead, who can then update the SDC constraints and distribute new files to all partition designers as required.

The `.sdc` should include clock creation and clock constraints for any clock used by more than one partition. These constraints are particularly important when working with complex clocking structures, such as the following:

- Cascaded clock multiplexers
- Cascaded PLLs
- Multiple independent clocks on the same clock pin

- Redundant clocking structures required for secure applications
- Virtual clocks and generated clocks that are consistently used for source synchronous interfaces
- Clock uncertainties

Additionally, the `.sdc` with project-wide constraints should contain all project-wide timing exception assignments, such as the following:

- Multicycle assignments, `set_multicycle_path`
- False path assignments, `set_false_path`
- Maximum delay assignments, `set_max_delay`
- Minimum delay assignments, `set_min_delay`

The project-wide `.sdc` can also contain any `set_input_delay` or `set_output_delay` constraints that are used for ports in separate Quartus II projects, because these represent delays external to a given partition. If the partition designer wants to set these constraints within the separate Quartus II projects, the team must ensure that the I/O port names are identical in all projects so that the assignments can be integrated successfully without changes.

Similarly, a constraint on a path that crosses a partition boundary should be in the project-wide `.sdc`, because it is not completely localized in a separate Quartus II project.

Example Step 1: Project Lead Produces `.sdc` with Project-Wide Constraints for Lower-Level Partitions

The device input `top_level_clk` in [Figure 12-24](#) drives the `input_clk` port of `module_A`. To make sure the clock constraint is passed correctly to the partition, the project lead creates an `.sdc` with project-wide constraints for `module_A` that contains the following command:

```
create_clock -name {clk} -period 3.000 -waveform { 0.000 1.500 }  
[get_ports {INPUT_CLK}]
```

The designer of `module_A` includes this `.sdc` as part of the separate Quartus II project.

Creating an `.sdc` with Partition-Specific Constraints

The `.sdc` with partition-specific constraints should contain all constraints that affect only the partition. For example, a `set_false_path` or `set_multicycle_path` constraint for a path entirely within the partition should be in the partition-specific `.sdc`. These constraints are required for correct compilation of the partition, but need not be present in any other separate Quartus II projects.

The partition-specific `.sdc` should be maintained by the partition designer; it is their responsibility to add any constraints required to properly compile and analyze their partition.

The partition-specific `.sdc` is used in the separate Quartus II project and must be exported back to the project lead for the top-level design. The project lead must use the partition-specific constraints to properly constrain the placement, routing, or both, if the partition logic is fit at the top level, and to ensure that final timing sign-off is accurate. Use the following guidelines in the partition-specific `.sdc` to simplify these export and integration steps:

- Create a hierarchy variable for the partition (such as `module_A_hierarchy`) and set it to an empty string because the partition is the top-level instance in the separate Quartus II project. The project lead modifies this variable for the top-level hierarchy, reducing the effort of translating constraints on lower-level design hierarchies into constraints that apply in the top-level hierarchy. Use the following Tcl command first to check if the variable is already defined in the project, so that the top-level design does not use this empty hierarchy path: `if {[info exists module_A_hierarchy]}`.
- Use the hierarchy variable in the partition-specific `.sdc` as a prefix for assignments in the project. For example, instead of naming a particular instance of a register `reg:inst`, use `${module_A_hierarchy}reg:inst`. Also use the hierarchy variable as a prefix to any wildcard characters (such as `"*"`).
- Pay attention to the location of the assignments to I/O ports of the partition. In most cases, these assignments should be specified in the `.sdc` with project-wide constraints because the partition interface depends on the top-level design. If you want to set I/O constraints within the partition, the team must ensure that the I/O port names are identical in all projects so that the assignments can be integrated successfully without changes.
- Be careful with the `derive_clocks` and `derive_pll_clocks` commands. In most cases, the `.sdc` with project-wide constraints should call these commands. Because these commands impact the entire design, integrating them unexpectedly into the top-level design might cause problems.

If the design team follows these recommendations, the project lead should be able to include the `.sdc` with the partition-specific constraints provided by the partition designer directly in the top-level design.

Example Step 2: Partition Designer Creates `.sdc` with Partition-Specific Constraints

The partition designer compiles the design with the `.sdc` with project-wide constraints and might want to add some additional constraints. In this example, the designer realizes that they must specify a false path between the register called `reg_in_1` and all destinations in this design block with the wildcard character (such as `"*"`). This constraint applies entirely within the partition and must be exported to the top-level design, so it qualifies for inclusion in the `.sdc` with partition-specific constraints. The designer first defines the `module_A_hierarchy` variable and uses it when writing the constraint as follows:

```
if {[info exists module_A_hierarchy]} {
    set module_A_hierarchy ""
}
set_false_path -from [get_registers ${module_A_hierarchy}reg_in_1] -to
[get_registers ${module_A_hierarchy}*]
```

Consolidating the `.sdc` in the Top-Level Design

When the partition designers complete their designs, they export the results to the project lead. The project lead receives the exported `.qxp` files and a copy of the `.sdc` with partition-specific constraints.

To set up the top-level .sdc constraint file to accept the .sdc files from the separate Quartus II projects, the top-level .sdc should define the hierarchy variables specified in the partition .sdc files. List the variable for each partition and set it to the hierarchy path, up to and including the instantiation of the partition in the top-level design, including the final hierarchy character "|".

To ensure that the .sdc files are used in the correct order, the project lead can use the Tcl Source command to load each .sdc.

Example Step 3: Project Lead Performs Final Timing Analysis and Sign-off

With these commands, the top-level .sdc file looks like the following example:

```
create_clock -name {clk} -period 3.000 -waveform { 0.000 1.500 }  
[get_ports {TOP_LEVEL_CLK}]  
# Include the lower-level SDC file  
set module_A_hierarchy "module_A:inst|" # Note the final '|' character  
source <partition-specific constraint file such as  
..\module_A\module_A_constraints>.sdc
```

When the project lead performs top-level timing analysis, the false path assignment from the lower-level module_A project expands to the following:

```
set_false_path -from module_A:inst|reg_in_1 -to module_A:inst|*
```

Adding the hierarchy path as a prefix to the SDC command makes the constraint legal in the top-level design, and ensures that the wildcard does not affect any nodes outside the partition that it was intended to target.

By following the guidelines in this section, constraint propagation between the separate Quartus II projects can be managed effectively.

Introduction to Design Floorplans

A floorplan represents the layout of the physical resources on the device. Creating a design floorplan, or floorplanning, describes describe the process of mapping the logical design hierarchy onto physical regions in the device.

In the Quartus II software, LogicLock regions can be used to constrain blocks of a design to a particular region of the device. LogicLock regions represent an area on the device with a user-defined or Fitter-defined size and location in the device layout.



For more information about design floorplans and LogicLock regions, refer to the *Analyzing and Optimizing the Design Floorplan* chapter in volume 2 of the *Quartus II Handbook*.

The Difference between Logical Partitions and Physical Regions

Design partitions are logical entities based on the design hierarchy. LogicLock regions are physical placement assignments that constrain logic to a particular region on the device.

It is a common misconception that logic from a design partition is always grouped together on the device when you use incremental compilation. Actually, logic from a partition can be placed anywhere in the device if it is not constrained to a LogicLock region, although the Fitter can pack related logic together to improve timing performance. A logical design partition does not refer to any physical area on the device and does not directly control *where* instances are placed on the device.

If you want to control the placement of logic from a design partition and isolate it to a particular part of the device, you can assign the logical design partition to a physical region in the device floorplan with a LogicLock region assignment. Altera recommends creating a design floorplan by assigning design partitions to LogicLock regions to improve the quality of results and avoid placement conflicts in some situations for incremental compilation. For more information, refer to [“Why Create a Floorplan?” on page 12-42](#).

Another misconception is that LogicLock assignments are used to preserve placement results for incremental compilation. Actually, LogicLock regions only *constrain* logic to a physical region on the device. Incremental compilation does not use LogicLock assignments or any location assignments to preserve the placement results; it simply reuses the results stored in the database netlist from a previous compilation.

Why Create a Floorplan?

Creating a design floorplan is usually required if you want to preserve placement for partitions that will be exported, to avoid resource conflicts between partitions in the top-level design. Floorplan location planning can be important for a design that uses incremental compilation, for the following reasons:

- To avoid resource conflicts between partitions, predominantly when integrating partitions exported from another Quartus II project.
- To ensure a good quality of results when recompiling individual timing-critical partitions.

Location assignments for each partition ensure that there are no placement conflicts between partitions. If there are no LogicLock region assignments, or if LogicLock regions are set to auto-size or floating location, no device resources are specifically allocated for the logic associated with the region. If you do not clearly define resource allocation, logic placement can conflict when you integrate the partitions in the top-level design if you reuse the placement information from the exported netlist.

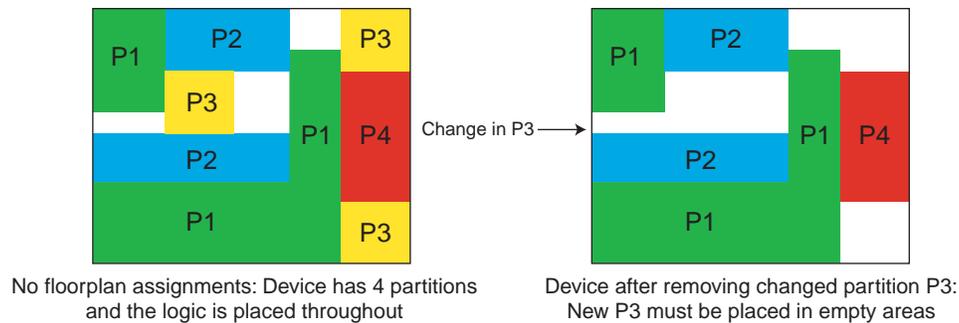
Creating a floorplan is also recommended for timing-critical partitions that have little timing margin to maintain good quality of results when the design changes.

Floorplan assignments are not required for non-critical partitions compiled in the same Quartus II project. The logic for partitions that are not timing-critical can be placed anywhere in the device on each recompilation if that is best for your design.

Design floorplan assignments prevent the situation in which the Fitter must place a partition in an area of the device where most resources are used by other partitions. A LogicLock region provides a reasonable region to re-place logic after a change, so the Fitter does not have to scatter logic throughout the available space in the device.

Figure 12-25 illustrates the problems that may be associated with refitting designs that do not have floorplan location assignments. The left floorplan shows the initial placement of a four-partition design (P1-P4) without any floorplan location assignments. The right floorplan shows the device if a change occurs to P3. After removing the logic for the changed partition, the Fitter must replace and reroute the new logic for P3 in the scattered white space shown in Figure 12-25. The placement of the post-fit netlists for other partitions forces the Fitter to implement P3 with the device resources that have not been used.

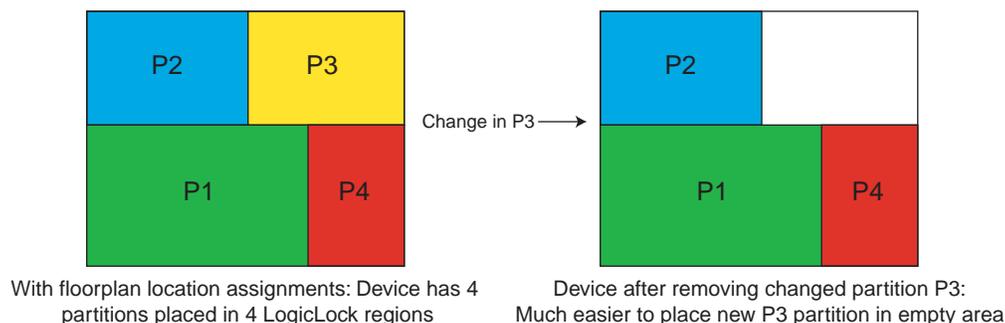
Figure 12-25. Representation of Device Floorplan without Location Assignments



The Fitter has a more difficult task because of more difficult physical constraints, and as a result, compilation time often increases. The Fitter might not be able to find any legal placement for the logic in partition P3, even if it could in the initial compilation. Additionally, if the Fitter can find a legal placement, the quality of results often decreases in these cases, sometimes dramatically, because the new partition is now scattered throughout the device.

Figure 12-26 shows the initial placement of a four-partition design with floorplan location assignments. Each partition is assigned to a LogicLock region. The second part of the figure shows the device after partition P3 is removed. This placement presents a much more reasonable task to the Fitter and yields better results.

Figure 12-26. Representation of Device Floorplan with Location Assignments



Altera recommends that you create a LogicLock floorplan assignment for timing-critical blocks with little timing margin that will be recompiled as you make changes to the design.

When to Create a Floorplan

It is important that you plan early to incorporate partitions into the design, and ensure that each partition follows partitioning guidelines. You can create floorplan assignments at different stages of the design flow, early or late in the flow. These guidelines help ensure better results as you begin creating floorplan location assignments.

Early Floorplan

An early floorplan is created before the design stage. You can plan an early floorplan at the top level of a design to allocate each partition a portion of the device resources. Doing so allows the designer for each block to create the logic for their design partition without conflicting with other logic. Each partition can be optimized in a separate Quartus II project if required, and the design can still be easily integrated in the top-level design. Even within one Quartus II project, each partition can be locked down with a post-fit netlist, and you can be sure there is space in the device floorplan for other partitions.

When you have compiled your complete design, or after you have integrated the first versions of partitions developed in separate Quartus II projects, you can use the design information and Quartus II features to tune and improve the floorplan, as described in the following section.

Late Floorplan

A late floorplan is created or modified after the design is created, when the code is close to complete and the design structure is likely to remain stable. Creating a late floorplan is typically necessary only if you are starting to use incremental compilation late in the design flow, or need to reserve space for a logic block that becomes timing-critical but still has HDL changes to be integrated. When the design is complete, you can take advantage of the Quartus II analysis features to check the floorplan quality. To adjust the floorplan, you can perform iterative compilations as required and assess the results of different assignments.



It may not be possible to create a good quality late floorplan if you do not create partitions in the early stages of the design.

Design Floorplan Placement Guidelines

The following guidelines are key to creating a good design floorplan:

- Capture correct resources in each region.
- Use good region placement to maintain design performance compared to flat compilation.

It is a common misconception that creating a floorplan enhances timing performance, as compared to a flat compilation with no location assignments. The Fitter does not usually require guidance to get optimal results for a full design.

Floorplan assignments can help maintain good performance when designs change incrementally, as described in “[Why Create a Floorplan?](#)” on page 12-42. However, poor placement assignments can often adversely affect performance results, as compared to a flat compilation, because the assignments limit the options for the Fitter. Investing time to find good region placement is required to match the performance of a full flat compilation.

Use the following general procedure to create a floorplan:

1. Divide the design into partitions.
2. Assign the partitions to LogicLock regions.
3. Compile the design.
4. Analyze the results.
5. Modify the placement and size of regions, as required.

You might have to iterate through these steps several times to find the best combination of design partitions and LogicLock regions that meet the resource and timing goals of the design.



For more information about performing these steps, refer to the [Quartus II Incremental Compilation for Hierarchical and Team-Based Design](#) chapter in volume 1 of the [Quartus II Handbook](#).

Assigning Partitions to LogicLock Regions

Before compiling a design with new LogicLock assignments, ensure that the partition netlist type is set to **Post-Synthesis** or **Source File**, so that the Fitter does not reuse previous placement results.

In most cases, you should include logic from one partition in each LogicLock region. This organization helps to prevent resource conflicts when partitions are exported and can lead to better performance preservation when locking down parts of a design in a single project.

The Quartus II software is flexible and allows exceptions to this rule. For example, you can place more than one partition in the same LogicLock region if the partitions are tightly connected, but you do not want to merge the partitions into one larger partition. For best results, ensure that you recompile all partitions in the LogicLock region every time the logic in one partition changes. Additionally, if a partition contains multiple lower-level entities, you can place those entities in different areas of the device with multiple LogicLock regions, even if they are defined in the same partition.

You can use the **Reserved** LogicLock option to ensure that you avoid conflicts with other logic that is not locked into a LogicLock region. This option prevents other logic from being placed in the region, and is useful if you have empty partitions at any point during your design flow, so that you can reserve space in the floorplan. Do not make reserved regions too large to prevent unused area because no other logic can be placed in a region with the **Reserved** LogicLock option.



For more information about LogicLock region properties, refer to the [LogicLock Region Properties Dialog Box](#) in Quartus II Help.

How to Size and Place Regions

In an early floorplan, assign physical locations based on design specifications. Use information about the connections between partitions, the partition size, and the type of device resources required.

In a late floorplan when the design is complete, you can use locations or regions chosen by the Fitter as a guideline. If you have compiled the full design, you can view the location of the partition logic in the Chip Planner. Refer to “[Checking Partition Quality](#)” on page 12-31 for information about viewing placement results for each partition in the device floorplan. You can use the natural grouping of each unconstrained partition as a starting point for a LogicLock region constraint. View the placement for each partition that requires a floorplan constraint, and create a new LogicLock region by drawing a box around the area on the floorplan then assign the partition to the region to constrain the partition placement.

- ② For step-by-step information on how to create a LogicLock region, refer to [Creating and Manipulating LogicLock Regions](#) in Quartus II Help.

Instead of creating regions based on the previous compilation results, you can start with the Fitter results for a default auto size and floating origin location for each new region when the design logic is complete. After compilation, lock the size and origin location. Instead of a full compilation, you can use the **Start Early Timing Estimate** command to perform a fast placement.

Alternatively, if the design logic is complete with auto sized or floating location regions, you can specify the size based on the synthesis results and use the locations chosen by the Fitter with the **Set to Estimated Size** command. Like the previous option, start with floating origin location. After compilation, lock the origin location. Again, instead of a full compilation, you can use the **Start Early Timing Estimate** command to perform a fast placement. You can also enable the **Fast Synthesis Effort** setting to reduce synthesis time.

After a compilation or early timing estimate, save the Fitter size and origin location of the Fitter with the **Set Size and Origin to Previous Fitter Results** command.

- ② For more information on LogicLock region, refer to [Creating and Manipulating LogicLock Regions](#) in Quartus II Help.

 It is important that you use the Fitter-chosen locations only as a starting point to give the regions a good fixed size and location. Ensure that all LogicLock regions in the design have a fixed size and have their origin locked to a specific location on the device. On average, regions with fixed size and location yield better timing performance than auto-sized regions.

Modifying Region Size and Origin

After saving the Fitter results from an initial compilation for a late floorplan, modify the regions using your knowledge of the design to set a specific size and location. If you have a good understanding of how the design fits together, you can often improve upon the regions placed in the initial compilation. In an early floorplan when the design has not yet been created, you can use the guidelines in this section to set the size and origin, even though there is no initial Fitter placement for a basis.

The easiest way to move and resize regions is to drag the region location and borders in the Chip Planner. Make sure that you select the **User-Defined** region in the floorplan (as opposed to the **Fitter-Placed** region from the last compilation) so that you can change the region.

Generally, you can keep the Fitter-determined relative placement of the regions, but make adjustments if required to meet timing performance. If you find that the early timing estimate did not result in good relative placements, try performing a full compilation so that the Fitter can optimize for a full placement and routing.

If two LogicLock regions have several connections between them, ensure they are placed near each other to improve timing performance. By placing connected regions near each other, the Fitter has more opportunity to optimize inter-region paths when both partitions are recompiled. Reducing the criticality of inter-region paths also allows the Fitter more flexibility when placing other logic in each region.

If resource utilization is low in the overall device, enlarge the regions. Doing so usually improves the final results because it gives the Fitter more freedom to place additional or modified logic added to the partition during subsequent incremental compilations. It also allows room for optimizations such as pipelining and physical synthesis logic duplication.

Try to have each region evenly full, with the same “fullness” that the complete design would have without LogicLock regions. As a very rough suggestion, try to have each region approximately 75% full.

Allow more area for regions that are densely populated, because overly congested regions can lead to poor results. Allow more empty space for timing-critical partitions to improve results. However, do not make regions too large for their logic. Regions that are too large can result in wasted resources and also lead to suboptimal results.

Ideally, almost the entire device should be covered by LogicLock regions if all partitions are assigned to regions.

Regions should not overlap in the device floorplan. If two partitions are allocated on an overlapping portion of the chip, each may independently claim common resources in this region. This leads to resource conflicts when integrating results into a top-level design. In a single project, overlapping regions give more difficult constraints to the Fitter and can lead to reduced quality of results.

You can create hierarchical LogicLock regions to ensure that the logic in a child partition is physically placed inside the LogicLock region for its parent partition. This can be useful when the parent partition does not contain registers at the boundary with the lower-level child partition and has a lot of signal connectivity. To create a hierarchical relationship between regions in the LogicLock Regions window, drag and drop the child region to the parent region.

I/O Connections

Consider I/O timing when placing regions. Using I/O registers can minimize I/O timing problems, and using boundary registers on partitions can minimize problems connecting regions or partitions. However, I/O timing might still be a concern. It is most important for flows where each partition is compiled independently, because the Fitter can optimize the placement for paths between partitions if the partitions are compiled at the same time.

Place regions close to the appropriate I/O, if necessary. For example, DDR memory interfaces have very strict placement rules to meet timing requirements. Incorporate any specific placement requirements into your floorplan as required. It is best to create LogicLock regions for internal logic only, and provide pin location assignments for external device I/O pins (instead of including the I/O cells in a LogicLock region to control placement).

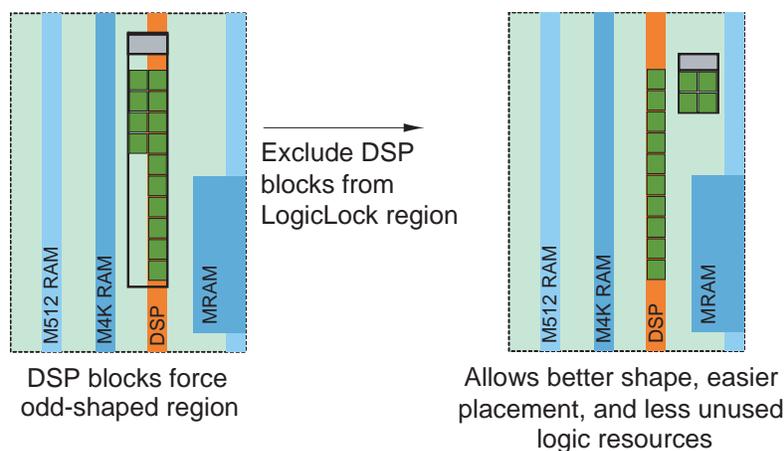
LogicLock Resource Exclusions

You can exclude certain resource types from a LogicLock region to manage the ratio of logic to dedicated DSP and RAM resources in the region.

If your design contains memory or Digital Signal Processing (DSP) elements, you may want to exclude these elements from the LogicLock region. LogicLock resource exceptions prevent certain types of elements from being assigned to a region. Therefore, those elements are not required to be placed inside the region boundaries. The option does not prevent them from being placed inside the region boundaries unless the **Reserved** property of the region is turned on.

Resource exceptions are useful in cases where it is difficult to place rectangular regions for design blocks that contain memory and DSP elements, due to their placement in columns throughout the device floorplan. Exclude RAMs, DSPs, or logic cells to give the Fitter more flexibility with region sizing and placement. Excluding RAM or DSP elements can help to resolve no-fit errors that are caused by regions spanning too many resources, especially for designs that are memory-intensive, DSP-intensive, or both. [Figure 12-27](#) shows an example of a design with an odd-shaped region to accommodate DSP blocks for a region that does not contain very much logic. The right side of the figure shows the result after excluding DSP blocks from the region. The region can be placed more easily without wasting logic resources.

Figure 12-27. LogicLock Resource Exclusion Example



To view any resource exceptions, right-click in the LogicLock Regions window, and then click **LogicLock Regions Properties**. In the **LogicLock Regions Properties** dialog box, select the design element (module or entity) in the **Members** box, and then click **Edit**. In the **Edit Node** dialog box, to set up a resource exception, click the **Edit** button next to the **Excluded element types** box, and then turn on the design element types to be excluded from the region. You can choose to exclude combinational logic or registers from logic cells, or any of the sizes of TriMatrix memory blocks, or DSP blocks.

If the excluded logic is in its own lower-level design entity (even if it is within the same design partition), you can assign the entity to a separate LogicLock region to constrain its placement in the device.

You can also use this feature with the LogicLock **Reserved** property to reserve specific resources for logic that will be added to the design.

Creating Floorplan Location Assignments With Tcl Commands—Excluding or Filtering Certain Device Elements (Such as RAM or DSP Blocks)

To assign a code block to a LogicLock region, with exclusions, use the following Tcl command:

```
set_logiclock_contents -region <LogicLock region name> -to <block>
-exceptions \"<keyword>:<keyword>\"
```

- *LogicLock region name*—The name of the LogicLock region to which the code block is assigned.
- *block*—A code block in a Quartus II project hierarchy, which can also be a design partition.
- *<keyword>*—The list of exceptions during assignment. For example, if DSP were in the keyword list, then the named block of code would be assigned to the LogicLock region, except for any DSP block within the code block. You can include the following exceptions in the `set_logiclock_contents` command:

Keyword variables:

- *REGISTER*—Any registers in the logic cells.
- *COMBINATIONAL*—Any combinational elements in the logic cells.
- *SMALL_MEM*—The small TriMatrix memory blocks (M512 or MLAB).
- *MEDIUM_MEM*—The medium TriMatrix memory blocks (M4K or M9K).
- *LARGE_MEM*—The large TriMatrix memory blocks (M-RAM or M144K).
- *DSP*—Any DSP blocks.
- *VIRTUAL_PIN*—Any virtual pins.



Resource filtering uses the optional Tcl argument `-exclude_resources` in the `set_logiclock_contents` function. If left unspecified, no resource filter is created. In the `.qsf`, resource filtering uses an extra LogicLock membership assignment called `LL_MEMBER_RESOURCE_EXCLUDE`. For example, the following line in the `.qsf` is used to specify a resource filter for the `alu:alu_unit` entity assigned to the ALU region.

```
set_instance_assignment -name LL_MEMBER_RESOURCE_EXCLUDE \
"DSP:SMALL_MEM" -to "alu:alu_unit" -section_id ALU
```

Creating Non-Rectangular Regions

To constrain placement to non-rectangular or non-contiguous areas of the device, you can connect multiple rectangular regions together using the **Merge** command.

For devices that do not support the **Merge** command (Arria™GX, Cyclone, Cyclone II, HardCopy II, MAX™ II, Stratix, Stratix II, Stratix II GX, and Stratix GX devices), you can limit entity placement to a sub-area of a LogicLock region to create non-rectangular constraints. In these devices, construct a LogicLock hierarchy by creating child regions inside of parent regions, and then use the **Reserved** option to control which logic can be placed inside these child regions. Setting the **Reserved** option for the region prevents the Fitter from placing nodes that are not assigned to the region inside the boundary of the region.

- ② For more information and examples of creating non-rectangular regions, refer to *Creating and Manipulating LogicLock Regions* in Quartus II Help.

Checking Floorplan Quality

This section provides an overview of tools that you can use as you create a floorplan in the Quartus II software. You can use these tools to assess your floorplan quality and use the information to improve your design or assignments as required to achieve the best results.

Incremental Compilation Advisor

You can use the Incremental Compilation Advisor to check that your design follows the recommendations for creating floorplan location assignments that are presented in this chapter. For more information, refer to *“Incremental Compilation Advisor”* on page 12-32.

LogicLock Region Resource Estimates

You can view resource estimates for a LogicLock region to determine the region's resource coverage, and use this estimate before compilation to check region size. Using this estimate helps to ensure adequate resources when you are sizing or moving regions.

- ② For information about how to view the properties of a LogicLock region, refer to *LogicLock Region Properties Dialog Box* in Quartus II Help.

LogicLock Region Properties Statistics Report

LogicLock region statistics are similar to design partition properties, but also include resource usage details after compilation.

The statistics report the number of resources used and the total resources covered by the region, and also lists the number of I/O connections and how many I/Os are registered (good), as well as the number of internal connections and the number of inter-region connections (bad).

- ❓ For information about the **Statistics** tab in the **LogicLock Region Properties** dialog box, refer to *LogicLock Region Properties Dialog Box* in Quartus II Help.

Locate the Quartus II TimeQuest Timing Analyzer Path in the Chip Planner

In the TimeQuest user interface, you can locate a specific path in the Chip Planner to view its placement and perform a report timing operation (for example, report timing for all paths with less than 0 ns slack).

- ❓ For information about how to locate paths between the TimeQuest Timing Analyzer and the Chip Planner, refer to *Locate Dialog Box* in Quartus II Help.

Inter-Region Connection Bundles

The Chip Planner can display bundles of connections between LogicLock regions, with filtering options that allow you to choose the relevant data for display. These bundles can help you to visualize how many connections there are between each LogicLock region to improve floorplan assignments or to change partition assignments, if required.

- ❓ For information about how to display bundles of connections between LogicLock regions, refer to *Generate Inter-region Bundles Dialog Box* in Quartus II Help.

Routing Utilization

The Chip Planner includes a feature to display a color map of routing congestion. This display helps identify areas of the chip that are too tightly packed.

In the Chip Planner, red LAB blocks indicate higher routing congestion. You can position the mouse pointer over a LAB to display a tooltip that reports the logic and routing utilization information.

- ❓ For information about how to view a color map of routing congestion in the Chip Planner, refer to *About the Chip Planner* in Quartus II Help.

Ensure Floorplan Assignments Do Not Significantly Impact Quality of Results

The end results of design partitioning and floorplan creation differ from design to design. However, it is important to evaluate your results to ensure that your scheme is successful. Compare your before and after results, and consider using another scheme if any of the following guidelines are not met:

- You should see only minor degradation in f_{MAX} after the design is partitioned and floorplan location assignments are created. There is some performance cost associated with setting up a design for incremental compilation; approximately 3% is typical.
- The area increase should be no more than 5% after the design is partitioned and floorplan location assignments are created.
- The time spent in the routing stage should not significantly increase.

The amount of compilation time spent in the routing stage is reported in the Messages window with an Info message that indicates the elapsed time for Fitter routing operations. If you notice a dramatic increase in routing time, the floorplan location assignments may be creating substantial routing congestion. In this case, decrease the number of LogicLock regions, which typically reduces the compilation time in subsequent incremental compilations and may also improve design performance.

Recommended Design Flows and Application Examples

This section provides design flows for partitioning and creating a design floorplan during common timing closure and team-based design scenarios. Each flow describes the situation in which it should be used, and provides a step-by-step description of the commands required to implement the flow.

Create a Floorplan for Major Design Blocks

Use this incremental compilation flow for designs when you want to assign a floorplan location for each major block in your design. A full floorplan ensures that partitions do not interact as they are changed and recompiled—each partition has its own area of the device floorplan.

To create a floorplan for major design blocks, follow this general methodology:

1. In the Design Partitions window, ensure that all partitions have their netlist type set to **Source File** or **Post-Synthesis**. If the netlist type is set to **Post-Fit**, floorplan location assignments are not used when recompiling the design.
2. Create a LogicLock region for each partition (including the top-level entity, which is set as a partition by default).
3. Run a full compilation of your design to view the initial Fitter-chosen placement of the LogicLock regions as a guideline.
4. In the Chip Planner, view the placement results of each partition and LogicLock on the device.
5. If required, modify the size and location of the regions in the Chip Planner. For example, enlarge the regions to fill up the device and allow for future logic changes.

You can also, if needed, create a new LogicLock region by drawing a box around an area on the floorplan.

6. Run an early timing estimate with the **Start Early Timing Estimate** command to estimate the timing performance of your design with the modified or new LogicLock regions.
7. Repeat steps 5 and 6 until you are satisfied with the quality of results for your design floorplan. Once you are satisfied with your results, run a full compilation of your design.

Create a Floorplan Assignment for One Design Block with Difficult Timing

Use this flow when you have one timing-critical design block that requires more optimization than the rest of your design. You can take advantage of incremental compilation to reduce your compilation time without creating a full design floorplan.

In this scenario, you do not want to create floorplan assignments for the entire design. Instead, you can create a region to constrain the location of your critical design block, and allow the rest of the logic to be placed anywhere on the device. To create a region for critical design block, follow these steps:

1. Divide up your design into partitions. Consider the guidelines in “[Design Partition Guidelines](#)” on page 12-10 to determine partition boundaries. Ensure that you isolate the timing-critical logic in a separate partition.
2. Define a LogicLock region for the timing-critical partition. Ensure that you capture the correct amount of device resources in the region. Turn on the **Reserved** property to prevent any other logic from being placed in the region.
 - If the design block is not complete, reserve space in the design floorplan based on your knowledge of the design specifications, connectivity between design blocks, and estimates of the size of the partition based on any initial implementation numbers.
 - If the critical design block has initial source code ready, compile the design to place the LogicLock region. Save the Fitter-determined size and origin, then enlarge the region to provide more flexibility and allow for future design changes.

As the rest of the design is completed, and the device fills up, the timing-critical region has a reserved area of the floorplan. When you make changes to the design block, the logic will be re-placed in the same part of the device, which helps ensure good quality of results.

Create a Floorplan as the Project Lead in a Team-Based Flow

Use this approach when you have several designs that will be implemented in separate Quartus II projects by different designers, or third-party IP designers who want to optimize their designs independently and pass the results to the project lead.

As the project lead in this scenario, follow these steps to prepare the top-level design for a successful team-based design methodology with early floorplan planning:

1. Create a new Quartus II project that will ultimately contain the full implementation of the entire design.
2. Create a “skeleton” or framework of the design that defines the hierarchy for the subdesigns that will be implemented by separate designers. Consider the partitioning guidelines in this chapter while determining the design hierarchy.
3. Make project-wide settings. Select the device, make global assignments for clocks and device I/O ports, and make any global signal constraints to specify which signals can use global routing resources.
4. Make design partition assignments for each major subdesign. Set the netlist type for each partition that will be implemented in a separate Quartus II project and later exported and integrated with the top-level design to **Empty**.

5. Create LogicLock regions for each partition to create a design floorplan. This floorplan should consider the connectivity between partitions and estimates of the size of each partition based on any initial implementation numbers and knowledge of the design specifications. Use the guidelines described in this chapter to choose a size and location for each LogicLock region.
 6. Provide the constraints from the top-level design to partition designers using one of the following procedures:
 - a. Create a copy of the top-level Quartus II project framework by checking out the appropriate files from a source control system, using the **Copy Project** command, or creating a project archive. Provide each partition designer with the copy of the project.
 - b. Provide the constraints with documentation or scripts.
- ② To use design partition scripts to pass constraints and generate separate Quartus II projects, refer to *Generating Design Partition Scripts for Project Management* in Quartus II Help.

Conclusion

Incremental compilation can significantly improve your design productivity, especially for large, complex designs. To take advantage of the feature, it is worth spending time to create quality partition and floorplan assignments.

Follow the guidelines to set up your design hierarchy and source code for incremental compilation. Keep partitions independent of each other and do not rely on any cross-boundary logic optimizations.

Floorplan location assignments are required when design blocks are developed independently, and are recommended for timing-critical partitions that are expected to change. Follow the guidelines to create and modify LogicLock regions to create good placement assignments for your design partitions.

Take advantage of the numerous Quartus II software tools to assess partition quality and analyze the floorplan to make good LogicLock location assignments. Remember that you do not have to follow all the guidelines exactly to implement an incremental compilation design flow, but following the guidelines can maximize your chances of success.

Document Revision History

Table 12-1 shows the revision history for this chapter.

Table 12-1. Document Revision History

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Changed to new document template. ■ Moved “Creating Floorplan Location Assignments With Tcl Commands—Excluding or Filtering Certain Device Elements (Such as RAM or DSP Blocks)” from the Quartus II Incremental Compilation for Hierarchical and Team-Based Design chapter in volume 1 of the <i>Quartus II Handbook</i>. ■ Consolidated Design Partition Planner and Incremental Compilation Advisor information between the Quartus II Incremental Compilation for Hierarchical and Team-Based Design and Best Practices for Incremental Compilation Partitions and Floorplan Assignments handbook chapters.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Removed the explanation of the “bottom-up design flow” where designers work completely independently, and replaced with Altera’s recommendations for team-based environments where partitions are developed in the same top-level project framework, plus an explanation of the bottom-up process for including independent partitions from third-party IP designers. ■ Expanded the Merge command explanation to explain how it now accommodates cross-partition boundary optimizations. ■ Restructured Altera recommendations for when to use a floorplan.
October 2009	9.1.0	<ul style="list-style-type: none"> ■ Redefined the bottom-up design flow as team-based and reorganized previous design flow examples to include steps on how to pass top-level design information to lower-level projects. ■ Added “Including SDC Constraints from Lower-Level Partitions for Third-Party IP Delivery” from the Quartus II Incremental Compilation for Hierarchical and Team-Based Design chapter in volume 1 of the <i>Quartus II Handbook</i>. ■ Reorganized the “Recommended Design Flows and Application Examples” section. ■ Removed HardCopy APEX and HardCopy Stratix Devices section.
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Added I/O register packing examples from <i>Incremental Compilation for Hierarchical and Team-Based Designs</i> chapter ■ Moved “Incremental Compilation Advisor” section ■ Added “Viewing Design Partition Planner and Floorplan Side-by-Side” section ■ Updated Figure 12-22 ■ Chapter 8 was previously Chapter 7 in software release 8.1.
November 2008	8.1.0	<ul style="list-style-type: none"> ■ Changed to 8-1/2 x 11 page size. No change to content.
May 2007	8.0.0	<ul style="list-style-type: none"> ■ Initial release.

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

As programmable logic devices become more complex and require increased performance, advanced design synthesis has become an important part of the design flow. In the Quartus® II software you can use the integrated Analysis and Synthesis module of the Compiler to synthesize your design files and create the project database for future stages of the compilation flow. You can also use other EDA synthesis tools to first synthesize your designs, and then generate EDIF netlist files or Verilog Quartus Mapping Files (.vqm) that you can use with the Quartus II software. The Quartus II netlist viewers allow you to visually analyze the design netlist at different stages of synthesis and compilation. This section explains the options that are available for each of these flows and how they are supported in the Quartus II software.

This section includes the following chapters:

- **Chapter 13, Quartus II Integrated Synthesis**

This chapter documents the integrated synthesis design flow and language support in the Quartus II software. It explains how you can improve synthesis results with Quartus II synthesis options and optimization techniques, and how you can control the inference of architecture-specific megafunctions. This chapter also explains some of the node-naming conventions used during synthesis to help you better understand your synthesized design and the messages issued during synthesis to improve your HDL code. Scripting techniques for applying all the options and settings described are also provided.

- **Chapter 14, Synopsys Synplify Support**

This chapter documents support for the Synopsys Synplify software in the Quartus II software, as well as key design flows, methodologies, and techniques for achieving good results in Altera® devices with the Synplify software.

- **Chapter 15, Mentor Graphics Precision Synthesis Support**

This chapter documents support for the Mentor Graphics® Precision Synthesis software in the Quartus II software, as well as key design flows, methodologies, and techniques for achieving good results in Altera® devices with the Precision Synthesis software.

- **Chapter 16, Mentor Graphics LeonardoSpectrum Support**

This chapter documents key design methodologies and techniques for Altera devices using the Mentor Graphics LeonardoSpectrum™ software and Quartus II design flow. The LeonardoSpectrum software is a mature synthesis tool supporting legacy devices and many current devices. Altera recommends using the advanced Precision Synthesis software for new designs in new device families.

- **Chapter 17, Analyzing Designs with Quartus II Netlist Viewers**

This chapter shows how to use the Quartus II netlist viewers to analyze your design at various stages of the design cycle. It also provides an introduction to the Quartus II design flow using netlist viewers, an overview of each viewer, and an explanation of the user interface.

This chapter documents the design flow of integrated synthesis and provides scripting techniques for applying all options and settings described in this chapter.

As programmable logic designs become more complex and require increased performance, advanced synthesis has become an important part of the design flow. The Altera® Quartus® II software includes advanced integrated synthesis that fully supports VHDL, Verilog HDL, and Altera-specific design entry languages, and provides options to control the synthesis process. With this synthesis support, the Quartus II software provides a complete, easy-to-use solution.

This chapter contains the following sections:

- “Design Flow” on page 13–1
- “Language Support” on page 13–3
- “Incremental Compilation” on page 13–20
- “Quartus II Synthesis Options” on page 13–22
- “Analyzing Synthesis Results” on page 13–72
- “Analyzing and Controlling Synthesis Messages” on page 13–72
- “Node-Naming Conventions in Quartus II Integrated Synthesis” on page 13–77
- “Scripting Support” on page 13–83



For examples of Verilog HDL and VHDL code synthesized for specific logic functions, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*. For information about coding with primitives that describe specific low-level functions in Altera devices, refer to the *Designing With Low-Level Primitives User Guide*.

Design Flow

The Quartus II Analysis & Synthesis stage of the compilation flow runs integrated synthesis, which fully supports Verilog HDL, VHDL, and Altera-specific languages, and major features of the SystemVerilog language. For more information, refer to “Language Support” on page 13–3.

In the synthesis stage of the compilation flow, the Quartus II software performs logic synthesis to optimize design logic and performs technology mapping to implement the design logic in device resources such as logic elements (LEs) or adaptive logic modules (ALMs), and other dedicated logic blocks. The synthesis stage generates a single project database that integrates all the design files in a project (including any netlists from third-party synthesis tools).

You can use Analysis & Synthesis to perform the following compilation processes:

- **Analyze Current File**—Parses the current design source file to check for syntax errors. This command does not report on many semantic errors that require further design synthesis. To perform this analysis, on the Processing menu, click **Analyze Current File**.
- **Analysis & Elaboration**—Checks a design for syntax and semantic errors and performs elaboration to identify the design hierarchy. To perform Analysis & Elaboration, on the Processing menu, point to **Start**, and then click **Start Analysis & Elaboration**.
- **Analysis & Synthesis**—Performs complete Analysis & Synthesis on a design, including technology mapping. To perform Analysis & Synthesis, on the Processing menu, point to **Start**, and then click **Start Analysis & Synthesis**.

The Quartus II integrated synthesis design and compilation flow consists of the following steps:

1. Create a project in the Quartus II software and specify the general project information, including the top-level design entity name.
2. Create design files in the Quartus II software or with a text editor.
3. On the Project menu, click **Add/Remove Files in Project** and add all design files to your Quartus II project using the **Files** page of the **Settings** dialog box.
4. Specify compiler settings that control the compilation and optimization of the design during synthesis and fitting. For synthesis settings, refer to “[Quartus II Synthesis Options](#)” on page 13-22.
5. Add timing constraints to specify the timing requirements.



To partition your design to reduce compilation time, refer to “[Incremental Compilation](#)” on page 13-20.

6. Compile the design. To synthesize the design, on the Processing menu, point to **Start**, and then click **Start Analysis & Synthesis**. To run a complete compilation flow including placement, routing, creation of a programming file, and timing analysis, click **Start Compilation** on the Processing menu.
7. After obtaining synthesis and placement and routing results that meet your requirements, program or configure your Altera device.

The Quartus II software generates netlists that enable you to perform functional simulation or gate-level timing simulation, timing analysis, and formal verification.



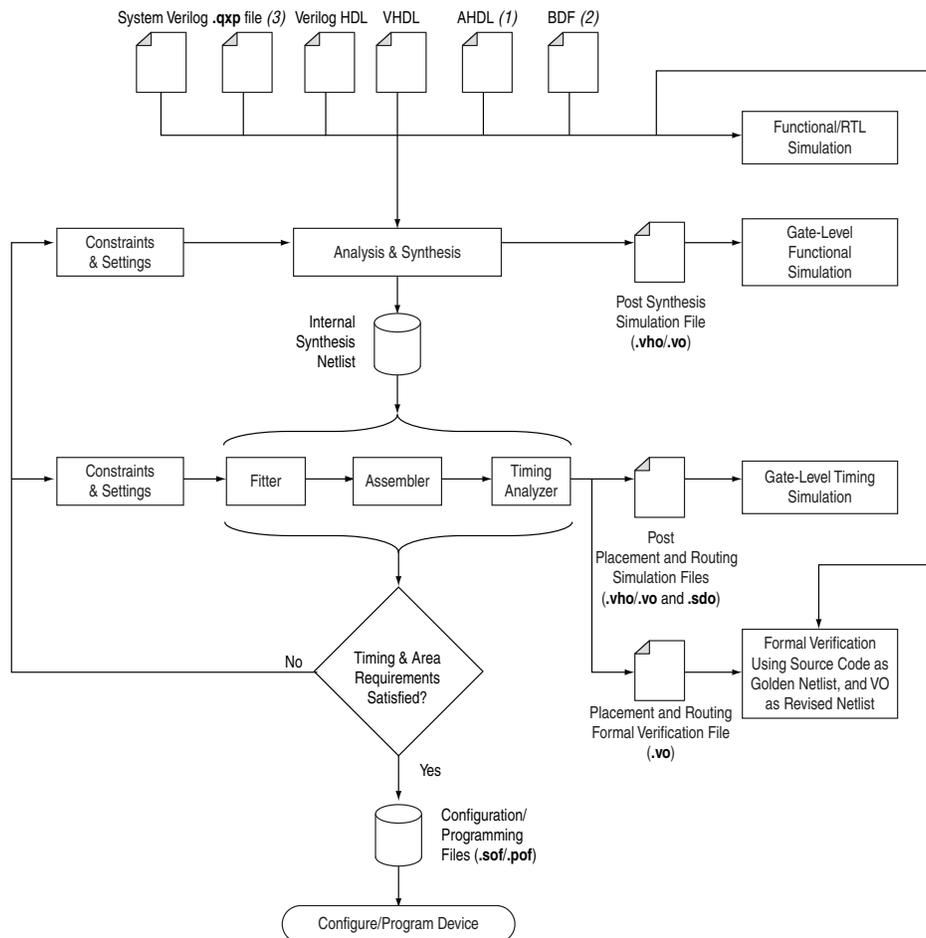
For an overall summary of features in the Quartus II software, refer to the [Introduction to the Quartus II Software](#) manual.



For more information about Quartus II projects and the compilation flow, refer to [Managing Files in a Project](#) and [About Compilation Flows](#) in Quartus II Help.

Figure 13-1 shows the basic design flow using Quartus II integrated synthesis.

Figure 13-1. Quartus II Design Flow Using Quartus II Integrated Synthesis



Notes to Figure 13-1:

- (1) AHDL stands for the Altera Hardware Description Language.
- (2) BDF stands for the Altera schematic Block Design File format (.bdf).
- (3) The Quartus II Exported Partition File (.qxp) is a precompiled netlist that you can use as a design source file. For more information about using .qxp as a design source file, refer to “Quartus II Exported Partition File as Source” on page 13-22.

Language Support

This section explains Quartus II integrated synthesis support for HDL, schematic design entry, graphical state machine entry, and how to specify the Verilog HDL or VHDL language version used in your design. This section also documents language features such as Verilog HDL macros, initial constructs and memory system tasks, and VHDL libraries. “Design Libraries” on page 13-12 describes how to compile and reference design units in custom libraries, and “Using Parameters/Generics” on page 13-16 describes how to use parameters or generics and pass them between languages.

To ensure that the Quartus II software reads all associated project files, add each file to your Quartus II project by clicking **Add/Remove Files in Project** on the Project menu. You can add design files to your project. You can mix all supported languages and netlists generated by third-party synthesis tools in a single Quartus II project.

Verilog HDL Support

The Quartus II Compiler's Analysis & Synthesis module supports the following Verilog HDL standards:

- Verilog-1995 (IEEE Standard 1364-1995)
- Verilog-2001 (IEEE Standard 1364-2001)
- SystemVerilog-2005 (IEEE Standard 1800-2005) (not all constructs are supported)



The Quartus II software does not support Verilog-2001 libraries and configurations.

- ② For more information about Verilog HDL, refer to *About Verilog HDL* in Quartus II Help.

The Verilog HDL code samples provided in this document follow the Verilog-2001 standard unless otherwise specified. The Quartus II Compiler uses the Verilog-2001 standard by default for files that have the extension `.v`, and the SystemVerilog standard for files that have the extension `.sv`.

- ② For more information about Quartus II Verilog HDL support, refer to *Quartus II Verilog HDL Support* in Quartus II Help.

You can specify a default Verilog HDL version for all files by performing the following steps:

1. On the Assignments menu, click **Settings**.
2. In the **Settings** dialog box, under **Category**, expand **Analysis & Synthesis Settings**, and select **Verilog HDL Input**.
3. On the **Verilog HDL Input** page, under **Verilog version**, select the appropriate Verilog HDL version and then click **OK**.

You can override the default Verilog HDL version for each Verilog HDL design file by following these steps:

1. On the Project menu, click **Add/Remove Files in Project**.
2. On the **Files** page, select the appropriate file in the list, and then click **Properties**.
3. In the **HDL Version** list, select **SystemVerilog_2005**, **Verilog_2001**, or **Verilog_1995**, and then click **OK**.

You can control the Verilog HDL version that compiles the design in a design file with the `VERILOG_INPUT_VERSION` synthesis directive, as shown in [Example 13-1](#). This directive overrides the default HDL version and any HDL version specified in the **File Properties** dialog box.

Example 13-1. Controlling the Verilog HDL Input Version with a Synthesis Directive

```
// synthesis VERILOG_INPUT_VERSION <language version>
```

The variable `<language version>` uses one of the following values:

- `VERILOG_1995`
- `VERILOG_2001`
- `SYSTEMVERILOG_2005`

When the Quartus II software reads a `VERILOG_INPUT_VERSION` synthesis directive, the current language version setting changes as specified until after the file or until the next `VERILOG_INPUT_VERSION` directive.

 You cannot change the language version in the middle of a Verilog HDL module.

For more information about specifying synthesis directives, refer to [“Synthesis Directives”](#) on page 13-27.

- ② For more information about Verilog HDL synthesis attributes and directives, refer to [Verilog HDL Synthesis Attributes and Directives](#) in Quartus II Help.

If you use scripts to add design files, you can use the `-HDL_VERSION` command to specify the HDL version for each design file. For more information, refer to [“Adding an HDL File to a Project and Setting the HDL Version”](#) on page 13-84.

The Quartus II software support for Verilog HDL is case-sensitive in accordance with the Verilog HDL standard. The Quartus II software supports the compiler directive ``define`, in accordance with the Verilog HDL standard.

The Quartus II software supports the `include` compiler directive to include files with absolute paths (with either `“/”` or `“\”` as the separator), or relative paths. When searching for a relative path, the Quartus II software initially searches relative to the project directory. If the Quartus II software cannot find the file, the software then searches relative to all user libraries, and finally, relative to the directory location of the current file.

SystemVerilog Support

The Quartus II software supports the following SystemVerilog constructs:

- Parameterized interfaces, generic interfaces, and `modport` constructs
- Packages
- Extern module declarations
- Built-in data types `logic`, `bit`, `byte`, `shortint`, `longint`, `int`
- Unsized integer literals ``0`, ``1`, ``x`, ``z`, ``X`, and ``Z`

- Structure data types using `struct`
- Ports and parameters with unrestricted data types
- Types you defined using `typedef`
- Global declarations of task/functions/parameters/types (does not support global variables)
- Coding constructs `always_comb`, `always_latch`, `always_ff`
- Continuous assignments to nodes other than `nets`, and procedural assignments to nodes other than `reg`
- Enumeration methods `First`, `Last`, `Next(n)`, `Prev(n)`, `Num`, and `Name`
- Assignment operators `+=`, `-=`, `*=`, `/=`, `%=`, `&=`, `|=`, `^=`, `<<=`, `>>=`, `<<<=`, and `>>>=`
- Increment `++` and decrement `--`
- Jump statements `return`, `break`, and `continue`
- Enhanced for loop (declare loop variables inside initial condition)



Quartus II integrated synthesis does not support multiple loop variable declaration.

- Do-while loop and local loop constructs
- Assignment patterns
- Keywords `unique` and `priority` in case statements
- Default values for function/task arguments
- Closing labels
- Extensions to directives ``define` and ``include`
- Expression size system function `$bits`
- Array query system functions `$dimensions`, `$unpacked_dimensions`, `$left`, `$right`, `$high`, `$low`, `$increment`, and `$size`
- Packed array (include multidimensional packed array and packed arrays of packed structures)
- Unpacked array (include single-valued range dimension)
- Implicit port connections with `.name` and `.*`
- Immediate assertions in procedural constructs
- System task such as `$fatal`, `$error`, `$warning`, `$info`, and `$display` statements
- Implicit port and argument direction for functions and tasks (Implicit port direction in modules is not supported)
- Declaration of time unit and time precision with `timeunit` and `timeprecision` in a module

Quartus II integrated synthesis also parses, but otherwise ignores the SystemVerilog assertions.

 Designs written to comply with the Verilog-2001 standard might not compile with the SystemVerilog setting because the SystemVerilog standard adds several new reserved keywords.

Initial Constructs and Memory System Tasks

The Quartus II software infers power-up conditions from Verilog HDL `initial` constructs. The Quartus II software also creates power-up settings for variables, including RAM blocks. If the Quartus II software encounters non-synthesizable constructs in an `initial` block, it generates an error. To avoid such errors, enclose non-synthesizable constructs (such as those intended only for simulation) in `translate_off` and `translate_on` synthesis directives, as described in “[Translate Off and On / Synthesis Off and On](#)” on page 13–63. Synthesis of initial constructs enables the power-up state of the synthesized design to match, as closely as possible, the power-up state of the original HDL code in simulation. For more information, refer to “[Power-Up Level](#)” on page 13–41.

 Initial blocks do not infer power-up conditions in some third-party EDA synthesis tools. If you convert between synthesis tools, you must set your power-up conditions correctly.

Quartus II integrated synthesis supports the `$readmemb` and `$readmemh` system tasks to initialize memories. [Example 13–2](#) shows an initial construct that initializes an inferred RAM with `$readmemb`.

Example 13–2. Verilog HDL Code: Initializing RAM with the `readmemb` Command

```
reg [7:0] ram[0:15];
initial
begin
  $readmemb("ram.txt", ram);
end
```

When creating a text file to use for memory initialization, specify the address using the format `@<location>` on a new line, and then specify the memory word such as 110101 or abcde on the next line. [Example 13–3](#) shows a portion of a Memory Initialization File (`.mif`) for the RAM in [Example 13–2](#).

Example 13–3. Text File Format: Initializing RAM with the `readmemb` Command

```
@0
00000000
@1
00000001
@2
00000010
...
@e
00001110
@f
00001111
```

Verilog HDL Macros

The Quartus II software fully supports Verilog HDL macros, which you can define with the `'define` compiler directive in your source code. You can also define macros in the GUI or on the command line.

Setting a Verilog HDL Macro Default Value in the GUI

To specify a macro in the GUI, follow these steps:

1. On the Assignments menu, click **Settings**.
2. In the **Category** list, expand **Analysis & Synthesis Settings** and select **Verilog HDL Input**.
3. Under **Verilog HDL macro**, type the macro name in the **Name** box and the value in the **Setting** box.
4. Click **Add**.

Setting a Verilog HDL Macro Default Value on the Command Line

To set a default value for a Verilog HDL macro on the command line, use the `--verilog_macro` option, as shown in [Example 13-4](#).

Example 13-4. Command Syntax for Specifying a Verilog HDL Macro

```
quartus_map <Design name> --verilog_macro= "<Macro name>=<Macro setting>" ←
```

The command in [Example 13-5](#) has the same effect as specifying ``define a 2` in the Verilog HDL source code.

Example 13-5. Specifying a Verilog HDL Macro `a = 2`

```
quartus_map my_design --verilog_macro="a=2" ←
```

To specify multiple macros, you can repeat the option more than once, as in [Example 13-6](#).

Example 13-6. Specifying Verilog HDL Macros `a = 2` and `b = 3`

```
quartus_map my_design --verilog_macro="a=2" --verilog_macro="b=3" ←
```

VHDL Support

The Quartus II Compiler's Analysis & Synthesis module supports the following VHDL standards:

- VHDL 1987 (IEEE Standard 1076-1987)
- VHDL 1993 (IEEE Standard 1076-1993)
- VHDL 2008 (IEEE Standard 1076-2008)

The Quartus II Compiler uses the VHDL 1993 standard by default for files that have the extension `.vhdl` or `.vhd`.

 The VHDL code samples provided in this chapter follow the VHDL 1993 standard.

To specify a default VHDL version for all files, follow these steps:

1. On the Assignments menu, click **Settings**.
2. In the **Category** list, expand **Analysis & Synthesis Settings** and select **VHDL Input**.
3. On the **VHDL Input** page, under **VHDL version**, select the appropriate version, and then click **OK**.

You can override the default VHDL version for each VHDL design file by performing the following steps:

1. On the Project menu, click **Add/Remove Files in Project**.
2. On the **Files** page, select the appropriate file in the list, and then click **Properties**.
3. In the HDL version list, select **VHDL_2008**, **VHDL_1993**, or **VHDL_1987**, and then click **OK**.

You can also specify the VHDL version that compiles the design for each design file with the `VHDL_INPUT_VERSION` synthesis directive, as shown in [Example 13-7](#). This directive overrides the default HDL version and any HDL version specified in the **File Properties** dialog box.

Example 13-7. Controlling the VHDL Input Version with a Synthesis Directive

```
--synthesis VHDL_INPUT_VERSION <language version>
```

Example 13-8. VHDL 2008—Controlling the VHDL Input Version with a Synthesis Directive

```
/* synthesis VHDL_INPUT_VERSION <language version> */
```

The variable `<language version>` requires one of the following values:

- `VHDL_1987`
- `VHDL_1993`
- `VHDL_2008`

When the Quartus II software reads a `VHDL_INPUT_VERSION` synthesis directive, it changes the current language version as specified until after the file or until it reaches the next `VHDL_INPUT_VERSION` directive.



You cannot change the language version in a VHDL design unit.

For more information about specifying synthesis directives, refer to [“Synthesis Directives”](#) on page 13-27.

If you use scripts to add design files, you can use the `-HDL_VERSION` command to specify the HDL version for each design file. For more information, refer to [“Adding an HDL File to a Project and Setting the HDL Version”](#) on page 13-84.

The Quartus II software reads default values for registered signals defined in the VHDL code and converts the default values into power-up level settings. This enables the power-up state of the synthesized design to match, as closely as possible, the power-up state of the original HDL code in simulation. For more information, refer to [“Power-Up Level”](#) on page 13-41.

VHDL-2008 Support

The Quartus II software supports the following VHDL 2008 constructs:

- Block comments
- Simplified sensitivity lists
- Extensions to generate
- Matching case statement
- Matching equality/inequality operators
- Enhanced bit string literals
- Unconstrained elements in arrays
- Condition operator (explicit and implicit)

VHDL Standard Libraries and Packages

The Quartus II software includes the standard IEEE libraries and several vendor-specific VHDL libraries. For information about organizing your own design units into custom libraries, refer to “[Design Libraries](#)” on page 13-12.

The **IEEE** library includes the standard VHDL packages `std_logic_1164`, `numeric_std`, `numeric_bit`, and `math_real`. The **STD** library is part of the VHDL language standard and includes the packages `standard` (included in every project by default) and `textio`. For compatibility with older designs, the Quartus II software also supports the following vendor-specific packages and libraries:

- Synopsys packages such as `std_logic_arith` and `std_logic_unsigned` in the **IEEE** library
- Mentor Graphics® packages such as `std_logic_arith` in the **ARITHMETIC** library
- Altera primitive packages `altera_primitives_components` (for primitives such as `GLOBAL` and `DFFE`) and `maxplus2` (for legacy support of MAX+PLUS® II primitives) in the **ALTERA** library
- Altera megafunction packages `altera_mf_components` and `stratixgx_mf_components` in the **ALTERA_MF** library (for Altera-specific megafunctions including `LCELL`), and `lpm_components` in the **LPM** library for library of parameterized modules (LPM) functions.



Altera recommends that you import component declarations for Altera primitives such as `GLOBAL` and `DFFE` from the `altera_primitives_components` package and not the `altera_mf_components` package.

VHDL wait Constructs

The Quartus II software supports one VHDL `wait until` statement per process block. However, the Quartus II software does not support other VHDL wait constructs, such as `wait for` and `wait on` statements, or processes with multiple wait statements.

Example 13-9 is a VHDL code example of a supported `wait until` construct.

Example 13-9. VHDL Code: Supported wait until Construct

```
architecture dff_arch of ls_dff is
begin
output: process begin
wait until (CLK'event and CLK='1');
Q <= D;
Qbar <= not D;
end process output;
end dff_arch;
```

AHDL Support

The Quartus II Compiler's Analysis & Synthesis module fully supports the Altera Hardware Description Language (AHDL).

AHDL designs use Text Design Files (**.tdf**). You can import AHDL Include Files (**.inc**) into a **.tdf** file with an AHDL `include` statement. Altera provides **.inc** files for all megafunctions shipped with the Quartus II software.

 The AHDL language does not support the synthesis directives or attributes described in this chapter.

 For more information about AHDL, refer to *About AHDL* in the Quartus II Help.

Schematic Design Entry Support

The Quartus II Compiler's Analysis & Synthesis module fully supports **.bdf** for schematic design entry.

You can use the Quartus II Block Editor to create and edit **.bdf** files and open Graphic Design Files (**.gdf**) imported from the MAX+PLUS II software. Use the Symbol Editor to create and edit Block Symbol Files (**.bsf**) and open MAX+PLUS II Symbol Files (**.sym**). You can read and edit these legacy MAX+PLUS II formats with the Quartus II Block and Symbol Editors; however, the Quartus II software saves them as **.bdf** or **.bsf** files.

 For information about creating and editing schematic designs, refer to the *About Schematic Design Entry* in Quartus II Help.

 Schematic entry methods do not support the synthesis directives or attributes in this chapter.

State Machine Editor

The Quartus II software supports graphical state machine entry. To create a new finite state machine (FSM) design, on the File menu, click **New**. In the **New** dialog box, expand the **Design Files** list, and then select **State Machine File**.

In the editor, you can use the State Machine Wizard to step you through the state machine creation. Click the **State Machine Wizard** icon. Specify the reset information, define the input ports, states, and transitions, and then define the output ports and output conditions. Click **Finish** to create the state machine diagram.

You can also create the state machine diagram with the editor GUI. Use the icons or right-click menu options to insert new input and output signals and create states in the schematic display. To specify transitions, select the **Transition Tool**, click on the source state, and then drag the mouse to the destination state. Double-click on a transition to specify the transition equation, with a syntax that conforms to Verilog HDL. Double-click on a state to open the **State Properties** dialog box, in which you can change the state name, specify whether it acts as the reset state, and change the incoming and outgoing transition equations.

To view and edit state machine information in a table format, click the **State Machine Table** icon.

The Quartus II software saves the state machine diagram as a State Machine File (.smf). After defining the state machine logic, you can create a Verilog HDL or VHDL design file by clicking the **Generate HDL File** icon. You can then instantiate the state machine in your design using any design entry language.

- ❓ For more information about creating and editing state machine diagrams, and a list of supported operators in transition equations syntax, refer to *Creating and Editing State Machines with the State Machine Editor* in Quartus II Help.

Design Libraries

By default, the Quartus II software compiles all design files into the work library. If you do not specify a design library, or if a file refers to a library that does not exist, or if the referenced library does not contain a referenced design unit, the Quartus II software searches the work library. This behavior allows the Quartus II software to compile most designs with minimal setup, while creating separate custom design libraries is optional.

To compile your design files into specific libraries (for example, when you have two or more functionally different design entities that share the same name), you can specify a destination library for each design file in various ways, as described in the following subsections:

- “Specifying a Destination Library Name in the Settings Dialog Box” on page 13-13
- “Specifying a Destination Library Name in the Quartus II Settings File or Using Tcl” on page 13-13

When the Quartus II Compiler analyzes the file, it stores the analyzed design units in the destination library of the file.



A design can contain two or more entities with the same name if the Quartus II software compiles the entities into separate libraries.

When compiling a design instance, the Quartus II software initially searches for the entity in the library associated with the instance (which is the work library if no other library is specified). If the Quartus II software could not locate the entity definition, the software searches for a unique entity definition in all design libraries. If the Quartus II software finds more than one entity with the same name, the software generates an error. If your design uses multiple entities with the same name, you must compile the entities into separate libraries.

In VHDL, you can associate an instance with a particular entity in several ways, as described in “[Mapping a VHDL Instance to an Entity in a Specific Library](#)” on [page 13-14](#). In Verilog HDL, BDF schematic entry, AHDL, VQM and EDIF netlists, you can use different libraries for each of the entities that have the same name, and compile the instantiation into the same library as the appropriate entity.

Specifying a Destination Library Name in the Settings Dialog Box

To specify a library name for one of your design files, follow these steps:

1. On the Assignments menu, click **Settings**.
2. In the **Category** list, select **Files**.
3. Select the file in the **File Name** list.
4. Click **Properties**.
5. In the **File Properties** dialog box, select the type of design file from the **Type** list.
6. Type the library name in the **Library** field.
7. Click **OK**.

Specifying a Destination Library Name in the Quartus II Settings File or Using Tcl

You can specify the library name with the `-library` option to the `<language type>_FILE` assignment in the Quartus II Settings File (`.qsf`) or with Tcl commands.

For example, the following assignments specify that the Quartus II software analyzes the `my_file.vhd` and stores its contents (design units) in the VHDL library `my_lib`, and then analyzes the Verilog HDL file `my_header_file.h` and stores its contents in a library called `another_lib`. Refer to [Example 13-10](#).

Example 13-10. Specifying a Destination Library Name

```
set_global_assignment -name VHDL_FILE my_file.vhd -library my_lib
set_global_assignment -name VERILOG_FILE my_header_file.h -library another_lib
```

For more information about Tcl scripting, refer to “[Scripting Support](#)” on [page 13-83](#).

Specifying a Destination Library Name in a VHDL File

You can use the `library` synthesis directive to specify a library name in your VHDL source file. This directive takes the name of the destination library as a single string argument. Specify the `library` directive in a VHDL comment before the context clause for a primary design unit (that is, a package declaration, an entity declaration, or a configuration), using one of the supported keywords for synthesis directives, that is, `altera`, `synthesis`, `pragma`, `synopsys`, or `exemplar`.

For more information about specifying synthesis directives, refer to “[Synthesis Directives](#)” on [page 13-27](#).

The `library` directive overrides the default library destination `work`, the library setting specified for the current file in the **Settings** dialog box, any existing QSF setting, any setting made through the Tcl interface, or any prior `library` directive in the current file. The directive remains effective until the end of the file or the next `library` synthesis directive.

Example 13-11 uses the `library` synthesis directive to create a library called `my_lib` that contains the design unit `my_entity`.

Example 13-11. Using the Library Synthesis Directive

```
-- synthesis library my_lib
library ieee;
use ieee.std_logic_1164.all;
entity my_entity(...)
end entity my_entity;
```



You can specify a single destination library for all the design units in a given source file by specifying the library name in the **Settings** dialog box, editing the `.qsf`, or using the Tcl interface. To organize the design units in a single file into different libraries, rather than just a single library, you can use the `library` directive to change the destination VHDL library in a source file.

The Quartus II software generates an error if you use the `library` directive in a design unit.

Mapping a VHDL Instance to an Entity in a Specific Library

The VHDL language provides several ways to map or bind an instance to an entity in a specific library, as described in the following subsections.

Direct Entity Instantiation

In the direct entity instantiation method, the instantiation refers to an entity in a specific library, as shown in **Example 13-12**.

Example 13-12. VHDL Code: Direct Entity Instantiation

```
entity entity1 is
port(...);
end entity entity1;

architecture arch of entity1 is
begin
inst: entity lib1.foo
port map(...);
end architecture arch;
```

Component Instantiation—Explicit Binding Instantiation

You can bind a component to an entity in several mechanisms. In an explicit binding indication, you bind a component instance to a specific entity, as shown in [Example 13-13](#).

Example 13-13. VHDL Code: Binding Instantiation

```
entity entity1 is
port(...);
end entity entity1;

package components is
component entity1 is
port map (...);
end component entity1;
end package components;

entity top_entity is
port(...);
end entity top_entity;

use lib1.components.all;
architecture arch of top_entity is
-- Explicitly bind instance I1 to entity1 from lib1
for I1: entity1 use entity lib1.entity1
port map(...);
end for;
begin
I1: entity1 port map(...);
end architecture arch;
```

Component Instantiation—Default Binding

If you do not provide an explicit binding indication, the Quartus II software binds a component instance to the nearest visible entity with the same name. If no such entity is visible in the current scope, the Quartus II software binds the instance to the entity in the library in which the component was declared. For example, if the component is declared in a package in library MY_LIB, an instance of the component binds to the entity in library MY_LIB. The portions of code in [Example 13-14](#) and [Example 13-15](#) show this instantiation method:

Example 13-14. VHDL Code: Default Binding to the Entity in the Same Library as the Component Declaration

```
use mylib.pkg.foo; -- import component declaration from package "pkg" in
-- library "mylib"
architecture rtl of top
...
begin
-- This instance will be bound to entity "foo" in library "mylib"
inst: foo
port map(...);
end architecture rtl;
```

Example 13-15. VHDL Code: Default Binding to the Directly Visible Entity

```
use mylib.foo; -- make entity "foo" in library "mylib" directly visible
architecture rtl of top
component foo is
generic (...)
port (...);
end component;
begin
-- This instance will be bound to entity "foo" in library "mylib"
inst: foo
port map(...);
end architecture rtl;
```

Using Parameters/Generics

This section describes how the Quartus II software supports parameters (known as generics in VHDL) and how you can pass these parameters between design languages.

You can enter default parameter values for your design in the **Default Parameters** page under the **Analysis & Synthesis Settings** page in the **Settings** dialog box. Default parameters enable you to add, change, and delete global parameters for the current assignment. In AHDL, the Quartus II software inherits parameters, so any default parameters apply to all AHDL instances in the design. You can also specify parameters for instantiated modules in a **.bdf**. To modify parameters in a **.bdf** instance, double-click on the parameter value box for the instance symbol, or right-click on the symbol and choose **Properties**, and then click the **Parameters** tab. For more information about the methods of GUI-based entry, the interpretation of parameter values, and format recommendations, refer to [“Setting Default Parameter Values and BDF Instance Parameter Values” on page 13-16](#).

You can specify parameters for instantiated modules in your design source files with the syntax provided for the language you choose. Some designs instantiate entities in a different language; for example, they might instantiate a VHDL entity from a Verilog HDL design file. You can pass parameters or generics between VHDL, Verilog HDL, AHDL, and BDF schematic entry, and from EDIF or VQM to any of these languages. You do not require an additional procedure to pass parameters from one language to another. However, sometimes you must specify the type of parameter you are passing. In those cases, you must follow certain guidelines to ensure that the Quartus II software correctly interprets the parameter value. For more information about parameter type rules, refer to [“Passing Parameters Between Two Design Languages” on page 13-18](#).

Setting Default Parameter Values and BDF Instance Parameter Values

Default parameter values and BDF instance parameter values do not have an explicitly declared type. Usually, the Quartus II software can correctly infer the type from the value without ambiguity. For example, the Quartus II software interprets “ABC” as a string, 123 as an integer, and 15.4 as a floating-point value. In other cases, such as when the instantiated subdesign language is VHDL, the Quartus II software uses the type of the parameter, generic, or both in the instantiated entity to determine how to interpret the value, so that the Quartus II software interprets a value of 123 as

a string if the VHDL parameter is of a type string. In addition, you can set the parameter value in a format that is legal in the language of the instantiated entity. For example, to pass an unsized bit literal value from **.bdf** to Verilog HDL, you can use '1 as the parameter value, and to pass a 4-bit binary vector from **.bdf** to Verilog HDL, you can use 4'b1111 as the parameter value.

In a few cases, the Quartus II software cannot infer the correct type of parameter value. To avoid ambiguity, specify the parameter value in a type-encoded format in which the first or first and second characters of the parameter indicate the type of the parameter, and the rest of the string indicates the value in a quoted sub-string. For example, to pass a binary string 1001 from **.bdf** to Verilog HDL, you cannot use the value 1001, because the Quartus II software interprets it as a decimal value. You also cannot use the string "1001" because the Quartus II software interprets it as an ASCII string. You must use the type-encoded string B"1001" for the Quartus II software to correctly interpret the parameter value. Table 13-1 provides a list of valid parameter strings, and shows how the Quartus II software interprets these parameter strings. Use the type-encoded format only when necessary to resolve ambiguity.

Table 13-1. Valid Parameter Strings and Interpretations

Parameter String	Quartus II Parameter Type, Format, and Value
S"abc", s"abc"	String value abc
"abc123", "123abc"	String value abc123 or 123abc
F"12.3", f"12.3"	Floating point number 12.3
-5.4	Floating point number -5.4
D"123", d"123"	Decimal number 123
123, -123	Decimal number 123, -123
X"ff", H"ff"	Hexadecimal value FF
Q"77", O"77"	Octal value 77
B"1010", b"1010"	Unsigned binary value 1010
SB"1010", sb"1010"	Signed binary value 1010
R"1", R"0", R"X", R"Z", r"1", r"0", r"X", r"Z"	Unsized bit literal
E"apple", e"apple"	Enumeration type, value name is apple
P"1 unit"	Physical literal, the value is (1, unit)
A(...), a(...)	Array type or record type, whose content is determined by the string (...)

You can select the parameter type for global parameters or global constants with the pull-down list in the **Parameter** tab of the **Symbol Properties** dialog box. If you do not specify the parameter type, the Quartus II software interprets the parameter value and defines the parameter type. You must specify parameter type with the pull-down list to avoid ambiguity.



If you open a **.bdf** in the Quartus II software, the software automatically updates the parameter types of old symbol blocks by interpreting the parameter value based on the language-independent format. If the Quartus II software does not recognize the parameter value type, the software sets the parameter type as **untyped**.

The Quartus II software supports the following parameter types:

- **Unsigned Integer**
- **Signed Integer**
- **Unsigned Binary**
- **Signed Binary**
- **Octal**
- **Hexadecimal**
- **Float**
- **Enum**
- **String**
- **Boolean**
- **Char**
- **Untyped/Auto**

Passing Parameters Between Two Design Languages

When passing a parameter between two different languages, a design block that is higher in the design hierarchy instantiates a lower-level subdesign block and provides parameter information. The subdesign language (the design entity that is instantiated) must correctly interpret the parameter. Based on the information provided by the higher-level design and the value format, and sometimes by the parameter type of the subdesign entity, the Quartus II software interprets the type and value of the passed parameter.

When passing a parameter whose value is an enumerated type value or literal from a language that does not support enumerated types to one that does (for example, from Verilog HDL to VHDL), you must ensure that the enumeration literal is in the correct spelling in the language of the higher-level design block (block that is higher in the hierarchy). The Quartus II software passes the parameter value as a string literal, and the language of the lower-level design correctly convert the string literal into the correct enumeration literal.

If the language of the lower-level entity is SystemVerilog, you must ensure that the enum value is in the correct case. In SystemVerilog, two enumeration literals differ in more than just case. For example, `enum {item, ITEM}` is not a good choice of item names because these names can create confusion and is more difficult to pass parameters from case-insensitive HDLs, such as VHDL.

Arrays have different support in different design languages. For details about the array parameter format, refer to the **Parameter** section in the Analysis & Synthesis Report of a design that contains array parameters or generics.

The following code shows examples of passing parameters from one design entry language to a subdesign written in another language. [Example 13-16](#) shows a VHDL subdesign that is instantiated in a top-level Verilog HDL design in [Example 13-17](#). [Example 13-18](#) shows a Verilog HDL subdesign that is instantiated in a top-level VHDL design in [Example 13-19](#).

Example 13-16. VHDL Parameterized Subdesign Entity

```
type fruit is (apple, orange, grape);
entity vhdl_sub is
generic (
name : string := "default",
width : integer := 8,
number_string : string := "123",
f : fruit := apple,
binary_vector : std_logic_vector(3 downto 0) := "0101",
signed_vector : signed (3 downto 0) := "1111");
```

Example 13-17. Verilog HDL Top-Level Design Instantiating and Passing Parameters to VHDL Entity from [Example 13-16](#)

```
vhdl_sub inst (...);
defparam inst.name = "lower";
defparam inst.width = 3;
defparam inst.num_string = "321";
defparam inst.f = "grape"; // Must exactly match enum value
defparam inst.binary_vector = 4'b1010;
    defparam inst.signed_vector = 4'sb1010;
```

Example 13-18. Verilog HDL Parameterized Subdesign Module

```
module veri_sub (... )
parameter name = "default";
parameter width = 8;
parameter number_string = "123";
parameter binary_vector = 4'b0101;
parameter signed_vector = 4'sb1111;
```

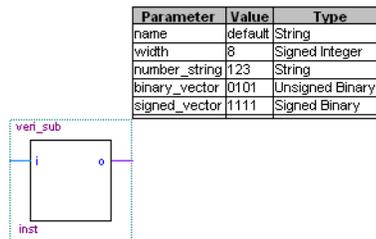
Example 13-19. VHDL Top-Level Design Instantiating and Passing Parameters to the Verilog HDL Module from [Example 13-18](#)

```
inst:veri_sub
generic map (
name => "lower",
width => 3,
number_string => "321"
binary_vector = "1010"
signed_vector = "1010")
```

To use an HDL subdesign such as the one shown in [Example 13-18](#) in a top-level .bdf design, you must generate a symbol for the HDL file, as shown in [Figure 13-2](#). Open the HDL file in the Quartus II software, and then, on the File menu, point to **Create/Update**, and then click **Create Symbol Files for Current File**.

To modify parameters on a **.bdf** instance, double-click on the parameter value box for the instance symbol, or right-click on the symbol and choose **Properties**, and then click the **Parameters** tab. Right-click on the symbol and choose **Update Design File from Selected Block** to pass the updated parameter to the HDL file.

Figure 13-2. BDF Top-Level Design Instantiating and Passing Parameters to the Verilog HDL Module from Example 13-18



Incremental Compilation

Incremental compilation manages a design hierarchy for incremental design by allowing you to divide the design into multiple partitions. Incremental compilation ensures that the Quartus II software resynthesizes only the updated partitions of the design during compilation, to reduce the compilation time and the runtime memory usage. The feature maintains node names during synthesis for all registered and combinational nodes in unchanged partitions. You can perform incremental synthesis by setting the Netlist Type for all design partitions to **Post-Synthesis**.

You can also preserve the placement and routing information for unchanged partitions. This feature allows you to preserve performance of unchanged blocks in your design and reduces the time required for placement and routing, which significantly reduces your design compilation time.

❓ For more information about incremental compilation, refer to *About Incremental Compilation* in Quartus II Help.

🔗 For more information about incremental compilation, refer to *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* and *Best Practices for Incremental Compilation Partitions and Floorplan Assignments* chapters in volume 1 of the *Quartus II Handbook*.

Partitions for Preserving Hierarchical Boundaries

A design partition represents a portion of the design that you want to synthesize and fit incrementally.

If you want to preserve the **Optimization Technique** and **Restructure Multiplexers** logic options set in any entity, you must create new partitions for the particular entity instead of using the **Preserve Hierarchical Boundary** logic option. If you have settings applied to specific existing design hierarchies, particularly those created in the Quartus II software versions before 9.0, you must create a design partition for the design hierarchy so that synthesis can optimize the design instance independently and preserve the hierarchical boundaries.

The **Preserve Hierarchical Boundary** logic option is available only in Quartus II software versions 8.1 and earlier. Altera recommends using design partitions if you want to preserve hierarchical boundaries through the synthesis and fitting process, because incremental compilation maintains the hierarchical boundaries of design partitions.

Parallel Synthesis

Parallel Synthesis one of the Analysis & Synthesis options that reduces compilation time for synthesis. The option enables the Quartus II software to use multiple processors to synthesize multiple partitions in parallel.

This option is available when you perform one or more of the following tasks:

- The number of processors allowed in a single machine is greater than 1. You can specify the maximum number of processors allowed under **Parallel Compilation** options in the **Compilation Process Settings** page of the **Settings** dialog box.
- Ensure that the incremental compilation feature is enabled.
- Your design has two or more partitions.
- You must turn on the **Parallel Synthesis** option.

By default, the Quartus II software enables the **Parallel Synthesis** option. To disable parallel synthesis, follow these steps:

1. On the Assignments menu, click **Settings**.
2. In the **Category** list, click **Analysis & Synthesis Settings**, and then click **More Settings** to select **Parallel Synthesis**.

You can also set the **Parallel Synthesis** option with the following Tcl command, as shown in [Example 13-20](#):

Example 13-20. Setting the Parallel Synthesis option with Tcl Command

```
set_global_assignment -name parallel_synthesis off
```

You can view all the interleaved messages from different partitions in the Messages window. The **Partition column** in the Messages window displays the partition ID of the partition referred to in the message. After compilation, you can sort the messages by partition.

- ④ For more information about displaying the **Partition** column, refer to [About the Messages Window](#) in Quartus II Help.

If you use the command line, you can differentiate among the interleaved messages by turning on the **Show partition that generated the message** option in the Messages page. This option shows the partition ID in parenthesis for each message.

Quartus II Exported Partition File as Source

You can use a **.qxp** as a source file in incremental compilation. The **.qxp** contains the precompiled design netlist exported as a partition from another Quartus II project, and fully defines the entity. Project team members or intellectual property (IP) providers can use a **.qxp** to send their design to the project lead, instead of sending the original HDL source code. The **.qxp** preserves the compilation results and instance-specific assignments. Not all global assignments can function in a different Quartus II project. You can override the assignments for the entity in the **.qxp** by applying assignments in the top-level design.

 For more information about **.qxp**, refer to *Quartus II Exported Partition File (.qxp)* in Quartus II Help.

 For more information about exporting design partitions and using **.qxp** files, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.

Quartus II Synthesis Options

The Quartus II software offers several options to help you control the synthesis process and achieve optimal results for your design. “[Setting Synthesis Options](#)” on [page 13-24](#) describes the **Analysis & Synthesis Settings** page of the **Settings** dialog box, in which you can set the most common global settings and options, and defines the following types of synthesis options: Quartus II logic options, synthesis attributes, and synthesis directives.

 When you apply a Quartus II Synthesis option globally or to an entity, the option affects all lower-level entities in the hierarchy path, including entities instantiated with Altera and third-party IP.

The following subsections describe the following common synthesis options in the Quartus II software, and provide HDL examples on how to use each option:

- Major Optimization Settings:
 - “[Optimization Technique](#)” on [page 13-28](#)
 - “[Auto Gated Clock Conversion](#)” on [page 13-28](#)
 - “[PowerPlay Power Optimization](#)” on [page 13-31](#)
 - “[Restructure Multiplexers](#)” on [page 13-33](#)
- Settings Related to Timing Constraints:
 - “[Timing-Driven Synthesis](#)” on [page 13-30](#)
 - “[Optimization Technique](#)” on [page 13-28](#)
 - “[Auto Gated Clock Conversion](#)” on [page 13-28](#)
 - “[SDC Constraint Protection](#)” on [page 13-31](#)

- State Machine Settings and Enumerated Types:
 - “State Machine Processing” on page 13-35
 - “Manually Specifying State Assignments Using the `syn_encoding` Attribute” on page 13-37
 - “Manually Specifying Enumerated Types Using the `enum_encoding` Attribute” on page 13-38
 - “Safe State Machines” on page 13-39
- Register Power-Up Settings:
 - “Power-Up Level” on page 13-41
 - “Power-Up Don’t Care” on page 13-42
- Controlling, Preserving, Removing, and Duplicating Logic and Registers:
 - “Limiting Resource Usage in Partitions” on page 13-32
 - “Remove Duplicate Registers” on page 13-42
 - “Preserve Registers” on page 13-43
 - “Disable Register Merging/Don’t Merge Register” on page 13-43
 - “Noprune Synthesis Attribute/Preserve Fan-out Free Register Node” on page 13-44
 - “Keep Combinational Node/Implement as Output of Logic Cell” on page 13-45
 - “Disabling Synthesis Netlist Optimizations with `dont_retime` Attribute” on page 13-46
 - “Disabling Synthesis Netlist Optimizations with `dont_replicate` Attribute” on page 13-47
 - “Maximum Fan-Out” on page 13-48
 - “Controlling Clock Enable Signals with Auto Clock Enable Replacement and `direct_enable`” on page 13-49
 - “Auto Gated Clock Conversion” on page 13-28
 - “Partitions for Preserving Hierarchical Boundaries” on page 13-20
- Megafunction Inference Options:
 - “Inferring Multiplier, DSP, and Memory Functions from HDL Code” on page 13-50
 - “RAM Style and ROM Style—for Inferred Memory” on page 13-53
 - “Turning Off the Add Pass-Through Logic to Inferred RAMs `no_rw_check` Attribute” on page 13-55
 - “RAM Initialization File—for Inferred Memory” on page 13-59
 - “Multiplier Style—for Inferred Multipliers” on page 13-59

- Controlling Synthesis with Other Synthesis Directives:
 - “Full Case Attribute” on page 13-61
 - “Parallel Case” on page 13-62
 - “Translate Off and On / Synthesis Off and On” on page 13-63
 - “Ignore translate_off and synthesis_off Directives” on page 13-64
 - “Read Comments as HDL” on page 13-65
- Specifying I/O-Related Assignments:
 - “Use I/O Flipflops” on page 13-66
 - “Specifying Pin Locations with chip_pin” on page 13-67
- Setting Quartus II Logic Options in Your HDL Source Code:
 - “Using altera_attribute to Set Quartus II Logic Options” on page 13-69
- Other Settings:
 - “Synthesis Effort” on page 13-35
 - “Synthesis Seed” on page 13-35

Setting Synthesis Options

You can set synthesis options in the **Settings** dialog box, or with logic options in the Quartus II software, or you can use synthesis attributes and directives in your HDL source code.

Analysis & Synthesis Settings Page of the Settings Dialog Box

The **Analysis & Synthesis Settings** page of the **Settings** dialog box allows you to set global synthesis options that apply to the entire project. You can also use a corresponding Tcl command.

The Quartus II software sets some of the advanced synthesis settings in the **Physical Synthesis Optimizations** page under **Compilation Process Settings**.



For more information about Physical Synthesis options, refer to the *Netlist Optimizations and Physical Synthesis* chapter in volume 2 of the *Quartus II Handbook*.

Quartus II Logic Options

The Quartus II logic options control many aspects of the synthesis and placement and routing process. To set logic options in the Quartus II GUI, on the Assignments menu, click **Assignment Editor**. You can also use a corresponding Tcl command to set global assignments. The Quartus II logic options enable you to set instance or node-specific assignments without editing the source HDL code.



For more information about using the Assignment Editor, refer to the *About the Assignment Editor* in Quartus II Help.

Synthesis Attributes

The Quartus II software supports synthesis attributes for Verilog HDL and VHDL, also commonly called pragmas. These attributes are not standard Verilog HDL or VHDL commands. Synthesis tools use attributes to control the synthesis process in a particular manner. The Quartus II software applies the attributes in the HDL source code, and attributes always apply to a specific design element. Some synthesis attributes are also available as Quartus II logic options via the Quartus II GUI or scripting. Each attribute description in this chapter indicates a corresponding setting or a logic option that you can set in the GUI. You can specify only some attributes with HDL synthesis attributes.

Attributes specified in your HDL code are not visible in the Assignment Editor or in the .qsf. Assignments or settings made with the Quartus II GUI, the .qsf, or the Tcl interface take precedence over assignments or settings made with synthesis attributes in your HDL code. The Quartus II software generates warning messages if the software finds invalid attributes, but does not generate an error or stop the compilation. This behavior is necessary because attributes are specific to various design tools, and attributes not recognized in the Quartus II software might be for a different EDA tool. The Quartus II software lists the attributes specified in your HDL code in the Source assignments table of the Analysis & Synthesis report.

The Verilog-2001, SystemVerilog, and VHDL language definitions provide specific syntax for specifying attributes, but in Verilog-1995, you must embed attribute assignments in comments. You can enter attributes in your code using the syntax in [Example 13-21](#) through [Example 13-27](#), in which *<attribute>*, *<attribute type>*, *<value>*, *<object>*, and *<object type>* are variables, and the entry in brackets is optional. The examples in this chapter demonstrate each syntax form.

 Verilog HDL is case-sensitive; therefore, synthesis attributes in Verilog HDL files are also case-sensitive.

Example 13-21. Specifying Synthesis Attributes in Verilog-1995

```
// synthesis <attribute> [ = <value> ]  
or  
/* synthesis <attribute> [ = <value> ] */
```

You must use Verilog-1995 comment-embedded attributes as a suffix to (that is, placed after) the declaration of an item and must appear before a semicolon, when a semicolon is necessary (refer to [Example 13-21](#)).

 You cannot use the open one-line comment in Verilog HDL when a semicolon is necessary after the line, because it is not clear to which HDL element that the attribute applies. For example, you cannot make an attribute assignment such as `reg r; // synthesis <attribute>` because the Quartus II software could read the attribute as part of the next line.

To apply multiple attributes to the same instance in Verilog-1995, separate the attributes with spaces, as shown in [Example 13-22](#):

Example 13-22. Applying Multiple Attributes to the Same Instance in Verilog-1996

```
//synthesis <attribute1> [ = <value> ] <attribute2> [ = <value> ]
```

For example, to set the `maxfan` attribute to 16 (for details, refer to “Maximum Fan-Out” on page 13-48) and set the `preserve` attribute (for details, refer to “Preserve Registers” on page 13-43) on a register called `my_reg`, use the following syntax as shown in Example 13-23:

Example 13-23. Setting `maxfan` and `preserve` Attribute on a Register

```
reg my_reg /* synthesis maxfan = 16 preserve */;
```

In addition to the `synthesis` keyword shown above, the Quartus II software supports the `pragma`, `synopsys`, and `exemplar` keywords for compatibility with other synthesis tools. The software also supports the `altera` keyword, which allows you to add synthesis attributes that the Quartus II integrated synthesis feature recognizes and not by other tools that recognize the same synthesis attribute.

 Because formal verification tools do not recognize the `exemplar`, `pragma`, and `altera` keywords, avoid using these attribute keywords when using formal verification.

Example 13-24. Specifying Synthesis Attributes in Verilog-2001 and SystemVerilog

```
(* <attribute> [ = <value> ] *)
```

You must use Verilog-2001 attributes as a prefix to (that is, placed before) a declaration, module item, statement, or port connection, and as a suffix to (that is, placed after) an operator or a Verilog HDL function name in an expression (refer to Example 13-24).

 Formal verification does not support the Verilog-2001 attribute syntax because the tools do not recognize the syntax.

To apply multiple attributes to the same instance in Verilog-2001 or SystemVerilog, separate the attributes with commas, as shown in Example 13-25:

Example 13-25. Applying Multiple Attributes

```
(* <attribute1> [ = <value1> ], <attribute2> [ = <value2> ] *)
```

For example, to set the `maxfan` attribute to 16 (refer to “Maximum Fan-Out” on page 13-48 for details) and set the `preserve` attribute (refer to “Preserve Registers” on page 13-43 for details) on a register called `my_reg`, use the following syntax as shown in Example 13-26:

Example 13-26.

```
(* maxfan = 16, preserve *) reg my_reg;
```

VHDL attributes, as shown in Example 13-27, declare and apply the attribute type to the object you specify.

Example 13-27. Synthesis Attributes in VHDL

```
attribute <attribute> : <attribute type> ;  
attribute <attribute> of <object> : <object type> is <value>;
```

The Quartus II software defines and applies each attribute separately to a given node. For VHDL designs, the software declares all supported synthesis attributes in the `altera_syn_attributes` package in the **Altera** library. You can call this library from your VHDL code to declare the synthesis attributes, as shown in [Example 13-28](#):

Example 13-28.

```
LIBRARY altera;  
USE altera.altera_syn_attributes.all;
```

Synthesis Directives

The Quartus II software supports synthesis directives, also commonly called compiler directives or pragmas. You can include synthesis directives in Verilog HDL or VHDL code as comments. These directives are not standard Verilog HDL or VHDL commands. Synthesis tools use directives to control the synthesis process in a particular manner. Directives do not apply to a specific design node, but change the behavior of the synthesis tool from the point in which they occur in the HDL source code. Other tools, such as simulators, ignore these directives and treat them as comments.

You can enter synthesis directives in your code using the syntax in [Example 13-29](#), [Example 13-30](#), and [Example 13-31](#), in which `<directive>` and `<value>` are variables, and the entry in brackets are optional. For synthesis directives, no equal sign before the value is necessary; this is different than the Verilog syntax for synthesis attributes. The examples in this chapter demonstrate each syntax form.



Verilog HDL is case sensitive; therefore, all synthesis directives are also case sensitive.

Example 13-29. Specifying Synthesis Directives with Verilog HDL

```
// synthesis <directive> [ <value> ]  
or  
/* synthesis <directive> [ <value> ] */
```

Example 13-30. Specifying Synthesis Directives with VHDL

```
-- synthesis <directive> [ <value> ]
```

Example 13-31. Specifying Synthesis Directives with VHDL-2008

```
/* synthesis <directive> [<value>] */
```

In addition to the `synthesis` keyword shown above, the software supports the `pragma`, `synopsys`, and `exemplar` keywords in Verilog HDL and VHDL for compatibility with other synthesis tools. The Quartus II software also supports the keyword `altera`, which allows you to add synthesis directives that only Quartus II integrated synthesis feature recognizes, and not by other tools that recognize the same synthesis directives.



Because formal verification tools ignore the `exemplar`, `pragma`, and `altera` keywords, Altera recommends that you avoid using these directive keywords when you use formal verification to prevent mismatches with the Quartus II results.

Optimization Technique

The **Optimization Technique** logic option specifies the goal for logic optimization during compilation; that is, whether to attempt to achieve maximum speed performance or minimum area usage, or a balance between the two. Table 13-2 lists the settings for this logic option, which you can apply only to a design entity. You can also set this logic option for your whole project in the **Settings** dialog box. If you want to set this logic option for an entity, you must create a design partition for the entity before setting the **Optimization Technique** logic option. The Quartus II software ignores this option when set on an entity that is not a design partition.

Table 13-2. Optimization Technique Settings

Setting	Description
Area	The Compiler makes the design as small as possible to minimize resource usage.
Speed	The Compiler chooses a design implementation that has the fastest f_{MAX} .
Balanced (1)	The Compiler maps part of the design for area and part for speed, providing better area utilization than optimizing for speed, with a slightly slower f_{MAX} than optimizing for speed.

Note to Table 13-2:

(1) Not all device families support the balanced optimization technique.

The default setting varies by device family. The software optimizes the default setting for best area or best speed trade-off.

Auto Gated Clock Conversion

Clock gating is a common optimization technique in ASIC designs to minimize power consumption. You can use the **Auto Gated Clock Conversion** logic option to optimize your prototype ASIC designs by converting gated clocks into clock enables when you use FPGAs in your ASIC prototyping. The automatic conversion of gated clocks to clock enables is more efficient than manually modifying source code. The **Auto Gated Clock Conversion** logic option automatically converts qualified gated clocks (base clocks as defined in the Synopsys Design Constraints [SDC]) to clock enables. To use **Auto Gated Clock Conversion**, you must select the option from the **More Analysis & Synthesis Settings** dialog box, in the **Analysis & Synthesis Settings** page.

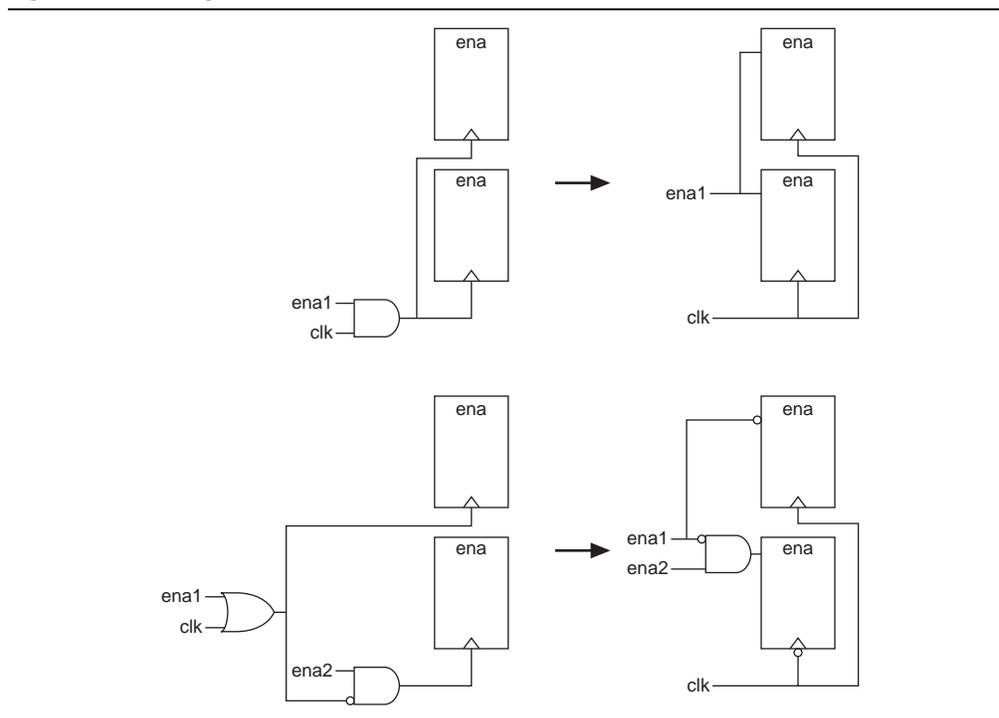
The gated clock conversion occurs when all these conditions are met:

- Only one base clock drives a gated-clock
- For one set of gating input values, the value output of the gated clock remains constant and does not change as the base clock changes
- For one value of the base clock, changes in the gating inputs do not change the value output for the gated clock

The option supports combinational gates in clock gating network.

Figure 13-3 shows example of gated clock conversions.

Figure 13-3. Example Gated Clock Conversion



 This option does not support registers in RAM, DSP blocks, or I/O related WYSIWYG primitives. Because the gated-clock conversion cannot trace the base clock from the gated clock, the gated clock conversion does not support multiple design partitions from incremental compilation in which the gated clock and base clock are not in the same hierarchical partition. A gated clock tree, instead of every gated clock, is the basis of each conversion. Therefore, if you cannot convert a gated clock from a root gated clock of a multiple cascaded gated clock, the conversion of the entire gated clock tree fails.

The **Info** tab in the Messages window lists all the converted gated clocks. You can view a list of converted and non-converted gated clocks from the Compilation Report under the **Optimization Results** of the Analysis & Synthesis Report. The **Gated Clock Conversion Details** table lists the reasons for non-converted gated clocks.

 This feature is available when using the TimeQuest analyzer.

 For more information about **Auto Gated Clock Conversion** logic option and a list of supported devices, refer to *Auto Gated Clock Conversion logic option* in Quartus II Help.

Timing-Driven Synthesis

The **Timing-Driven Synthesis** logic option specifies whether Analysis & Synthesis should use the design's SDC timing constraints to better optimize the circuit. When you turn on this option, Analysis & Synthesis runs timing analysis to obtain timing information about the netlist, and then considers the SDC timing constraints to focus on critical portions of the design when optimizing for performance, while optimizing non-critical portions for area. When you turn on this option, Analysis & Synthesis also protects SDC constraints by not merging duplicate registers that have incompatible timing constraints. For more information, refer to “[SDC Constraint Protection](#)” on page 13-31.

When you turn on the **Timing-Driven Synthesis** logic option, Analysis & Synthesis increases performance by improving logic depth on critical portions of the design, and improving area on non-critical portions of the design. The increased performance affects the amount of area used, specifically adaptive look-up tables (ALUTs) and registers in the design. Depending on how much of the design is timing critical, overall area can increase or decrease when you turn on the **Timing-Driven Synthesis** logic option. Runtime and peak memory use increases slightly if you turn on the **Timing-Driven Synthesis** logic option.

When you turn on the **Timing-Driven Synthesis** logic option, the **Optimization Technique** logic option has the following affect. With **Optimization Technique Speed**, **Timing-Driven Synthesis** optimizes timing-critical portions of the design for performance at the cost of increasing area (logic and register utilization). With an **Optimization Technique** of **Balanced**, **Timing-Driven Synthesis** also optimizes the timing-critical portions of the design for performance, but it only allows limited area increase. With **Optimization Technique Area**, **Timing-Driven Synthesis** only optimizes the design for area. **Timing-Driven Synthesis** prevents registers with incompatible timing constraints from merging for any **Optimization Technique** setting. If your design contains multiple partitions, you can select **Timing-Driven Synthesis** unique options for each partition. If you use a .qxp as a source file, or if your design uses partitions developed in separate Quartus II projects, the software cannot properly compute timing of paths that cross the partition boundaries.

Even with the **Optimization Technique** logic option set to **Speed**, the **Timing-Driven Synthesis** option still considers the resource usage in your design when increasing area to improve timing. For example, the **Timing-Driven Synthesis** option checks if a device has enough registers before deciding to implement the shift registers in logics instead of RAM for better timing performance.

When using incremental compilation, integrated synthesis allows each partition to use up all the registers in a device. You can use the **Maximum Number of LABs** settings to specify the number of LABs that every partition can use. If your design has only one partition, you can also use the **Maximum Number of LABs** settings to limit the number of resources that your design can use. This is useful when you add more logic to your design.

To change the **Timing-Driven Synthesis** logic option, follow these steps:

1. On the Assignment menu, click **Settings**.
2. In the **Category** list, select **Analysis & Synthesis Settings**. In the **Analysis & Synthesis Settings** page, select or unselect **Timing-Driven Synthesis**.

-  The option is available when using the TimeQuest analyzer. Altera recommends that you select a specific device for timing-driven synthesis to have the most accurate timing information. When you select auto device, timing-driven synthesis uses the smallest device for the selected family to obtain timing information.
-  Arria® II GX, Cyclone® III, Cyclone IV, Hardcopy® III, Hardcopy IV, Stratix® III, and Stratix IV devices turn on this feature by default.
-  For more information about **Timing-Driven Synthesis** logic option and a list of supported devices, refer to *Timing-Driven Synthesis logic option* in Quartus II Help.

SDC Constraint Protection

The **SDC Constraint Protection** option specifies whether Analysis & Synthesis should protect registers from merging when they have incompatible timing constraints. For example, two registers that are duplicates of each other but have different multicycle constraints on them are not merged when you turn on this option. When you turn on the **Timing-Driven Synthesis** option, the software detects registers with incompatible constraints, and you do not need to turn on **SDC Constraint Protection**. To use the **SDC constraint protection** option, you must turn on the option in the **More Analysis & Synthesis Settings** dialog box, which is found in the **Analysis & Synthesis Settings** page.

This feature supports Arria II GX, Arria GX, Cyclone II, Cyclone III, Cyclone IV, HardCopy series, MAX II, Stratix II, Stratix II GX, Stratix III, and Stratix IV device families.

PowerPlay Power Optimization

The **PowerPlay Power Optimization** logic option controls the power-driven compilation setting of Analysis & Synthesis and determines how aggressively Analysis & Synthesis optimizes the design for power.

To display the **Analysis & Synthesis Settings** page, follow these steps:

1. On the Assignments menu, click **Settings**.
2. In the **Category** list, select **Analysis & Synthesis Settings**.

-  For more information about the available settings for the **PowerPlay power optimization** logic option and a list of supported devices, refer to *PowerPlay power optimization (Analysis & Synthesis Settings Page) logic option* in Quartus II Help.
-  For more information about optimizing your design for power utilization, refer to the *Power Optimization* chapter in volume 2 of the *Quartus II Handbook*. For information about analyzing your power results, refer to the *PowerPlay Power Analysis* chapter in volume 3 of the *Quartus II Handbook*.

Limiting Resource Usage in Partitions

Resource balancing is important when performing Analysis & Synthesis. During resource balancing, Quartus II integrated synthesis considers the amount of used and available DSP and RAM blocks in the device, and tries to balance these resources to prevent no-fit errors.

For DSP blocks, resource balancing converts the remaining DSP blocks to equivalent logic if there are more DSP blocks in the design that can be placed in the device. For RAM blocks, resource balancing converts RAM blocks to different types of RAM blocks if there are not enough blocks of a certain type available in the device; however, Quartus II integrated synthesis does not convert RAM blocks to logic.



The RAM balancing feature does not support Stratix V devices because Stratix V has only M20K memory blocks.

By default, Quartus II integrated synthesis considers the information in the targeted device to identify the number of available DSP or RAM blocks. However, in incremental compilation, each partition considers the information in the device independently and consequently assumes that the partition has all the DSP and RAM blocks in the device available for use, resulting in over allocation of DSP or RAM blocks in the design, which means that the total number of DSP or RAM blocks used by all the partitions is greater than the number of DSP or RAM blocks available in the device. This can eventually lead to a no-fit error during the fitting process.

The following sections describe the methods to prevent a no-fit error during the fitting process:

- [“Creating LogicLock Regions” on page 13-32](#)
- [“Using Assignments to Limit the Number of RAM and DSP Blocks” on page 13-33](#)

Creating LogicLock Regions

The floorplan-aware synthesis feature allows you to use LogicLock regions to define resource allocation for DSP blocks and RAM blocks. For example, if a certain partition is assigned to a certain LogicLock region, resource balancing takes into account that all the DSP and RAM blocks in that partition need to fit in this LogicLock region. Resource balancing then balances the DSP and RAM blocks accordingly.

Because floorplan-aware balancing step considers only one partition at a time, it does not know that nodes from another partition may be using the same resources. When using this feature, Altera recommends that you do not manually assign nodes from different partitions to the same LogicLock region.

If you do not want the software to consider the LogicLock floorplan constraints when performing DSP and RAM balancing, you can turn off the floorplan-aware synthesis feature. You can turn off the **Use LogicLock Constraints During Resource Balancing** option in the **Analysis & Synthesis Settings** page by clicking **More Settings**.



For more information about using LogicLock regions to create a floorplan for incremental compilation, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.

Using Assignments to Limit the Number of RAM and DSP Blocks

For DSP and RAM block balancing, you can use assignments to limit the maximum number of blocks that the balancer allows. You can set these assignments globally or on individual partitions. For DSP block balancing, the **Maximum DSP Block Usage** logic option allows you to specify the maximum number of DSP blocks that the DSP block balancer assumes are available for the current partition. For RAM blocks, the floorplan-aware logic option allows you to specify maximum resources for different RAM types, such as **Maximum Number of M4K/M9K Memory Blocks**, **Maximum Number of M512 Memory Blocks**, **Maximum Number of M-RAM/M144K Memory Blocks**, or **Maximum Number of LABs**.

The partition-specific assignment overrides the global assignment, if any. However, each partition that does not have a partition-specific assignment uses the value set by the global assignment, or the value derived from the device size if no global assignment exists. This can also lead to over-allocation. Therefore, Altera recommends that you always set the assignment on each partition individually.

To select the **Maximum Number <block type> Memory Blocks** option or the **Maximum DSP Block Usage** option globally, follow these steps:

1. On the Assignment menu, click **Settings**.
2. Under Category, click **Analysis & Synthesis Settings**.
3. In the **Analysis & Synthesis Settings** dialog box, click **More Settings**.
4. In the **Name** pull-down list, select the required option and set the desired value.

You can use the Assignment Editor to set this assignment on a partition by selecting the assignment, and setting it on the root entity of a partition. You can set any positive integer as the value of this assignment. If you set this assignment on a name other than a partition root, then Analysis & Synthesis gives an error.

- ④ For more information about the logic options, including a list of supported device families, refer to *Maximum DSP Block Usage logic option*, *Maximum Number of M4K/M9K Memory Blocks logic option*, *Maximum Number of M512 Memory Blocks logic option*, *Maximum Number of M-RAM/144K Memory Blocks logic option*, and *Maximum Number of LABs logic option* in Quartus II Help.

Restructure Multiplexers

The **Restructure Multiplexers** logic option restructures multiplexers to make more efficient use of area, allowing the design to implement multiplexers with a reduced number of LEs or ALMs.

When multiplexers from one part of your design feed multiplexers in another part of your design, trees of multiplexers form. Multiplexers may arise in different parts of the design through Verilog HDL or VHDL constructs such as the “if,” “case,” or “?:” statements. Multiplexer buses occur most often as a result of multiplexing together arrays in Verilog HDL, or STD_LOGIC_VECTOR signals in VHDL. The **Restructure Multiplexers** logic option identifies buses of multiplexer trees that have a similar structure. This logic option optimizes the structure of each multiplexer bus for the target device to reduce the overall amount of logic in your design.

Results of the multiplexer optimizations are design dependent, but area reductions as high as 20% are possible. The option can negatively affect your design's f_{MAX} .

Table 13-3 lists the settings for the logic option, which you can apply to an individual node or to an entity that is a design partition. You can also specify this option for your whole project on the **Analysis & Synthesis Settings** page of the **Settings** dialog box by clicking **More Settings** and setting the option value.

Table 13-3. Restructure Multiplexer Settings

Setting	Description
On	Enables multiplexer restructuring to minimize your design area. This setting can reduce the f_{MAX} .
Off	Disables multiplexer restructuring to avoid possible reductions in f_{MAX} .
Auto (Default)	Allows the Compiler to determine whether to enable the option based on your other Quartus II synthesis settings. When the Optimization Technique option is set to Area or Balanced , Quartus II integrated synthesis restructures all multiplexers. When the Optimization Technique option is set to Speed , Quartus II integrated synthesis attempts to restructure the multiplexers selectively and makes a good trade-off between area and f_{MAX} .

After compilation, you can view multiplexer restructuring information in the **Multiplexer Restructuring Statistics** report in the **Multiplexer Statistics** folder under **Analysis & Synthesis Optimization Results** in the **Analysis & Synthesis** section of the Compilation Report.

Table 13-4 describes the information in the **Multiplexer Restructuring Statistics** report table for each bus of multiplexers.

Table 13-4. Multiplexer Information in the Multiplexer Restructuring Statistics Report

Heading	Description
Multiplexer Inputs	The number of different inputs that are multiplexed together.
Bus Width	The width of the bus in bits.
Baseline Area	An estimate of how many logic cells are necessary to implement the bus of multiplexers (before any multiplexer restructuring). You can use this estimate to identify any large multiplexers in the design.
Area if Restructured	An estimate of how many logic cells are necessary to implement the bus of multiplexers when applying Multiplexer Restructuring .
Saving if Restructured	An estimate of logic cells saving when applying Multiplexer Restructuring .
Registered	An indication of whether registers are present on the multiplexer outputs. Multiplexer Restructuring uses the secondary control signals of a register (such as synchronous clear and synchronous load) to further reduce the amount of logic required to implement the bus of multiplexers.
Example Multiplexer Output	The name of one of the multiplexer outputs. This name can help determine the source location of the multiplexer bus in your design.

 For more information about **Restructure Multiplexers** logic option, including a list of supported device families, refer to *Restructure Multiplexers logic option* in Quartus II Help.

 For more information about optimizing for multiplexers, refer to the *Multiplexers* section of the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*.

Synthesis Effort

The **Synthesis Effort** logic option specifies the overall synthesis effort level in the Quartus II software. You can set the effort level to either **Fast** or **Auto**.

The **Auto** setting indicates standard synthesis effort. The Quartus II software attempts to optimize your design as much as possible.

When you set the effort level to **Fast**, Quartus II integrated synthesis skips several steps to ensure that synthesis runs much faster (at the cost of performance and resource utilization). You can use the **Fast** synthesis effort level with the Fitter early timing estimate feature. The early timing estimate feature gives you preliminary timing estimates before running a full compilation, which results in a quicker iteration time; therefore, you can save significant compilation time to get a good estimation of the final timing of your design. When you use the **Fast** synthesis effort level as part of a full compilation, Fitter runtime might increase because fast synthesis generates a netlist that is slightly more difficult for the Fitter to route when compared to the netlist from a normal synthesis. When you set the **Synthesis Effort** option to **Fast**, **Timing-Driven Synthesis** turns off.

To set the **Synthesis Effort** option from the Quartus II GUI, on the **Analysis & Synthesis Settings** page, click **More Settings**. Select **Auto** or **Fast** from the pull-down menu in the **Synthesis Effort** option, and then click **OK** to close the **Settings** dialog box.

- ❓ For more information about **Synthesis Effort** logic option, including a list of supported device families, refer to *Synthesis Effort logic option* in Quartus II Help.

Synthesis Seed

The **Synthesis Seed** option specifies the seed that Synthesis uses to randomly run synthesis in a slightly different way. You can use this seed when your design is close to meeting requirements, to get a slightly different result. The seeds that produce the best result for a design may change if the design changes.

To set the **Synthesis Seed** option from the Quartus II GUI, on the **Analysis & Synthesis Settings** page, click **More Settings**. The default value is 1. You can specify a positive integer value.

State Machine Processing

The **State Machine Processing** logic option specifies the processing style to synthesize a state machine. [Table 13-5](#) lists the settings for this logic option, which you can apply to a state machine name or to a design entity that contains a state machine. You can also set this option for your whole project on the **Analysis & Synthesis Settings** page in the **Settings** dialog box.

Table 13-5. State Machine Processing Settings (Part 1 of 2)

Setting	Description
Auto (Default)	Allows the Compiler to choose what it determines to be the best encoding for the state machine.
Minimal Bits	Uses the least number of bits to encode the state machine.
One-Hot	Encodes the state machine in one-hot style. See the example below for details.

Table 13-5. State Machine Processing Settings (Part 2 of 2)

Setting	Description
User-Encoded	Encodes the state machine in the manner that you specify.
Sequential	Uses a binary encoding in which the first enumeration literal in the Enumeration Type has encoding 0 and the second 1.
Gray	Uses an encoding in which the encodings for adjacent enumeration literals differ by one bit. An N -bit gray code can represent 2^N values.
Johnson	Uses an encoding similar to a gray code, in which each state has only one bit different from its neighboring states. To generate each state, you can shift the bits of the previous state to the right by 1. The MSB of each state is the negation of the LSB of the previous state. An N -bit Johnson code can represent at most 2^N states, but requires less logic than a gray encoding.

The default state machine encoding, which is **Auto**, uses one-hot encoding for FPGA devices and minimal-bits encoding for CPLDs. These settings achieve the best results on average, but another encoding style might be more appropriate for your design, so this option allows you to control the state machine encoding.



For guidelines on how to correctly infer and encode your state machine, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*.

For one-hot encoding, the Quartus II software does not guarantee that each state has one bit set to one and all other bits set to zero. Quartus II integrated synthesis creates one-hot register encoding with standard one-hot encoding and then inverts the first bit. This results in an initial state with all zero values, and the remaining states have two 1 values. Quartus II integrated synthesis encodes the initial state with all zeros for the state machine power-up because all device registers power up to a low value. This encoding has the same properties as true one-hot encoding: the software recognizes each state by the value of one bit. For example, in a one-hot-encoded state machine with five states, including an initial or reset state, the software uses the register encoding shown in [Example 13-32](#):

Example 13-32. Register Encoding

```
State 0    0 0 0 0 0
State 1    0 0 0 1 1
State 2    0 0 1 0 1
State 3    0 1 0 0 1
State 4    1 0 0 0 1
```

If you set the **State Machine Processing** logic option to **User-Encoded** in a Verilog HDL design, the software starts with the original design values for the state constants. For example, a Verilog HDL design can contain a declaration such as shown in [Example 13-33](#):

Example 13-33.

```
parameter S0 = 4'b1010, S1 = 4'b0101, ...
```

If the software infers the states S_0 , S_1 , ... the software uses the encoding $4'b1010$, $4'b0101$, ... If necessary, the software inverts bits in a user-encoded state machine to ensure that all bits of the reset state of the state machine are zero.

 You can view the state machine encoding from the Compilation Report under the State Machines of the Analysis & Synthesis Report. The State Machine Viewer only displays a graphical representation of the state machines as interpreted from your design.

 For more information about the State Machine Viewer, refer to the State Machine Viewer section of the *Analyzing Designs with Quartus II Netlist Viewers* chapter in volume 1 of the *Quartus II Handbook*.

To assign your own state encoding with the **User-Encoded** setting of the **State Machine Processing** option in a VHDL design, you must apply specific binary encoding to the elements of an enumerated type because enumeration literals have no numeric values in VHDL. Use the `syn_encoding` synthesis attribute to apply your encoding values. For more information, refer to [“Manually Specifying State Assignments Using the `syn_encoding` Attribute”](#).

 For information about the **State Machine Processing** logic option, including supported devices, refer to *State Machine Processing logic option* in Quartus II Help.

Manually Specifying State Assignments Using the `syn_encoding` Attribute

The Quartus II software infers state machines from enumerated types and automatically assigns state encoding based on [“State Machine Processing” on page 13–35](#). With this logic option, you can choose the value **User-Encoded** to use the encoding from your HDL code. However, in standard VHDL code, you cannot specify user encoding in the state machine description because enumeration literals have no numeric values in VHDL.

To assign your own state encoding for the **User-Encoded State Machine Processing** setting, use the `syn_encoding` synthesis attribute to apply specific binary encodings to the elements of an enumerated type or to specify an encoding style. The Quartus II software can implement Enumeration Types with different encoding styles shown in [Table 13–6](#).

Table 13–6. `syn_encoding` Attribute Values

Attribute Value	Enumeration Types
"default"	Use an encoding based on the number of enumeration literals in the Enumeration Type. If the number of literals is less than five, use the "sequential" encoding. If the number of literals is more than five, but fewer than 50, use a "one-hot" encoding. Otherwise, use a "gray" encoding.
"sequential"	Use a binary encoding in which the first enumeration literal in the Enumeration Type has encoding 0 and the second 1
"gray"	Use an encoding in which the encodings for adjacent enumeration literals differ by exactly one bit. An N -bit gray code can represent 2^N values.
"johnson"	Use an encoding similar to a gray code. An N -bit Johnson code can represent at most 2^N states, but requires less logic than a gray encoding.
"one-hot"	The default encoding style requiring N bits, in which N is the number of enumeration literals in the Enumeration Type.
"compact"	Use an encoding with the fewest bits.
"user"	Encode each state using its value in the Verilog source. By changing the values of your state constants, you can change the encoding of your state machine.

The `syn_encoding` attribute must follow the enumeration type definition, but precede its use.

Manually Specifying Enumerated Types Using the `enum_encoding` Attribute

By default, the Quartus II software one-hot encodes all enumerated types you defined. With the `enum_encoding` attribute, you can specify the logic encoding for an enumerated type and override the default one-hot encoding to improve the logic efficiency.



If an enumerated type represents the states of a state machine, using the `enum_encoding` attribute to specify a manual state encoding prevents the Compiler from recognizing state machines based on the enumerated type. Instead, the Compiler processes these state machines as regular logic with the encoding specified by the attribute, and the Report window for the project does not list these states machines as state machines. If you want to control the encoding for a recognized state machine, use the **State Machine Processing** logic option and the `syn_encoding` synthesis attribute.

To use the `enum_encoding` attribute in a VHDL design file, associate the attribute with the enumeration type whose encoding you want to control. The `enum_encoding` attribute must follow the enumeration type definition, but precede its use. In addition, the attribute value should be a string literal that specifies either an arbitrary user encoding or an encoding style of "default", "sequential", "gray", "johnson", or "one-hot".

An arbitrary user encoding consists of a space-delimited list of encodings. The list must contain as many encodings as the number of enumeration literals in your enumeration type. In addition, the encodings should have the same length, and each encoding must consist solely of values from the `std_ulogic` type declared by the `std_logic_1164` package in the IEEE library. In [Example 13-34](#), the `enum_encoding` attribute specifies an arbitrary user encoding for the enumeration type `fruit`.

Example 13-34. Specifying an Arbitrary User Encoding for Enumerated Type

```
type fruit is (apple, orange, pear, mango);
attribute enum_encoding : string;
attribute enum_encoding of fruit : type is "11 01 10 00";
```

[Example 13-35](#) shows the encoded enumeration literals:

Example 13-35.

```
apple   = "11"
orange  = "01"
pear    = "10"
mango   = "00"
```

Altera recommends that you specify an encoding style, rather than a manual user encoding, especially when the enumeration type has a large number of enumeration literals. The Quartus II software can implement enumeration types with the different encoding styles as shown in [Table 13-7](#).

Table 13-7. enum_encoding Attribute Values

Attribute Value	Enumeration Types
"default"	Use an encoding based on the number of enumeration literals in the enumeration type. If the number of literals are fewer than five, use the "sequential" encoding. If the number of literals are more than five, but fewer than 50 literals, use a "one-hot" encoding. Otherwise, use a "gray" encoding.
"sequential"	Use a binary encoding in which the first enumeration literal in the enumeration type has encoding 0 and the second 1.
"gray"	Use an encoding in which the encodings for adjacent enumeration literals differ by exactly one bit. An N -bit gray code can represent 2^N values.
"johnson"	Use an encoding similar to a gray code. An N -bit Johnson code can represent at most 2^N states, but requires less logic than a gray encoding.
"one-hot"	The default encoding style requiring N bits, in which N is the number of enumeration literals in the enumeration type.

In [Example 13-34](#), the `enum_encoding` attribute manually specified a gray encoding for the enumeration type `fruit`. You can concisely write this example by specifying the "gray" encoding style instead of a manual encoding, as shown in [Example 13-36](#).

Example 13-36. Specifying the “gray” Encoding Style or Enumeration Type

```
type fruit is (apple, orange, pear, mango);
attribute enum_encoding : string;
attribute enum_encoding of fruit : type is "gray";
```

Safe State Machines

The **Safe State Machine** logic option and corresponding `syn_encoding` attribute value `safe` specify that the software must insert extra logic to detect an illegal state, and force the transition of the state machine to the reset state.

A finite state machine can enter an illegal state—meaning the state registers contain a value that does not correspond to any defined state. By default, the behavior of the state machine that enters an illegal state is undefined. However, you can set the `syn_encoding` attribute to `safe` or use the **Safe State Machine** logic option if you want the state machine to recover deterministically from an illegal state. The software inserts extra logic to detect an illegal state, and forces the transition of the state machine to the reset state. You can use this logic option when the state machine enters an illegal state. The most common cause of an illegal state is a state machine that has control inputs that come from another clock domain, such as the control logic for a clock-crossing FIFO, because the state machine must have inputs from another clock domain. This option protects only state machines (and not other registers) by forcing them into the reset state. You can use this option if your design has asynchronous inputs. However, Altera recommends using a synchronization register chain instead of relying on the safe state machine option.

The safe state machine value does not use any user-defined default logic from your HDL code that corresponds to unreachable states. Verilog HDL and VHDL enable you to specify a behavior for all states in the state machine explicitly, including unreachable states. However, synthesis tools detect if state machine logic is unreachable and minimize or remove the logic. Synthesis tools also remove any flag signals or logic that indicate such an illegal state. If the state machine is implemented as safe, the recovery logic added by Quartus II integrated synthesis forces its transition from an illegal state to the reset state.

You can set the **Safe State Machine** logic option globally, or on individual state machines. To set this logic option, on the **Analysis & Synthesis Settings** page, select **More Settings**. In the **Existing option settings** list, select **Safe State Machine**, and turn on this option in the **Setting** list.

You can set the `syn_encoding` safe attribute on a state machine in HDL, as shown in [Example 13-37](#) through [Example 13-39](#).

Example 13-37. Verilog HDL Code: a Safe State Machine Attribute

```
reg [2:0] my_fsm /* synthesis syn_encoding = "safe" */;
```

Example 13-38. Verilog-2001 and SystemVerilog Code: a Safe State Machine Attribute

```
(* syn_encoding = "safe" *) reg [2:0] my_fsm;
```

Example 13-39. VHDL Code: a Safe State Machine Attribute

```
ATTRIBUTE syn_encoding OF my_fsm : TYPE IS "safe";
```

If you specify an encoding style (refer to [“Manually Specifying State Assignments Using the `syn_encoding` Attribute”](#) on page 13-37), separate the encoding style value in the quotation marks with the `safe` value with a comma, as follows: `"safe, one-hot"` or `"safe, gray"`.

Safe state machine implementation can result in a noticeable area increase for the design. Therefore, Altera recommends that you set this option only on the critical state machines in the design in which the safe mode is necessary, such as a state machine that uses inputs from asynchronous clock domains. You can also reduce the necessity of this option by correctly synchronizing inputs coming from other clock domains.

-  If you make the safe state machine assignment on an instance that the software fails to recognize as a state machine, or an entity that contains a state machine, the software takes no action. You must restructure the code, so that the software recognizes and infers the instance as a state machine.
-  For more information about the **Safe State Machine** logic option, refer to [Safe State Machine logic option](#) in Quartus II Help.
-  For guidelines to ensure that the software correctly infers your state machine, refer to the [Recommended HDL Coding Styles](#) chapter in volume 1 of the *Quartus II Handbook*.

Power-Up Level

This logic option causes a register (flipflop) to power up with the specified logic level, either **High** (1) or **Low** (0). Registers in the core hardware power up to 0 in all Altera devices. For the register to power up with a logic level **High**, the Compiler performs an optimization referred to as **NOT-gate push back** on the register. **NOT-gate push back** adds an inverter to the input and the output of the register, so that the reset and power-up conditions appear to be high and the device operates as expected. The register itself still powers up **Low**, but the register output inverts so the signal arriving at all destinations is **High**.

The **Power-Up Level** option supports wildcard characters, and you can apply this option to any register, registered logic cell WYSIWYG primitive, or to a design entity containing registers if you want to set the power level for all registers in the design entity. If you assign this option to a registered logic cell WYSIWYG primitive, such as an atom primitive from a third-party synthesis tool, you must turn on the **Perform WYSIWYG Primitive Resynthesis** logic option for the option to take effect. You can also apply the option to a pin with the logic configurations described in the following list:

- If you turn on this option for an input pin, the option transfers to the register that the pin drives, if all of these conditions are present:
 - No logic, other than inversion, between the pin and the register.
 - The input pin drives the data input of the register.
 - The input pin does not fan-out to any other logic.
- If you turn on this option for an output or bidirectional pin, the option transfers to the register that feeds the pin, if all of these conditions are present:
 - No logic, other than inversion, between the register and the pin.
 - The register does not fan-out to any other logic.

❓ For more information about the **Power-Up Level** logic option, including information on the supported device families, refer to *Power-Up Level logic option* in Quartus II Help.

Inferred Power-Up Levels

Quartus II integrated synthesis reads default values for registered signals defined in Verilog HDL and VHDL code, and converts the default values into Power-Up Level settings. The software also synthesizes variables with assigned values in Verilog HDL initial blocks into power-up conditions. Synthesis of these default and initial constructs allows synthesized behavior of your design to match, as closely as possible, the power-up state of the HDL code during a functional simulation.

The following register declarations all set a power-up level of V_{CC} or a logic value “1”, as shown in [Example 13-40](#):

Example 13-40.

```
signal q : std_logic = '1'; -- power-up to VCC

reg q = 1'b1; // power-up to VCC

reg q;
initial begin q = 1'b1; end // power-up to VCC
```



For more information about NOT-gate push back, the power-up states for Altera devices, and how set and reset control signals affect the power-up level, refer to the [Recommended HDL Coding Styles](#) chapter in volume 1 of the *Quartus II Handbook*.

Power-Up Don't Care

This logic option allows the Compiler to optimize registers in the design that do not have a defined power-up condition.

For example, your design might have a register with its D input tied to V_{CC} , and with no clear signal or other secondary signals. If you turn on this option, the Compiler can choose for the register to power up to V_{CC} . Therefore, the output of the register is always V_{CC} . The Compiler can remove the register and connect its output to V_{CC} . If you turn this option off or if you set a **Power-Up Level** assignment of **Low** for this register, the register transitions from GND to V_{CC} when the design starts up on the first clock signal. Thus, the register is at V_{CC} and you cannot remove the register. Similarly, if the register has a clear signal, it is not removed because after the clear is asserted, the register transitions again to GND and back to V_{CC} .

If the Compiler performs a **Power-Up Don't Care** optimization that allows it to remove a register, it issues a message to indicate that it is doing so.

This project-wide option does not apply to registers that have the **Power-Up Level** logic option set to either **High** or **Low**.

-  For more information about **Power-Up Don't Care** logic option and a list of supported devices, refer to [Power-Up Don't Care logic option](#) in Quartus II Help.

Remove Duplicate Registers

When you turn on the **Remove Duplicate Registers** logic option, the Compiler removes registers that are identical to other registers. If two registers generate the same logic, the Compiler removes the second register, and the first register fans out to the destinations of the second register. Also, if the deleted register has different logic option assignments, the Compiler ignores them.



Altera recommends using this option only if you want to prevent the Compiler from removing duplicate registers. Use this option only with the **Off** setting. You can apply this option to an individual register or a design entity that contains registers.

-  For more information about **Remove Duplicate Registers** logic option and the supported devices, refer to [Remove Duplicate Registers logic option](#) in Quartus II Help.

Preserve Registers

This attribute and logic option directs the Compiler not to minimize or remove a specified register during synthesis optimizations or register netlist optimizations. Optimizations can eliminate redundant registers and registers with constant drivers; this option prevents a register from being reduced to a constant or merged with a duplicate register. This option can preserve a register so you can observe the register during simulation or with the SignalTap® II Logic Analyzer. Additionally, it can preserve registers if you create a preliminary version of the design in which you have not specified the secondary signals. You can also use the attribute to preserve a duplicate of an I/O register so that you can place one copy of the I/O register in an I/O cell and the second in the core.

 This option cannot preserve registers that have no fan-out. To prevent the removal of registers with no fan-out, refer to “[Noprune Synthesis Attribute/Preserve Fan-out Free Register Node](#)” on page 13-44.

The **Preserve Registers** logic option prevents the software from inferring a register as a state machine.

You can set the **Preserve Registers** logic option in the Quartus II GUI, or you can set the preserve attribute in your HDL code, as shown in [Example 13-41](#) through [Example 13-43](#). In these examples, the `my_reg` register is preserved.

 In addition to `preserve`, the Quartus II software supports the `syn_preserve` attribute name for compatibility with other synthesis tools.

Example 13-41. Verilog HDL Code: `syn_preserve` Attribute

```
reg my_reg /* synthesis syn_preserve = 1 */;
```

Example 13-42. Verilog-2001 Code: `syn_preserve` Attribute

```
(* syn_preserve = 1 *) reg my_reg;
```

 The `= 1` after the `preserve` in [Example 13-41](#) and [Example 13-42](#) is optional, because the assignment uses a default value of 1 when you specify the assignment.

Example 13-43. VHDL Code: `preserve` Attribute

```
signal my_reg : stdlogic;  
attribute preserve : boolean;  
attribute preserve of my_reg : signal is true;
```

 For more information about the **Preserve Registers** logic option and the supported devices, refer to [Preserve Registers logic option](#) in Quartus II Help.

Disable Register Merging/Don't Merge Register

This logic option and attribute prevents the specified register from merging with other registers and prevents other registers from merging with the specified register. When applied to a design entity, it applies to all registers in the entity.

You can use this option to instruct the Compiler to correctly use your timing constraints for the register during synthesis. For example, if the register has a multicycle constraint, this option prevents the Compiler from merging other registers into the specified register, avoiding unintended timing effects and functional differences.

This option differs from the **Preserve Register** logic option because it does not prevent the removal of a register with constant drivers or a redundant register.

You can set the **Disable Register Merging** logic option in the Quartus II GUI, or you can set the `dont_merge` attribute in your HDL code, as shown in [Example 13-44](#) through [Example 13-46](#). In these examples, the `my_reg` register is prevented from merging.

Example 13-44. Verilog HDL Code: `dont_merge` Attribute

```
reg my_reg /* synthesis dont_merge */;
```

Example 13-45. Verilog-2001 and SystemVerilog Code: `dont_merge` Attribute

```
(* dont_merge *) reg my_reg;
```

Example 13-46. VHDL Code: `dont_merge` Attribute

```
signal my_reg : stdlogic;
attribute dont_merge : boolean;
attribute dont_merge of my_reg : signal is true;
```

-  For more information about the **Disable Register Merging** logic option and the supported devices, refer to [Disable Register Merging logic option](#) in Quartus II Help.

Noprune Synthesis Attribute/Preserve Fan-out Free Register Node

This synthesis attribute and corresponding logic option direct the Compiler to preserve a fan-out-free register through the entire compilation flow. This is different from the **Preserve Registers** option, which prevents a register from being reduced to a constant or merged with a duplicate register. Standard synthesis optimizations remove nodes that do not directly or indirectly feed a top-level output pin. This option can retain a register so you can observe it in the Simulator or the SignalTap II Logic Analyzer. Additionally, it can retain registers if you create a preliminary version of the design in which you have not specified the fan-out logic of the register.

You can set the **Preserve Fan-out Free Register Node** logic option in the Quartus II GUI, or you can set the `noprune` attribute in your HDL code, as shown in [Example 13-47](#) through [Example 13-49](#). In these examples, the `my_reg` register is preserved.

-  You must use the `noprune` attribute instead of the logic option if the register has no immediate fan-out in its module or entity. If you do not use the synthesis attribute, the software removes (or “prunes”) registers with no fan-out during Analysis & Elaboration before the logic synthesis stage applies any logic options. If the register has no fan-out in the full design, but has fan-out in its module or entity, you can use the logic option to retain the register through compilation.

The software supports the attribute name `syn_noprune` for compatibility with other synthesis tools.

Example 13-47. Verilog HDL Code: `syn_noprune` Attribute

```
reg my_reg /* synthesis syn_noprune */;
```

Example 13-48. Verilog-2001 and SystemVerilog Code: `noprune` Attribute

```
(* noprune *) reg my_reg;
```

Example 13-49. VHDL Code: `noprune` Attribute

```
signal my_reg : stdlogic;  
attribute noprune: boolean;  
attribute noprune of my_reg : signal is true;
```

-  For more information about **Preserve Fan-out Free Register Node** logic option and a list of supported devices, refer to *Preserve Fan-out Free Register logic option* in Quartus II Help.

Keep Combinational Node/Implement as Output of Logic Cell

This synthesis attribute and corresponding logic option direct the Compiler to keep a wire or combinational node through logic synthesis minimizations and netlist optimizations. A wire that has a `keep` attribute or a node that has the **Implement as Output of Logic Cell** logic option applied becomes the output of a logic cell in the final synthesis netlist, and the name of the logic cell remains the same as the name of the wire or node. You can use this directive to make combinational nodes visible to the SignalTap II Logic Analyzer.

-  The option cannot keep nodes that have no fan-out. You cannot maintain node names for wires with tri-state drivers, or if the signal feeds a top-level pin of the same name (in this case, the node name is changed to a name, such as `<net name>~buf0`).

You can use the **Ignore LCELL Buffers** logic option to direct Analysis & Synthesis to ignore logic cell buffers created by the **Implement as Output of Logic Cell** logic option or the `LCELL` primitive. If you apply this logic option to an entity, it affects all lower-level entities in the hierarchy path.

-  To avoid unintended design optimizations, make sure the **Ignore LCELL Buffers** logic option is not inherited by an entity instantiated with Altera or third-party IP that relies on logic cell buffers for correct behavior. For example, if an IP core uses logic cell buffers to manage high fan-out signals and inherits the **Ignore LCELL Buffers** logic option, the target device may no longer function properly.

You can turn off the **Ignore LCELL Buffers** logic option for a specific entity to override any assignments inherited from higher-level entities in the hierarchy path if logic cell buffers created by the **Implement as Output of Logic Cell** logic option or the `LCELL` primitive are required for correct behavior.

You can set the **Implement as Output of Logic Cell** logic option in the Quartus II GUI, or you can set the `keep` attribute in your HDL code, as shown in [Example 13-50](#) through [Example 13-52](#). In these examples, the Compiler maintains the node name `my_wire`.



In addition to `keep`, the Quartus II software supports the `syn_keep` attribute name for compatibility with other synthesis tools.

Example 13-50. Verilog HDL Code: keep Attribute

```
wire my_wire /* synthesis keep = 1 */;
```

Example 13-51. Verilog-2001 Code: keep Attribute

```
(* keep = 1 *) wire my_wire;
```

Example 13-52. VHDL Code: syn_keep Attribute

```
signal my_wire: bit;  
attribute syn_keep: boolean;  
attribute syn_keep of my_wire: signal is true;
```

- ② For more information about the **Implement as Output of Logic Cell** logic option and the supported devices, refer to *Implement as Output of Logic Cell logic option* in Quartus II Help.

Disabling Synthesis Netlist Optimizations with `dont_retime` Attribute

This attribute disables synthesis retiming optimizations on the register you specify. When applied to a design entity, it applies to all registers in the entity.

You can turn off retiming optimizations with this option and prevent node name changes, so that the Compiler can correctly use your timing constraints for the register.

You can set the **Netlist Optimizations** logic option to **Never Allow** in the Quartus II GUI to disable retiming along with other synthesis netlist optimizations, or you can set the `dont_retime` attribute in your HDL code, as shown in [Example 13-53](#) through [Example 13-55](#). In these examples, the code prevents `my_reg` register from being retimed.

Example 13-53. Verilog HDL Code: dont_retime Attribute

```
reg my_reg /* synthesis dont_retime */;
```

Example 13-54. Verilog-2001 and SystemVerilog Code: dont_retime Attribute

```
(* dont_retime *) reg my_reg;
```

Example 13-55. VHDL Code: dont_retime Attribute

```
signal my_reg : std_logic;  
attribute dont_retime : boolean;  
attribute dont_retime of my_reg : signal is true;
```



For compatibility with third-party synthesis tools, Quartus II integrated synthesis also supports the attribute `syn_allow_retiming`. To disable retiming, set `syn_allow_retiming` to 0 (Verilog HDL) or `false` (VHDL). This attribute does not have any effect when you set the attribute to 1 or `true`.

Disabling Synthesis Netlist Optimizations with dont_replicate Attribute

This attribute disables synthesis replication optimizations on the register you specify. When applied to a design entity, it applies to all registers in the entity.

You can turn off register replication (or duplication) optimizations with this option, so that the Compiler uses your timing constraints for the register.

You can set the **Netlist Optimizations** logic option to **Never Allow** in the Quartus II GUI to disable replication along with other synthesis netlist optimizations, or you can set the `dont_replicate` attribute in your HDL code, as shown in [Example 13-56](#) through [Example 13-58](#). In these examples, the code prevents `my_reg` register from being replicated.

Example 13-56. Verilog HDL Code: dont_replicate Attribute

```
reg my_reg /* synthesis dont_replicate */;
```

Example 13-57. Verilog-2001 and SystemVerilog Code: dont_replicate Attribute

```
(* dont_replicate *) reg my_reg;
```

Example 13-58. VHDL Code: dont_replicate Attribute

```
signal my_reg : std_logic;  
attribute dont_replicate : boolean;  
attribute dont_replicate of my_reg : signal is true;
```

 For compatibility with third-party synthesis tools, Quartus II integrated synthesis also supports the attribute `syn_replicate`. To disable replication, set `syn_replicate` to 0 (Verilog HDL) or `false` (VHDL). This attribute does not have any effect when you set the attribute to 1 or `true`.

Maximum Fan-Out

This **Maximum Fan-Out** attribute and logic option directs the Compiler to control the number of destinations that a node feeds. The Compiler duplicates a node and splits its fan-out until the individual fan-out of each copy falls below the maximum fan-out restriction. You can apply this option to a register or a logic cell buffer, or to a design entity that contains these elements. You can use this option to reduce the load of critical signals, which can improve performance. You can use the option to instruct the Compiler to duplicate a register that feeds nodes in different locations on the target device. Duplicating the register can enable the Fitter to place these new registers closer to their destination logic to minimize routing delay.

To turn off the option for a given node if you set the option at a higher level of the design hierarchy, in the **Netlist Optimizations** logic option, select **Never Allow**. If not disabled by the **Netlist Optimizations** option, the software acknowledges the maximum fan-out constraint as long as the following conditions are met:

- The node is not part of a cascade, carry, or register cascade chain.
- The node does not feed itself.
- The node feeds other logic cells, DSP blocks, RAM blocks, and/or pins through data, address, clock enable, and other ports, but not through any asynchronous control ports (such as asynchronous clear).

The software does not create duplicate nodes in these cases, because there is no clear way to duplicate the node, or to avoid the small differences in timing which could produce functional differences in the implementation (in the third condition above in which asynchronous control signals are involved). If the constraint cannot be applied because one of these conditions is not met, the Quartus II software issues a message to indicate that it ignored the maximum fan-out assignment. To instruct the software not to check node destinations for possible problems such as the third condition, you can set the **Netlist Optimizations** logic option to **Always Allow** for a given node.

 If you have enabled any of the Quartus II netlist optimizations that affect registers, add the `preserve` attribute to any registers to which you have set a `maxfan` attribute. The `preserve` attribute ensures that the netlist optimization algorithms, such as register retiming, do not affect the registers.

 For details about netlist optimizations, refer to the *Netlist Optimizations and Physical Synthesis* chapter in volume 2 of the *Quartus II Handbook*.

You can set the **Maximum Fan-Out** logic option in the Quartus II GUI. This option supports wildcard characters. You can also set the `maxfan` attribute in your HDL code, as shown in [Example 13-59](#) through [Example 13-61](#). In these examples, the Compiler duplicates the `clk_gen` register, so its fan-out is not greater than 50.



In addition to `maxfan`, the Quartus II software supports the `syn_maxfan` attribute for compatibility with other synthesis tools.

Example 13-59. Verilog HDL Code: `syn_maxfan` Attribute

```
reg clk_gen /* synthesis syn_maxfan = 50 */;
```

Example 13-60. Verilog-2001 Code: `maxfan` Attribute

```
(* maxfan = 50 *) reg clk_gen;
```

Example 13-61. VHDL Code: `maxfan` Attribute

```
signal clk_gen : stdlogic;  
attribute maxfan : signal ;  
attribute maxfan of clk_gen : signal is 50;
```

- ❓ For more information about the **Maximum Fan-Out** logic option and the supported devices, refer to *Maximum Fan-Out logic option* in Quartus II Help.

Controlling Clock Enable Signals with Auto Clock Enable Replacement and `direct_enable`

The **Auto Clock Enable Replacement** logic option allows the software to find logic that feeds a register and move the logic to the register's clock enable input port. To solve fitting or performance issues with designs that have many clock enables, you can turn off this option for individual registers or design entities. Turning the option off prevents the software from using the register's clock enable port. The software implements the clock enable functionality using multiplexers in logic cells.

If the specific logic is not automatically moved to a clock enable input with the **Auto Clock Enable Replacement** logic option, you can instruct the software to use a direct clock enable signal. Applying the `direct_enable` attribute to a specific signal instructs the software to use the clock enable port of a register to implement the signal. The attribute ensures that the signal directly drives the clock enable port, and the signal is not optimized or combined with any other logic.

Example 13-62 through *Example 13-64* show how to set this attribute to ensure that the signal is preserved and used directly as a clock enable.



In addition to `direct_enable`, the Quartus II software supports the `syn_direct_enable` attribute name for compatibility with other synthesis tools.

Example 13-62. Verilog HDL Code: `direct_enable` attribute

```
wire my_enable /* synthesis direct_enable = 1 */ ;
```

Example 13-63. Verilog-2001 and SystemVerilog Code: `syn_direct_enable` attribute

```
(* syn_direct_enable *) wire my_enable;
```

Example 13-64. VHDL Code: `direct_enable` attribute

```
attribute direct_enable: boolean;
attribute direct_enable of my_enable: signal is true;
```



For more information about the **Auto Clock Enable Replacement** logic option and the supported devices, refer to *Auto Clock Enable Replacement logic option* in Quartus II Help.

Inferring Multiplier, DSP, and Memory Functions from HDL Code

The Quartus II Compiler automatically recognizes multipliers, multiply-accumulators, multiply-adders, or memory functions described in HDL code, and converts the HDL code into respective megafunction, or, may map them directly to device atoms or memory atoms. If the software converts the HDL code into megafunction, the software uses the Altera megafunction code when compiling your design, even when you do not specifically instantiate the megafunction. The software infers megafunctions to take advantage of logic that is optimized for Altera devices. The area and performance of such logic can be better than the results from inferring generic logic from the same HDL code.

Additionally, you must use megafunctions to access certain architecture-specific features, such as RAM, DSP blocks, and shift registers that provide improved performance compared with basic logic cells.



For details about coding style recommendations when targeting megafunctions in Altera devices, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*.

The Quartus II software provides options to control the inference of certain types of megafunctions, as described in the following subsections:

- “Multiply-Accumulators and Multiply-Adders”
- “Shift Registers” on page 13-51
- “RAM and ROM” on page 13-51
- “Resource Aware RAM, ROM, and Shift-Register Inference” on page 13-52
- “Auto RAM to Logic Cell Conversion” on page 13-53

Multiply-Accumulators and Multiply-Adders

Use the **Auto DSP Block Replacement** logic option to control DSP block inference for multiply-accumulations and multiply-adders. To disable inference, turn off this option for the entire project on the **Analysis & Synthesis Settings** page of the **Settings** dialog box, or turn off the option for a specific block with the Assignment Editor. By default, the software enables this logic option for Stratix V devices.

- For more information about the **Auto DSP Block Replacement** logic option and the supported devices, refer to *Auto DSP Block Replacement logic option* in Quartus II Help.

Shift Registers

Use the **Auto Shift Register Replacement** logic option to control shift register inference. This option has three settings: **Off**, **Auto** and **Always**. **Auto** is the default setting in which Quartus II integrated synthesis decides which shift registers to replace or leave in registers. Placing shift registers in memory saves logic area, but can have a negative effect on f_{max} . Quartus II integrated synthesis uses the optimization technique setting, logic and RAM utilization of the design, and timing information from **Timing-Driven Synthesis** to determine which shift registers are located in memory and which are located in registers. To disable inference, turn off this option for the entire project on the **Analysis & Synthesis Settings** page of the **Settings** dialog box by clicking **More Settings** and setting the option to **Off**. You can also disable the option for a specific block with the Assignment Editor. Even if you set the logic option to **On** or **Auto**, the software might not infer small shift registers because small shift registers do not benefit from implementation in dedicated memory. However, you can use the **Allow Any Shift Register Size for Recognition** logic option to instruct synthesis to infer a shift register even when its size is too small.

- The registers that the software maps to the ALTSHIFT_TAPS megafunction and places in RAM are not available in the Simulator because their node names do not exist after synthesis.

The software turns off the **Auto Shift Register Replacement** logic option when you select a formal verification tool on the **EDA Tool Settings** page. The software issues a warning and lists shift registers that would have been inferred if no formal verification tool was selected in the compilation report. To enable a megafunction for the shift register in the formal verification flow, you can either instantiate a shift register explicitly with the MegaWizard™ Plug-In Manager or make the shift register into a black box in a separate entity or module.

- For more information about the **Auto Shift Register Replacement** logic option and the supported devices, refer to *Auto Shift Register Replacement logic option* in Quartus II Help.

RAM and ROM

Use the **Auto RAM Replacement** and **Auto ROM Replacement** logic options to control RAM and ROM inference, respectively. To disable inference, turn off the appropriate option for the entire project on the **Analysis & Synthesis Settings** page of the **Settings** dialog box by clicking **More Settings** and setting the option to **Off**. You can also disable the option for a specific block with the Assignment Editor.

 Although inferred shift registers are implemented in RAM blocks, you cannot turn off the **Auto RAM Replacement** option to disable shift register replacement. Use the **Auto Shift Register Replacement** option (refer to “[Shift Registers](#)”).

The software might not infer very small RAM or ROM blocks because very small memory blocks can be implemented more efficiently with the registers in the logic. However, you can use the **Allow Any RAM Size for Recognition** and **Allow Any ROM Size for Recognition** logic options to instruct synthesis to infer a memory block even when its size is too small.

 The software turns off the **Auto ROM Replacement** logic option when you select a formal verification tool in the **EDA Tool Settings** page. If you do not select a formal verification tool, the software issues a warning and displays a report panel that lists ROMs that would have been inferred. To enable a megafunction for the shift register in the formal verification flow, you can either instantiate a ROM explicitly using the MegaWizard Plug-In Manager or create a black box for the ROM in a separate entity or module.

Although formal verification tools do not support inferred RAM blocks, because of the importance of inferring RAM in many designs, the software turns on the **Auto RAM Replacement** logic option when you select a formal verification tool in the **EDA Tool Settings** page. The software automatically performs black box instance for any module or entity that contains an inferred RAM block. The software issues a warning and lists the black box created in the compilation report. This black box allows formal verification tools to proceed; however, the formal verification tool cannot verify the entire module or entire entity that contains the RAM. Altera recommends that you explicitly instantiate RAM blocks in separate modules or entities so that the formal verification tool can verify as much logic as possible.

 For more information about the **Auto RAM Replacement** and **Auto ROM Replacement** logic options and their supported devices, refer to [Auto RAM Replacement logic option](#) and [Auto ROM Replacement logic option](#) in Quartus II Help.

Resource Aware RAM, ROM, and Shift-Register Inference

The Quartus II integrated synthesis considers resource usage when inferring RAM, ROM, and shift registers. During RAM, ROM, and shift register inferencing, synthesis looks at the number of memories available in the current device and does not infer more memory than is available to avoid a no-fit error. Synthesis tries to select the memories that are not inferred in a way that aims at the smallest increase in logic and registers.

Resource aware RAM, ROM and shift register inference is controlled by the **Resource Aware Inference for Block RAM** option. You can disable this option for the entire project in the **More Analysis & Synthesis Settings** dialog box, or per partition in the Assignment Editor.

 Stratix V devices do not support this option.

When you select the **Auto** setting, resource aware RAM, ROM, and shift register inference use the resource counts from the largest device.

For designs with multiple partitions, Quartus II integrated synthesis considers one partition at a time. Therefore, for each partition, it assumes that all RAM blocks are available to that partition. If this causes a no-fit error, you can limit the number of RAM blocks available per partition with the **Maximum Number of M512 Memory Blocks**, **Maximum Number of M4K/M9K Memory Blocks**, **Maximum Number of M-RAM/M144K Memory Blocks** and **Maximum Number of LABs** settings in the Assignment Editor. The balancer also uses these options. For more information, refer to [“Limiting Resource Usage in Partitions”](#) on page 13-32.

Auto RAM to Logic Cell Conversion

The **Auto RAM to Logic Cell Conversion** logic option allows Quartus II integrated synthesis to convert small RAM blocks to logic cells if the logic cell implementation gives better quality of results. The software converts only single-port or simple-dual port RAMs with no initialization files to logic cells. You can set this option globally or apply it to individual RAM nodes. You can enable this option by turning on the appropriate option for the entire project in the **More Analysis & Synthesis Settings** dialog box.

For Arria GX and Stratix family of devices, the software uses the following rules to determine the placement of a RAM, either in logic cells or a dedicated RAM block:

- If the number of words is less than 16, use a RAM block if the total number of bits is greater than or equal to 64.
- If the number of words is greater than or equal to 16, use a RAM block if the total number of bits is greater than or equal to 32.
- Otherwise, implement the RAM in logic cells.

For the Cyclone family of devices, the software uses the following rules:

- If the number of words is greater than or equal to 64, use a RAM block.
- If the number of words is greater than or equal to 16 and less than 64, use a RAM block if the total number of bits is greater than or equal to 128.
- Otherwise, implement the RAM in logic cells.

- ❓ For more information about the **Auto RAM to Logic Cell Conversion** logic options and the supported devices, refer to [Auto RAM to Logic Cell Conversion logic option](#) in Quartus II Help.

RAM Style and ROM Style—for Inferred Memory

These attributes specify the implementation for an inferred RAM or ROM block. You can specify the type of TriMatrix embedded memory block, or specify the use of standard logic cells (LEs or ALMs). The attributes are supported only for device families with TriMatrix embedded memory blocks.

The `ramstyle` and `romstyle` attributes take a single string value. The `M512`, `M4K`, `M-RAM`, `MLAB`, `M9K`, `M144K` and `M20K` values (as applicable for the target device family) indicate the type of memory block to use for the inferred RAM or ROM. If you set the attribute to a block type that does not exist in the target device family, the software generates a warning and ignores the assignment. The value `logic` indicates that the RAM or ROM should be implemented in regular logic rather than dedicated memory blocks. You

can set the attribute on a module or entity, in which case it specifies the default implementation style for all inferred memory blocks in the immediate hierarchy. You can also set the attribute on a specific signal (VHDL) or variable (Verilog HDL) declaration, in which case it specifies the preferred implementation style for that specific memory, overriding the default implementation style.



If you specify a value of `logic`, the memory still appears as a RAM or ROM block in the RTL Viewer, but it is converted to regular logic during synthesis.

In addition to `ramstyle` and `romstyle`, the Quartus II software supports the `syn_ramstyle` attribute name for compatibility with other synthesis tools.

[Example 13-65](#) through [Example 13-67](#) specify that all memory in the module or entity `my_memory_blocks` should be implemented using a specific type of block.

Example 13-65. Verilog-1995 Code: Applying a `romstyle` Attribute to a Module Declaration

```
module my_memory_blocks (...) /* synthesis romstyle = "M4K" */;
```

Example 13-66. Verilog-2001 and SystemVerilog Code: Applying a `ramstyle` Attribute to a Module Declaration

```
(* ramstyle = "M512" *) module my_memory_blocks (...);
```

Example 13-67. VHDL Code: Applying a `romstyle` Attribute to an Architecture

```
architecture rtl of my_my_memory_blocks is
attribute romstyle : string;
attribute romstyle of rtl : architecture is "M-RAM";
begin
```

[Example 13-68](#) through [Example 13-70](#) specify that the inferred memory `my_ram` or `my_rom` should be implemented using regular logic instead of a TriMatrix memory block.

Example 13-68. Verilog-1995 Code: Applying a `syn_ramstyle` Attribute to a Variable Declaration

```
reg [0:7] my_ram[0:63] /* synthesis syn_ramstyle = "logic" */;
```

Example 13-69. Verilog-2001 and SystemVerilog Code: Applying a `romstyle` Attribute to a Variable Declaration

```
(* romstyle = "logic" *) reg [0:7] my_rom[0:63];
```

Example 13-70. VHDL Code: Applying a `ramstyle` Attribute to a Signal Declaration

```
type memory_t is array (0 to 63) of std_logic_vector (0 to 7);
signal my_ram : memory_t;
attribute ramstyle : string;
attribute ramstyle of my_ram : signal is "logic";
```

You can control the depth of an inferred memory block with the `max_depth` attribute. You can also optimize the usage of the memory block with this attribute. [Example 13-71](#) through [Example 13-73](#) specify the depth of the inferred memory `mem` using the `max_depth` synthesis attribute.

Example 13-71. Verilog-1995 Code: Applying a `max_depth` Attribute to a Variable Declaration

```
reg [7:0] mem [127:0] /* synthesis max_depth = 2048 */
```

Example 13-72. Verilog-2001 and SystemVerilog Code: Applying a `max_depth` Attribute to a Variable Declaration

```
(* max_depth = 2048*) reg [7:0] mem [127:0];
```

Example 13-73. VHDL Code: Applying a `max_depth` Attribute to a Variable Declaration

```
type ram_block is array (0 to 31) of std_logic_vector (2 downto 0);  
signal mem : ram_block;  
attribute max_depth : natural;  
attribute max_depth OF mem : signal is 2048;
```

The syntax for setting these attributes in HDL is the same as the syntax for other synthesis attributes, as shown in [“Synthesis Attributes” on page 13-25](#).

Turning Off the Add Pass-Through Logic to Inferred RAMs `no_rw_check` Attribute

Setting the `no_rw_check` value for the `ramstyle` attribute, or turning off the corresponding global **Add Pass-Through Logic to Inferred RAMs** logic option indicates that your design does not depend on the behavior of the inferred RAM when there are reads and writes to the same address in the same clock cycle. If you specify the attribute or turn off the logic option, the Quartus II software can choose a read-during-write behavior instead of using the read-during-write behavior of your HDL source code.

In some cases, an inferred RAM should be mapped into regular logic cells because it has a read-during-write behavior that is not supported by the TriMatrix memory blocks in your target device. In other cases, the Quartus II software must insert extra logic to mimic read-during-write behavior of the HDL source to increase the area of your design and potentially reduce its performance. In some of these cases, you can use the attribute to specify that the software can implement the RAM directly in a TriMatrix memory block without using logic. You can also use the attribute to prevent a warning message for dual-clock RAMs in the case that the inferred behavior in the device does not exactly match the read-during-write conditions described in the HDL code.



For more information about recommended styles for inferring RAM and some of the issues involved with different read-during-write conditions, refer to the [Recommended HDL Coding Styles](#) chapter in volume 1 of the *Quartus II Handbook*.

To set the **Add Pass-Through Logic to Inferred RAMs** logic option with the Quartus II GUI, click **More Settings** on the **Analysis & Synthesis Settings** page of the **Settings** dialog box. [Example 13-74](#) and [Example 13-75](#) use two addresses and normally require extra logic after the RAM to ensure that the read-during-write conditions in the device match the HDL code. If your design does not require a defined read-during-write condition, the extra logic is not necessary. With the `no_rw_check` attribute, Quartus II integrated synthesis does not generate the extra logic.

Example 13-74. Verilog HDL Inferred RAM Using `no_rw_check` Attribute

```
module ram_infer (q, wa, ra, d, we, clk);
    output [7:0] q;
    input [7:0] d;
    input [6:0] wa;
    input [6:0] ra;
    input we, clk;
    reg [6:0] read_add;
    (* ramstyle = "no_rw_check" *) reg [7:0] mem [127:0];
    always @ (posedge clk) begin
        if (we)
            mem[wa] <= d;
            read_add <= ra;
        end
        assign q = mem[read_add];
    endmodule
```

Example 13-75. VHDL Inferred RAM Using `no_rw_check` Attribute

```
LIBRARY ieee;
USE ieee.std_logic_1164.ALL;

ENTITY ram IS
    PORT (
        clock: IN STD_LOGIC;
        data: IN STD_LOGIC_VECTOR (2 DOWNTO 0);
        write_address: IN INTEGER RANGE 0 to 31;
        read_address: IN INTEGER RANGE 0 to 31;
        we: IN STD_LOGIC;
        q: OUT STD_LOGIC_VECTOR (2 DOWNTO 0));
END ram;

ARCHITECTURE rtl OF ram IS
    TYPE MEM IS ARRAY(0 TO 31) OF STD_LOGIC_VECTOR(2 DOWNTO 0);
    SIGNAL ram_block: MEM;
    ATTRIBUTE ramstyle : string;
    ATTRIBUTE ramstyle of ram_block : signal is "no_rw_check";
    SIGNAL read_address_reg: INTEGER RANGE 0 to 31;
BEGIN
    PROCESS (clock)
    BEGIN
        IF (clock'event AND clock = '1') THEN
            IF (we = '1') THEN
                ram_block(write_address) <= data;
            END IF;
            read_address_reg <= read_address;
        END IF;
    END PROCESS;
    q <= ram_block(read_address_reg);
END rtl;
```

You can use a `ramstyle` attribute with the `MLAB` value, so that the Quartus II software can infer a small RAM block and place it in an MLAB.



You can use this attribute in cases in which some asynchronous RAM blocks might be coded with read-during-write behavior that does not match the Stratix III, Stratix IV, and Stratix V architectures. Thus, the device behavior would not exactly match the behavior that the code describes. If the difference in behavior is acceptable in your design, use the `ramstyle` attribute with the `no_rw_check` value to specify that the software should not check the read-during-write behavior when inferring the RAM. When you set this attribute, Quartus II integrated synthesis allows the behavior of the output to differ when the asynchronous read occurs on an address that had a write on the most recent clock edge. That is, the functional HDL simulation results do not match the hardware behavior if you write to an address that is being read. To include these attributes, set the value of the `ramstyle` attribute to `MLAB, no_rw_check`.

Example 13-76 and **Example 13-77** show the method of setting two values to the `ramstyle` attribute with a small asynchronous RAM block, with the `ramstyle` synthesis attribute set, so that the software can implement the memory in the MLAB memory block and so that the read-during-write behavior is not important. Without the attribute, this design requires 512 registers and 240 ALUTs. With the attribute, the design requires eight memory ALUTs and only 15 registers.

Example 13-76. Verilog HDL Inferred RAM Using `no_rw_check` and MLAB Attributes

```

module async_ram (
    input  [5:0] addr,
    input  [7:0] data_in,
    input          clk,
    input          write,
    output [7:0] data_out );

    (* ramstyle = "MLAB, no_rw_check" *) reg [7:0] mem[0:63];

    assign data_out = mem[addr];

    always @ (posedge clk)
    begin
        if (write)
            mem[addr] = data_in;
    end
endmodule

```

Example 13-77. VHDL Inferred RAM Using `no_rw_check` and MLAB Attributes

```

LIBRARY ieee;
USE ieee.std_logic_1164.ALL;
ENTITY ram IS
    PORT (
        clock: IN STD_LOGIC;
        data: IN STD_LOGIC_VECTOR (2 DOWNTO 0);
        write_address: IN INTEGER RANGE 0 to 31;
        read_address: IN INTEGER RANGE 0 to 31;
        we: IN STD_LOGIC;
        q: OUT STD_LOGIC_VECTOR (2 DOWNTO 0));
END ram;

ARCHITECTURE rtl OF ram IS
    TYPE MEM IS ARRAY(0 TO 31) OF STD_LOGIC_VECTOR(2 DOWNTO 0);
    SIGNAL ram_block: MEM;
    ATTRIBUTE ramstyle : string;
    ATTRIBUTE ramstyle of ram_block : signal is "MLAB , no_rw_check";
    SIGNAL read_address_reg: INTEGER RANGE 0 to 31;
BEGIN
    PROCESS (clock)
    BEGIN
        IF (clock'event AND clock = '1') THEN
            IF (we = '1') THEN
                ram_block(write_address) <= data;
            END IF;
            read_address_reg <= read_address;
        END IF;
    END PROCESS;
    q <= ram_block(read_address_reg);
END rtl;

```

-  For more information about the **Add Pass-Through Logic to Inferred RAMs** logic option and the supported devices, refer to *Add Pass-Through Logic to Inferred RAMs logic option* in Quartus II Help.

RAM Initialization File—for Inferred Memory

The `ram_init_file` attribute specifies the initial contents of an inferred memory in the form of a `.mif`. The attribute takes a string value containing the name of the RAM initialization file.

Example 13-78. Verilog-1995 Code: Applying a `ram_init_file` Attribute

```
reg [7:0] mem[0:255] /* synthesis ram_init_file  
= " my_init_file.mif" */;
```

Example 13-79. Verilog-2001 Code: Applying a `ram_init_file` Attribute

```
(* ram_init_file = "my_init_file.mif" *) reg [7:0] mem[0:255];
```

Example 13-80. VHDL Code: Applying a `ram_init_file` Attribute

```
type mem_t is array(0 to 255) of unsigned(7 downto 0);  
signal ram : mem_t;  
attribute ram_init_file : string;  
attribute ram_init_file of ram :  
signal is "my_init_file.mif";
```

-  In VHDL, you can also initialize the contents of an inferred memory by specifying a default value for the corresponding signal. In Verilog HDL, you can use an initial block to specify the memory contents. Quartus II integrated synthesis automatically converts the default value into a `.mif` for the inferred RAM.

Multiplier Style—for Inferred Multipliers

The `multstyle` attribute specifies the implementation style for multiplication operations (*) in your HDL source code. You can use this attribute to specify whether you prefer the Compiler to implement a multiplication operation in general logic or dedicated hardware, if available in the target device.

The `multstyle` attribute takes a string value of "logic" or "dsp", indicating a preferred implementation in logic or in dedicated hardware, respectively. In Verilog HDL, apply the attribute to a module declaration, a variable declaration, or a specific binary expression that contains the * operator. In VHDL, apply the synthesis attribute to a signal, variable, entity, or architecture.

-  Specifying a `multstyle` of "dsp" does not guarantee that the Quartus II software can implement a multiplication in dedicated DSP hardware. The final implementation depends on several conditions, including the availability of dedicated hardware in the target device, the size of the operands, and whether or not one or both operands are constant.

In addition to `multstyle`, the Quartus II software supports the `syn_multstyle` attribute name for compatibility with other synthesis tools.

When applied to a Verilog HDL module declaration, the attribute specifies the default implementation style for all instances of the `*` operator in the module. For example, in the following code examples, the `multstyle` attribute directs the Quartus II software to implement all multiplications inside module `my_module` in the dedicated multiplication hardware.

Example 13-81. Verilog-1995 Code: Applying a `multstyle` Attribute to a Module Declaration

```
module my_module (...) /* synthesis multstyle = "dsp" */;
```

Example 13-82. Verilog-2001 Code: Applying a `multstyle` Attribute to a Module Declaration

```
(* multstyle = "dsp" *) module my_module(...);
```

When applied to a Verilog HDL variable declaration, the attribute specifies the implementation style for a multiplication operator, which has a result directly assigned to the variable. The attribute overrides the `multstyle` attribute with the enclosing module, if present. In [Example 13-83](#) and [Example 13-84](#), the `multstyle` attribute applied to variable `result` directs the Quartus II software to implement `a * b` in logic rather than the dedicated hardware.

Example 13-83. Verilog-2001 Code: Applying a `multstyle` Attribute to a Variable Declaration

```
wire [8:0] a, b;
(* multstyle = "logic" *) wire [17:0] result;
assign result = a * b; //Multiplication must be
                       //directly assigned to result
```

Example 13-84. Verilog-1995 Code: Applying a `multstyle` Attribute to a Variable Declaration

```
wire [8:0] a, b;
wire [17:0] result /* synthesis multstyle = "logic" */;
assign result = a * b; //Multiplication must be
                       //directly assigned to result
```

When applied directly to a binary expression that contains the `*` operator, the attribute specifies the implementation style for that specific operator alone and overrides any `multstyle` attribute with the target variable or enclosing module. In [Example 13-85](#), the `multstyle` attribute indicates that `a * b` should be implemented in the dedicated hardware.

Example 13-85. Verilog-2001 Code: Applying a `multstyle` Attribute to a Binary Expression

```
wire [8:0] a, b;
wire [17:0] result;
assign result = a * (* multstyle = "dsp" *) b;
```



You cannot use Verilog-1995 attribute syntax to apply the `multstyle` attribute to a binary expression.

When applied to a VHDL entity or architecture, the attribute specifies the default implementation style for all instances of the * operator in the entity or architecture. In [Example 13-86](#), the `multstyle` attribute directs the Quartus II software to use dedicated hardware, if possible, for all multiplications inside architecture `rtl` of entity `my_entity`.

Example 13-86. VHDL Code: Applying a multstyle Attribute to an Architecture

```
architecture rtl of my_entity is
    attribute multstyle : string;
    attribute multstyle of rtl : architecture is "dsp";
begin
```

When applied to a VHDL signal or variable, the attribute specifies the implementation style for all instances of the * operator, which has a result directly assigned to the signal or variable. The attribute overrides the `multstyle` attribute with the enclosing entity or architecture, if present. In [Example 13-87](#), the `multstyle` attribute associated with signal `result` directs the Quartus II software to implement a * `b` in logic rather than the dedicated hardware.

Example 13-87. VHDL Code: Applying a multstyle Attribute to a Signal or Variable

```
signal a, b : unsigned(8 downto 0);
signal result : unsigned(17 downto 0);

attribute multstyle : string;
attribute multstyle of result : signal is "logic";
result <= a * b;
```

Full Case Attribute

A Verilog HDL case statement is full when its case items cover all possible binary values of the case expression or when a default case statement is present. A `full_case` attribute attached to a case statement header that is not full forces synthesis to treat the unspecified states as a don't care value. VHDL case statements should be full, so the attribute does not apply to VHDL.



Using this attribute on a case statement that is not full avoids the latch inference problems discussed in the *Design Recommendations for Altera Devices and the Quartus II Design Assistant* chapter in volume 1 of the *Quartus II Handbook*.



Latches have limited support in formal verification tools. Make sure that you do not infer latches unintentionally; for example, through an incomplete case statement when using formal verification. Formal verification tools do support the `full_case` synthesis attribute (with limited support for attribute syntax, as described in *"Synthesis Attributes"* on page 13-25).

Using the `full_case` attribute might cause a simulation mismatch between the Verilog HDL functional and the post-Quartus II simulation because unknown case statement cases can still function as latches during functional simulation. For example, a simulation mismatch can occur with the code in [Example 13-88](#) when `sel` is `2'b11` because a functional HDL simulation output behaves as a latch and the Quartus II simulation output behaves as a don't care value.

 Altera recommends making the case statement “full” in your regular HDL code, instead of using the `full_case` attribute.

The case statement in [Example 13-88](#) is not full because not all binary values for `sel` are specified. Because the `full_case` attribute is used, synthesis treats the output as “don’t care” when the `sel` input is `2'b11`.

Example 13-88. Verilog HDL Code: a `full_case` Attribute

```
module full_case (a, sel, y);
    input [3:0] a;
    input [1:0] sel;
    output y;
    reg y;
    always @ (a or sel)
        case (sel) // synthesis full_case
            2'b00: y=a[0];
            2'b01: y=a[1];
            2'b10: y=a[2];
        endcase
endmodule
```

Verilog-2001 syntax also accepts the statements in [Example 13-89](#) in the case header instead of the comment form as shown in [Example 13-88](#).

Example 13-89. Verilog-2001 Syntax for the `full_case` Attribute

```
(* full_case *) case (sel)
```

Parallel Case

The `parallel_case` attribute indicates that a Verilog HDL case statement should be considered parallel; that is, only one case item can be matched at a time. Case items in Verilog HDL case statements might overlap. To resolve multiple matching case items, the Verilog HDL language defines a priority among case items in which the case statement always executes the first case item that matches the case expression value. By default, the Quartus II software implements the extra logic necessary to satisfy this priority relationship.

Attaching a `parallel_case` attribute to a case statement header allows the Quartus II software to consider its case items as inherently parallel; that is, at most one case item matches the case expression value. Parallel case items simplify the generated logic.

In VHDL, the individual choices in a case statement might not overlap, so they are always parallel and this attribute does not apply.

Altera recommends that you only use this attribute when the case statement is truly parallel. If you use the attribute in any other situation, the generated logic does not match the functional simulation behavior of the Verilog HDL.

 Altera recommends that you avoid using the `parallel_case` attribute, because you may mismatch the Verilog HDL functional and the post-Quartus II simulation.

If you specify SystemVerilog-2005 as the supported Verilog HDL version for your design, you can use the SystemVerilog keyword `unique` to achieve the same result as the `parallel_case` directive without causing simulation mismatches.

Example 13-90 shows a `casez` statement with overlapping case items. In functional HDL simulation, the three case items are prioritized by the bits in `sel`. For example, `sel[2]` takes priority over `sel[1]`, which takes priority over `sel[0]`. However, the synthesized design can simulate differently because the `parallel_case` attribute eliminates this priority. If more than one bit of `sel` is high, more than one output (`a`, `b`, or `c`) is high as well, a situation that cannot occur in functional HDL simulation.

Example 13-90. Verilog HDL Code: a `parallel_case` Attribute

```
module parallel_case (sel, a, b, c);
    input [2:0] sel;
    output a, b, c;
    reg a, b, c;
    always @ (sel)
    begin
        {a, b, c} = 3'b0;
        casez (sel) // synthesis parallel_case
            3'b1??: a = 1'b1;
            3'b?1?: b = 1'b1;
            3'b??1: c = 1'b1;
        endcase
    end
endmodule
```

Verilog-2001 syntax also accepts the statements as shown in **Example 13-91** in the `case` (or `casez`) header instead of the comment form, as shown in **Example 13-90**.

Example 13-91. Verilog-2001 Syntax

```
(* parallel_case *) casez (sel)
```

Translate Off and On / Synthesis Off and On

The `translate_off` and `translate_on` synthesis directives indicate whether the Quartus II software or a third-party synthesis tool should compile a portion of HDL code that is not relevant for synthesis. The `translate_off` directive marks the beginning of code that the synthesis tool should ignore; the `translate_on` directive indicates that synthesis should resume. You can also use the `synthesis_on` and `synthesis_off` directives as a synonym for `translate on` and `off`.

You can use these directives to indicate a portion of code for simulation only. The synthesis tool reads synthesis-specific directives and processes them during synthesis; however, third-party simulation tools read the directives as comments and ignore them. [Example 13-92](#), [Example 13-93](#), and [Example 13-94](#) show these directives.

Example 13-92. Verilog HDL Code: Translate Off and On

```
// synthesis translate_off
parameter tpd = 2;    // Delay for simulation
#tpd;
// synthesis translate_on
```

Example 13-93. VHDL Code: Translate Off and On

```
-- synthesis translate_off
use std.textio.all;
-- synthesis translate_on
```

Example 13-94. VHDL 2008 Code: Translate Off and On

```
/* synthesis translate_off */
use std.textio.all;
/* synthesis translate_on */
```

If you want to ignore only a portion of code in Quartus II integrated synthesis, you can use the Altera-specific attribute keyword `altera`. For example, use the `// altera translate_off` and `// altera translate_on` directives to direct Quartus II integrated synthesis to ignore a portion of code that you intend only for other synthesis tools.

Ignore `translate_off` and `synthesis_off` Directives

The **Ignore `translate_off` and `synthesis_off` Directives** logic option directs Quartus II integrated synthesis to ignore the `translate_off` and `synthesis_off` directives described in the previous section. Turning on this logic option allows you to compile code that you intended to be ignored by third-party synthesis tools; for example, megafunction declarations that were treated as black boxes in other tools but can be compiled in the Quartus II software. To set the **Ignore `translate_off` and `synthesis_off` Directives** logic option, click **More Settings** on the **Analysis & Synthesis Settings** page of the **Settings** dialog box.

- ④ For more information about the **Ignore `translate_off` and `synthesis_off` Directives** logic option and the supported devices, refer to *Ignore `translate_off` and `synthesis_off` Directives logic option* in Quartus II Help.

Read Comments as HDL

The `read_comments_as_HDL` synthesis directive indicates that the Quartus II software should compile a portion of HDL code that you commented out. This directive allows you to comment out portions of HDL source code that are not relevant for simulation, while instructing the Quartus II software to read and synthesize that same source code. Setting the `read_comments_as_HDL` directive to `on` indicates the beginning of commented code that the synthesis tool should read; setting the `read_comments_as_HDL` directive to `off` indicates the end of the code.

 You can use this directive with `translate_off` and `translate_on` to create one HDL source file that includes a megafunction instantiation for synthesis and a behavioral description for simulation.

Because formal verification tools do not recognize the `read_comments_as_HDL` directive, the directive is not supported when you use formal verification.

In [Example 13-95](#), [Example 13-96](#), and [Example 13-97](#), the commented code enclosed by `read_comments_as_HDL` is visible to the Quartus II Compiler and is synthesized. VHDL 2008 allows block comments, which is also supported for synthesis directives.

 Because synthesis directives are case-sensitive in Verilog HDL, you must match the case of the directive, as shown in the following examples.

Example 13-95. Verilog HDL Code: Read Comments as HDL

```
// synthesis read_comments_as_HDL on
// my_rom lpm_rom (.address (address),
//                .data      (data));
// synthesis read_comments_as_HDL off
```

Example 13-96. VHDL Code: Read Comments as HDL

```
-- synthesis read_comments_as_HDL on
-- my_rom : entity lpm_rom
--   port map (
--     address => address,
--     data    => data,    );
-- synthesis read_comments_as_HDL off
```

Example 13-97. VHDL 2008 Code: Read Block Comments as HDL

```
/* synthesis read_comments_as_HDL on */
/* my_rom : entity lpm_rom
   port map (
     address => address,
     data => data, ); */
synthesis read_comments_as_HDL off */
```

Use I/O Flipflops

The `useioff` attribute directs the Quartus II software to implement input, output, and output enable flipflops (or registers) in I/O cells that have fast, direct connections to an I/O pin, when possible. To improve I/O performance by minimizing setup, clock-to-output, and clock-to-output enable times, you can apply the `useioff` synthesis attribute. This synthesis attribute is supported using the **Fast Input Register**, **Fast Output Register**, and **Fast Output Enable Register** logic options that can also be set in the Assignment Editor.

The `useioff` synthesis attribute takes a boolean value and can only be applied to the port declarations of a top-level Verilog HDL module or VHDL entity (it is ignored if applied elsewhere). Setting the value to 1 (Verilog HDL) or `TRUE` (VHDL) instructs the Quartus II software to pack registers into I/O cells. Setting the value to 0 (Verilog HDL) or `FALSE` (VHDL) prevents register packing into I/O cells.

In [Example 13-98](#) and [Example 13-99](#), the `useioff` synthesis attribute directs the Quartus II software to implement the registers `a_reg`, `b_reg`, and `o_reg` in the I/O cells corresponding to the ports `a`, `b`, and `o`, respectively.

Example 13-98. Verilog HDL Code: the `useioff` Attribute

```
module top_level(clk, a, b, o);
    input clk;
    input [1:0] a, b /* synthesis useioff = 1 */;
    output [2:0] o /* synthesis useioff = 1 */;
    reg [1:0] a_reg, b_reg;
    reg [2:0] o_reg;
    always @ (posedge clk)
    begin
        a_reg <= a;
        b_reg <= b;
        o_reg <= a_reg + b_reg;
    end
    assign o = o_reg;
endmodule
```

Example 13-99 and Example 13-100 show that the Verilog-2001 syntax also accepts the type of statements instead of the comment form in Example 13-98.

Example 13-99. Verilog-2001 Code: the useioff Attribute

```
(* useioff = 1 *)   input [1:0] a, b;  
(* useioff = 1 *)   output [2:0] o;
```

Example 13-100. VHDL Code: the useioff Attribute

```
library ieee;  
use ieee.std_logic_1164.all;  
use ieee.numeric_std.all;  
entity useioff_example is  
    port (  
        clk : in  std_logic;  
        a, b : in  unsigned(1 downto 0);  
        o   : out unsigned(1 downto 0));  
    attribute useioff : boolean;  
    attribute useioff of a : signal is true;  
    attribute useioff of b : signal is true;  
    attribute useioff of o : signal is true;  
end useioff_example;  
architecture rtl of useioff_example is  
    signal o_reg, a_reg, b_reg : unsigned(1 downto 0);  
begin  
    process(clk)  
    begin  
        if (clk = '1' AND clk'event) then  
            a_reg <= a;  
            b_reg <= b;  
            o_reg <= a_reg + b_reg;  
        end if;  
    end process;  
    o <= o_reg;  
end rtl;
```

Specifying Pin Locations with chip_pin

The `chip_pin` attribute allows you to assign pin locations in your HDL source. You can only use the attribute on the ports of the top-level entity or module in your design. You can assign pins only to single-bit or one-dimensional bus ports in your design.

For single-bit ports, the value of the `chip_pin` attribute is the name of the pin on the target device, as specified by the pin table of the device.



In addition to the `chip_pin` attribute, the Quartus II software supports the `altera_chip_pin_lc` attribute name for compatibility with other synthesis tools. When using this attribute in other synthesis tools, some older device families require an “@” symbol in front of each pin assignment. In the Quartus II software, the “@” is optional.

[Example 13-101](#) through [Example 13-103](#) show different ways of assigning `my_pin1` to Pin C1 and `my_pin2` to Pin 4 on a different target device.

Example 13-101. Verilog-1995 Code: Applying Chip Pin to a Single Pin

```
input my_pin1 /* synthesis chip_pin = "C1" */;
input my_pin2 /* synthesis altera_chip_pin_lc = "@4" */;
```

Example 13-102. Verilog-2001 Code: Applying Chip Pin to a Single Pin

```
(* chip_pin = "C1" *) input my_pin1;
(* altera_chip_pin_lc = "@4" *) input my_pin2;
```

Example 13-103. VHDL Code: Applying Chip Pin to a Single Pin

```
entity my_entity is
port(my_pin1: in std_logic; my_pin2: in std_logic;...);
end my_entity;
attribute chip_pin : string;
attribute altera_chip_pin_lc : string;
attribute chip_pin of my_pin1 : signal is "C1";
attribute altera_chip_pin_lc of my_pin2 : signal is "@4";
```

For bus I/O ports, the value of the chip pin attribute is a comma-delimited list of pin assignments. The order in which you declare the range of the port determines the mapping of assignments to individual bits in the port. To leave a particular bit unassigned, leave its corresponding pin assignment blank.

[Example 13-104](#) assigns `my_pin[2]` to Pin_4, `my_pin[1]` to Pin_5, and `my_pin[0]` to Pin_6.

Example 13-104. Verilog-1995 Code: Applying Chip Pin to a Bus of Pins

```
input [2:0] my_pin /* synthesis chip_pin = "4, 5, 6" */;
```

[Example 13-105](#) reverses the order of the signals in the bus, assigning `my_pin[0]` to Pin_4 and `my_pin[2]` to Pin_6 but leaves `my_pin[1]` unassigned.

Example 13-105. Verilog-1995 Code: Applying Chip Pin to Part of a Bus

```
input [0:2] my_pin /* synthesis chip_pin = "4, ,6" */;
```

[Example 13-106](#) assigns `my_pin[2]` to Pin 4 and `my_pin[0]` to Pin 6, but leaves `my_pin[1]` unassigned.

Example 13-106. VHDL Code: Applying Chip Pin to Part of a Bus of Pins

```
entity my_entity is
port(my_pin: in std_logic_vector(2 downto 0);...);
end my_entity;

attribute chip_pin of my_pin: signal is "4, , 6";
```

Using altera_attribute to Set Quartus II Logic Options

altera_attribute allows you to apply Quartus II logic options and assignments to an object in your HDL source code. You can set this attribute on an entity, architecture, instance, register, RAM block, or I/O pin. You cannot set it on an arbitrary combinational node such as a net. With altera_attribute, you can control synthesis options from your HDL source even when the options lack a specific HDL synthesis attribute (such as many of the logic options presented earlier in this chapter). You can also use this attribute to pass entity-level settings and assignments to phases of the Compiler flow that follow Analysis & Synthesis, such as Fitting.

Assignments or settings made through the Quartus II GUI, the .qsf, or the Tcl interface take precedence over assignments or settings made with the altera_attribute synthesis attribute in your HDL code.

The syntax for setting this attribute in HDL is the same as the syntax for other synthesis attributes, as shown in “Synthesis Attributes” on page 13-25.

The attribute value is a single string containing a list of .qsf variable assignments separated by semicolons, as shown in Example 13-107.

Example 13-107. Variable Assignments Separated by Semicolons

```
-name <variable_1> <value_1>;-name <variable_2> <value_2>[;...]
```

If the Quartus II option or assignment includes a target, source, and section tag, you must use the syntax in Example 13-108 for each .qsf variable assignment:

Example 13-108. Syntax for Each .qsf Variable Assignment

```
-name <variable> <value>  
-from <source> -to <target> -section_id <section>
```

The syntax for the full attribute value, including the optional target, source, and section tags for two different .qsf assignments, is shown in Example 13-109.

Example 13-109. Syntax for Fill Attribute Value

```
" -name <variable_1> <value_1> [-from <source_1>] [-to <target_1>] [-section_id \  
<section_1>]; -name <variable_2> <value_2> [-from <source_2>] [-to <target_2>] \  
[-section_id <section_2>] "
```

If the assigned value of a variable is a string of text, you must use escaped quotes around the value in Verilog HDL (as shown in Example 13-110), or double-quotes in VHDL (as shown in Example 13-111):

Example 13-110. Assigned Value of a Variable in Verilog HDL (With Nonexistent Variable and Value Terms)

```
"VARIABLE_NAME \"STRING_VALUE\" "
```

Example 13-111. Assigned Value of a Variable in VHDL (With Nonexistent Variable and Value Terms)

```
"VARIABLE_NAME "STRING_VALUE" "
```

To find the `.qsf` variable name or value corresponding to a specific Quartus II option or assignment, you can make the option setting or assignment in the Quartus II GUI and then note the changes in the `.qsf`. You can also refer to the *Quartus II Settings File Manual*, which documents all variable names.

[Example 13-112](#) through [Example 13-114](#) use `altera_attribute` to set the power-up level of an inferred register.



For inferred instances, you cannot apply the attribute to the instance directly. Therefore, you must apply the attribute to one of the output nets of the instance. The Quartus II software automatically moves the attribute to the inferred instance.

Example 13-112. Verilog-1995 Code: Applying altera_attribute to an Instance

```
reg my_reg /* synthesis altera_attribute = "-name POWER_UP_LEVEL HIGH"
*/;
```

Example 13-113. Verilog-2001 Code: Applying altera_attribute to an Instance

```
(* altera_attribute = "-name POWER_UP_LEVEL HIGH" *) reg my_reg;
```

Example 13-114. VHDL Code: Applying altera_attribute to an Instance

```
signal my_reg : std_logic;
attribute altera_attribute : string;
attribute altera_attribute of my_reg: signal is "-name POWER_UP_LEVEL
HIGH";
```

Example 13–115 through Example 13–117 use the `altera_attribute` to disable the **Auto Shift Register Replacement** synthesis option for an entity. To apply the Altera Attribute to a VHDL entity, you must set the attribute on its architecture rather than on the entity itself.

Example 13–115. Verilog-1995 Code: Applying `altera_attribute` to an Entity

```
module my_entity(...) /* synthesis altera_attribute = "-name  
AUTO_SHIFT_REGISTER_RECOGNITION OFF" */;
```

Example 13–116. Verilog-2001 Code: Applying `altera_attribute` to an Entity

```
(* altera_attribute = "-name AUTO_SHIFT_REGISTER_RECOGNITION OFF" *)  
module my_entity(...) ;
```

Example 13–117. VHDL Code: Applying `altera_attribute` to an Entity

```
entity my_entity is  
-- Declare generics and ports  
end my_entity;  
architecture rtl of my_entity is  
attribute altera_attribute : string;  
-- Attribute set on architecture, not entity  
attribute altera_attribute of rtl: architecture is "-name  
AUTO_SHIFT_REGISTER_RECOGNITION OFF";  
begin  
-- The architecture body  
end rtl;
```

You can also use `altera_attribute` for more complex assignments that have more than one instance. In Example 13–119 through Example 13–121, the `altera_attribute` cuts all timing paths from `reg1` to `reg2`, equivalent to this Tcl or QSF command, as shown in Example 13–118:

Example 13–118.

```
set_instance_assignment -name CUT ON -from reg1 -to reg2 ←
```

Example 13–119. Verilog-1995 Code: Applying `altera_attribute` with the `-to` Option

```
reg reg2;  
reg reg1 /* synthesis altera_attribute = "-name CUT ON -to reg2" */;
```

Example 13–120. Verilog-2001 and SystemVerilog Code: Applying `altera_attribute` with the `-to` Option

```
reg reg2;  
(* altera_attribute = "-name CUT ON -to reg2" *) reg reg1;
```

Example 13–121. VHDL Code: Applying `altera_attribute` with the `-to` Option

```
signal reg1, reg2 : std_logic;  
attribute altera_attribute: string;  
attribute altera_attribute of reg1 : signal is "-name CUT ON -to reg2";
```

You can specify either the `-to` option or the `-from` option in a single `altera_attribute`; integrated synthesis automatically sets the remaining option to the target of the `altera_attribute`. You can also specify wildcards for either option. For example, if you specify `"*"` for the `-to` option instead of `reg2` in these examples, the Quartus II software cuts all timing paths from `reg1` to every other register in this design entity.

You can use the `altera_attribute` only for entity-level settings, and the assignments (including wildcards) apply only to the current entity.

Analyzing Synthesis Results

After performing synthesis, you can check your synthesis results in the **Analysis & Synthesis** section of the Compilation Report and the Project Navigator.

Analysis & Synthesis Section of the Compilation Report

The Compilation Report, which provides a summary of results for the project, appears after a successful compilation. After Analysis & Synthesis, the Summary section of the Compilation Report provides a summary of utilization based on synthesis data, before Fitter optimizations have occurred. Synthesis-specific information is listed in the **Analysis & Synthesis** section.

Analysis & Synthesis includes various report sections, including a list of the source files read for the project, the resource utilization by entity after synthesis, and information about state machines, latches, optimization results, and parameter settings.

- ④ For more information about each report section, refer to the *Analysis & Synthesis Summary Reports* in Quartus II Help.

Project Navigator

The **Hierarchy** tab of the Project Navigator provides a summary of resource information about the entities in the project. After Analysis & Synthesis, before the Fitter begins, the Project Navigator provides a summary of utilization based on synthesis data, before Fitter optimizations have occurred.

If you hold your mouse pointer over one of the entities in the **Hierarchy** tab, a tooltip appears that shows parameter information for each instance.

Analyzing and Controlling Synthesis Messages

This section provides information about the messages generated during synthesis, and how you can control which messages appear during compilation.

Quartus II Messages

The messages that appear during Analysis & Synthesis describe many of the optimizations that the software performs during the synthesis stage, and provide information about how the design is interpreted. Altera recommends checking the messages to analyze **Critical Warnings** and **Warnings**, because these messages can relate to important design problems. You should also read the information messages **Info** and **Extra Info** to get more information about how the software processes your design.

The **Info**, **Extra Info**, **Warning**, **Critical Warning**, and **Error** tabs display messages grouped by type.

 For more information about the Messages window and message suppression, refer to *About the Messages Window* and *About Message Suppression* in Quartus II Help.

 For more information about the Messages, refer to *Managing Quartus II Projects* chapter in volume 2 of the *Quartus II Handbook*.

You can specify the type of Analysis & Synthesis messages that you want to view by selecting the **Analysis & Synthesis Message Level** option. You can specify the display level by performing the following steps:

1. On the Assignments menu, click **Settings**.
2. In the **Category** list, click **Analysis & Synthesis Settings**.
3. Click **More Settings**. Select the level for the **Analysis & Synthesis Message Level** option.

VHDL and Verilog HDL Messages

The Quartus II software issues a variety of messages when it is analyzing and elaborating the Verilog HDL and VHDL files in your design. These HDL messages are a subset of all Quartus II messages that help you identify potential problems early in the design process.

HDL messages fall into the following categories:

- **Info message**—Lists a property of your design.
- **Warning message**—Indicates a potential problem in your design. Potential problems come from a variety of sources, including typos, inappropriate design practices, or the functional limitations of your target device. Though HDL warning messages do not always identify actual problems, Altera recommends investigating code that generates an HDL warning. Otherwise, the synthesized behavior of your design might not match your original intent or its simulated behavior.
- **Error message**—Indicates an actual problem with your design. Your HDL code can be invalid due to a syntax or semantic error, or it might not be synthesizable as written.

In [Example 13-122](#), the sensitivity list contains multiple copies of the variable `i`. While the Verilog HDL language does not prohibit duplicate entries in a sensitivity list, it is clear that this design has a typing error: Variable `j` should be listed on the sensitivity list to avoid a possible simulation or synthesis mismatch.

Example 13-122. Generating an HDL Warning Message

```
//dup.v
module dup(input i, input j, output reg o);
always @ (i or i)
    o = i & j;
endmodule
```

When processing the HDL code, the Quartus II software generates the following warning message, as shown in [Example 13-123](#):

Example 13-123.

```
Warning: (10276) Verilog HDL sensitivity list warning at dup.v(2):
sensitivity list contains multiple entries for "i".
```

In Verilog HDL, variable names are case-sensitive, so the variables `my_reg` and `MY_REG` in [Example 13-124](#) are two different variables. However, declaring variables which have names in different cases is potentially confusing, especially if you use VHDL, in which variables are not case-sensitive.

Example 13-124. Generating HDL Info Messages

```
// namecase.v
module namecase (input i, output o);
    reg my_reg;
    reg MY_REG;
    assign o = i;
endmodule
```

When processing the HDL code, the Quartus II software generates the following informational message, as shown in [Example 13-125](#):

Example 13-125.

```
Info: (10281) Verilog HDL information at namecase.v(3): variable name
"MY_REG" and variable name "my_reg" should not differ only in case.
```

In addition, the Quartus II software generates additional HDL info messages to inform you that neither `my_reg` nor `MY_REG` are used in this small design, as shown in [Example 13-126](#):

Example 13-126.

```
Info: (10035) Verilog HDL or VHDL information at namecase.v(3): object
"my_reg" declared but not used
Info: (10035) Verilog HDL or VHDL information at namecase.v(4): object
"MY_REG" declared but not used
```

The Quartus II software allows you to control how many HDL messages you can view during the Analysis & Elaboration of your design files. You can set the HDL Message Level to enable or disable groups of HDL messages, or you can enable or disable specific messages, as described in the following sections.

For more information about synthesis directives and their syntax, refer to “[Synthesis Directives](#)” on page 13-27.

Setting the HDL Message Level

The HDL Message Level specifies the types of messages that the Quartus II software displays when it is analyzing and elaborating your design files. [Table 13-8](#) describes the HDL message levels.

Table 13-8. HDL Info Message Level

Level	Purpose	Description
Level1	High-severity messages only	If you only want to view the HDL messages that identify likely problems with your design, select Level1. When Level1 is selected, the Quartus II software issues a message only if there is a high probability that it points to an actual problem with your design.
Level2	High-severity and medium-severity messages	If you want to view additional HDL messages that identify possible problems with your design, select Level2. This is the default setting.
Level3	All messages, including low-severity messages	If you want to view all HDL info and warning messages, select Level3. This level includes extra “LINT” messages that suggest changes to improve the style of your HDL code or make it easier to understand.

You must address all issues reported at the **Level1** setting. The default HDL message level is **Level2**.

To set the HDL Message Level in the GUI, follow these steps:

1. On the Assignments menu, click **Settings**.
2. In the **Category** list, click **Analysis & Synthesis Settings**.
3. Set the desired message level from the pull-down menu in the **HDL Message Level** list, and then click **OK**.

You can override this default setting in a source file with the `message_level` synthesis directive, which takes the values `level1`, `level2`, and `level3`, as shown in [Example 13-127](#) and [Example 13-128](#).

Example 13-127. Verilog HDL Examples of `message_level` Directive

```
// altera message_level level1
or
/* altera message_level level3 */
```

Example 13-128. VHDL Code: `message_level` Directive

```
-- altera message_level level2
```

A `message_level` synthesis directive remains effective until the end of a file or until the next `message_level` directive. In VHDL, you can use the `message_level` synthesis directive to set the HDL Message Level for entities and architectures, but not for other design units. An HDL Message Level for an entity applies to its architectures, unless overridden by another `message_level` directive. In Verilog HDL, you can use the `message_level` directive to set the HDL Message Level for a module.

Enabling or Disabling Specific HDL Messages by Module/Entity

You can enable or disable a specific HDL info or warning message with its Message ID, which is displayed in parentheses at the beginning of the message. Enabling or disabling a specific message overrides its HDL Message Level. This method is different from the message suppression in the Messages window because you can disable messages for a specific module or entity. This method applies only to the HDL messages, and if you disable a message with this method, the message is listed as a suppressed message in the Quartus II GUI.

To disable specific HDL messages in the GUI, follow these steps:

1. On the Assignments menu, click **Settings**.
2. In the **Category** list, expand **Analysis & Synthesis Settings** and select **Advanced**.
3. In the **Advanced Message Settings** dialog box, add the Message IDs you want to enable or disable.

To enable or disable specific HDL messages in your HDL, use the `message_on` and `message_off` synthesis directives. These directives require a space-separated list of Message IDs. You can enable or disable messages with these synthesis directives immediately before Verilog HDL modules, VHDL entities, or VHDL architectures. You cannot enable or disable a message during an HDL construct.

A message enabled or disabled via a `message_on` or `message_off` synthesis directive overrides its HDL Message Level or any `message_level` synthesis directive. The message remains disabled until the end of the source file or until its status is changed by another `message_on` or `message_off` directive.

Example 13-129. Verilog HDL `message_off` Directive for Message with ID 10000

```
// altera message_off 10000  
or  
/* altera message_off 10000 */
```

Example 13-130. VHDL `message_off` Directive for Message with ID 10000

```
-- altera message_off 10000
```

Node-Naming Conventions in Quartus II Integrated Synthesis

This section provides an overview of the conventions used by the Quartus II software during synthesis to name the nodes created from your HDL design. This information is useful for finding logic node names during verification and debugging of a design. This section focuses on the conventions for Verilog HDL and VHDL code, but AHDL and .bdf files are discussed when appropriate.

Whenever possible, Quartus II integrated synthesis uses wire or signal names from your source code to name nodes such as LEs or ALMs. Some nodes, such as registers, have predictable names that do not change when a design is resynthesized, although certain optimizations can affect register names. The names of other nodes, particularly LEs or ALMs that contain only combinational logic, can change due to logic optimizations that the software performs.

This section discusses the following topics:

- “Hierarchical Node-Naming Conventions”
- “Node-Naming Conventions for Registers (DFF or D Flipflop Atoms)”
- “Register Changes During Synthesis” on page 13-79
- “Preserving Register Names” on page 13-81
- “Node-Naming Conventions for Combinational Logic Cells” on page 13-81
- “Preserving Combinational Logic Names” on page 13-82

Hierarchical Node-Naming Conventions

To make each name in the design unique, the Quartus II software adds the hierarchy path to the beginning of each name. The “|” separator indicates a level of hierarchy. For each instance in the hierarchy, the software adds the entity name and the instance name of that entity, using the “:” separator between each entity name and its instance name. For example, if a design defines entity A with the name `my_A_inst`, the hierarchy path of that entity would be `A:my_A_inst`. The full name of any node is obtained by starting with the hierarchical instance path, followed by a “|”, and ending with the node name inside that entity, using the convention shown in [Example 13-131](#):

Example 13-131.

```
<entity 0> : <instance_name 0> | <entity 1> : <instance_name 1> | . . . | <instance_name n> | <node_name>
```

For example, if entity A contains a register (DFF atom) called `my_dff`, its full hierarchy name would be `A:my_A_inst|my_dff`.

To instruct the Compiler to generate node names that do not contain entity names, on the **Compilation Process Settings** page of the **Settings** dialog box, click **More Settings**, and then turn off **Display entity name for node name**. With this option turned off, the node names use the convention in shown in [Example 13-132](#):

Example 13-132.

```
<instance_name 0> | <instance_name 1> | . . . | <instance_name n> | <node_name>
```

Node-Naming Conventions for Registers (DFF or D Flipflop Atoms)

In Verilog HDL and VHDL, inferred registers are named after the reg or signal connected to the output.

[Example 13-133](#) is an example of a register in Verilog HDL that creates a DFF primitive called `my_dff_out`:

Example 13-133. Verilog HDL Register

```
wire dff_in, my_dff_out, clk;

always @ (posedge clk)
my_dff_out <= dff_in;
```

Similarly, [Example 13-134](#) is an example of a register in VHDL that creates a DFF primitive called `my_dff_out`.

Example 13-134. VHDL Register

```
signal dff_in, my_dff_out, clk;
process (clk)
begin
if (rising_edge(clk)) then
my_dff_out <= dff_in;
end if;
end process;
```

In AHDL designs, DFF registers are declared explicitly rather than inferred, so the software uses the user-declared name for the register.

For schematic designs using a `.bdf`, all elements are given a name when they are instantiated in the design, so the software uses the name you defined for the register or DFF.

In the special case that a wire or signal (such as `my_dff_out` in the preceding examples) is also an output pin of your top-level design, the Quartus II software cannot use that name for the register (for example, cannot use `my_dff_out`) because the software requires that all logic and I/O cells have unique names. In this case, Quartus II integrated synthesis appends `~reg0` to the register name.

For example, the Verilog HDL code in [Example 13-135](#) generates a register called `q~reg0`:

Example 13-135. Verilog HDL Register Feeding Output Pin

```
module my_dff (input clk, input d, output q);
always @ (posedge clk)
q <= d;
endmodule
```

This situation occurs only for registers driving top-level pins. If a register drives a port of a lower level of the hierarchy, the port is removed during hierarchy flattening and the register retains its original name, in this case, `q`.

Register Changes During Synthesis

On some occasions, you might not find registers that you expect to view in the synthesis netlist. Registers might be removed by logic optimization, or their names might be changed due to synthesis optimization. Common optimizations include inference of a state machine, counter, adder-subtractor, or shift register from registers and surrounding logic. Other common register changes occur when registers are packed into dedicated hardware on the FPGA, such as a DSP block or a RAM block.

This section describes the following factors that can affect register names:

- “Synthesis and Fitting Optimizations”
- “State Machines”
- “Inferred Adder-Subtractors, Shift Registers, Memory, and DSP Functions” on page 13–80
- “Packed Input and Output Registers of RAM and DSP Blocks” on page 13–81
- “Preserving Register Names” on page 13–81
- “Preserving Combinational Logic Names” on page 13–82

Synthesis and Fitting Optimizations

Registers might be removed by logic optimization during synthesis if they are not connected to inputs or outputs in the design, or if the logic can be simplified due to constant signal values. Register names might also be changed due to synthesis optimizations, such as when duplicate registers are merged together to reduce resource utilization.

NOT-gate push back optimizations can affect registers that use preset signals. This type of optimization can impact your timing assignments when registers are used as clock dividers. If this situation occurs in your design, change the clock settings to work on the new register name.

Synthesis netlist optimizations often change node names because registers can be combined or duplicated to optimize the design.



For more information about the type of optimizations performed by synthesis netlist optimizations, refer to the *Netlist Optimizations and Physical Synthesis* chapter in volume 2 of the *Quartus II Handbook*.

The Quartus II Compilation Report provides a list of registers that are removed during synthesis optimizations, and a brief reason for the removal. To generate the Quartus II Compilation Report, follow these steps:

1. In the **Analysis & Synthesis** folder, open **Optimization Results**.
2. Open **Register Statistics**, and then click the **Registers Removed During Synthesis** report.
3. Click **Removed Registers Triggering Further Register Optimizations**.

The second report contains a list of registers that caused other registers to be removed from the design. The report provides a brief reason for the removal, and a list of registers that were removed due to the removal of the initial register.

Quartus II integrated synthesis creates synonyms for registers duplicated with the **Maximum Fan-Out** option (or `maxfan` attribute). Therefore, timing assignments applied to nodes that are duplicated with this option are applied to the new nodes as well.

The Quartus II Fitter can also change node names after synthesis (for example, when the Fitter uses register packing to pack a register into an I/O element, or when logic is modified by physical synthesis). The Fitter creates synonyms for duplicated registers so timing analysis can use the existing node name when applying assignments.

You can instruct the Quartus II software to preserve certain nodes throughout compilation so you can use them for verification or making assignments. For more information, refer to [“Preserving Register Names” on page 13-81](#).

State Machines

If a state machine is inferred from your HDL code, the registers that represent the states are mapped into a new set of registers that implement the state machine. Most commonly, the software converts the state machine into a one-hot form where each state is represented by one register. In this case, for Verilog HDL or VHDL designs, the registers are named according to the name of the state register and the states.

For example, consider a Verilog HDL state machine in which the states are parameter `state0 = 1`, `state1 = 2`, `state2 = 3`, and in which the state machine register is declared as `reg [1:0] my_fsm`. In this example, the three one-hot state registers are named `my_fsm.state0`, `my_fsm.state1`, and `my_fsm.state2`.

In AHDL, state machines are explicitly specified with a state machine name. State machine registers are given synthesized names based on the state machine name, but not the state names. For example, if a state machine is called `my_fsm` and has four state bits, they might be synthesized with names such as `my_fsm~12`, `my_fsm~13`, `my_fsm~14`, and `my_fsm~15`.

Inferred Adder-Subtractors, Shift Registers, Memory, and DSP Functions

The Quartus II software infers megafunctions from Verilog HDL and VHDL code for logic that forms adder-subtractors, shift registers, RAM, ROM, and arithmetic functions that are placed in DSP blocks.

 For information about inferring megafunctions, refer to the [Recommended HDL Coding Styles](#) chapter in volume 1 of the *Quartus II Handbook*.

Because adder-subtractors are part of a megafunction instead of generic logic, the combinational logic exists in the design with different names. For shift registers, memory, and DSP functions, the registers and logic are implemented inside the dedicated RAM or DSP blocks in the device. Thus, the registers are not visible as separate LEs or ALMs.

Packed Input and Output Registers of RAM and DSP Blocks

Registers are packed into the input registers and output registers of RAM and DSP blocks, so that they are not visible as separate registers in LEs or ALMs.



For information about packing registers into RAM and DSP megafunctions, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*.

Preserving Register Names

Altera recommends that you preserve certain register names for verification or debugging, or to ensure that timing assignments are applied correctly. Quartus II integrated synthesis preserves certain nodes automatically if they might be used in a timing constraint.

Use the `preserve` attribute to instruct the Compiler not to minimize or remove a specified register during synthesis optimizations or register netlist optimizations. For details, refer to “[Preserve Registers](#)” on page 13-43.

Use the `noprune` attribute to preserve a fan-out-free register through the entire compilation flow. For details, refer to “[Noprune Synthesis Attribute/Preserve Fan-out Free Register Node](#)” on page 13-44.

Use the synthesis attribute `syn_dont_merge` to ensure that the registers are not merged with other registers, and other registers are not merged with them. For details, refer to “[Disable Register Merging/Don't Merge Register](#)” on page 13-43.

Node-Naming Conventions for Combinational Logic Cells

Whenever possible for Verilog HDL, VHDL, and AHDL code, the Quartus II software uses wire names that are the targets of assignments, but can change the node names due to synthesis optimizations.

For example, consider the Verilog HDL code in [Example 13-136](#). Quartus II integrated synthesis uses the names `c`, `d`, `e`, and `f` for the generated combinational logic cells.

Example 13-136. Naming Nodes for Combinational Logic Cells in Verilog HDL

```
wire c;
reg d, e, f;

assign c = a | b;
always @ (a or b)
d = a & b;
always @ (a or b) begin : my_label
e = a ^ b;
end

always @ (a or b)
f = ~(a | b);
```

For schematic designs using a `.bdf`, all elements are given a name when they are instantiated in the design and the software uses the name you defined when possible.

If logic cells, such as those created in [Example 13-136](#), are packed with registers in device architectures such as the Stratix and Cyclone device families, those names might not appear in the netlist after fitting. In other devices, such as newer families in the Stratix and Cyclone series device families, the register and combinational nodes are kept separate throughout the compilation, so these names are more often maintained through fitting.

When logic optimizations occur during synthesis, it is not always possible to retain the initial names as described. In some cases, synthesized names are used, which are the wire names with a tilde (~) and a number appended. For example, if a complex expression is assigned to wire *w* and that expression generates several logic cells, those cells can have names such as *w*, *w~1*, and *w~2*. Sometimes the original wire name *w* is removed, and an arbitrary name such as *rt1~123* is created. Quartus II integrated synthesis attempts to retain user names whenever possible. Any node name ending with *~<number>* is a name created during synthesis, which can change if the design is changed and re-synthesized. Knowing these naming conventions helps you understand your post-synthesis results, helping you to debug your design or create assignments.

The software maintains combinational clock logic by ensuring nodes that might be a clock are not changed during synthesis. The software also maintains or protects multiplexers in clock trees, so that the TimeQuest analyzer has information about which paths are unate, to allow complete and correct analysis of combinational clocks. Multiplexers often occur in clock trees when the software selects between different clocks. To help with the analysis of clock trees, the software ensures that each multiplexer encountered in a clock tree is broken into 2:1 multiplexers, and each of those 2:1 multiplexers is mapped into one lookup table (independent of the device family). This optimization might result in a slight increase in area, and for some designs a decrease in timing performance. You can turn off this multiplexer protection with the option **Clock MUX Protection** under **More Settings** on the **Analysis & Synthesis Settings** page of the **Settings** dialog box.

- ❓ For more information about **Clock MUX Protection** logic option and a list of supported devices, refer to [Clock MUX Protection logic option](#) in Quartus II Help.

Preserving Combinational Logic Names

You can preserve certain combinational logic node names for verification or debugging, or to ensure that timing assignments are applied correctly.

Use the `keep` attribute to keep a wire name or combinational node name through logic synthesis minimizations and netlist optimizations. For details, refer to [“Keep Combinational Node/Implement as Output of Logic Cell”](#) on page 13-45.

For any internal node in your design clock network, use `keep` to protect the name so that you can apply correct clock settings. Also, set the attribute for combinational logic involved in cut and -through assignments.

- 👉 Setting the `keep` attribute for combinational logic can increase the area utilization and increase the delay of the final mapped logic because the attribute requires the insertion of extra combinational logic. Use the attribute only when necessary.

Scripting Support

You can run procedures and make settings in a Tcl script. You can also run some procedures at a command prompt. For detailed information about scripting command options, refer to the Quartus II Command-Line and Tcl API Help browser.

To run the Help browser, type the command at the command prompt shown in [Example 13-137](#):

Example 13-137.

```
quartus_sh --qhelp ←
```

 For more information about Tcl scripting, refer to the [Tcl Scripting](#) chapter in volume 2 of the *Quartus II Handbook*. For more information about all settings and constraints in the Quartus II software, refer to the [Quartus II Settings File Manual](#). For more information about command-line scripting, refer to the [Command-Line Scripting](#) chapter in volume 2 of the *Quartus II Handbook*.

 For more information about Tcl scripting, refer to [API Functions for Tcl](#) in Quartus II Help.

You can specify many of the options described in this section either on an instance, at the global level, or both.

To make a global assignment, use the Tcl command shown in [Example 13-138](#):

Example 13-138.

```
set_global_assignment -name <QSF Variable Name> <Value> ←
```

To make an instance assignment, use the Tcl command shown in [Example 13-139](#):

Example 13-139.

```
set_instance_assignment -name <QSF Variable Name> <Value>\ -to  
<Instance Name> ←
```

To set the **Synthesis Effort** option at the command line, use the `--effort` option with the `quartus_map` executable, as shown in [Example 13-140](#).

Example 13-140. Command Syntax for Specifying Synthesis Effort Option

```
quartus_map <Design name> --effort= "auto | fast" ←
```

The early timing estimate feature gives you preliminary timing estimates before running a full compilation, which results in a quicker iteration time; therefore, you can save significant compilation time to get a good estimation of the final timing of your design.

If you want to run fast synthesis with the Fitter **Early Timing Estimate** option, use the command shown in [Example 13-141](#). This command runs the full flow with timing analysis:

Example 13-141. Command Syntax for Running Fast Synthesis with Early Timing Estimate Option

```
quartus_sh --flow early_timing_estimate_with_synthesis <Design name> ←
```

For more information, refer to “[Synthesis Effort](#)” on page 13-35.

Adding an HDL File to a Project and Setting the HDL Version

To add an HDL or schematic entry design file to your project, use the Tcl assignments shown in [Example 13-142](#):

Example 13-142.

```
set_global_assignment -name VERILOG_FILE <file name>.<v|sv>
set_global_assignment -name SYSTEMVERILOG_FILE <file name>.sv
set_global_assignment -name VHDL_FILE <file name>.<vhd|vhdl>
set_global_assignment -name AHDL_FILE <file name>.tdf
set_global_assignment -name BDF_FILE <file name>.bdf
```



You can use any file extension for design files, as long as you specify the correct language when adding the design file. For example, you can use **.h** for Verilog HDL header files.

To specify the Verilog HDL or VHDL version, use the option shown in [Example 13-143](#), at the end of the `VERILOG_FILE` or `VHDL_FILE` command:

Example 13-143.

```
- HDL_VERSION <language version> ←
```

The variable *<language version>* takes one of the following values:

- VERILOG_1995
- VERILOG_2001
- SYSTEMVERILOG_2005
- VHDL_1987
- VHDL_1993
- VHDL_2008

For example, to add a Verilog HDL file called **my_file.v** written in Verilog-1995, use the command shown in [Example 13-144](#):

Example 13-144.

```
set_global_assignment -name VERILOG_FILE my_file.v -HDL_VERSION \
VERILOG_1995
```

In [Example 13-145](#), the `syn_encoding` attribute associates a binary encoding with the states in the enumerated type `count_state`. In this example, the states are encoded with the following values: zero = "11", one = "01", two = "10", three = "00".

Example 13-145. Specifying User-Encoded States with the `syn_encoding` Attribute in VHDL

```
ARCHITECTURE rtl OF my_fsm IS
    TYPE count_state IS (zero, one, two, three);
    ATTRIBUTE syn_encoding : STRING;
    ATTRIBUTE syn_encoding OF count_state : TYPE IS "11 01 10 00";
    SIGNAL present_state, next_state : count_state;
BEGIN
```

You can also use the `syn_encoding` attribute in Verilog HDL to direct the synthesis tool to use the encoding from your HDL code, instead of using the **State Machine Processing** option.

The `syn_encoding` value "user" instructs the Quartus II software to encode each state with its corresponding value from the Verilog HDL source code. By changing the values of your state constants, you can change the encoding of your state machine.

In [Example 13-146](#), the states are encoded as follows:

```
init = "00"
last = "11"
next = "01"
later = "10"
```

Example 13-146. Verilog-2001 and SystemVerilog Code: Specifying User-Encoded States with the `syn_encoding` Attribute

```
(* syn_encoding = "user" *) reg [1:0] state;
parameter init = 0, last = 3, next = 1, later = 2;
always @ (state) begin
    case (state)
        init:
            out = 2'b01;
        next:
            out = 2'b10;
        later:
            out = 2'b11;
        last:
            out = 2'b00;
    endcase
end
```

Without the `syn_encoding` attribute, the Quartus II software encodes the state machine based on the current value of the **State Machine Processing** logic option.

If you also specify a safe state machine (as described in ["Safe State Machines" on page 13-39](#)), separate the encoding style value in the quotation marks from the safe value with a comma, as follows: "safe, one-hot" or "safe, gray".

For more information, refer to ["Manually Specifying State Assignments Using the `syn_encoding` Attribute" on page 13-37](#).

Assigning a Pin

To assign a signal to a pin or device location, use the Tcl command shown in [Example 13-147](#):

Example 13-147.

```
set_location_assignment -to <signal name> <location>
```

Valid locations are pin location names. Some device families also support edge and I/O bank locations. Edge locations are `EDGE_BOTTOM`, `EDGE_LEFT`, `EDGE_TOP`, and `EDGE_RIGHT`. I/O bank locations include `IOBANK_1` to `IOBANK_n`, in which `n` is the number of I/O banks in a particular device.

Creating Design Partitions for Incremental Compilation

To create a partition, use the command shown in [Example 13-148](#):

Example 13-148.

```
set_instance_assignment -name PARTITION_HIERARCHY \  
<file name> -to <destination> -section_id <partition name>
```

The *<file name>* is the name used for internally generated netlist files during incremental compilation. If you create the partition in the Quartus II GUI, netlist files are named automatically by the Quartus II software based on the instance name. If you use Tcl to create your partitions, you must assign a custom file name that is unique across all partitions. For the top-level partition, the specified file name is ignored, and you can use any dummy value. To ensure the names are safe and platform independent, file names should be unique, regardless of case. For example, if a partition uses the file name `my_file`, no other partition can use the file name `MY_FILE`. To make file naming simple, Altera recommends that you base each file name on the corresponding instance name for the partition.

The *<destination>* is the short hierarchy path of the entity. A *short* hierarchy path is the full hierarchy path without the top-level name, for example: `"ram:ram_unit|altsyncram:altsyncram_component"` (with quotation marks). For the top-level partition, you can use the pipe (`|`) symbol to represent the top-level entity.

For more information about hierarchical naming conventions, refer to [“Node-Naming Conventions in Quartus II Integrated Synthesis”](#) on page 13-77.

The *<partition name>* is the partition name you designate, which should be unique and less than 1024 characters long. The name may only consist of alphanumeric characters, as well as pipe (`|`), colon (`:`), and underscore (`_`) characters. Altera recommends enclosing the name in double quotation marks (`" "`).

Quartus II Synthesis Options

Table 13-9 lists the .qsf variable names and applicable values for the settings discussed in this chapter. Use the .qsf variable name in the Tcl assignment to make the setting along with the appropriate value.

 Applying a Quartus II Synthesis option globally or to an entity affects all lower-level entities in the hierarchy path, including entities instantiated with Altera and third-party IP.

Table 13-9. Quartus II Synthesis Options (Part 1 of 3) *(Note 1)*

Setting Name	Quartus II Settings File Variable	Values
Add Pass-Through Logic to Inferred RAMs	ADD_PASS_THROUGH_LOGIC_TO_INFERRED_RAMs	On/Off
Allow Any RAM Size for Recognition	ALLOW_ANY_RAM_SIZE_FOR_RECOGNITION	On/Off
Allow Any ROM Size for Recognition	ALLOW_ANY_ROM_SIZE_FOR_RECOGNITION	On/Off
Allow Any Shift Register Size for Recognition	ALLOW_ANY_SHIFT_REGISTER_SIZE_FOR_RECOGNITION	On/Off
Allow Asynchronous Clear Usage For Shift Register Replacement	ALLOW_ACLR_FOR_SHIFT_REGISTER_RECOGNITION	On/Off
Allow Synchronous Control Signals	ALLOW_SYNCH_CTRL_USAGE	On/Off
Analysis & Synthesis Message Level	SYNTH_MESSAGE_LEVEL	Low/Medium/High
Auto Carry Chains	AUTO_CARRY_CHAINS	On/Off
Auto Clock Enable Replacement	AUTO_CLOCK_ENABLE_RECOGNITION	On/Off
Auto DSP Block Replacement	AUTO_DSP_RECOGNITION	On/Off
Auto Gated Clock Conversion	SYNTH_GATED_CLOCK_CONVERSION	On/Off
Auto Open-Drain Pins	AUTO_OPEN_DRAIN_PINS	On/Off
Auto RAM Block Balancing	AUTO_RAM_BLOCK_BALANCING	On/Off
Auto RAM to Logic Cell Conversion	AUTO_RAM_TO_LCELL_CONVERSION	On/Off
Auto RAM Replacement	AUTO_RAM_RECOGNITION	On/Off
Auto Resource Sharing	AUTO_RESOURCE_SHARING	On/Off
Auto ROM Replacement	AUTO_ROM_RECOGNITION	On/Off
Auto Shift-Register Replacement	AUTO_SHIFT_REGISTER_RECOGNITION	Always/Auto/Off
Block Design Naming	BLOCK_DESIGN_NAMING	Auto/Max+Plus II/ Quartus II
Carry Chain Length	<device name>_CARRY_CHAIN_LENGTH	<Maximum allowable length of a chain>
Clock MUX Protection	SYNTH_CLOCK_MUX_PROTECTION	On/Off
Create Debugging Nodes for IP Cores	ENABLE_IP_DEBUG	On/Off

Table 13-9. Quartus II Synthesis Options (Part 2 of 3) (Note 1)

Setting Name	Quartus II Settings File Variable	Values
DSP Block Balancing	DSP_BLOCK_BALANCING	Auto/DSP Blocks/ Logic Elements/ Off/Simple 18-bit Multipliers/ Simple Multipliers/Width 18-bit Multipliers
Extract Verilog State Machines	EXTRACT_VERILOG_STATE_MACHINES	On/Off
Extract VHDL State Machines	EXTRACT_VHDL_STATE_MACHINES	On/Off
Force Use of Synchronous Clear Signals	FORCE_SYNCH_CLEAR	On/Off
HDL Message Level	HDL_MESSAGE_LEVEL	Level1/Level2/ Level3
Ignore CARRY Buffers	IGNORE_CARRY_BUFFERS	On/Off
Ignore CASCADE Buffers	IGNORE_CASCADE_BUFFERS	On/Off
Ignore GLOBAL Buffers	IGNORE_GLOBAL_BUFFERS	On/Off
Ignore Maximum Fan-Out Assignments	IGNORE_MAX_FANOUT_ASSIGNMENTS	On/Off
Ignore ROW GLOBAL Buffers	IGNORE_ROW_GLOBAL_BUFFERS	On/Off
Ignore translate_off and synthesis_off directives	IGNORE_TRANSLATE_OFF_AND_SYNTHESIS_OFF	On/Off
Ignore Verilog Initial Constructs	IGNORE_VERILOG_INITIAL_CONSTRUCTS	On/Off
Iteration limit for constant Verilog loops	VERILOG_CONSTANT_LOOP_LIMIT	<Maximum limit to infinite loops before exhaustion of memory>
Iteration limit for non-constant Verilog loops	VERILOG_NON_CONSTANT_LOOP_LIMIT	<Maximum limit to infinite loops before exhaustion of memory>
Limit AHDL Integers to 32 Bits	LIMIT_AHDL_INTEGERS_TO_32_BITS	On/Off
Maximum DSP Block Usage (2)	MAX_BALANCING_DSP_BLOCKS	<Maximum DSP Block Usage Value>
Maximum Number of M4K/M9K Memory Blocks	MAX_RAM_BLOCKS_M4K	<Maximum memory blocks usage>
Maximum Number of M512 Memory Blocks	MAX_RAM_BLOCKS_M512	<Maximum memory blocks usage>
Maximum Number of M-RAM/M144K Memory Blocks	MAX_RAM_BLOCKS_MRAM	<Maximum memory blocks usage>
Maximum Number of LABs	MAX_LABS	<Maximum number of LAB usage>
NOT Gate Push-Back	NOT_GATE_PUSH_BACK	On/Off
Number of Inverted Registers Reported in Synthesis Report	NUMBER_OF_INVERTED_REGISTERS_REPORTED	<Maximum number of inverted registers>
Number of Removed Registers Reported in Synthesis Report	NUMBER_OF_REMOVED_REGISTERS_REPORTED	<Maximum number of inverted registers>
Optimization Technique	<device family>_OPTIMIZATION_TECHNIQUE	Area/Speed/ Balanced
Parallel Synthesis	PARALLEL_SYNTHESIS	On/Off

Table 13–9. Quartus II Synthesis Options (Part 3 of 3) (Note 1)

Setting Name	Quartus II Settings File Variable	Values
Perform WYSIWYG Primitive Resynthesis	ADV_NETLIST_OPT_SYNTH_WYSIWYG_REMAP	On/Off
PowerPlay Power Optimization	OPTIMIZE_POWER_DURING_SYNTHESIS	Normal compilation/ Extra effort/Off
Power-Up Don't Care (2)	ALLOW_POWER_UP_DONT_CARE	On/Off
Remove Duplicate Registers	REMOVE_DUPLICATE_REGISTERS	On/Off
Remove Redundant Logic Cells (2)	REMOVE_REDUNDANT_LOGIC_CELLS	On/Off
Restructure Multiplexers	MUX_RESTRUCTURE	On/Off/Auto
Resource Aware Inference for Block RAM	SYNTH_RESOURCE_AWARE_INFERENCE_FOR_BLOCK_RAM	On/Off
Safe State Machine	SAFE_STATE_MACHINE	On/Off
SDC Constraint Protection	SYNTH_PROTECT_SDC_CONSTRAINT	On/Off
Show Parameter Settings Tables in Synthesis Report	SHOW_PARAMETER_SETTINGS_TABLES_IN_SYNTHESIS_REPORT	On/Off
State Machine Processing	STATE_MACHINE_PROCESSING	Auto/One-Hot/ Gray/Johnson/ Minimal Bits/ Sequential/ User-Encoded
Strict RAM Replacement	STRICT_RAM_RECOGNITION	On/Off
Synthesis Effort (2)	SYNTHESIS_EFFORT	Auto/Fast
Synthesis Seed (2)	SYNTHESIS_SEED	<Non-negative integer>
Timing Driven Synthesis	SYNTH_TIMING_DRIVEN_SYNTHESIS	On/Off
Use LogicLock Constraints during Resource Balancing	USE_LOGICLOCK_CONSTRAINTS_IN_BALANCING	On/Off

Notes to Table 13–9:

- (1) These settings are supported as Global and Instance settings, unless specified.
(2) This setting is only a Global setting.

Conclusion

The Quartus II software includes Verilog HDL, SystemVerilog, and VHDL language support, as well as support for Altera-specific languages, making the synthesis feature an easy-to-use, standalone solution for Altera designs. You can use the synthesis options in the software or in your HDL code to better control the way your design is synthesized, helping you improve your synthesis results. Use Quartus II reports and messages to analyze your compilation results.

Document Revision History

Table 13-10 shows the revision history for this document.

Table 13-10. Document Revision History (Part 1 of 2)

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Updated “Verilog HDL Support” on page 13-4 to include Verilog-2001 support. ■ Updated “VHDL-2008 Support” on page 13-10 to include the condition operator (explicit and implicit) support. ■ Rewrote “Limiting Resource Usage in Partitions” on page 13-32. ■ Added “Creating LogicLock Regions” on page 13-32 and “Using Assignments to Limit the Number of RAM and DSP Blocks” on page 13-33. ■ Updated “Turning Off the Add Pass-Through Logic to Inferred RAMs no_rw_check Attribute” on page 13-55. ■ Updated “Auto Gated Clock Conversion” on page 13-28. ■ Added links to Quartus II Help.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Removed Referenced Documents section. ■ Added “Synthesis Seed” on page 9-36 section. ■ Updated the following sections: <ul style="list-style-type: none"> ■ “SystemVerilog Support” on page 9-5 ■ “VHDL-2008 Support” on page 9-10 ■ “Using Parameters/Generics” on page 9-16 ■ “Parallel Synthesis” on page 9-21 ■ “Limiting Resource Usage in Partitions” on page 9-32 ■ “Synthesis Effort” on page 9-35 ■ “Synthesis Attributes” on page 9-25 ■ “Synthesis Directives” on page 9-27 ■ “Auto Gated Clock Conversion” on page 9-29 ■ “State Machine Processing” on page 9-36 ■ “Multiply-Accumulators and Multiply-Adders” on page 9-50 ■ “Resource Aware RAM, ROM, and Shift-Register Inference” on page 9-52 ■ “RAM Style and ROM Style—for Inferred Memory” on page 9-53 ■ “Turning Off the Add Pass-Through Logic to Inferred RAMs no_rw_check Attribute” on page 9-55 ■ “Using altera_attribute to Set Quartus II Logic Options” on page 9-68 ■ “Adding an HDL File to a Project and Setting the HDL Version” on page 9-83 ■ “Creating Design Partitions for Incremental Compilation” on page 9-85 ■ “Inferring Multiplier, DSP, and Memory Functions from HDL Code” on page 9-50 ■ Updated Table 9-9 on page 9-86.
December 2009	9.1.1	<ul style="list-style-type: none"> ■ Added information clarifying inheritance of Synthesis settings by lower-level entities, including Altera and third-party IP ■ Updated “Keep Combinational Node/Implement as Output of Logic Cell” on page 9-46

Table 13-10. Document Revision History (Part 2 of 2)

Date	Version	Changes
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Updated the following sections: <ul style="list-style-type: none"> ■ “Initial Constructs and Memory System Tasks” on page 9-7 ■ “VHDL Support” on page 9-9 ■ “Parallel Synthesis” on page 9-21 ■ “Synthesis Directives” on page 9-27 ■ “Timing-Driven Synthesis” on page 9-31 ■ “Safe State Machines” on page 9-40 ■ “RAM Style and ROM Style—for Inferred Memory” on page 9-53 ■ “Translate Off and On / Synthesis Off and On” on page 9-62 ■ “Read Comments as HDL” on page 9-63 ■ “Adding an HDL File to a Project and Setting the HDL Version” on page 9-81 ■ Removed “Remove Redundant Logic Cells” section ■ Added “Resource Aware RAM, ROM, and Shift-Register Inference” section ■ Updated Table 9-9 on page 9-83
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Updated Table 9-9. ■ Updated the following sections: <ul style="list-style-type: none"> ■ “Partitions for Preserving Hierarchical Boundaries” on page 9-20 ■ “Analysis & Synthesis Settings Page of the Settings Dialog Box” on page 9-24 ■ “Timing-Driven Synthesis” on page 9-30 ■ “Turning Off Add Pass-Through Logic to Inferred RAMs/ no_rw_check Attribute Setting” on page 9-54 ■ Added “Parallel Synthesis” on page 9-21 ■ Chapter 9 was previously Chapter 8 in software version 8.1

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter documents support for the Synopsys Synplify software in the Quartus® II software, as well as key design flows, methodologies, and techniques for achieving optimal results in Altera® devices.

This chapter includes the following topics:

- General design flow with the Synplify and Quartus II software
- Synplify software optimization strategies, including timing-driven compilation settings, optimization options, and Altera-specific attributes
- Exporting designs and constraints to the Quartus II software using NativeLink integration
- Guidelines for Altera megafunctions and library of parameterized module (LPM) functions, instantiating them with the MegaWizard™ Plug-In Manager, and tips for inferring them from hardware description language (HDL) code
- Incremental compilation and block-based design, including the MultiPoint flow in the Synplify Pro and Synplify Premier software

The content in this chapter applies to the Synplify, Synplify Pro, and Synplify Premier software unless otherwise specified. This chapter includes the following sections:

- “Altera Device Family Support”
- “Design Flow” on page 14–2
- “Synplify Optimization Strategies” on page 14–6
- “Exporting Designs to the Quartus II Software Using NativeLink Integration” on page 14–13
- “Guidelines for Altera Megafunctions and Architecture-Specific Features” on page 14–16
- “Incremental Compilation and Block-Based Design” on page 14–29



This chapter assumes that you have set up, licensed, and are familiar with the Synplify software.

Altera Device Family Support

The Synplify software supports active devices available in the current version of the Quartus II software. Support for newly released device families may require an overlay. Contact Synopsys at www.synopsys.com for more information.

The Synplify software also supports the FLEX 8000 and MAX 9000 legacy devices that are supported only in the Altera MAX+PLUS® II software, as well as ACEX® 1K, APEX™ II, APEX 20K, APEX 20KC, APEX 20KE, FLEX® 10K, and FLEX 6000 legacy devices that are supported by the Quartus II software version 9.0 and earlier.

Design Flow

The following steps describe a basic Quartus II software design flow using the Synplify software:

1. Create Verilog HDL or VHDL design files.
2. Set up a project in the Synplify software and add the HDL design files for synthesis.
3. Select a target device and add timing constraints and compiler directives in the Synplify software to help optimize the design during synthesis.
4. Synthesize the project in the Synplify software.
5. Create a Quartus II project and import the following files generated by the Synplify software into the Quartus II software. Use the following files for placement and routing, and for performance evaluation:
 - The technology-specific Verilog Quartus Mapping File (**.vqm**) netlist or EDIF Input File (**.edf**) netlist for legacy devices also supported by the MAX+PLUS II software
 - The Synopsys Constraints Format (**.scf**) file for TimeQuest Timing Analyzer constraints

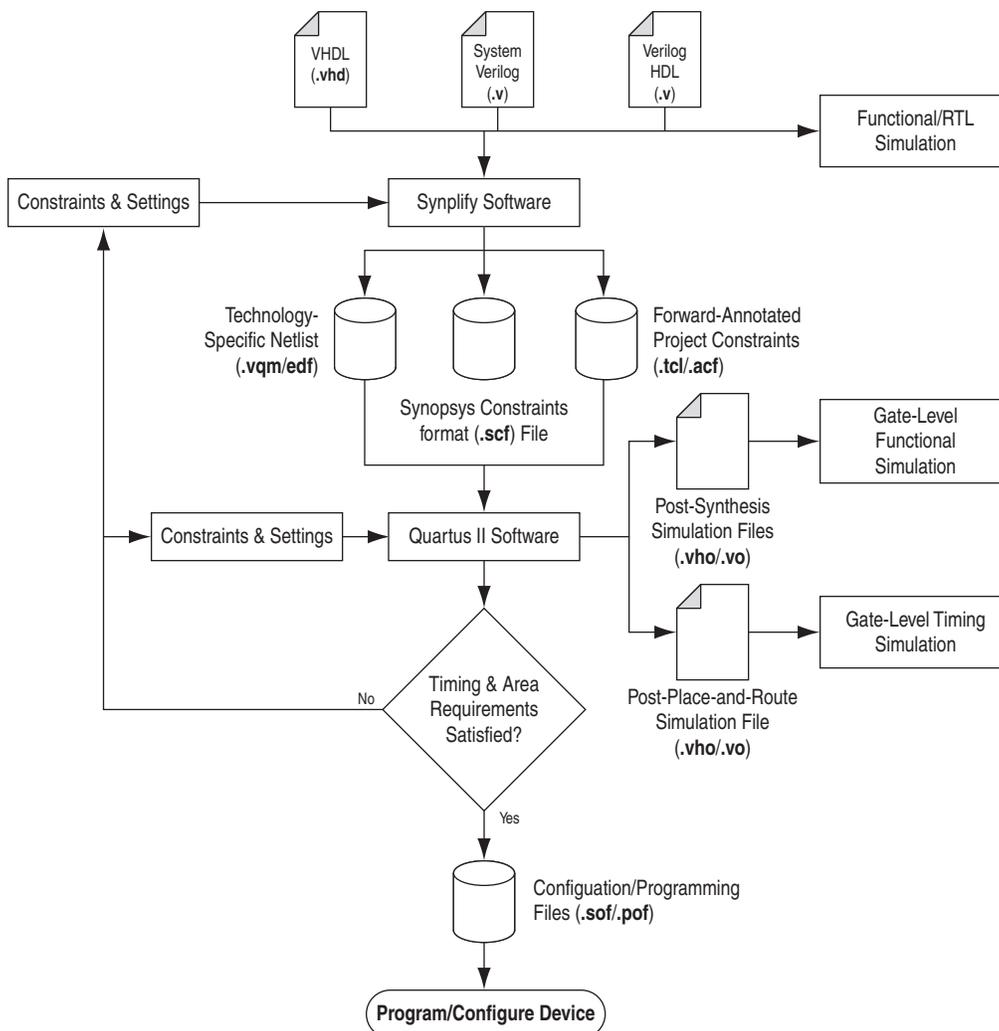
 If your design uses the Classic Timing Analyzer for timing analysis in the Quartus II software versions 10.0 and earlier, the Synplify software generates timing constraints in the Tcl Constraints File (**.tcl**). If you are using the Quartus II software versions 10.1 and later, you must use the TimeQuest Timing Analyzer for timing analysis.

 - The **.tcl** to set up your Quartus II project and pass constraints

Alternatively, you can run the Quartus II software from within the Synplify software. For more information, refer to [“Running the Quartus II Software from within the Synplify Software”](#) on page 14–14.
6. After obtaining place-and-route results that meet your requirements, configure or program the Altera device.

Figure 14-1 shows the recommended design flow using the Synplify and the Quartus II software.

Figure 14-1. Recommended Design Flow



The Synplify software supports VHDL, Verilog HDL, and SystemVerilog source files. However, only the Synplify Pro and Premier software support mixed synthesis, allowing a combination of VHDL and Verilog HDL or SystemVerilog format source files.

Specify timing constraints and attributes for a design in a SCOPE Design Constraints File (.sdc) with the SCOPE window in the Synplify software using the standard Synopsys Design Constraint (SDC) format, or directly in the HDL source file. You can also define compiler directives in the HDL source file. Many of these constraints are forward-annotated for use by the Quartus II software. See Table 14-1 on page 14-4 for a list of the files generated by Synplify.

The HDL Analyst that is included in the Synplify software is a graphical tool for generating schematic views of the technology-independent RTL view netlist (.srs) and technology-view netlist (.srm) files. You can use the Synplify HDL Analyst to analyze and debug your design visually. The HDL Analyst supports cross-probing between the RTL and Technology views, the HDL source code, the Finite State Machine (FSM) viewer, and between the technology view and the timing report file in the Quartus II software.

 A separate license file is required to enable the HDL Analyst in the Synplify software. The Synplify Pro and Premier software include the HDL Analyst.

After synthesis is complete, import the .vqm or .edf netlist to the Quartus II software for place-and-route. Use the .tcl file generated by the Synplify software to forward-annotate your project constraints including device selection, called the generated .scf file to forward-annotate TimeQuest Timing Analyzer timing constraints, and optionally to set up your project in the Quartus II software.

If area and timing requirements are satisfied, use the files generated by the Quartus II software to program or configure the Altera device. As shown in [Figure 14-1](#), if your area or timing requirements are not met, you can change the constraints in the Synplify software or the Quartus II software and rerun synthesis. Altera recommends that you provide timing constraints in the Synplify software and any placement constraints in the Quartus II software. Repeat the process until area and timing requirements are met.

You can perform simulation and formal verification at various stages in the design process. You can perform final timing analysis after placement and routing is complete.

 For more information about how the Synplify software supports formal verification, refer to [Section V. Formal Verification](#) in volume 3 of the *Quartus II Handbook*.

You can also use other options and techniques in the Quartus II software to meet area and timing requirements, such as WYSIWYG Primitive Resynthesis, which can perform optimizations on your .vqm netlist within the Quartus II software.

 For information about netlist optimizations, refer to the [Netlist Optimizations and Physical Synthesis](#) chapter in volume 3 of the *Quartus II Handbook*.

 In some cases, you might be required to modify the source code if the area and timing requirements cannot be met using options in the Synplify and Quartus II software.

During synthesis, the Synplify software produces several intermediate and output files, which are listed and described in [Table 14-1](#).

Table 14-1. Synplify Intermediate and Output Files (Part 1 of 2)

File Extensions	File Description
.srs	Technology-independent RTL netlist file that can be read only by the Synplify software.
.srm	Technology view netlist file.
.srr (1)	Synthesis Report file.

Table 14–1. Synplify Intermediate and Output Files (Part 2 of 2)

File Extensions	File Description
.vqm/.edf	Technology-specific netlist in .vqm or .edf file format. A .vqm file is created for all Altera device families supported by the Quartus II software. An .edf file is created for devices supported by the MAX+PLUS II software.
.tcl	Forward-annotated constraints file containing constraints and assignments. A .tcl file for the Quartus II software is created for all devices. The .tcl file contains the appropriate Tcl commands to create and set up a Quartus II project and pass placement constraints.
.acf	Assignment and Configurations file for backward compatibility with the MAX+PLUS II software. For devices supported by the MAX+PLUS II software, the MAX+PLUS II assignments are imported from the MAX+PLUS II .acf file.
.scf	Synopsys Constraint Format file containing timing constraints for the TimeQuest Timing Analyzer.

Note to Table 14–1:

- (1) This report file includes performance estimates that are often based on pre-place-and-route information. Use the f_{MAX} reported by the Quartus II software after place-and-route—it is the only reliable source of timing information. This report file includes post-synthesis device resource utilization statistics that might inaccurately predict resource usage after place-and-route. The Synplify software does not account for black box functions nor for logic usage reduction achieved through register packing performed by the Quartus II software. Register packing combines a single register and look-up table (LUT) into a single logic cell, reducing logic cell utilization below the Synplify software estimate. Use the device utilization reported by the Quartus II software after place-and-route.

Specifying the Output Netlist File Name and Result Format

The **Result Format** list specifies an **.edf** or **.vqm** netlist, depending on your device family. The software creates an **.edf** output netlist file only for devices supported by the MAX+PLUS II software. For current Altera devices, the software generates a **.vqm**-formatted netlist.

To specify the output netlist directory location, name, and format for the Synplify software, perform the following steps:

1. In the Synplify software, on the Project menu, click **Implementation Options**.
2. Click the **Implementation Results** tab.
3. In the **Results Directory** box, type your output netlist file directory location.
4. In the **Result File Name** box, type your output netlist file name.
5. In the **Results Format** list, specify either **.vqm** for devices using the Quartus II software or **.edf** for devices using the MAX+PLUS II software.

Specifying the Quartus II Software Version

You can specify your version of the Quartus II software in the **Implementation Results** tab of the **Implementation Options** dialog box. This option ensures that the netlist is compatible with the software version and supports the newest features. Altera recommends using the latest version of the Quartus II software whenever possible. If your Quartus II software version is newer than the versions available in the **Quartus Version** list, check if there is a newer version of the Synplify software available that supports the current Quartus II software version. Otherwise, select the latest version in the list for the best compatibility.



The **Quartus Version** list is available only after selecting an Altera device.

To specify the Quartus II software version used in the Synplify software, perform the following steps:

1. In the Synplify software, on the Project menu, click **Implementation Options**.
2. Click the **Implementation Results** tab.
3. Specify your version of the Quartus II software in the **Quartus Version** list.

Alternatively, type the following command at the command line:

```
set_option -quartus_version <version number> ←
```

Synplify Optimization Strategies

Combining Synplify software constraints with VHDL and Verilog HDL coding techniques and Quartus II software options can help you obtain the results that you require. This section provides an overview of some of the techniques you can use to help improve the quality of your results.

-  For additional design and optimization techniques, refer to the *Design Recommendations for Altera Devices and the Quartus II Design Assistant* chapter in volume 1 and the *Area and Timing Optimization* chapter in volume 2 of the *Quartus II Handbook*.
-  For more information about applying the attributes discussed in this section, refer to the *Altera Constraints, Attributes, and Options* chapter of the *Synopsys FPGA Synthesis Reference Manual*.

Using Synplify Premier to Optimize Your Design

Compared to other Synplify products, the Synplify Premier software offers additional physical synthesis optimizations. After typical logic synthesis, the Synplify Premier software places and routes the design and attempts to restructure the netlist based on the physical location of the logic in the Altera device. The Synplify Premier software forward-annotates the design netlist to the Quartus II software to perform the final placement and routing. In the default flow, the Synplify Premier software also forward-annotates placement information for the critical path(s) in the design, which can improve the compilation time in the Quartus II software.

The physical location annotation file is called `<design name>_plc.tcl`. If you open the Quartus II software from the Synplify Premier software user interface, the Quartus II software automatically uses this file for the placement information.

The Physical Analyst allows you to examine the placed netlist from the Synplify Premier software, which is similar to the HDL Analyst for a logical netlist. You can use this display to analyze and diagnose potential problems.

Using Implementations in Synplify Pro or Premier

To create different synthesis results without overwriting the existing results, in the Synplify Pro or Premier software, on the Project menu, click **New Implementation**. For each implementation, specify the target device, synthesis options, and constraint files. Each implementation generates its own subdirectory that contains all the resulting files, including `.vqm/.edf`, `.scf`, and `.tcl` files, from a compilation of the particular implementation. You can then compare the results of the different implementations to find the optimal set of synthesis options and constraints for a design.

Timing-Driven Synthesis Settings

The Synplify software supports timing-driven synthesis with user-assigned timing constraints to optimize the performance of the design. This section explains important timing constraints in the Synplify software.

The Quartus II NativeLink feature allows timing constraints that are applied in the Synplify software to be forward-annotated for the Quartus II software with an `.scf` file for timing-driven place and route. Refer to [“Passing TimeQuest SDC Timing Constraints to the Quartus II Software”](#) on page 14–15 for more details about how constraints such as clock frequencies, false paths, and multicycle paths are forward-annotated.



The Synplify Synthesis Report File (`.srr`) contains timing reports of estimated place-and-route delays. The Quartus II software can perform further optimizations on a post-synthesis netlist from third-party synthesis tools. In addition, designs might contain black boxes or intellectual property (IP) functions that have not been optimized by the third-party synthesis software. Actual timing results are obtained only after the design has been fully placed and routed in the Quartus II software. For these reasons, the Quartus II post place-and-route timing reports provide a more accurate representation of the design. Use the statistics in these reports to evaluate design performance.

Clock Frequencies

For single-clock designs, you can specify a global frequency when using the push-button flow. While this flow is simple and provides good results, it often does not meet the performance requirements for more advanced designs. You can use timing constraints, compiler directives, and other attributes to help optimize the performance of a design. You can enter these attributes and directives directly in the HDL code. Alternatively, you can enter attributes (not directives) into an `.sdc` file with the SCOPE window in the Synplify software.

Use the SCOPE window to set global frequency requirements for the entire design and individual clock settings. Use the **Clocks** tab in the SCOPE window to specify frequency (or period), rise times, fall times, duty cycle, and other settings. Assigning individual clock settings, rather than over-constraining the global frequency, helps the Quartus II software and the Synplify software achieve the fastest clock frequency for the overall design. The `define_clock` attribute assigns clock constraints.

Multiple Clock Domains

The Synplify software can perform timing analysis on unrelated clock domains. Each clock group is a different clock domain and is treated as unrelated to the clocks in all other clock groups. All clocks in a single clock group are assumed to be related, and the Synplify software automatically calculates the relationship between the clocks. You can assign clocks to a new clock group or put related clocks in the same clock group with the **Clocks** tab in the SCOPE window, or with the `define_clock` attribute.

Input and Output Delays

Specify the input and output delays for the ports of a design in the **Input/Output** tab of the SCOPE window, or with the `define_input_delay` and `define_output_delay` attributes. The Synplify software does not allow you to assign the t_{CO} and t_{SU} values directly to inputs and outputs. However, a t_{CO} value can be inferred by setting an external output delay; a t_{SU} value can be inferred by setting an external input delay.

Equation 14-1 illustrates the relationship between t_{CO} and the output delay:

Equation 14-1.

$$t_{CO} = \text{clock period} - \text{external output delay}$$

Equation 14-2 illustrates the relationship between t_{SU} and the input delay:

Equation 14-2.

$$t_{SU} = \text{clock period} - \text{external input delay}$$

When the `syn_forward_io_constraints` attribute is set to 1, the Synplify software passes the external input and output delays to the Quartus II software using NativeLink integration. The Quartus II software then uses the external delays to calculate the maximum system frequency.

Multicycle Paths

A multicycle path is a path that requires more than one clock cycle to propagate. Specify any multicycle paths in the design in the **Multi-Cycle Paths** tab of the SCOPE window, or with the `define_multicycle_path` attribute. You should specify which paths are multicycle to prevent the Quartus II and the Synplify compilers from working excessively on a non-critical path. Not specifying these paths can also result in an inaccurate critical path reported during timing analysis.

False Paths

False paths are paths that should be ignored during timing analysis, or be assigned low (or no) priority during optimization. Some examples of false paths include slow asynchronous resets, and test logic that has been added to the design. Set these paths in the **False Paths** tab of the SCOPE window, or use the `define_false_path` attribute.

FSM Compiler

If the FSM Compiler is turned on, the compiler automatically detects state machines in a design, which are then extracted and optimized. The FSM Compiler analyzes state machines and implements sequential, gray, or one-hot encoding, based on the number of states. The compiler also performs unused-state analysis, optimization of unreachable states, and minimization of transition logic. Implementation is based on the number of states, regardless of the coding style in the HDL code.

If the FSM Compiler is turned off, the compiler does not optimize logic as state machines. The state machines are implemented as coded in the HDL code. Thus, if the coding style for a state machine is sequential, the implementation is also sequential.

Use the `syn_state_machine` compiler directive to specify or prevent a state machine from being extracted and optimized. To override the default encoding of the FSM Compiler, use the `syn_encoding` directive.

The values for the `syn_encoding` directive are described in [Table 14-2](#).

Table 14-2. `syn_encoding` Directive Values

Value	Description
Sequential	Generates state machines with the fewest possible flipflops. Sequential, also called binary, state machines are useful for area-critical designs when timing is not the primary concern.
Gray	Generates state machines where only one flipflop changes during each transition. Gray-encoded state machines tend to be free of glitches.
One-hot	Generates state machines containing one flipflop for each state. One-hot state machines typically provide the best performance and shortest clock-to-output delays. However, one-hot implementations are usually larger than sequential implementations.
Safe	Generates extra control logic to force the state machine to the reset state if an invalid state is reached. You can use the safe value in conjunction with any of the other three values, which results in the state machine being implemented with the requested encoding scheme and the generation of the reset logic.

[Example 14-1](#) shows sample VHDL code for applying the `syn_encoding` directive.

Example 14-1. Sample VHDL Code for `syn_encoding`

```
SIGNAL current_state : STD_LOGIC_VECTOR(7 DOWNTO 0);  
ATTRIBUTE syn_encoding : STRING;  
ATTRIBUTE syn_encoding OF current_state : SIGNAL IS "sequential";
```

By default, the state machine logic is optimized for speed and area, which may be potentially undesirable for critical systems. The safe value generates extra control logic to force the state machine to the reset state if an invalid state is reached.

FSM Explorer in Synplify Pro and Premier

The Synplify Pro and Premier software use the FSM Explorer to explore different encoding styles for a state machine automatically, and then implement the best encoding based on the overall design constraints. The FSM Explorer uses the FSM Compiler to identify and extract state machines from a design. However, unlike the FSM Compiler, which chooses the encoding style based on the number of states, the FSM Explorer attempts several different encoding styles before choosing a specific one. The trade-off is that the compilation requires more time to analyze the state machine, but finds an optimal encoding scheme for the state machine.

Optimization Attributes and Options

The following sections describe more attributes and options that you can modify in the Synplify software to improve your design performance.

Retiming in Synplify Pro and Premier

The Synplify Pro and Premier software can retime a design, which can improve the timing performance of sequential circuits by moving registers (register balancing) across combinational elements. Be aware that retimed registers incur name changes. To retime your design, turn on **Retiming** in the **Device** tab in the **Implementation Options** section, or use the `syn_allow_retiming` attribute.

Maximum Fan-Out

When your design has critical path nets with high fan-out, use the `syn_maxfan` attribute to control the fan-out of the net. Setting this attribute for a specific net results in the replication of the driver of the net to reduce overall fan-out. The `syn_maxfan` attribute takes an integer value and applies it to inputs or registers. The `syn_maxfan` attribute cannot be used to duplicate control signals. The minimum allowed value of the attribute is 4. Using this attribute might result in increased logic resource utilization, thus straining routing resources, which can lead to long compilation times and difficult fitting.

If you must duplicate an output register or an output enable register, you can create a register for each output pin by using the `syn_useioff` attribute. Refer to [“Register Packing” on page 14-10](#).

Preserving Nets

During synthesis, the compiler maintains ports, registers, and instantiated components. However, some nets cannot be maintained to create an optimized circuit. Applying the `syn_keep` directive overrides the optimization of the compiler and preserves the net during synthesis. The `syn_keep` directive is a Boolean data type value and can be applied to wires (Verilog HDL) and signals (VHDL). Setting the value to **true** preserves the net through synthesis.

Register Packing

Altera devices allow register packing into I/O cells. Altera recommends allowing the Quartus II software to make the I/O register assignments. However, you can control register packing with the `syn_useioff` attribute. The `syn_useioff` attribute is a Boolean data type value that can be applied to ports or entire modules. Setting the value to **1** instructs the compiler to pack the register into an I/O cell. Setting the value to **0** prevents register packing in both the Synplify and Quartus II software.

Resource Sharing

The Synplify software uses resource sharing techniques during synthesis, by default, to reduce area. Turning off the **Resource Sharing** option on the **Options** tab of the **Implementation Options** dialog box improves performance results for some designs. You can also turn off the option for a specific module with the `syn_sharing` attribute. If you turn off this option, be sure to check the results to verify improvement in timing performance. If there is no improvement, turn on **Resource Sharing**.

Preserving Hierarchy

The Synplify software performs cross-boundary optimization by default, which causes the design to flatten to allow optimization. You can use the `syn_hier` attribute to override the default compiler settings. The `syn_hier` attribute applies a string value to modules, architectures, or both. Setting the value to **hard** maintains the boundaries of a module, architecture, or both, but allows constant propagation. Setting the value to **locked** prevents all cross-boundary optimizations. Use the **locked** setting with the `partition` setting to create separate design blocks and multiple output netlists for incremental compilation, as described in “Using MultiPoint Synthesis with Incremental Compilation” on page 14-31.

By default, the Synplify software generates a hierarchical `.vqm` file. To flatten the file, set the `syn_netlist_hierarchy` attribute to 0.

Register Input and Output Delays

Two advanced options, `define_reg_input_delay` and `define_reg_output_delay`, can speed up paths feeding a register, or coming from a register, by a specific number of nanoseconds. The Synplify software attempts to meet the global clock frequency goals for a design as well as the individual clock frequency goals (set with the `define_clock` attribute). You can use these attributes to add a delay to paths feeding into or out of registers to further constrain critical paths. You can slow down a path that is too highly optimized by setting this attributes to a negative number.

The `define_reg_input_delay` and `define_reg_output_delay` options are useful to close timing if your design does not meet timing goals, because the routing delay after placement and routing exceeds the delay predicted by the Synplify software. Rerun synthesis using these options, specifying the actual routing delay (from place-and-route results) so that the tool can meet the required clock frequency. Synopsys recommends that for best results, do not make these assignments too aggressively. For example, you can increase the routing delay value, but do not also use the full routing delay from the last compilation.

In the SCOPE constraint window, the registers panel contains the following options:

- **Register**—Specifies the name of the register. If you have initialized a compiled design, select the name from the list.
- **Type**—Specifies whether the delay is an input or output delay.
- **Route**—Shrinks the effective period for the constrained registers by the specified value without affecting the clock period that is forward-annotated to the Quartus II software.

Use the following Tcl command syntax to specify an input or output register delay in nanoseconds.

Example 14-2. Specifying an Input or Output Register Delay Using Tcl Command Syntax

```
define_reg_input_delay {<register>} -route <delay in ns>  
define_reg_output_delay {<register>} -route <delay in ns>
```

syn_direct_enable

This attribute controls the assignment of a clock-enable net to the dedicated enable pin of a register. With this attribute, you can direct the Synplify mapper to use a particular net as the only clock enable when the design has multiple clock enable candidates.

To use this attribute as a compiler directive to infer registers with clock enables, enter the `syn_direct_enable` directive in your source code, instead of the SCOPE spreadsheet.

The `syn_direct_enable` data type is Boolean. A value of **1** or **true** enables net assignment to the clock-enable pin. The following is the syntax for Verilog HDL:

```
object /* synthesis syn_direct_enable = 1 */ ;
```

I/O Standard

For certain Altera devices, specify the I/O standard type for an I/O pad in the design with the **I/O Standard** panel in the Synplify SCOPE window.

Example 14-3 shows the Synplify SDC syntax for the `define_io_standard` constraint, in which the `delay_type` must be either `input_delay` or `output_delay`.

Example 14-3. Synplify SDC Syntax for the `define_io_standard` Constraint

```
define_io_standard [-disable|-enable] {<objectName>} -delay_type \  
[input_delay|output_delay] <columnTc1Name>{<value>} [<columnTc1Name>{<value>}...]
```

 For details about supported I/O standards, refer to the *Altera I/O Standards* section in the *Synopsys FPGA Synthesis Reference Manual*.

Altera-Specific Attributes

You can use the attributes described in this section with specific Altera device features, which are forward-annotated to the Quartus II project, and are used during place-and-route.

altera_chip_pin_lc

Use the `altera_chip_pin_lc` attribute to make pin assignments. This attribute applies a string value to inputs and outputs. Use the attribute only on the ports of the top-level entity in the design. Do not use this attribute to assign pin locations from entities at lower levels of the design hierarchy.

 The `altera_chip_pin_lc` attribute is not supported for any MAX series device.

In the SCOPE window, set the value of the `altera_chip_pin_lc` attribute to a pin number or a list of pin numbers.

Example 14-4 shows VHDL code for making location assignments for supported Altera devices. Pin location assignments for these devices are written to the output `.tcl` file.

Example 14-4. Making Location Assignments in VHDL

```
ENTITY sample (data_in : IN STD_LOGIC_VECTOR (3 DOWNTO 0);  
              data_out: OUT STD_LOGIC_VECTOR (3 DOWNTO 0));  
ATTRIBUTE altera_chip_pin_lc : STRING;  
ATTRIBUTE altera_chip_pin_lc OF data_out : SIGNAL IS "14, 5, 16, 15";
```

 The `data_out` signal is a 4-bit signal; `data_out[3]` is assigned to pin 14 and `data_out[0]` is assigned to pin 15.

altera_io_powerup

Use the `altera_io_powerup` attribute to define the power-up value of an I/O register that has no set or reset. This attribute applies a string value (**high** | **low**) to ports with I/O registers. By default, the power-up value of the I/O register is set to **low**.

altera_io_opendrain

Use the `altera_io_opendrain` attribute to specify open-drain mode I/O ports. This attribute applies a boolean data type value to outputs or bidirectional ports for devices that support open-drain mode.

Exporting Designs to the Quartus II Software Using NativeLink Integration

The NativeLink feature in the Quartus II software facilitates the seamless transfer of information between the Quartus II software and EDA tools, and allows you to run other EDA design entry or synthesis, simulation, and timing analysis tools automatically from within the Quartus II software. After a design is synthesized in the Synplify software, a `.vqm` or `.edf` netlist file, an `.scf` file for TimeQuest Timing Analyzer timing constraints, and `.tcl` files are used to import the design into the Quartus II software for place-and-route. You can run the Quartus II software from within the Synplify software or as a stand-alone application. After you import the design into the Quartus II software, you can specify different options to further optimize the design.

 When you are using NativeLink integration, the path to your project must not contain empty spaces. The Synplify software uses Tcl scripts to communicate with the Quartus II software, and the Tcl language does not accept arguments with empty spaces in the path.

Use NativeLink integration to integrate the Synplify software and Quartus II software with a single GUI for both synthesis and place-and-route operations. NativeLink integration allows you to run the Quartus II software from within the Synplify software GUI, or to run the Synplify software from within the Quartus II software GUI.

This section explains the different NativeLink flows and provides details about how constraints are passed to the Quartus II software, and describes the following topics:

- “Running the Quartus II Software from within the Synplify Software”
- “Using the Quartus II Software to Run the Synplify Software” on page 14–15
- “Running the Quartus II Software Manually With the Synplify-Generated Tcl Script” on page 14–15
- “Passing TimeQuest SDC Timing Constraints to the Quartus II Software” on page 14–15

Running the Quartus II Software from within the Synplify Software

To run the Quartus II software from within the Synplify software, you must set the `QUARTUS_ROOTDIR` environment variable to the Quartus II software installation directory located in the `<Quartus II system directory>\altera\ <version number>\quartus`. You must set this environment variable to use the Synplify and Quartus II software together. Synplify also uses this variable to open the Quartus II software in the background and obtain detailed information for Altera megafunctions used in the design.

In the Windows operating system, set the environment variable using the Control Panel, System options. In the **Advanced** tab, click **Environment Variables**, and create a `QUARTUS_ROOTDIR` system variable.

In the Linux operating system, create an environment variable `QUARTUS_ROOTDIR` that points to the `<home directory>/altera <version number>` location.

You can create new place and route implementations with the **New P&R** button in the Synplify software GUI. Under each implementation, the Synplify Pro software creates a place-and-route implementation called `pr_<number>` **Altera Place and Route**. To run the Quartus II software in command-line mode after each synthesis run, use the text box to turn on the place-and-route implementation. The results of the place-and-route are written to a log file in the `pr_<number>` directory under the current implementation directory.

You can also use the commands in the Quartus II menu to run the Quartus II software at any time following a successful completion of synthesis. In the Synplify software, on the Options menu, click **Quartus II** and then choose one of the following commands:

- **Launch Quartus**—Opens the Quartus II software GUI and creates a Quartus II project with the synthesized output file, forward-annotated timing constraints, and pin assignments. Use this command to configure options for the project and to execute any Quartus II commands.
- **Run Background Compile**—Runs the Quartus II software in command-line mode with the project settings from the synthesis run. The results of the place-and-route are written to a log file.

The `<project_name>_cons.tcl` file is used to set up the Quartus II project and directs the `<project_name>.tcl` file to pass constraints from the Synplify software to the Quartus II software. By default, the `<project_name>.tcl` file contains device, timing, and location assignments. The `<project_name>.tcl` file contains the command to use the Synplify-generated `.scf` constraints file with the TimeQuest Timing Analyzer.

Using the Quartus II Software to Run the Synplify Software

You can set up the Quartus II software to run the Synplify software for synthesis with NativeLink integration. This feature allows you to use the Synplify software to quickly synthesize a design as part of a standard compilation in the Quartus II software. When you use this feature, the Synplify software does not use any timing constraints or assignments, such as incremental compilation partitions, that you have set in the Quartus II software.



For best results, Synopsys recommends that you set constraints in the Synplify software and use a Tcl script to pass these constraints to the Quartus II software, instead of opening the Synplify software from within the Quartus II software.

To set up the Synplify software in the Quartus II software, on the Tools menu, click **Options**. In the **Options** dialog box, click **EDA Tool Options** and specify the path of the Synplify or Synplify Pro software under **Location of Executable**.

- For more information about using NativeLink integration with the Synplify software in the Quartus II software, refer to [About Using the Synplify Software with the Quartus II Software](#) in Quartus II Help.

Running the Synplify software with NativeLink integration is supported on both floating network and node-locked fixed PC licenses. Both types of licenses support batch mode compilation.

Running the Quartus II Software Manually With the Synplify-Generated Tcl Script

You can also run the Quartus II software with a Synplify-generated Tcl script. To run the Tcl scrip to set up your project assignments, perform the following steps:

- Ensure the **.vqm/.edf**, **.scf**, and **.tcl** files are located in the same directory.
- In the Quartus II software, on the View menu, point to **Utility Windows** and click **Tcl Console**. The Quartus II Tcl Console opens.
- At the Tcl Console command prompt, type the following:

```
source <path>/<project name>_cons.tcl ←
```

Passing TimeQuest SDC Timing Constraints to the Quartus II Software

The TimeQuest Timing Analyzer is a powerful ASIC-style timing analysis tool that validates the timing performance of all logic in your design using an industry standard constraints format, Synopsys Design Constraints (SDC). This section explains how timing constraints set in the Synplify software are passed to the Quartus II software for use with the TimeQuest Timing Analyzer.

The Synplify-generated **.tcl** file contains constraints for the Quartus II software, such as the device specification and any location constraints. Timing constraints are forward-annotated in the Synopsys Constraints Format (**.scf**) file.



For additional information about the TimeQuest Timing Analyzer, refer to the [Quartus II TimeQuest Timing Analyzer](#) chapter in volume 3 of the *Quartus II Handbook*.

 Synopsys recommends that you modify constraints using the SCOPE constraint editor window, rather than using the generated `.sdc`, `.scf`, or `.tcl` file.

The following list of Synplify constraints are converted to the equivalent Quartus II SDC commands and are forward-annotated to the Quartus II software in the `.scf` file:

- `define_clock`
- `define_input_delay`
- `define_output_delay`
- `define_multicycle_path`
- `define_false_path`

All Synplify constraints described in the following sections are mapped to SDC commands for the TimeQuest Timing Analyzer.

-  For syntax and arguments for these commands, refer to the applicable subsection or refer to Synplify Help. For a list of corresponding commands in the Quartus II software, refer to the Quartus II Help.

Individual Clocks and Frequencies

Specify clock frequencies for individual clocks in the Synplify software with the `define_clock` command. This command is passed to the Quartus II software with the `create_clock` command.

Input and Output Delay

Specify input delay and output delay constraints in the Synplify software with the `define_input_delay` and `define_output_delay` commands, respectively. These commands are passed to the Quartus II software with the `set_input_delay` and `set_output_delay` commands.

Multicycle Path

Specify a multicycle path constraint in the Synplify software with the `define_multicycle_path` command. This command is passed to the Quartus II software with the `set_multicycle_path` command.

False Path

Specify a false path constraint in the Synplify software with the `define_false_path` command. This command is passed to the Quartus II software with the `set_false_path` command.

Guidelines for Altera Megafunctions and Architecture-Specific Features

Altera provides parameterizable megafunctions, including LPMs, device-specific Altera megafunctions, IP available as Altera MegaCore[®] functions, and IP available through the Altera Megafunction Partners Program (AMPPSM). You can use megafunctions and IP functions by instantiating them in your HDL code, or by inferring certain megafunctions from generic HDL code.

You can instantiate a megafunction in your HDL code with the MegaWizard Plug-In Manager to parameterize the function, or instantiate the function using the port and parameter definition. The MegaWizard Plug-In Manager provides a graphical interface within the Quartus II software for customizing and parameterizing any available megafunction for the design. For more information about the MegaWizard Plug-In Manager flow with the Synplify software, refer to “Instantiating Altera Megafunctions With the MegaWizard Plug-In Manager” on page 14-17 and “Instantiating Intellectual Property With the MegaWizard Plug-In Manager and IP Toolbench” on page 14-19.

-  For more information about specific Altera megafunctions, refer to the Quartus II Help. For more information about IP functions, refer to the appropriate IP documentation.

The Synplify software also automatically recognizes certain types of HDL code, and infers the appropriate megafunction when a megafunction provides optimal results. The Synplify software provides options to control inference of certain types of megafunctions, as described in “Inferring Altera Megafunctions from HDL Code” on page 14-22.

-  For more information about instantiating versus inferring megafunctions, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*. This chapter also provides details about using the MegaWizard Plug-In Manager in the Quartus II software and explains the files generated by the wizard, as well as coding style recommendations and HDL examples for inferring megafunctions in Altera devices.

Instantiating Altera Megafunctions With the MegaWizard Plug-In Manager

This section describes how to instantiate Altera megafunctions with the MegaWizard Plug-In Manager.

When you use the MegaWizard Plug-In Manager to set up and parameterize a megafunction, the MegaWizard Plug-In Manager creates a VHDL or Verilog HDL wrapper file `<output file>.v | vhd` that instantiates the megafunction.

The Synplify software uses the Quartus II timing and resource estimation netlist feature to report more accurate resource utilization and timing performance estimates, and leverages timing-driven optimization, instead of treating the megafunction as a “black box.” Including the MegaWizard-generated megafunction variation wrapper file in your Synplify project, gives the Synplify software complete information about the megafunction.

-  There is an option in the MegaWizard Plug-In Manager to generate a netlist for resource and timing estimation. This option is not recommended for the Synplify software because the software automatically generates this information in the background without a separate netlist. If you do create a separate netlist `<output file>_syn.v` and use that file in your synthesis project, you must also include the `<output file>.v | vhd` file in your Quartus II project.

Verify that the correct Quartus II version is specified in the Synplify software before compiling the MegaWizard-generated file to ensure that the software uses the correct library definitions for the megafunction. The **Quartus Version** setting must match the version of the Quartus II software used to generate the customized megafunction in the MegaWizard Plug-In Manager.

For details about how to set the Quartus II version in the Synplify software, refer to [“Specifying the Quartus II Software Version” on page 14-5](#).

In addition, ensure that the `QUARTUS_ROOTDIR` environment variable specifies the installation directory location of the correct Quartus II version. The Synplify software uses this information to launch the Quartus II software in the background. The environment variable setting must match the version of the Quartus II software used to generate the customized megafunction in the MegaWizard Plug-In Manager. Refer to [“Using the Quartus II Software to Run the Synplify Software” on page 14-15](#) for more details.

Instantiating Megafunctions With MegaWizard Plug-In Manager-Generated Verilog HDL Files

If you turn on the `<output file>_inst.v` option on the last page of the MegaWizard interface, the MegaWizard Plug-In Manager generates a Verilog HDL instantiation template file for use in your Synplify design. The instantiation template file, `<output file>_inst.v`, helps to instantiate the megafunction variation wrapper file, `<output file>.v`, in your top-level design. Include the megafunction variation wrapper file `<output file>.v` in your Synplify project. The Synplify software includes the megafunction information in the output `.vqm` netlist file. You do not need to include the MegaWizard-generated megafunction variation wrapper file in your Quartus II project.

Instantiating Megafunctions with MegaWizard Plug-In Manager-Generated VHDL Files

If you turn on the `<output file>.cmp` and `<output file>_inst.vhd` options on the last page of the MegaWizard interface, the MegaWizard Plug-In Manager generates a VHDL component declaration file and a VHDL instantiation template file for use in your Synplify design. These files can help you instantiate the megafunction variation wrapper file, `<output file>.vhd`, in your top-level design. Include the `<output file>.vhd` in your Synplify project. The Synplify software includes the megafunction information in the output `.vqm` netlist file. You do not need to include the MegaWizard-generated megafunction variation wrapper file in your Quartus II project.

Changing Synplify’s Default Behavior for Instantiated Altera Megafunctions

By default, the Synplify software automatically opens the Quartus II software in the background to generate a resource and timing estimation netlist for megafunctions, as described in the previous sections.

You may want to change this behavior to reduce run times in the Synplify software, because generating the netlist files can take several minutes for large designs, or if the Synplify software cannot access your Quartus II software installation to generate the files. Changing this behavior might speed up the compilation time in the Synplify software, but the Quality of Results (QoR) might be reduced.

The Synplify software directs the Quartus II software to generate information in two ways:

- Some megafunctions provide a “clear box” model—the Synplify software fully synthesizes this model and includes the device architecture-specific primitives in the output **.vqm** netlist file.
- Other megafunctions provide a “grey box” model—the Synplify software reads the resource information, but the netlist does not contain all the logic functionality.

For these functions, the Synplify software uses the logic information for resource and timing estimation and optimization, and then instantiates the megafunction in the output **.vqm** netlist file so the Quartus II software can implement the appropriate device primitives. By default, the Synplify software uses the clear box model when available, and otherwise uses the grey box model. To change this behavior, perform the following steps:

1. In the Synplify software, click **Implementation Options**.
2. On the **Device** tab, specify one of the following values for the **Altera Models** option:
 - **On**—uses the clearbox model when available and the grey box model when the clearbox model is unavailable
 - **clearbox_only**—enables the clear box model, but not the grey box model
 - **Off**—turns off the feature entirely

Instantiating Intellectual Property With the MegaWizard Plug-In Manager and IP Toolbench

Many Altera IP functions include a resource and timing estimation netlist that the Synplify software uses to report more accurate resource utilization and timing performance estimates, and leverage timing-driven optimization rather than a black box function.

To create this netlist file, perform the following steps:

1. Select the IP function in the MegaWizard Plug-In Manager.
2. Click **Next** to open the IP Toolbench.
3. Click **Set Up Simulation**, which sets up all the EDA options.
4. Turn on the **Generate netlist** option to generate a netlist for resource and timing estimation and click **OK**.
5. Click **Generate** to generate the netlist file.

The Quartus II software generates a file *<output file>_syn.v*. This netlist contains the grey box information for resource and timing estimation, but does not contain the actual implementation. Include this netlist file in your Synplify project. Next, include the megafunction variation wrapper file *<output file>.v | vhd* in the Quartus II project along with your Synplify **.vqm** output netlist.

If your IP function does not include a resource and timing estimation netlist, the Synplify software must treat the IP function as a black box. In this case, refer to the following subsections for details about creating black boxes.

For information about including Quartus II-specific files in your Synplify project so they are automatically passed to the Quartus II software along with the output .vqm file, refer to [“Including Files for Quartus II Placement and Routing Only”](#) on page 14-22.

Instantiating Black Box IP Functions With Generated Verilog HDL Files

Use the `syn_black_box` compiler directive to declare a module as a black box. The top-level design files must contain the IP port-mapping and a hollow-body module declaration. Apply the `syn_black_box` directive to the module declaration in the top-level file or a separate file included in the project so that the Synplify software recognizes the module is a black box. The software compiles successfully without this directive, but reports an additional warning message. Using this directive allows you to add other directives, as discussed in [“Other Synplify Software Attributes for Creating Black Boxes”](#) on page 14-21.

[Example 14-5](#) shows a sample top-level file that instantiates `my_verilogIP.v`, which is a simplified customized variation generated by the MegaWizard Plug-In Manager and the IP Toolbench.

Example 14-5. Sample Top-Level Verilog HDL Code with Black Box Instantiation of IP

```
module top (clk, count);
    input clk;
    output[7:0] count;
    my_verilogIP verilogIP_inst (.clock (clk), .q (count));
endmodule
// Module declaration
// The following attribute is added to create a
// black box for this module.
module my_verilogIP (clock, q) /* synthesis syn_black_box */;
    input clock;
    output[7:0] q;
endmodule
```

Instantiating Black Box IP Functions With Generated VHDL Files

Use the `syn_black_box` compiler directive to declare a component as a black box. The top-level design files must contain the megafunction variation component declaration and port-mapping. Apply the `syn_black_box` directive to the component declaration in the top-level file. The software compiles successfully without this directive, but reports an additional warning message. Using this directive allows you to add other directives, such as the ones in the [“Other Synplify Software Attributes for Creating Black Boxes”](#) section.

Example 14-6 shows a sample top-level file that instantiates `my_vhdlIP.vhd`, which is a simplified customized variation generated by the MegaWizard Plug-In Manager and the IP Toolbench.

Example 14-6. Sample Top-Level VHDL Code with Black Box Instantiation of IP

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;

ENTITY top IS
  PORT (
    clk: IN STD_LOGIC ;
    count: OUT STD_LOGIC_VECTOR (7 DOWNTO 0)
  );
END top;

ARCHITECTURE rtl OF top IS
  COMPONENT my_vhdlIP
  PORT (
    clock: IN STD_LOGIC ;
    q: OUT STD_LOGIC_VECTOR (7 DOWNTO 0)
  );
end COMPONENT;
attribute syn_black_box : boolean;
attribute syn_black_box of my_vhdlIP: component is true;
BEGIN
  vhdlIP_inst : my_vhdlIP PORT MAP (
    clock => clk,
    q => count
  );
END rtl;
```

Other Synplify Software Attributes for Creating Black Boxes

Instantiating a function as a black box does not provide visibility into the function module for the synthesis tool. Thus, it does not take full advantage of the synthesis tool's timing-driven optimization. For better timing optimization, especially if the black box does not have registered inputs and outputs, add timing models to black boxes by adding the `syn_tpd`, `syn_tsu`, and `syn_tco` attributes. Refer to **Example 14-7** for a Verilog HDL example.

Example 14-7. Adding Timing Models to Black Boxes in Verilog HDL

```
module ram32x4(z,d,addr,we,clk);
  /* synthesis syn_black_box syn_tco1="clk->z[3:0]=4.0"
    syn_tpd1="addr[3:0]->z[3:0]=8.0"
    syn_tsu1="addr[3:0]->clk=2.0"
    syn_tsu2="we->clk=3.0" */
  output[3:0]z;
  input[3:0]d;
  input[3:0]addr;
  input we;
  input clk;
endmodule
```

The following additional attributes are supported by the Synplify software to communicate details about the characteristics of the black box module within the HDL code:

- `syn_resources`—Specifies the resources used in a particular black box

- `black_box_pad_pin`—Prevents mapping to I/O cells
- `black_box_tri_pin`—Indicates a tri-stated signal



For more information about applying these attributes, refer to the *Altera Constraints, Attributes, and Options* chapter of the *Synopsys FPGA Synthesis Reference Manual*.

Including Files for Quartus II Placement and Routing Only

In the Synplify software, you can add files to your project that are used only during placement and routing in the Quartus II software. This can be useful if you have grey or black boxes for Synplify synthesis that require the full design files to be compiled in the Quartus II software.

To include the files for Quartus II place-and-route only, perform the following steps:

1. Add the files to the Synplify project as source files.
2. Right-click the file, and on the shortcut menu, click **File options**.
3. Turn on **Use for Place and Route Only**. You can also set the option in a script using the `-job_owner par` option.

For example, the commands in [Example 14-8](#) define files for a Synplify project that includes a top-level design file, a grey box netlist file, an IP wrapper file, and an encrypted IP file. With these files, the Synplify software writes an empty instantiation of “core” in the `.vqm` file and uses the grey box netlist for resource and timing estimation. The files `core.v` and `core_enc8b10b.v` are not compiled by the Synplify software, but are copied into the place-and-route directory. The Quartus II software compiles these files to implement the “core” IP block.

Example 14-8. Commands to Define Files for a Synplify Project

```
add_file -verilog -job_owner par "core_enc8b10b.v"  
add_file -verilog -job_owner par "core.v"  
add_file -verilog "core_gb.v"  
add_file -verilog "top.v"
```

Inferring Altera Megafunctions from HDL Code

The Synplify software uses Behavior Extraction Synthesis Technology (BEST) algorithms to infer high-level structures such as RAMs, ROMs, operators, FSMs, and DSP multiplication operations. Then, the Synplify software keeps the structures abstract for as long as possible in the synthesis process. This allows the use of technology-specific resources to implement these structures by inferring the appropriate Altera megafunction when a megafunction provides optimal results. The following sections outline some of the Synplify-specific details when inferring Altera megafunctions. The Synplify software provides options to control inference of certain types of megafunctions, which is also described in the following sections.

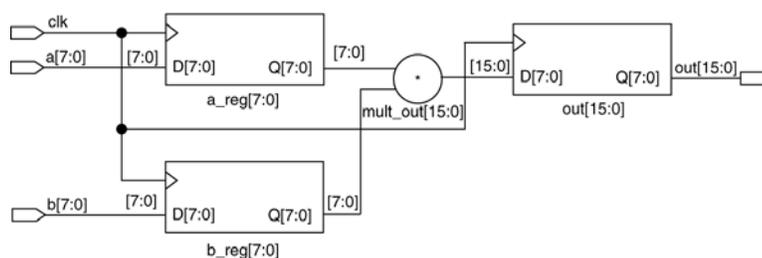


For coding style recommendations and examples for inferring megafunctions in Altera devices, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*.

Inferring Multipliers

Figure 14-2 shows the HDL Analyst view of an unsigned 8×8 multiplier with two pipeline stages after synthesis in the Synplify software. This multiplier is converted into an ALTMULT_ADD or ALTMULT_ACCUM megafunction. For devices with DSP blocks, the software might implement the function in a DSP block instead of regular logic, depending on device utilization. For some devices, the software maps directly to DSP block device primitives instead of instantiating a megafunction in the .vqm file.

Figure 14-2. HDL Analyst View of LPM_MULT Megafunction (Unsigned 8×8 Multiplier with Pipeline=2)



Resource Balancing

While mapping multipliers to DSP blocks, the Synplify software performs resource balancing for optimum performance.

Altera devices have a fixed number of DSP blocks, which includes a fixed number of embedded multipliers. If the design uses more multipliers than are available, the Synplify software automatically maps the extra multipliers to logic elements (LEs), or adaptive logic modules (ALMs).

If a design uses more multipliers than are available in the DSP blocks, the Synplify software maps the multipliers in the critical paths to DSP blocks. Next, any wide multipliers, which might or might not be in the critical paths, are mapped to DSP blocks. Smaller multipliers and multipliers that are not in the critical paths might then be implemented in the logic (LEs or ALMs). This ensures that the design fits successfully in the device.

Controlling the DSP Block Inference

You can implement multipliers in DSP blocks or in logic in Altera devices that contain DSP blocks. You can control this implementation through attribute settings in the Synplify software.

Signal Level Attribute

You can control the implementation of individual multipliers by using the `syn_multstyle` attribute as shown in the following Verilog HDL code:

```
<signal_name> /* synthesis syn_multstyle = "logic" */;
```

where `<signal_name>` is the name of the signal.



The `syn_multstyle` attribute applies to wires only; it cannot be applied to registers.

Table 14-3 describes the signal level attribute values that control the implementation of the multipliers in the DSP blocks or LEs in the Synplify software.

Table 14-3. DSP Block Attribute Settings in the Synplify Software

Attribute Name	Value	Description
syn_multstyle	lpm_mult	LPM function inferred and multipliers implemented in DSP blocks
	logic	LPM function not inferred and multipliers implemented LEs by the Synplify software
	block_mult	DSP megafunction is inferred and multipliers are mapped directly to DSP block device primitives (for supported devices)

Example 14-9 and Example 14-10 show simple Verilog HDL and VHDL code using the syn_multstyle attribute.

Example 14-9. Signal Attributes for Controlling DSP Block Inference in Verilog HDL Code

```

module mult(a,b,c,r,en);
  input [7:0] a,b;
  output [15:0] r;
  input [15:0] c;
  input en;
  wire [15:0] temp /* synthesis syn_multstyle="logic" */;

  assign temp = a*b;
  assign r = en ? temp : c;
endmodule

```

Example 14-10. Signal Attributes for Controlling DSP Block Inference in VHDL Code

```

library ieee;
use ieee.std_logic_1164.all;
use ieee.std_logic_unsigned.all;

entity onereg is port (
  r : out std_logic_vector(15 downto 0);
  en : in std_logic;
  a : in std_logic_vector(7 downto 0);
  b : in std_logic_vector(7 downto 0);
  c : in std_logic_vector(15 downto 0)
);
end onereg;

architecture beh of onereg is
  signal temp : std_logic_vector(15 downto 0);
  attribute syn_multstyle : string;
  attribute syn_multstyle of temp : signal is "logic";

begin
  temp <= a * b;
  r <= temp when en='1' else c;
end beh;

```

Inferring RAM

When a RAM block is inferred from an HDL design, the Synplify software uses an Altera megafunction to target the device memory architecture. For some devices, the Synplify software maps directly to memory block device primitives instead of instantiating a megafunction in the `.vqm` file.

Follow these guidelines for the Synplify software to successfully infer RAM in a design:

- The address line must be at least two bits wide.
- Resets on the memory are not supported. Refer to the device family documentation for information about whether read and write ports must be synchronous.
- Some Verilog HDL statements with blocking assignments might not be mapped to RAM blocks, so avoid blocking statements when modeling RAMs in Verilog HDL.

For some device families, the `syn_ramstyle` attribute specifies the implementation to use for an inferred RAM. You can apply the `syn_ramstyle` attribute globally, to a module, or to a RAM instance, to specify `registers` or `block_ram` values. To turn off RAM inference, set the attribute value to `registers`.

When inferring RAM for some Altera device families, the Synplify software generates additional bypass logic. This logic is generated to resolve a half-cycle read/write behavior difference between the RTL and post-synthesis simulations. The RTL simulation shows the memory being updated on the positive edge of the clock; the post-synthesis simulation shows the memory being updated on the negative edge of the clock. To eliminate bypass logic, the output of the RAM must be registered. By adding this register, the output of the RAM is seen after a full clock cycle, by which time the update has occurred, thus eliminating the need for bypass logic.

For devices with TriMatrix memory blocks, disable the creation of glue logic by setting the `syn_ramstyle` value to `no_rw_check`. Set `syn_ramstyle` to `no_rw_check` to disable the creation of glue logic in dual-port mode.

Example 14-11 shows sample VHDL code for inferring dual-port RAM.

Example 14-11. VHDL Code for Inferred Dual-Port RAM

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.std_logic_signed.all;

ENTITY dualport_ram IS
PORT ( data_out: OUT STD_LOGIC_VECTOR (7 DOWNTO 0);
      data_in: IN STD_LOGIC_VECTOR (7 DOWNTO 0);
      wr_addr, rd_addr: IN STD_LOGIC_VECTOR (6 DOWNTO 0);
      we: IN STD_LOGIC;
      clk: IN STD_LOGIC);
END dualport_ram;

ARCHITECTURE ram_infer OF dualport_ram IS
TYPE Mem_Type IS ARRAY (127 DOWNTO 0) OF STD_LOGIC_VECTOR (7 DOWNTO 0);
SIGNAL mem: Mem_Type;
SIGNAL addr_reg: STD_LOGIC_VECTOR (6 DOWNTO 0);

BEGIN
  data_out <= mem (CONV_INTEGER(rd_addr));
  PROCESS (clk, we, data_in) BEGIN
    IF (clk='1' AND clk'EVENT) THEN
      IF (we='1') THEN
        mem(CONV_INTEGER(wr_addr)) <= data_in;
      END IF;
    END IF;
  END PROCESS;
END ram_infer;
```

Example 14-12 shows an example of the VHDL code preventing bypass logic for inferring dual-port RAM. The extra latency behavior stems from the inferring methodology and is not required when instantiating a megafunction.

Example 14-12. VHDL Code for Inferred Dual-Port RAM Preventing Bypass Logic

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.std_logic_signed.all;

ENTITY dualport_ram IS
PORT ( data_out: OUT STD_LOGIC_VECTOR (7 DOWNTO 0);
      data_in : IN STD_LOGIC_VECTOR (7 DOWNTO 0);
      wr_addr, rd_addr : IN STD_LOGIC_VECTOR (6 DOWNTO 0);
      we : IN STD_LOGIC;
      clk : IN STD_LOGIC);
END dualport_ram;

ARCHITECTURE ram_infer OF dualport_ram IS
TYPE Mem_Type IS ARRAY (127 DOWNTO 0) OF STD_LOGIC_VECTOR (7 DOWNTO 0);
SIGNAL mem : Mem_Type;
SIGNAL addr_reg : STD_LOGIC_VECTOR (6 DOWNTO 0);
SIGNAL tmp_out : STD_LOGIC_VECTOR(7 DOWNTO 0); --output register

BEGIN
  tmp_out <= mem (CONV_INTEGER(rd_addr));
  PROCESS (clk, we, data_in) BEGIN
    IF (clk='1' AND clk'EVENT) THEN
      IF (we='1') THEN
        mem(CONV_INTEGER(wr_addr)) <= data_in;
      END IF;
      data_out <= tmp_out; --registers output preventing
                          -- bypass logic generation.
    END IF;
  END PROCESS;
END ram_infer;
```

RAM Initialization

Use the Verilog HDL `$readmemb` or `$readmemh` system tasks in your HDL code to initialize RAM memories. The Synplify compiler forward-annotates the initialization values in the `.srs` (technology-independent RTL netlist) file and the mapper generates a corresponding hexadecimal memory initialization (`.hex`) file. One `.hex` file is created for each of the `altsyncram` megafunctions that are inferred in the design. The `.hex` file is associated with the `altsyncram` instance in the `.vqm` file using the `init_file` attribute.

[Example 14-13](#) and [Example 14-14](#) illustrate how RAM memories can be initialized through HDL code, and how the corresponding .hex file is generated using Verilog HDL.

Example 14-13. Using \$readmemb System Task to Initialize an Inferred RAM in Verilog HDL Code

```
initial
begin
    $readmemb("mem.ini", mem);
end

always @(posedge clk)
begin
    raddr_reg <= raddr;
    if(we)
        mem[waddr] <= data;
end
```

Example 14-14. Sample of .vqm Instance Containing Memory Initialization File from [Example 14-13](#)

```
altsyncram mem_hex( .wren_a(we), .wren_b(GND),...);

defparam mem_hex.lpm_type = "altsyncram";
defparam mem_hex.operation_mode = "Dual_Port";
...
defparam mem_hex.init_file = "mem_hex.hex";
```

Inferring ROM

When a ROM block is inferred from an HDL design, the Synplify software uses an Altera megafunction to target the device memory architecture. For some devices, the Synplify software maps directly to memory block device atoms instead of instantiating a megafunction in the .vqm file. Follow these guidelines for the Synplify software to successfully infer ROM in a design:

- The address line must be at least two bits wide.
- The ROM must be at least half full.
- A CASE or IF statement must make 16 or more assignments using constant values of the same width.

Inferring Shift Registers

The Synplify software infers shift registers for sequential shift components so that they can be placed in dedicated memory blocks in supported device architectures using the ALTSHIFT_TAPS megafunction.

If necessary, set the implementation style with the `syn_srlstyle` attribute. If you do not want the components automatically mapped to shift registers, set the value to `registers`. You can set the value globally, or on individual modules or registers.

For some designs, turning off shift register inference improves the design performance.

Incremental Compilation and Block-Based Design

As designs become more complex and designers work in teams, a block-based incremental design flow is often an effective design approach. In an incremental compilation flow, you can make changes to part of the design while maintaining the placement and performance of unchanged parts of the design. Design iterations are made dramatically faster by focusing new compilations on particular design partitions and merging results with previous compilation results of other partitions. You can perform optimization on individual subblocks and then preserve the results before you integrate the blocks into a final design and optimize it at the top-level.

MultiPoint synthesis, which is available for certain device technologies in the Synplify Pro and Premier software, provides an automated block-based incremental synthesis flow. The MultiPoint feature manages a design hierarchy to let you design incrementally and synthesize designs that take too long for synthesis of the entire project. MultiPoint synthesis allows different netlist files to be created for different sections of a design hierarchy and supports the Quartus II incremental compilation methodology. This feature also ensures that only those sections of a design that have been updated are resynthesized when the design is compiled, reducing synthesis run time and preserving the results for the unchanged blocks. You can change and resynthesize one section of a design without affecting other sections.

You can also partition your design and create different netlist files manually with the Synplify software by creating a separate project for the logic in each partition of the design. Creating different netlist files for each partition of the design also means that each partition can be independent of the others.

Hierarchical design methodologies can improve the efficiency of your design process, providing better design reuse opportunities and fewer integration problems when working in a team environment. When you use these incremental synthesis methodologies, you can take advantage of incremental compilation in the Quartus II software. You can perform placement and routing on only the changed partitions of the design, which reduces place-and-route time and preserves your fitting results. Follow the guidelines in this section to help you optimize results with these methodologies.

The following steps describe the general incremental compilation flow when using these features of the Quartus II software:

1. Create Verilog HDL or VHDL design files.
2. Determine which hierarchical blocks you want to treat as separate partitions in your design.
3. Set up your design using the MultiPoint synthesis feature or separate projects so that a separate netlist file is created for each design partition.
4. If using separate projects, disable I/O pad insertion in the implementations for lower-level partitions.
5. Compile and map each partition in the Synplify software, making constraints as you would in a non-incremental design flow.
6. Import the **.vqm** netlist and **.tcl** file for each partition into the Quartus II software and set up the Quartus II project(s) for incremental compilation.

7. Compile your design in the Quartus II software and preserve the compilation results with the post-fit netlist in incremental compilation.
8. When you make design or synthesis optimization changes to part of your design, resynthesize only the partition you modified to generate a new netlist and `.tcl` file. Do not regenerate netlist files for the unmodified partitions.
9. Import the new netlist and `.tcl` file into the Quartus II software and recompile the design in the Quartus II software with incremental compilation.

 For more information about creating partitions and using the incremental compilation in the Quartus II software, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.

Creating a Design with Separate Netlist Files for Incremental Compilation

The first stage of a hierarchical or incremental design flow is to ensure that different parts of your design do not affect each other. Ensure that you have separate netlists for each partition in your design so you can take advantage of incremental compilation in the Quartus II software. If the entire design is in one netlist file, changes in one partition might affect other partitions because of possible node name changes when you resynthesize the design.

To ensure proper functionality of the synthesis flow, create separate netlist files only for modules and entities. In addition, each module or entity requires its own design file. If two different modules are in the same design file, but are defined as being part of different partitions, incremental compilation cannot be maintained since both partitions must be recompiled when one module is changed.

Altera recommends that you register all inputs and outputs of each partition. This makes logic synchronous, and avoids any delay penalty on signals that cross partition boundaries.

If you use boundary tri-states in a lower-level block, the Synplify software pushes, or bubbles, the tri-states through the hierarchy to the top-level to use the tri-state drivers on output pins of Altera devices. Because bubbling tri-states requires optimizing through hierarchies, lower-level tri-states are not supported with a block-based compilation methodology. Use tri-state drivers only at the external output pins of the device and in the top-level block in the hierarchy.

 For more detailed recommendations about designing your hierarchy and creating partitions, refer to the *Best Practices for Incremental Compilation Partitions and Floorplan Assignments* chapter in volume 1 of the *Quartus II Handbook*.

You can generate multiple `.vqm` netlist files with the MultiPoint synthesis flow in the Synplify Pro and Premier software, or by manually creating separate Synplify projects and creating a black box for each block that you want to designate as a separate design partition.

In the MultiPoint synthesis flow in the Synplify Pro and Premier software, you create multiple `.vqm` netlist files from one easy-to-manage, top-level synthesis project. By using the manual black box method, you have multiple synthesis projects, which might be required for certain team-based or bottom-up designs where a single top-level project is not desired.

After you have created multiple `.vqm` files using one of these two methods, you must create the appropriate Quartus II projects to place-and-route the design.

Using MultiPoint Synthesis with Incremental Compilation

This section describes how to generate multiple `.vqm` files using the Synplify Pro and Premier software MultiPoint synthesis flow. You must first set up your constraint file and Synplify options, then apply the appropriate Compile Point settings to write multiple `.vqm` files and create design partition assignments for incremental compilation.

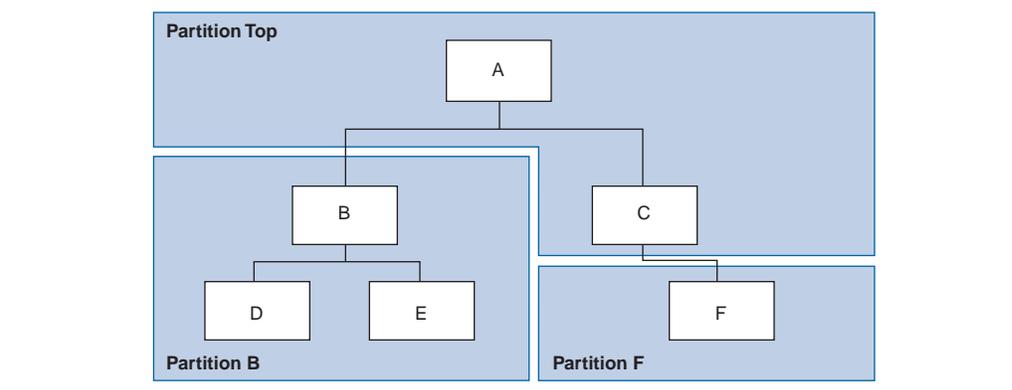
Set Compile Points and Create Constraint Files

The MultiPoint flow lets you segment a design into smaller synthesis units, called Compile Points. The synthesis software treats each Compile Point as a partition for incremental mapping, which allows you to isolate and work on each Compile Point module as independent segments of the larger design without impacting other design modules. A design can have any number of Compile Points, and Compile Points can be nested. The top-level module is always treated as a Compile Point.

Compile Points are optimized in isolation from their parent, which can be another Compile Point or a top-level design. Each block created with a Compile Point is unaffected by critical paths or constraints on its parent or other blocks. A Compile Point is independent, with its own individual constraints. During synthesis, any Compile Points that have not yet been synthesized are synthesized before the top level. Nested Compile Points are synthesized before the parent Compile Points in which they are contained. When you apply the appropriate setting for the Compile Point, a separate netlist is created for that Compile Point, isolating that logic from any other logic in the design.

Figure 14-3 shows an example of a design hierarchy that is split into multiple partitions. The top-level block of each partition can be synthesized as a separate Compile Point.

Figure 14-3. Partitions in a Hierarchical Design



In this case, modules A, B, and F are Compile Points. The top-level Compile Point consists of the top-level block in the design (that is, block A in this example), including the logic that is not defined under another Compile Point. In this example, the design for top-level Compile Point A also includes the logic in one of its subblocks, C. Because block F is defined as its own Compile Point, it is not treated as part of the top-level Compile Point A. Another separate Compile Point B contains the logic in blocks B, D, and E. One netlist is created for the top-level module A and submodule C, another netlist is created for B and its submodules D and E, while a third netlist is created for F.

Apply Compile Points to the module, or to the architecture in the Synplify Pro SCOPE spreadsheet, or to the .sdc file. You cannot set a Compile Point in the Verilog HDL or VHDL source code. You can set the constraints manually using Tcl or by editing the .sdc file, or you can use one of two methods in the GUI, as described in the following subsections.

Defining Compile Points With .tcl or .sdc Files

To set Compile Points with a .tcl or .sdc file, use the `define_compile_point` command, as shown in [Example 14-15](#).

Example 14-15. The `define_compile_point` Command

```
define_compile_point [-disable] {<objname>} -type {locked, partition}
```

In [Example 14-15](#), `<objname>` represents any module in the design. The Compile Point type `{locked, partition}` indicates that the Compile Point represents a partition for the Quartus II incremental compilation flow.

Each Compile Point has a set of constraint files that begin with the `define_current_design` command to set up the SCOPE environment, as follows:

```
define_current_design {<my_module>}
```

Defining Compile Points in the Top-Level SCOPE Window

The following method requires that you create separate constraint files for the top-level and lower-level Compile Points:

1. In the top-level SCOPE window, select the **Compile Points** tab.
2. Select the modules that you want to define as Compile Points and set **Type** to **locked, partition**.
3. Manually create a constraint file for each module to set constraints for each Compile Point.

Defining Compile Points by Creating a New SCOPE File

When you use the following method, the lower-level constraint file is created automatically:

1. On the File menu, click **New** and select **Constraint File**.
2. On the **Select File Type** tab of the **Create a New SCOPE File** dialog box, select **Compile Point**.

3. Select the module you want to designate as a Compile Point and click **OK**. The software automatically sets the Compile Points in the top-level constraint file and creates a lower-level constraint file for each Compile Point.

Additional Considerations for Compile Points

To ensure that changes to a Compile Point do not affect the top-level parent module, turn off the **Update Compile Point Timing Data** option in the **Implementation Options** dialog box. If this option is turned on, updates to a child module can impact the top-level module.

You can apply the `syn_allowed_resources` attribute to any Compile Point view to restrict the number of resources for a particular module.

When using Compile Points with incremental compilation, be aware of the following restrictions:

- To use Compile Points effectively, you must provide timing constraints (timing budgeting) for each Compile Point; the more accurate the constraints, the better your results are. Constraints are not automatically budgeted, so manual time budgeting is essential. Altera recommends that you register all inputs and outputs of each partition. This avoids any logic delay penalty on signals that cross-partition boundaries.
- When using the Synplify attribute `syn_useioff` to pack registers in the I/O Elements (IOEs) of Altera devices, these registers must be in the top-level module. Otherwise, you must direct the Quartus II software to perform I/O register packing instead of the `syn_useioff` attribute. You can use the **Fast Input Register** or **Fast Output Register** options, or set I/O timing constraints and turn on **Optimize I/O cell register placement for timing** on the **Fitter Settings** page of the **Settings** dialog box in the Quartus II software.
- There is no incremental synthesis support for top-level logic; any logic in the top-level is resynthesized during every compilation in the Synplify software.

 For more information about using Compile Points and setting Synplify attributes and constraints for both top-level and lower-level Compile Points, refer to the *Synopsys FPGA Synthesis User Guide* and the *Synopsys FPGA Synthesis Reference Manual* in the Synplify software.

Creating a Quartus II Project for Compile Points and Multiple .vqm Files

During compilation, the Synplify Pro and Premier software creates a `<top-level project>.tcl` file that provides the Quartus II software with the appropriate constraints and design partition assignments, creating a partition for each `.vqm` file along with the information to set up a Quartus II project. For details about using this Tcl script to set up your Quartus II project and pass your constraints, refer to [“Running the Quartus II Software Manually With the Synplify-Generated Tcl Script”](#) on page 14–15.

Depending on your design methodology, you can create one Quartus II project for all netlists or a separate Quartus II project for each netlist. In the standard incremental compilation design flow, you create design partition assignments and optional LogicLock™ floorplan location assignments for each partition in the design within a single Quartus II project. This methodology allows for the best quality of results and performance preservation during incremental changes to your design.

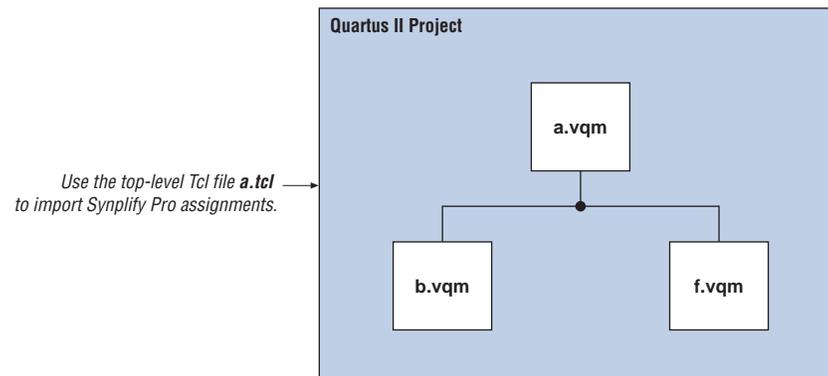
You might require a bottom-up design flow if each partition must be optimized separately, such as for third-party IP delivery. If you use this flow, Altera recommends you create a design floorplan to avoid placement conflicts between each partition. To follow this design flow in the Quartus II software, create separate Quartus II projects, export each design partition and incorporate it into a top-level design using the incremental compilation features to maintain placement results.

The following sections describe how to create the Quartus II projects for these two design flows.

Creating a Single Quartus II Project for a Standard Incremental Compilation Flow

Use the `<top-level project>.tcl` file that contains the Synplify assignments for all partitions within the project. This method allows you to import all the partitions into one Quartus II project and optimize all modules within the project at once, while taking advantage of the performance preservation and compilation-time reduction that incremental compilation offers. Figure 14-4 shows a visual representation of the design flow for the example design in Figure 14-3 on page 14-31.

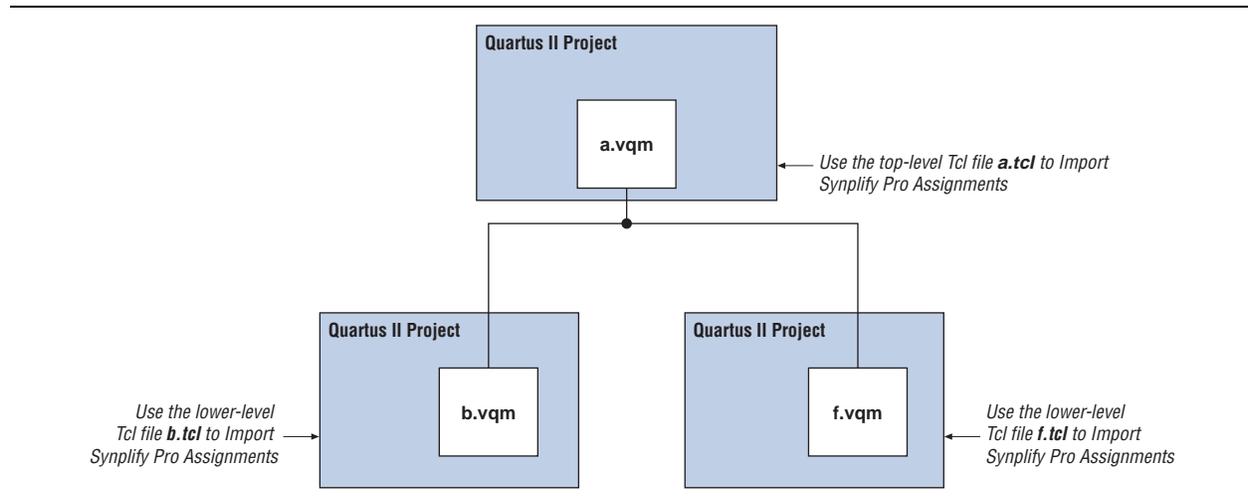
Figure 14-4. Design Flow Using Multiple .vqm Files with One Quartus II Project



Creating Multiple Quartus II Projects for a Bottom-Up Incremental Compilation Flow

Use the `<lower-level compile point>.tcl` files that contain the Synplify assignments for each Compile Point. Generate multiple Quartus II projects, one for each partition and netlist in the design. The designers in the project can optimize their own partitions separately within the Quartus II software and export the results for their own partitions. Figure 14-5 shows a visual representation of the design flow for the example design in Figure 14-3 on page 14-31. You can export the optimized subdesigns and then import them into one top-level Quartus II project using incremental compilation to complete the design.

Figure 14-5. Design Flow Using Multiple .vqm Files with Multiple Quartus II Projects



Creating Multiple .vqm Files for an Incremental Compilation Flow With Separate Synplify Projects

This section describes how to manually generate multiple .vqm files for an incremental compilation flow with black boxes and separate Synplify projects for each design partition. This manual flow is supported by versions of the Synplify software without the MultiPoint Synthesis feature.

Manually Creating Multiple .vqm Files With Black Boxes

To create multiple .vqm files manually in the Synplify software, create a separate project for each lower-level module and top-level design that you want to maintain as a separate .vqm file for an incremental compilation partition. Implement black box instantiations of lower-level partitions in your top-level project.

When synthesizing the projects for the lower-level modules, perform the following steps:

1. In the **Implementation Options** dialog box, turn on **Disable I/O Insertion** for the target technology.
2. Read the HDL files for the modules.



Modules might include black box instantiations of lower-level modules that are also maintained as separate .vqm files.

3. Add constraints with the SCOPE constraint window.
4. Enter the clock frequency to ensure that the sub-design is correctly optimized.
5. In the **Attributes** tab, set `syn_netlist_hierarchy` to 0.

When synthesizing the top-level design project, perform the following steps:

1. In the **Implementation Options** dialog box, turn off **Disable I/O Insertion** for the target technology.
2. Read the HDL files for top-level designs.

3. Create black boxes using lower-level modules in the top-level design.
4. Add constraints with the SCOPE constraint window.
5. Enter the clock frequency to ensure that the design is correctly optimized.
6. In the **Attributes** tab, set `syn_netlist_hierarchy` to 0.

The following sections describe an example of black box implementation to create separate `.vqm` files. [Figure 14-3 on page 14-31](#) shows an example of a design hierarchy that is split into multiple partitions.

The partition top contains the top-level block in the design (block A) and the logic that is not defined as part of another partition. In this example, the partition for top-level block A also includes the logic in one of its sub-blocks, block C. Because block F is contained in its own partition, it is not treated as part of the top-level partition A. Another separate partition, partition B, contains the logic in blocks B, D, and E. In a team-based design, engineers can work independently on the logic in different partitions. One netlist is created for the top-level module A and its submodule C, another netlist is created for module B and its submodules D and E, while a third netlist is created for module F.

To create multiple `.vqm` files for this design, follow these steps:

1. Generate a `.vqm` file for module B. Use `B.v.vhd`, `D.v.vhd`, and `E.v.vhd` as the source files.
2. Generate a `.vqm` file for module F. Use `F.v.vhd` as the source files.
3. Generate a top-level `.vqm` file for module A. Use `A.v.vhd` and `C.v.vhd` as the source files. Ensure that you use black box modules B and F, which were optimized separately in the previous steps.

Creating Black Boxes in Verilog HDL

Any design block that is not defined in the project, or included in the list of files to be read for a project, is treated as a black box by the software. Use the `syn_black_box` attribute to indicate that you intend to create a black box for the module. In Verilog HDL, you must provide an empty module declaration for a module that is treated as a black box.

Example 14-16 shows an example of the **A.v** top-level file. Follow the same procedure for lower-level files that also contain a black box for any module beneath the current level hierarchy.

Example 14-16. Verilog HDL Black Box for Top-Level File A.v

```

module A (data_in, clk, e, ld, data_out);
    input data_in, clk, e, ld;
    output [15:0] data_out;

    wire [15:0] cnt_out;

    B U1 (.data_in (data_in), .clk(clk), .ld (ld), .data_out(cnt_out));
    F U2 (.d(cnt_out), .clk(clk), .e(e), .q(data_out));

    // Any other code in A.v goes here.
endmodule

// Empty Module Declarations of Sub-Blocks B and F follow here.
// These module declarations (including ports) are required for black
// boxes.

module B (data_in, clk, ld, data_out) /* synthesis syn_black_box */ ;
    input data_in, clk, ld;
    output [15:0] data_out;
endmodule

module F (d, clk, e, q) /* synthesis syn_black_box */ ;
    input [15:0] d;
    input clk, e;
    output [15:0] q;
endmodule

```

Creating Black Boxes in VHDL

Any design block that is not defined in the project, or included in the list of files to be read for a project, is treated as a black box by the software. Use the `syn_black_box` attribute to indicate that you intend to treat the component as a black box. In VHDL, you must have a component declaration for the black box.

 Although VHDL is not case-sensitive, a **.vqm** (a subset of Verilog HDL) file is case-sensitive. Entity names and their port declarations are forwarded to the **.vqm** file. Black box names and port declarations are also passed to the **.vqm** file. To prevent case-based mismatches, use the same capitalization for black box and entity declarations in VHDL designs.

Example 14-17 shows an example of the **A.vhd** top-level file. Follow this same procedure for any lower-level files that contain a black box for any block beneath the current level of hierarchy.

Example 14-17. VHDL Black Box for Top-Level File A.vhd

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;
LIBRARY synplify;
USE synplify.attributes.all;

ENTITY A IS
PORT ( data_in : IN INTEGER RANGE 0 TO 15;
      clk, e, ld : IN STD_LOGIC;
      data_out : OUT INTEGER RANGE 0 TO 15 );
END A;

ARCHITECTURE a_arch OF A IS

COMPONENT B PORT(
  data_in : IN INTEGER RANGE 0 TO 15;
  clk, ld : IN STD_LOGIC;
  d_out : OUT INTEGER RANGE 0 TO 15 );
END COMPONENT;

COMPONENT F PORT(
  d : IN INTEGER RANGE 0 TO 15;
  clk, e: IN STD_LOGIC;
  q : OUT INTEGER RANGE 0 TO 15 );
END COMPONENT;

attribute syn_black_box of B: component is true;
attribute syn_black_box of F: component is true;

-- Other component declarations in A.vhd go here
signal cnt_out : INTEGER RANGE 0 TO 15;

BEGIN

U1 : B
PORT MAP (
  data_in => data_in,
  clk => clk,
  ld => ld,
  d_out => cnt_out );

U2 : F
PORT MAP (
  d => cnt_out,
  clk => clk,
  e => e,
  q => data_out );

-- Any other code in A.vhd goes here

END a_arch;

```

After you complete the steps described in this section, you have a netlist file for each partition of the design. These files are ready for use with the incremental compilation flow in the Quartus II software.

Creating a Quartus II Project for Multiple .vqm Files

The Synplify software creates a .tcl file for each .vqm file that provides the Quartus II software with the appropriate constraints and information to set up a project. For details about using the Tcl script generated by the Synplify software to set up your Quartus II project and pass your constraints, refer to “Running the Quartus II Software Manually With the Synplify-Generated Tcl Script” on page 14-15.

Depending on your design methodology, you can create one Quartus II project for all netlists or a separate Quartus II project for each netlist. In the standard incremental compilation design flow, you create design partition assignments and optional LogicLock floorplan location assignments for each partition in the design within a single Quartus II project. This methodology allows for the best quality of results and performance preservation during incremental changes to your design. You might require a bottom-up design flow where each partition must be optimized separately, such as for third-party IP delivery.

To perform follow this design flow in the Quartus II software, create separate Quartus II projects, export each design partition and incorporate it into a top-level design using the incremental compilation features to maintain the results.

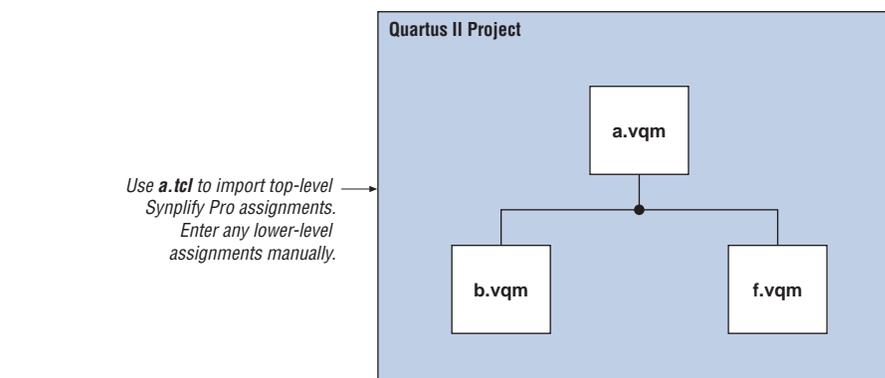
The following sections describe how to create the Quartus II projects for these two design flows.

Creating a Single Quartus II Project for a Standard Incremental Compilation Flow

Use the <top-level project>.tcl file that contains the Synplify assignments for the top-level design. This method allows you to import all of the partitions into one Quartus II project and optimize all modules within the project at once, taking advantage of the performance preservation and compilation time reduction offered by incremental compilation. Figure 14-6 shows a visual representation of the design flow for the example design in Figure 14-3 on page 14-31.

All of the constraints from the top-level project are passed to the Quartus II software in the top-level .tcl file, but constraints made in the lower-level projects within the Synplify software are not forward-annotated. Enter these constraints manually in your Quartus II project.

Figure 14-6. Design Flow Using Multiple .vqm Files with One Quartus II Project

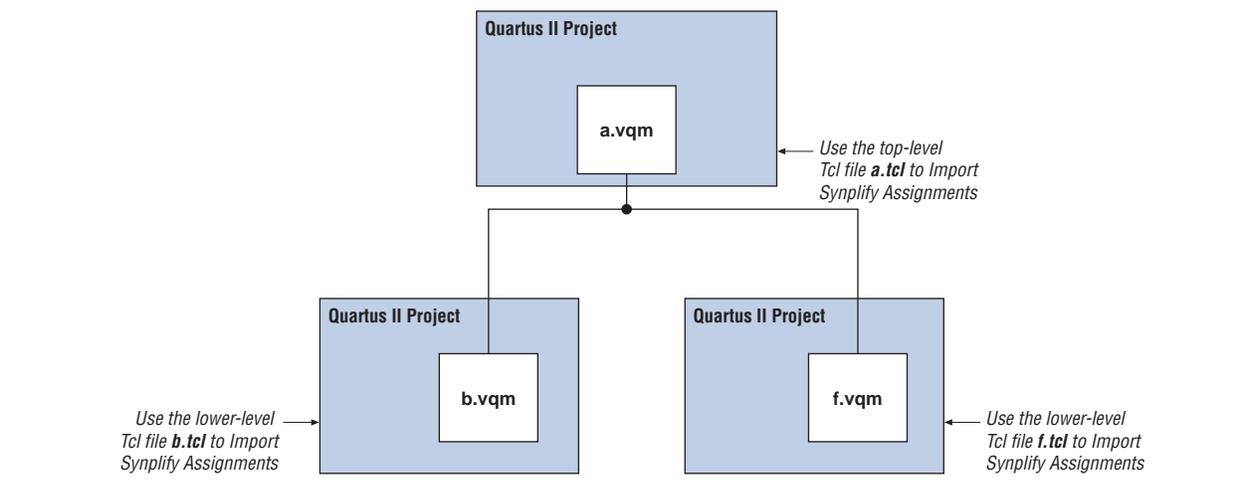


Creating Multiple Quartus II Projects for a Bottom-Up Incremental Compilation Flow

Use the `.tcl` file that is created for each `.vqm` file by the Synplify software for each Synplify project. This method generates multiple Quartus II projects, one for each block in the design. The designers in the project can optimize their own blocks separately within the Quartus II software and export the placement of their own blocks. Figure 14-7 shows a visual representation of the design flow for the example in Figure 14-3 on page 14-31.

Designers should create a LogicLock region to create a design floorplan for each block to avoid conflicts between partitions. The top-level designer then imports all the blocks and assignments into the top-level project. This method allows each block in the design to be optimized separately and then imported into one top-level project.

Figure 14-7. Design Flow Using Multiple Synplify Projects and Multiple Quartus II Projects



Performing Incremental Compilation in the Quartus II Software

In a standard design flow using Multipoint Synthesis, the Synplify software uses the Quartus II top-level `.tcl` file to ensure that the two tools databases stay synchronized. The Tcl file creates, changes, or deletes partition assignments in the Quartus II software for Compile Points that you create, change, or delete in the Synplify software. However, if you create, change, or delete a partition in the Quartus II software, the Synplify software does not change your Compile Point settings. Make any corresponding change in your Synplify project to ensure that you create the correct `.vqm` files.

 If you use the NativeLink integration feature described in “Using the Quartus II Software to Run the Synplify Software” on page 14-15, the Synplify software does not use any information about design partition assignments that you have set in the Quartus II software.

If you create netlist files with multiple Synplify projects, or if you do not use the Synplify Pro or Premier-generated `.tcl` files to update constraints in your Quartus II project, you must ensure that your Synplify `.vqm` netlists align with your Quartus II partition settings.

After you have set up your Quartus II project with `.vqm` netlist files as separate design partitions, set the appropriate Quartus II options to preserve your compilation results. On the Assignments menu, click **Design Partitions Window**. Change the **Netlist Type** to **Post-Fit** to preserve the previous compilation's post-fit placement results. If you do not make these settings, the Quartus II software does not reuse the placement or routing results from the previous compilation.

You can take advantage of incremental compilation with your Synplify design to reduce compilation time in the Quartus II software and preserve the results for unchanged design blocks.

 For more information about using Quartus II incremental compilation, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.

Conclusion

Taking advantage of the Synopsys Synplify and Altera Quartus II design flows allow you to control how your design files are prepared for the Quartus II place-and-route process, as well as improve performance and optimize a design for use with Altera devices.

Document Revision History

Table 14-4. Document Revision History (Part 1 of 2)

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Changed to new document template. ■ Removed Classic Timing Analyzer support. ■ Removed the “altera_implement_in_esb or altera_implement_in_eab” section. ■ Edited the “Creating a Quartus II Project for Compile Points and Multiple <code>.vqm</code> Files” on page 14–33 section for changes with the incremental compilation flow. ■ Edited the “Creating a Quartus II Project for Multiple <code>.vqm</code> Files” on page 14–39 section for changes with the incremental compilation flow. ■ Editorial changes.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Minor updates for the Quartus II software version 10.0 release.
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Minor updates for the Quartus II software version 9.1 release.
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Added new section “Exporting Designs to the Quartus II Software Using NativeLink Integration” on page 14–14. ■ Minor updates for the Quartus II software version 9.0 release. ■ Chapter 10 was previously Chapter 9 in software version 8.1.

Table 14-4. Document Revision History (Part 2 of 2)

Date	Version	Changes
November 2008	8.1.0	<ul style="list-style-type: none"> ■ Changed to 8-1/2 x 11 page size ■ Changed the chapter title from “Synplicity Synplify & Synplify Pro Support” to “Synopsys Synplify Support” ■ Replaced references to Synplicity with references to Synopsys ■ Added information about Synplify Premier ■ Updated supported device list ■ Added SystemVerilog information to Figure 14-1
May 2008	8.0.0	<ul style="list-style-type: none"> ■ Updated supported device list ■ Updated constraint annotation information for the TimeQuest Timing Analyzer ■ Updated RAM and MAC constraint limitations ■ Revised Table 9-1 ■ Added new section “Changing Synplify’s Default Behavior for Instantiated Altera Megafunctions” ■ Added new section “Instantiating Intellectual Property Using the MegaWizard Plug-In Manager and IP Toolbench” ■ Added new section “Including Files for Quartus II Placement and Routing Only” ■ Added new section “Additional Considerations for Compile Points” ■ Removed section “Apply the LogicLock Attributes” ■ Modified Figure 9-4, 9-43, 9-47. and 9-48 ■ Added new section “Performing Incremental Compilation in the Quartus II Software” ■ Numerous text changes and additions throughout the chapter ■ Renamed several sections ■ Updated “Referenced Documents” section

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this chapter.

This chapter documents support for the Mentor Graphics® Precision RTL Synthesis and Precision RTL Plus Synthesis software in the Quartus® II software design flow, as well as key design methodologies and techniques for improving your results for Altera® devices.

The topics discussed in this chapter include:

- “Altera Device Family Support”
- “Design Flow” on page 15–2
- “Creating and Compiling a Project in the Precision Synthesis Software” on page 15–4
- “Mapping the Precision Synthesis Design” on page 15–5
- “Synthesizing the Design and Evaluating the Results” on page 15–9
- “Exporting Designs to the Quartus II Software Using NativeLink Integration” on page 15–10
- “Guidelines for Altera Megafunctions and Architecture-Specific Features” on page 15–15
- “Incremental Compilation and Block-Based Design” on page 15–24

This chapter assumes that you have set up, licensed, and installed the Precision Synthesis software and the Quartus II software. You must set up, license, and install the Precision RTL Plus Synthesis software if you want to use the incremental synthesis feature for incremental compilation and block-based design.

-  To obtain and license the Precision Synthesis software, refer to the Mentor Graphics website at www.mentor.com. To install and run the Precision Synthesis software and to set up your work environment, refer to the *Precision Synthesis Installation Guide* in the Precision Manuals Bookcase. To access the Manuals Bookcase in the Precision Synthesis software, click **Help** and select **Open Manuals Bookcase**.

Altera Device Family Support

The Precision Synthesis software supports active devices available in the current version of the Quartus II software. Support for newly released device families may require an overlay. Contact Mentor Graphics for more information.

The Precision Synthesis software also supports the FLEX 8000 and MAX 9000 legacy devices that are supported only in the Altera MAX+PLUS® II software, as well as ACEX® 1K, APEX™ II, APEX 20K, APEX 20KC, APEX 20KE, FLEX® 10K, and FLEX 6000 legacy devices that are supported by the Quartus II software version 9.0 and earlier.

Design Flow

The following steps describe a basic Quartus II design flow using the Precision Synthesis software:

1. Create Verilog HDL or VHDL design files.
2. Create a project in the Precision Synthesis software that contains the HDL files for your design, select your target device, and set global constraints. Refer to [“Creating and Compiling a Project in the Precision Synthesis Software”](#) on page 15-4 for details.
3. Compile the project in the Precision Synthesis software.
4. Add specific timing constraints, optimization attributes, and compiler directives to optimize the design during synthesis.



For best results, Mentor Graphics recommends specifying constraints that are as close as possible to actual operating requirements. Properly setting clock and I/O constraints, assigning clock domains, and indicating false and multicycle paths guide the synthesis algorithms more accurately toward a suitable solution in the shortest synthesis time.

5. Synthesize the project in the Precision Synthesis software. With the design analysis and cross-probing capabilities of the Precision Synthesis software, you can identify and improve circuit area and performance issues using prelayout timing estimates.
6. Create a Quartus II project and import the following files generated by the Precision Synthesis software into the Quartus II project:
 - The technology-specific EDIF (.edf) netlist or Verilog Quartus Mapping File (.vqm) netlist
 - Synopsys Design Constraints File (.sdc) for TimeQuest Timing Analyzer constraints



If your design uses the Classic Timing Analyzer for timing analysis in the Quartus II software versions 10.0 and earlier, the Synplify software generates timing constraints in the Tcl Constraints File (.tcl). If you are using the Quartus II software versions 10.1 and later, you must use the TimeQuest Timing Analyzer for timing analysis.

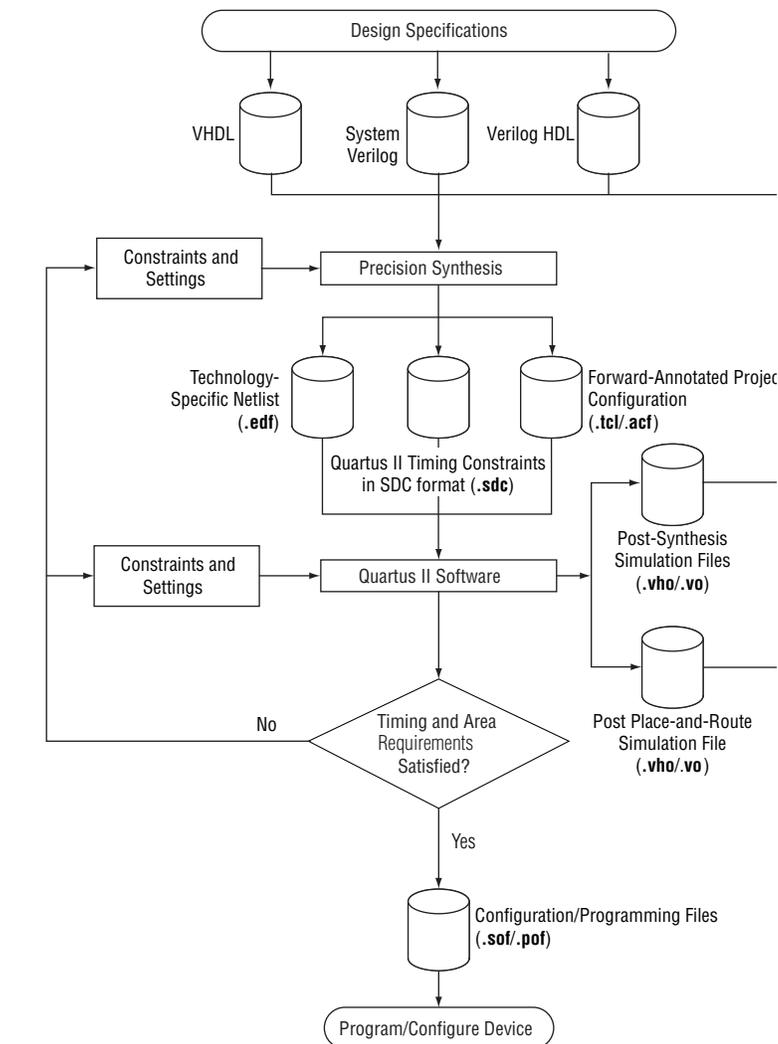
- Tcl Script Files (.tcl) to set up your Quartus II project and pass constraints

You can run the Quartus II software from within the Precision Synthesis software, or run the Precision Synthesis software using the Quartus II software. Refer to [“Running the Quartus II Software from within the Precision Synthesis Software”](#) on page 15-10 and [“Using the Quartus II Software to Run the Precision Synthesis Software”](#) on page 15-12 for more information.

7. After obtaining place-and-route results that meet your requirements, configure or program the Altera device.

Figure 15-1 shows the Quartus II design flow using the Precision Synthesis software as described in these steps, which are further described in detail in this chapter.

Figure 15-1. Design Flow Using the Precision Synthesis Software and Quartus II Software



If your area or timing requirements are not met, you can change the constraints and resynthesize the design in the Precision Synthesis software, or you can change the constraints to optimize the design during place-and-route in the Quartus II software. Repeat the process until the area and timing requirements are met.

You can use other options and techniques in the Quartus II software to meet area and timing requirements. For example, the **WYSIWYG Primitive Resynthesis** option can perform optimizations on your EDIF netlist in the Quartus II software.

 For more information about netlist optimizations, refer to the *Netlist Optimizations and Physical Synthesis* chapter in volume 2 of the *Quartus II Handbook*. For more recommendations about how to optimize your design, refer to the *Area and Timing Optimization* chapter in volume 2 of the *Quartus II Handbook*.

While simulation and analysis can be performed at various points in the design process, final timing analysis should be performed after placement and routing is complete.

During synthesis, the Precision Synthesis software produces several intermediate and output files, which are described in [Table 15-1](#).

Table 15-1. Precision Synthesis Software Intermediate and Output Files

File Extension	File Description
.psp	Precision Synthesis Project File.
.xdb	Mentor Graphics Design Database File.
.rep (1)	Synthesis Area and Timing Report File.
.vqm/.edf (2)	Technology-specific netlist in .vqm or .edf file format. By default, the Precision Synthesis software creates .vqm files for Arria series, Cyclone series, and Stratix series devices, and creates .edf files for ACEX, APEX, FLEX, and MAX series devices. The Precision Synthesis software can create .edf files for all Altera devices supported by the Quartus II software, but defaults to creating .vqm files when the device is supported.
.tcl	Forward-annotated Tcl assignments and constraints file. The <i><project name>.tcl</i> file is generated for all devices. The .tcl file acts as the Quartus II Project Configuration file and is used to make basic project and placement assignments, and to create and compile a Quartus II project.
.acf	Assignment and Configurations file for backward compatibility with the MAX+PLUS II software. For devices supported by the MAX+PLUS II software, the MAX+PLUS II assignments are imported from the MAX+PLUS II .acf file.
.sdc	Quartus II timing constraints file in Synopsys Design Constraints format. This file is generated automatically if the device uses the TimeQuest Timing Analyzer by default in the Quartus II software, and has the naming convention <i><project name>_pnr_constraints.sdc</i> . For more information about generating a TimeQuest constraint file, refer to “Exporting Designs to the Quartus II Software Using NativeLink Integration” on page 15-10.

Notes to Table 15-1:

- (1) The timing report file includes performance estimates that are based on pre-place-and-route information. Use the f_{MAX} reported by the Quartus II software after place-and-route for accurate post-place-and-route timing information. The area report file includes post-synthesis device resource utilization statistics that can differ from the resource usage after place-and-route due to black boxes or further optimizations performed during placement and routing. Use the device utilization reported by the Quartus II software after place-and-route for final resource utilization results. See [“Synthesizing the Design and Evaluating the Results”](#) on page 15-9 for details.
- (2) The Precision Synthesis software-generated VQM file is supported by the Quartus II software version 10.1 and later.

Creating and Compiling a Project in the Precision Synthesis Software

After creating your design files, create a project in the Precision Synthesis software that contains the basic settings for compiling the design.

To create a project, follow these steps:

1. In the Precision Synthesis software, click **New Project** in the **Design Bar** on the left side of the GUI.
2. Specify the **Project Name** and the **Project Folder**. The implementation name of the design corresponds to this project name.
3. Add input files to the project by clicking **Add Input Files** in the **Design Bar**. The Precision Synthesis software automatically detects the top-level module/entity of the design and uses it to name the current implementation directory, logs, reports, and netlist files.

4. In the **Design Bar**, click **Setup Design**.
5. To specify a target device family, expand **Altera** and select the target device and speed grade.
6. If you want, you can set a global design frequency and/or default input and output delays. This constrains all clock paths and I/O pins in your design. Modify the settings for individual paths or pins that do not require such a setting.
7. On the **Design Center** tab, right-click the **Output Files** folder and click **Output Options**.
8. To generate additional HDL netlists for post-synthesis simulation, select the desired output format. The Precision Synthesis software generates a separate file for each selected type of file: EDIF and Verilog HDL or VHDL.
9. To compile the design into a technology-independent implementation, in the **Design Bar**, click **Compile**.

Mapping the Precision Synthesis Design

In the next steps, you set constraints and map the design to technology-specific cells. The Precision Synthesis software maps the design by default to the fastest possible implementation that meets your timing constraints. To accomplish this, you must specify timing requirements for the automatically determined clock sources. With this information, the Precision Synthesis software performs static timing analysis to determine the location of the critical timing paths. The Precision Synthesis software achieves the best results for your design when you set as many realistic constraints as possible. Be sure to set constraints for timing, mapping, false paths, multicycle paths, and other factors that control the structure of the implemented design.

Mentor Graphics recommends creating an **.sdc** file and adding this file to the **Constraint Files** section of the **Project Files** list. You can create this file with a text editor, by issuing command-line constraint parameters, or by directing the Precision Synthesis software to generate the file automatically the first time you synthesize your design. To create a constraint file with the user interface, set constraints on design objects (such as clocks, design blocks, or pins) in the Design Hierarchy browser. By default, the Precision Synthesis software saves all timing constraints and attributes in two files: **precision_rtl.sdc** and **precision_tech.sdc**. The **precision_rtl.sdc** file contains constraints set on the RTL-level database (post-compilation) and the **precision_tech.sdc** file contains constraints set on the gate-level database (post-synthesis) located in the current implementation directory.

You can also enter constraints at the command line. After adding constraints at the command line, update the **.sdc** file with the **update constraint file** command. You can add constraints that change infrequently directly to the HDL source files with HDL attributes or pragmas.



The Precision **.sdc** file contains all the constraints for the Precision Synthesis project. For the Quartus II software, placement constraints are written in a **.tcl** file and timing constraints for the TimeQuest Timing Analyzer are written in the Quartus II **.sdc** file.

-  For details about the syntax of Synopsys Design Constraint commands, refer to the *Precision RTL Synthesis User's Manual* and the *Precision Synthesis Reference Manual*. For more details and examples of attributes, refer to the *Attributes* chapter in the *Precision Synthesis Reference Manual*. To access these manuals in the Precision Synthesis software, click **Help** and select **Open Manuals Bookcase**.

Setting Timing Constraints

The Precision Synthesis software uses timing constraints, based on the industry-standard `.sdc` file format, to deliver optimal results. Missing timing constraints can result in incomplete timing analysis and might prevent timing errors from being detected. The Precision Synthesis software provides constraint analysis prior to synthesis to ensure that designs are fully and accurately constrained. The `<project name>_pnr_constraints.sdc` file, which contains timing constraints in SDC format, is generated in the Quartus II software.

-  Because the `.sdc` file format requires that timing constraints be set relative to defined clocks, you must specify your clock constraints before applying any other timing constraints.

You also can use multicycle path and false path assignments to relax requirements or exclude nodes from timing requirements, which can improve area utilization and allow the software optimizations to focus on the most critical parts of the design.

-  For details about the syntax of Synopsys Design Constraint commands, refer to the *Precision RTL Synthesis User's Manual* and the *Precision Synthesis Reference Manual*. To access these manuals in the Precision Synthesis software, click **Help** and select **Open Manuals Bookcase**.

Setting Mapping Constraints

Mapping constraints affect how your design is mapped into the target Altera device. You can set mapping constraints in the user interface, in HDL code, or with the `set_attribute` command in the constraint file.

Assigning Pin Numbers and I/O Settings

The Precision Synthesis software supports assigning device pin numbers, I/O standards, drive strengths, and slew-rate settings to top-level ports of the design. You can set these timing constraints with the `set_attribute` command, the GUI, or by specifying synthesis attributes in your HDL code. These constraints are forward-annotated in the `<project name>.tcl` file that is read by the Quartus II software during place-and-route and do not affect synthesis.

You can use the `set_attribute` command in the Precision Synthesis software `.sdc` file format to specify pin number constraints, I/O standards, drive strengths, and slow slew-rate settings. Table 15-2 describes the format to use for entries in the Precision Synthesis software constraint file.

Table 15-2. Constraint File Settings

Constraint	Entry Format for Precision Constraint File
Pin number	<code>set_attribute -name PIN_NUMBER -value "<pin number>" -port <port name></code>
I/O standard	<code>set_attribute -name IOSTANDARD -value "<I/O Standard>" -port <port name></code>
Drive strength	<code>set_attribute -name DRIVE -value "<drive strength in mA>" -port <port name></code>
Slew rate	<code>set_attribute -name SLEW -value "TRUE FALSE" -port <port name></code>

You can also specify these options in the GUI. To specify a pin number or other I/O setting in the Precision Synthesis GUI, follow these steps:

1. After compiling the design, expand **Ports** in the Design Hierarchy Browser.
2. Under **Ports**, expand **Inputs** or **Outputs**.



You also can assign I/O settings by right-clicking the pin in the Schematic Viewer.

3. Right-click the desired pin name and select **Set Input Constraints** under **Inputs** or **Set Output Constraints** under **Outputs**.
4. Type the desired pin number on the Altera device in the **Pin Number** box in the **Port Constraints** dialog box.
5. Select the I/O standard from the **IO_STANDARD** list.
6. For output pins, you can also select a drive strength setting and slew rate setting using the **DRIVE** and **SLOW SLEW** lists.

You also can use synthesis attributes or pragmas in your HDL code to make these assignments. Example 15-1 and Example 15-2 show code samples that make a pin assignment in your HDL code.

Example 15-1. Verilog HDL Pin Assignment

```
//pragma attribute clk pin_number P10;
```

Example 15-2. VHDL Pin Assignment

```
attribute pin_number : string
attribute pin_number of clk : signal is "P10";
```

You can use the same syntax to assign the I/O standard using the **IOSTANDARD** attribute, drive strength using the attribute **DRIVE**, and slew rate using the **SLEW** attribute.



For more details about attributes and how to set these attributes in your HDL code, refer to the *Precision Synthesis Reference Manual*. To access this manual, in the Precision Synthesis software, click **Help** and select **Open Manuals Bookcase**.

Assigning I/O Registers

The Precision Synthesis software performs timing-driven I/O register mapping by default. You can force a register to the device's IO element (IOE) using the Complex I/O constraint. This option does not apply if you turn off **I/O pad insertion**. Refer to [“Disabling I/O Pad Insertion” on page 15-8](#) for more information.

To force an I/O register into the device's IOE using the GUI, follow these steps:

1. After compiling the design, expand **Ports** in the Design Hierarchy browser.
2. Under **Ports**, expand **Inputs** or **Outputs**.
3. Under **Inputs** or **Outputs**, right-click the desired pin name, point to **Map Input Register to IO** or **Map Output Register to IO**, for input or output respectively, and click **True**.



You also can make the assignment by right-clicking on the pin in the Schematic Viewer.

For the Stratix series, Cyclone series, and the MAX II device families, the Precision Synthesis software can move an internal register to an I/O register without any restrictions on design hierarchy.

For more mature devices, the Precision Synthesis software can move an internal register to an I/O register only when the register exists in the top-level of the hierarchy. If the register is buried in the hierarchy, you must flatten the hierarchy so that the buried registers are moved to the top-level of the design.

Disabling I/O Pad Insertion

The Precision Synthesis software assigns I/O pad atoms (device primitives used to represent the I/O pins and I/O registers) to all ports in the top-level of a design by default. In certain situations, you might not want the software to add I/O pads to all I/O pins in the design. The Quartus II software can compile a design without I/O pads; however, including I/O pads provides the Precision Synthesis software with more information about the top-level pins in the design.

Preventing the Precision Synthesis Software from Adding I/O Pads

If you are compiling a subdesign as a separate project, I/O pins cannot be primary inputs or outputs of the device; therefore, the I/O pins should not have an I/O pad associated with them. To prevent the Precision Synthesis software from adding I/O pads, perform the following steps:

1. On the Tools menu, click **Set Options**. The **Options** dialog box appears.
2. On the **Optimization** page, turn off **Add IO Pads**.
3. Click **Apply**.

These steps add the following command to the project file:

```
setup_design -addio=false
```

Preventing the Precision Synthesis Software from Adding an I/O Pad on an Individual Pin

To prevent I/O pad insertion on an individual pin when you are using a black box, such as DDR or a phase-locked loop (PLL), at the external ports of the design, perform the following steps:

1. After compiling the design, in the Design Hierarchy browser, expand **Ports**.
2. Under **Ports**, expand **Inputs** or **Outputs**.
3. Under **Inputs** or **Outputs**, right-click the desired pin name and click **Set Input Constraints**.
4. In the **Port Constraints** dialog box for the selected pin name, turn off **Insert Pad**.



You also can make this assignment by right-clicking the pin in the Schematic Viewer or by attaching the nopad attribute to the port in the HDL source code.

Controlling Fan-Out on Data Nets

Fan-out is defined as the number of nodes driven by an instance or top-level port. High fan-out nets can cause significant delays that result in an unroutable net. On a critical path, high fan-out nets can cause longer delays in a single net segment that result in the timing constraints not being met. To prevent this behavior, each device family has a global fan-out value set in the Precision Synthesis software library. In addition, the Quartus II software automatically routes high fan-out signals on global routing lines in the Altera device whenever possible.

To eliminate routability and timing issues associated with high fan-out nets, the Precision Synthesis software also allows you to override the library default value on a global or individual net basis. You can override the library value by setting a `max_fanout` attribute on the net.

Synthesizing the Design and Evaluating the Results

To synthesize the design for the target device, click **Synthesize** in the **Precision Synthesis Design Bar**. During synthesis, the Precision Synthesis software optimizes the compiled design, and then writes out netlists and reports to the implementation subdirectory of your working directory after the implementation is saved, using the following naming convention:

```
<project name>_impl_<number>
```



After synthesis is complete, you can evaluate the results for area and timing. The *Precision RTL Synthesis User's Manual* on the Mentor Graphics website describes different results that can be evaluated in the software.

There are several schematic viewers available in the Precision Synthesis software: RTL schematic, Technology-mapped schematic, and Critical Path schematic. These analysis tools allow you to quickly and easily isolate the source of timing or area issues, and to make additional constraint or code changes to optimize the design.

Obtaining Accurate Logic Utilization and Timing Analysis Reports

Historically, designers have relied on post-synthesis logic utilization and timing reports to determine the amount of logic their design requires, the size of the device required, and how fast the design runs. However, today's FPGA devices provide a wide variety of advanced features in addition to basic registers and look-up tables (LUTs). The Quartus II software has advanced algorithms to take advantage of these features, as well as optimization techniques to increase performance and reduce the amount of logic required for a given design. In addition, designs can contain black boxes and functions that take advantage of specific device features. Because of these advances, synthesis tool reports provide post-synthesis area and timing estimates, but you should use the place-and-route software to obtain final logic utilization and timing reports.

Exporting Designs to the Quartus II Software Using NativeLink Integration

The NativeLink feature in the Quartus II software facilitates the seamless transfer of information between the Quartus II software and EDA tools, which allows you to run other EDA design entry/synthesis, simulation, and timing analysis tools automatically from within the Quartus II software.

After a design is synthesized in the Precision Synthesis software, the technology-mapped design is written to the current implementation directory as an EDIF netlist file, along with a Quartus II Project Configuration File and a place-and-route constraints file. You can use the Project Configuration script, `<project name>.tcl`, to create and compile a Quartus II project for your EDIF or VQM netlist. This script makes basic project assignments, such as assigning the target device specified in the Precision Synthesis software. If you select an Arria GX, Stratix III, Cyclone III, or newer device, the constraints are written in SDC format to the `<project name>_pnr_constraints.sdc` file by default, which is used by the Fitter and the TimeQuest Timing Analyzer in the Quartus II software.

Use the following Precision Synthesis software command before compilation to generate the `<project name>_pnr_constraints.sdc`:

```
setup_design -timequest_sdc
```

With this command, the file is generated after the synthesis.

Running the Quartus II Software from within the Precision Synthesis Software

The Precision Synthesis software also has a built-in place-and-route environment that allows you to run the Quartus II Fitter and view the results in the Precision Synthesis GUI. This feature is useful when performing an initial compilation of your design to view post-place-and-route timing and device utilization results, but not all the advanced Quartus II options that control the compilation process are available.

After you specify an Altera device as the target, set the options for the Quartus II software. On the Tools menu, click **Set Options**. On the **Integrated Place and Route** page, under **Quartus II Modular**, specify the path to the Quartus II executables in the **Path to Quartus II installation tree** box.

To automate the place-and-route process, click **Run Quartus II** in the **Quartus II Modular** window of the Precision Synthesis toolbar. The Quartus II software uses the current implementation directory as the Quartus II project directory and runs a full compilation in the background (that is, the user interface does not appear).

Two primary Precision Synthesis software commands control the place-and-route process. Use the `setup_place_and_route` command to set the place-and-route options. Start the process with the `place_and_route` command.

Precision Synthesis software uses individual Quartus II executables, such as analysis and synthesis (`quartus_map`), Fitter (`quartus_fit`), and the TimeQuest Timing Analyzer (`quartus_sta`) for improved runtime and memory utilization during place and route. This flow is referred to as the **Quartus II Modular** flow option in the Precision Synthesis software. By default, the Precision Synthesis software generates a Quartus II Project Configuration File (`.tcl` file) for current device families. Timing constraints that you set during synthesis are exported to the Quartus II place-and-route constraints file `<project name>_pnr_constraints.sdc`.

After you compile the design in the Quartus II software from within the Precision Synthesis software, you can invoke the Quartus II GUI manually and then open the project using the generated Quartus II project file. You can view reports, run analysis tools, specify options, and run the various processing flows available in the Quartus II software.



For more information about running the Quartus II software from within the Precision Synthesis software, refer to the *Altera Quartus II Integration* chapter in the *Precision Synthesis Reference Manual*. To access this manual in the Precision Synthesis software, click **Help** and select **Open Manuals Bookcase**.

Running the Quartus II Software Manually Using the Precision Synthesis-Generated Tcl Script

You can run the Quartus II software using a Tcl script generated by the Precision Synthesis software. To run the Tcl script generated by the Precision Synthesis software to set up your project and start a full compilation, perform the following steps:

1. Ensure the `.edf` or `.vqm` file, `.tcl` files, and `.sdc` file are located in the same directory. The files should be located in the implementation directory by default.
2. In the Quartus II software, on the View menu, point to **Utility Windows** and click **Tcl Console**.
3. At the Tcl Console command prompt, type the command:

```
source <path>/<project name>.tcl ←
```
4. On the File menu, click **Open Project**. Browse to the project name and click **Open**.
5. Compile the project in the Quartus II software.

Using the Quartus II Software to Run the Precision Synthesis Software

With NativeLink integration, you can set up the Quartus II software to run the Precision Synthesis software. This feature allows you to use the Precision Synthesis software to synthesize a design as part of a standard compilation. When you use this feature, the Precision Synthesis software does not use any timing constraints or assignments, such as incremental compilation partitions, that you have set in the Quartus II software.

- ❓ For detailed information about using NativeLink integration with the Precision Synthesis software, refer to *Using the NativeLink Feature with Other EDA Tools* in the Quartus II Help.

Passing Constraints to the Quartus II Software

The place-and-route constraints script forward-annotates timing constraints that you made in the Precision Synthesis software. This integration allows you to enter these constraints once in the Precision Synthesis software, and then pass them automatically to the Quartus II software.

Refer to the introductory text in the section “[Exporting Designs to the Quartus II Software Using NativeLink Integration](#)” on page 15-10 for information on how to ensure the Precision Synthesis software targets the TimeQuest Timing Analyzer.

The following constraints are translated by the Precision Synthesis software and are applicable to the TimeQuest Timing Analyzer:

- `create_clock`
- `set_input_delay`
- `set_output_delay`
- `set_max_delay`
- `set_min_delay`
- `set_false_path`
- `set_multicycle_path`

create_clock

You can specify a clock in the Precision Synthesis software, as shown in [Example 15-3](#).

Example 15-3. Specifying a Clock using `create_clock`

```
create_clock -name <clock_name> -period <period in ns> -waveform {<edge_list>} -domain \  
<ClockDomain> <pin>
```

The period is specified in units of nanoseconds (ns). If no clock domain is specified, the clock belongs to a default clock domain `main`. All clocks in the same clock domain are treated as synchronous (related) clocks. If no `<clock_name>` is provided, the default name `virtual_default` is used. The `<edge_list>` sets the rise and fall edges of the clock signal over an entire clock period. The first value in the list is a rising transition, typically the first rising transition after time zero. The waveform can contain any even number of alternating edges, and the edges listed should alternate between rising and falling. The position of any edge can be equal to or greater than zero but must be equal to or less than the clock period.

If `-waveform <edge_list>` is not specified and `-period <period in ns>` is specified, the default waveform has a rising edge of 0.0 and a falling edge of `<period_value>/2`.

The Precision Synthesis software maps the clock constraint to the TimeQuest `create_clock` setting in the Quartus II software.

The Quartus II software supports only clock waveforms with two edges in a clock cycle. If the Precision Synthesis software finds a multi-edge clock, it issues an error message when you synthesize your design in the Precision Synthesis software.

set_input_delay

This port-specific input delay constraint is specified in the Precision Synthesis software, as shown in [Example 15-4](#).

Example 15-4. Specifying set_input_delay

```
set_input_delay {<delay_value> <port_pin_list>} -clock <clock_name> -rise -fall -add_delay
```

This constraint is mapped to the `set_input_delay` setting in the Quartus II software.

When the reference clock `<clock_name>` is not specified, all clocks are assumed to be the reference clocks for this assignment. The input pin name for the assignment can be an input pin name of a time group. The software can use the `clock_fall` option to specify delay relative to the falling edge of the clock.



Although the Precision Synthesis software allows you to set input delays on pins inside the design, these constraints are not sent to the Quartus II software, and a message is displayed.

set_output_delay

This port-specific output delay constraint is specified in the Precision Synthesis software, as shown in [Example 15-5](#).

Example 15-5. Using the set_output_delay Constraint

```
set_output_delay {<delay_value> <port_pin_list>} -clock <clock_name> -rise -fall -add_delay
```

This constraint is mapped to the `set_output_delay` setting in the Quartus II software.

When the reference clock `<clock_name>` is not specified, all clocks are assumed to be the reference clocks for this assignment. The output pin name for the assignment can be an output pin name of a time group.



Although the Precision Synthesis software allows you to set output delays on pins inside the design, these constraints are not sent to the Quartus II software.

set_max_delay and set_min_delay

The maximum delay for a point-to-point timing path constraint is specified in the Precision Synthesis software, as shown in [Example 15-6](#). The minimum delay for a point-to-point timing path constraint is shown in [Example 15-7](#).

Example 15-6. Using the set_max_delay Constraint

```
set_max_delay -from {<from_node_list>} -to {<to_node_list>} <delay_value>
```

Example 15-7. Using the set_min_delay Constraint

```
set_min_delay -from {<from_node_list>} -to {<to_node_list>} <delay_value>
```

The `set_max_delay` and `set_min_delay` commands specify that the maximum and minimum respectively, required delay for any start point in `<from_node_list>` to any endpoint in `<to_node_list>` must be less than or greater than `<delay_value>`. Typically, you use these commands to override the default setup constraint for any path with a specific maximum or minimum time value for the path.

The node lists can contain a collection of clocks, registers, ports, pins, or cells. The `-from` and `-to` parameters specify the source (start point) and the destination (endpoint) of the timing path, respectively. The source list (`<from_node_list>`) cannot include output ports, and the destination list (`<to_node_list>`) cannot include input ports. If you include more than one node on a list, you must enclose the nodes in quotes or in braces ({}).

If you specify a clock in the source list, you must specify a clock in the destination list. Applying `set_max_delay` or `set_min_delay` setting between clocks applies the exception from all registers or ports driven by the source clock to all registers or ports driven by the destination clock. Applying exceptions between clocks is more efficient than applying them for specific node-to-node, or node-to-clock paths. If you want to specify pin names in the list, the source must be a clock pin and the destination must be any non-clock input pin to a register. Assignments from clock pins, or to and from cells, apply to all registers in the cell or for those driven by the clock pin.

set_false_path

The false path constraint is specified in the Precision Synthesis software, as shown in [Example 15-8](#).

Example 15-8. Using the set_false_path Constraint

```
set_false_path -to <to_node_list> -from <from_node_list> -reset_path
```

The node lists can be a list of clocks, ports, instances, and pins. Multiple elements in the list can be represented using wildcards such as `*` and `?`.

In a place-and-route Tcl constraints file, this false path setting in the Precision Synthesis software is mapped to a `set_false_path` setting. The Quartus II software supports `setup`, `hold`, `rise`, or `fall` options for this assignment.

The node lists for this assignment represents top-level ports and/or nets connected to instances (end points of timing assignments).

Any false path setting in the Precision Synthesis software can be mapped to a setting in the Quartus II software with a through path specification.

set_multicycle_path

This multicycle path constraint is specified in the Precision Synthesis software, as shown in [Example 15-9](#).

Example 15-9. Using the set_multicycle_path Constraint

```
set_multicycle_path <multiplier_value> [-start] [-end] -to <to_node_list> -from <from_node_list> \  
-reset_path
```

The node list can contain clocks, ports, instances, and pins. Multiple elements in the list can be represented using wildcards such as * and ?. Paths without multicycle path definitions are identical to paths with multipliers of 1. To add one additional cycle to the datapath, use a multiplier value of 2. The option `start` indicates that source clock cycles should be considered for the multiplier. The option `end` indicates that destination clock cycles should be considered for the multiplier. The default is to reference the end clock.

In the place-and-route Tcl constraints file, the multicycle path setting in the Precision Synthesis software is mapped to a `set_multicycle_path` setting. The Quartus II software supports the `rise` or `fall` options on this assignment.

The node lists represent top-level ports and/or nets connected to instances (end points of timing assignments). The node lists can contain wildcards (such as *); the Quartus II software automatically expands all wildcards.

Any multicycle path setting in Precision Synthesis software can be mapped to a setting in the Quartus II software with a `-through` specification.

Guidelines for Altera Megafunctions and Architecture-Specific Features

Altera provides parameterizable megafunctions, including the LPMs, device-specific Altera megafunctions, IP available as Altera MegaCore functions, and IP available through the Altera Megafunction Partners Program (AMPPSM). You can use megafunctions and IP functions by instantiating them in your HDL code or by inferring certain megafunctions from generic HDL code.

If you want to instantiate a megafunction such as a PLL in your HDL code, you can instantiate and parameterize the function using the port and parameter definitions, or you can customize a function with the MegaWizardTM Plug-In Manager. Altera recommends using the MegaWizard Plug-In Manager, which provides a graphical interface within the Quartus II software for customizing and parameterizing any available megafunction for the design. [“Instantiating Altera Megafunctions Using the MegaWizard Plug-In Manager”](#) and [“Instantiating Intellectual Property With the MegaWizard Plug-In Manager and IP Toolbench”](#) on page 15-17 describe the MegaWizard Plug-In Manager flow with the Precision Synthesis software.

- For more information about specific Altera megafunctions and IP functions, refer to the Altera [IP and Megafunctions](#) website.

The Precision Synthesis software automatically recognizes certain types of HDL code and infers the appropriate function. The Precision Synthesis software provides options to control inference of certain types of megafunctions, as described in [“Inferring Altera Megafunctions from HDL Code”](#) on page 15-19.

- For a detailed discussion about instantiating functions versus inferring functions to target Altera architecture-specific features, refer to the [Recommended HDL Coding Styles](#) chapter in volume 1 of the *Quartus II Handbook*. This chapter also provides details on using the MegaWizard Plug-In Manager in the Quartus II software and explains the files generated by the wizard, as well as coding style recommendations and HDL examples for inferring functions in Altera devices.

Instantiating Altera Megafunctions Using the MegaWizard Plug-In Manager

This section describes how to instantiate Altera megafunctions with the MegaWizard Plug-In Manager, and how to generate the files that are included in the Precision Synthesis project for synthesis.

You can run the stand-alone version of the MegaWizard Plug-In Manager by typing the following command at a command prompt:

```
qmegawiz ←
```

Instantiating Megafunctions With MegaWizard Plug-In Manager-Generated Verilog HDL Files

The MegaWizard Plug-In Manager generates a Verilog HDL instantiation template file `<output file>_inst.v` and a hollow-body black box module declaration `<output file>_bb.v` for use in your Precision Synthesis design. Incorporate the instantiation template file, `<output file>_inst.v`, into your top-level design to instantiate the megafunction wrapper file, `<output file>.v`.

Include the hollow-body black box module declaration `<output file>_bb.v` in your Precision Synthesis project to describe the port connections of the black box. Adding the megafunction wrapper file `<output file>.v` in your Precision Synthesis project is optional, but you must add it to your Quartus II project along with the Precision Synthesis-generated EDIF or VQM netlist.

Alternatively, you can include the megafunction wrapper file `<output file>.v` in your Precision Synthesis project and then right-click the file in the input file list, and select **Properties**. In the **Input file properties** dialog box, turn on **Exclude file from Compile Phase** and click **OK**. When this option is turned on, the Precision Synthesis software excludes the file from compilation and copies the file to the appropriate directory for use by the Quartus II software during place-and-route.

Instantiating Megafunctions With MegaWizard Plug-In Manager-Generated VHDL Files

The MegaWizard Plug-In Manager generates a VHDL component declaration file `<output file>.cmp` and a VHDL instantiation template file `<output file>_inst.vhd` for use in your Precision Synthesis design. Incorporate the component declaration and instantiation template into your top-level design to instantiate the megafunction wrapper file, `<output file>.vhd`.

Adding the megafunction wrapper file `<output file>.vhd` in your Precision Synthesis project is optional, but you must add the file to your Quartus II project along with the Precision Synthesis-generated EDIF or VQM netlist.

Alternatively, you can include the megafunction wrapper file `<output file>.vhd` in your Precision Synthesis project and then right-click the file in the input file list, and select **Properties**. In the **Input file properties** dialog box, turn on **Exclude file from Compile Phase** and click **OK**. When this option is turned on, the Precision Synthesis software excludes the file from compilation and copies the file to the appropriate directory for use by the Quartus II software during place-and-route.

Instantiating Intellectual Property With the MegaWizard Plug-In Manager and IP Toolbench

Many Altera IP functions include a resource and timing estimation netlist that the Precision Synthesis software can use to synthesize and optimize logic around the IP efficiently. As a result, the Precision Synthesis software provides better timing correlation, area estimates, and Quality of Results (QoR) than a black box approach.

To create this netlist file, perform the following steps:

1. Select the IP function in the MegaWizard Plug-In Manager.
2. Click **Next** to open the IP Toolbench.
3. Click **Set Up Simulation**, which sets up all the EDA options.
4. Turn on the **Generate netlist** option to generate a netlist for resource and timing estimation and click **OK**.
5. Click **Generate** to generate the netlist file.

The Quartus II software generates a file `<output file>_syn.v`. This netlist contains the “grey box” information for resource and timing estimation, but does not contain the actual implementation. Include this netlist file into your Precision Synthesis project as an input file. Then include the megafunction wrapper file `<output file>.v | vhd` in the Quartus II project along with your EDIF or VQM output netlist.

 The generated “grey box” netlist file, `<output file>_syn.v`, is always in Verilog HDL format, even if you select VHDL as the output file format.

 There is currently no grey box support for SOPC Builder systems in the MegaWizard Plug-In Manager. For information about creating a grey box netlist file from the command line, search Altera's Knowledge Database. Alternatively, you can use a black box approach as described in [“Instantiating Black Box IP Functions With Generated Verilog HDL Files”](#).

Instantiating Black Box IP Functions With Generated Verilog HDL Files

You can use the `syn_black_box` or `black_box` compiler directives to declare a module as a black box. The top-level design files must contain the IP port mapping and a hollow-body module declaration. You can apply the directive to the module declaration in the top-level file or a separate file included in the project so that the Precision Synthesis software recognizes the module is a black box.



The `syn_black_box` and `black_box` directives are supported only on module or entity definitions.

[Example 15-10](#) shows a sample top-level file that instantiates `my_verilogIP.v`, which is a simplified customized variation generated by the MegaWizard Plug-In Manager and IP Toolbench.

Example 15-10. Top-Level Verilog HDL Code with Black Box Instantiation of IP

```
module top (clk, count);
    input clk;
    output[7:0] count;

    my_verilogIP verilogIP_inst (.clock (clk), .q (count));
endmodule

// Module declaration
// The following attribute is added to create a
// black box for this module.
module my_verilogIP (clock, q) /* synthesis syn_black_box */;
    input clock;
    output[7:0] q;
endmodule
```

Instantiating Black Box IP Functions With Generated VHDL Files

You can use the `syn_black_box` or `black_box` compiler directives to declare a component as a black box. The top-level design files must contain the megafunction variation component declaration and port mapping. Apply the directive to the component declaration in the top-level file.



The `syn_black_box` and `black_box` directives are supported only on module or entity definitions.

Example 15-11 shows a sample top-level file that instantiates `my_vhdlIP.vhd`, which is a simplified customized variation generated by the MegaWizard Plug-In Manager and IP Toolbench.

Example 15-11. Top-Level VHDL Code with Black Box Instantiation of IP

```
LIBRARY ieee;
USE ieee.std_logic_1164.all;
ENTITY top IS
  PORT (
    clk: IN STD_LOGIC ;
    count: OUT STD_LOGIC_VECTOR (7 DOWNT0 0)
  );
END top;

ARCHITECTURE rtl OF top IS
  COMPONENT my_vhdlIP
  PORT (
    clock: IN STD_LOGIC ;
    q: OUT STD_LOGIC_VECTOR (7 DOWNT0 0)
  );
  end COMPONENT;
  attribute syn_black_box : boolean;
  attribute syn_black_box of my_vhdlIP: component is true;
  BEGIN
    vhdlIP_inst : my_vhdlIP PORT MAP (
      clock => clk,
      q => count
    );
  END rtl;
```

Inferring Altera Megafunctions from HDL Code

The Precision Synthesis software automatically recognizes certain types of HDL code and maps arithmetical and relational operators, and memory (RAM and ROM), to efficient technology-specific implementations. This functionality allows technology-specific resources to implement these structures by inferring the appropriate Altera function to provide optimal results. In some cases, the Precision Synthesis software has options that you can use to disable or control inference.

-  For coding style recommendations and examples for inferring technology-specific architecture in Altera devices, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*, and the *Precision Synthesis Style Guide* in the Precision Manuals Bookcase. To access these manuals, in the Precision Synthesis software, click **Help** and select **Open Manuals Bookcase**.

Multipliers

The Precision Synthesis software detects multipliers in HDL code and maps them directly to device atoms to implement the multiplier in the appropriate type of logic. The Precision Synthesis software also allows you to control the device resources that are used to implement individual multipliers, as described in the following section.

Controlling DSP Block Inference for Multipliers

By default, the Precision Synthesis software uses DSP blocks available in Stratix series devices to implement multipliers. The default setting is **AUTO**, which allows the Precision Synthesis software to map to logic look-up tables (LUTs) or DSP blocks, depending on the size of the multiplier. You can use the Precision Synthesis GUI or HDL attributes to direct mapping to only logic elements or to only DSP blocks.

The options for multiplier mapping in the Precision Synthesis software are described in [Table 15-3](#).

Table 15-3. Options for dedicated_mult Parameter to Control Multiplier Implementation in Precision Synthesis

Value	Description
ON	Use only DSP blocks to implement multipliers, regardless of the size of the multiplier.
OFF	Use only logic (LUTs) to implement multipliers, regardless of the size of the multiplier.
AUTO	Use logic (LUTs) or DSP blocks to implement multipliers, depending on the size of the multipliers.

Setting the Use Dedicated Multiplier Option

To set the Use Dedicated Multiplier option in the Precision Synthesis GUI, compile the design, and then in the Design Hierarchy browser, right-click the operator for the desired multiplier and click **Use Dedicated Multiplier**.

Setting the dedicated_mult Attribute

To control the implementation of a multiplier in your HDL code, use the `dedicated_mult` attribute with the appropriate value from [Table 15-3](#), as shown in [Example 15-12](#) and [Example 15-13](#).

Example 15-12. Setting the dedicated_mult Attribute in Verilog HDL

```
//synthesis attribute <signal name> dedicated_mult <value>
```

Example 15-13. Setting the dedicated_mult Attribute in VHDL

```
ATTRIBUTE dedicated_mult: STRING;
ATTRIBUTE dedicated_mult OF <signal name>: SIGNAL IS <value>;
```

The `dedicated_mult` attribute can be applied to signals and wires; it does not work when applied to a register. This attribute can be applied only to simple multiplier code, such as `a = b * c`.

Some signals for which the `dedicated_mult` attribute is set can be removed during synthesis by the Precision Synthesis software for design optimization. In such cases, if you want to force the implementation, you should preserve the signal by setting the `preserve_signal` attribute to `TRUE`, as shown in [Example 15-14](#) and [Example 15-15](#).

Example 15-14. Setting the preserve_signal Attribute in Verilog HDL

```
//synthesis attribute <signal name> preserve_signal TRUE
```

Example 15-15. Setting the `preserve_signal` Attribute in VHDL

```
ATTRIBUTE preserve_signal: BOOLEAN;  
ATTRIBUTE preserve_signal OF <signal name>: SIGNAL IS TRUE;
```

[Example 15-16](#) and [Example 15-17](#) are examples, in Verilog HDL and VHDL, of using the `dedicated_mult` attribute to implement the given multiplier in regular logic in the Quartus II software.

Example 15-16. Verilog HDL Multiplier Implemented in Logic

```
module unsigned_mult (result, a, b);  
    output [15:0] result;  
    input [7:0] a;  
    input [7:0] b;  
    assign result = a * b;  
    //synthesis attribute result dedicated_mult OFF  
endmodule
```

Example 15-17. VHDL Multiplier Implemented in Logic

```
LIBRARY ieee;  
USE ieee.std_logic_1164.ALL;  
USE ieee.std_logic_arith.ALL;  
USE ieee.std_logic_unsigned.ALL;  
  
ENTITY unsigned_mult IS  
    PORT(  
        a: IN std_logic_vector (7 DOWNTO 0);  
        b: IN std_logic_vector (7 DOWNTO 0);  
        result: OUT std_logic_vector (15 DOWNTO 0));  
    ATTRIBUTE dedicated_mult: STRING;  
END unsigned_mult;  
  
ARCHITECTURE rtl OF unsigned_mult IS  
    SIGNAL a_int, b_int: UNSIGNED (7 downto 0);  
    SIGNAL pdt_int: UNSIGNED (15 downto 0);  
    ATTRIBUTE dedicated_mult OF pdt_int: SIGNAL IS "OFF";  
BEGIN  
    a_int <= UNSIGNED (a);  
    b_int <= UNSIGNED (b);  
    pdt_int <= a_int * b_int;  
    result <= std_logic_vector(pdt_int);  
END rtl;
```

Multiplier-Accumulators and Multiplier-Adders

The Precision Synthesis software also allows you to control the device resources used to implement multiply-accumulators or multiply-adders in your project or in a particular module.

The Precision Synthesis software detects multiply-accumulators or multiply-adders in HDL code and infers an `ALTMULT_ACCUM` or `ALTMULT_ADD` megafunction so that the logic can be placed in DSP blocks, or the software maps these functions directly to device atoms to implement the multiplier in the appropriate type of logic.

-  The Precision Synthesis software supports inference for these functions only if the target device family has dedicated DSP blocks. Refer to “[Controlling DSP Block Inference](#)” for more information.
-  For more information about DSP blocks in Altera devices, refer to the appropriate Altera device family handbook and device-specific documentation. For details about which functions a given DSP block can implement, refer to the DSP Solutions Center on the Altera website at www.altera.com.
-  For more information about inferring multiply-accumulator and multiply-adder megafunctions in HDL code, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*, and the *Precision Synthesis Style Guide* in the Precision Synthesis Manuals Bookcase.

Controlling DSP Block Inference

By default, the Precision Synthesis software infers the ALTMULT_ADD or ALTMULT_ACCUM megafunction appropriately in your design. These megafunctions allow the Quartus II software to select either logic or DSP blocks, depending on the device utilization and the size of the function.

You can use the `extract_mac` attribute to prevent inference of an ALTMULT_ADD or ALTMULT_ACCUM megafunction in a certain module or entity. The options for this attribute are described in [Table 15-4](#).

Table 15-4. Options for `extract_mac` Attribute Controlling DSP Implementation

Value	Description
TRUE	The ALTMULT_ADD or ALTMULT_ACCUM megafunction is inferred.
FALSE	The ALTMULT_ADD or ALTMULT_ACCUM megafunction is not inferred.

To control inference, use the `extract_mac` attribute with the appropriate value from [Table 15-4](#) in your HDL code, as shown in [Example 15-18](#) and [Example 15-19](#).

Example 15-18. Setting the `extract_mac` Attribute in Verilog HDL

```
//synthesis attribute <module name> extract_mac <value>
```

Example 15-19. Setting the `extract_mac` Attribute in VHDL

```
ATTRIBUTE extract_mac: BOOLEAN;
ATTRIBUTE extract_mac OF <entity name>: ENTITY IS <value>;
```

To control the implementation of the multiplier portion of a multiply-accumulator or multiply-adder, you must use the `dedicated_mult` attribute.

[Example 15-20](#) and [Example 15-21](#) on [page 15-23](#) use the `extract_mac`, `dedicated_mult`, and `preserve_signal` attributes (in Verilog HDL and VHDL) to implement the given DSP function in logic in the Quartus II software.

Example 15-20. Using `extract_mac`, `dedicated_mult` and `preserve_signal` in Verilog HDL

```

module unsig_altmult_accum1 (dataout, dataa, datab, clk, aclr, clken);
    input [7:0] dataa, datab;
    input clk, aclr, clken;
    output [31:0] dataout;

    reg [31:0] dataout;
    wire [15:0] multa;
    wire [31:0] adder_out;

    assign multa = dataa * datab;

    //synthesis attribute multa preserve_signal TRUE
    //synthesis attribute multa dedicated_mult OFF
    assign adder_out = multa + dataout;

    always @ (posedge clk or posedge aclr)
    begin
        if (aclr)
            dataout <= 0;
        else if (clken)
            dataout <= adder_out;
    end

    //synthesis attribute unsig_altmult_accum1 extract_mac FALSE
endmodule
    
```

Example 15-21. Using `extract_mac`, `dedicated_mult`, and `preserve_signal` in VHDL

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;
USE ieee.std_logic_arith.all;
USE ieee.std_logic_signed.all;
ENTITY signedmult_add IS
    PORT(
        a, b, c, d: IN STD_LOGIC_VECTOR (7 DOWNTO 0);
        result: OUT STD_LOGIC_VECTOR (15 DOWNTO 0));
    ATTRIBUTE preserve_signal: BOOLEAN;
    ATTRIBUTE dedicated_mult: STRING;
    ATTRIBUTE extract_mac: BOOLEAN;
    ATTRIBUTE extract_mac OF signedmult_add: ENTITY IS FALSE;
END signedmult_add;
ARCHITECTURE rtl OF signedmult_add IS
    SIGNAL a_int, b_int, c_int, d_int : signed (7 DOWNTO 0);
    SIGNAL pdt_int, pdt2_int : signed (15 DOWNTO 0);
    SIGNAL result_int: signed (15 DOWNTO 0);
    ATTRIBUTE preserve_signal OF pdt_int: SIGNAL IS TRUE;
    ATTRIBUTE dedicated_mult OF pdt_int: SIGNAL IS "OFF";
    ATTRIBUTE preserve_signal OF pdt2_int: SIGNAL IS TRUE;
    ATTRIBUTE dedicated_mult OF pdt2_int: SIGNAL IS "OFF";
BEGIN
    a_int <= signed (a);
    b_int <= signed (b);
    c_int <= signed (c);
    d_int <= signed (d);
    pdt_int <= a_int * b_int;
    pdt2_int <= c_int * d_int;
    result_int <= pdt_int + pdt2_int;
    result <= STD_LOGIC_VECTOR(result_int);
END rtl;
    
```

RAM and ROM

The Precision Synthesis software detects memory structures in HDL code and converts them to an operator that infers an ALTSYNCRAM or LPM_RAM_DP megafunction, depending on the device family. The software then places these functions in memory blocks.

The software supports inference for these functions only if the target device family has dedicated memory blocks.

- For more information about inferring RAM and ROM megafunctions in HDL code, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*, and the *Precision Synthesis Style Guide* in the Precision Synthesis Manuals Bookcase. To access these manuals, in the Precision Synthesis software, click **Help** and select **Open Manuals Bookcase**.

Incremental Compilation and Block-Based Design

As designs become more complex and designers work in teams, a block-based incremental design flow is often an effective design approach. In an incremental compilation flow, you can make changes to one part of the design while maintaining the placement and performance of unchanged parts of the design. Design iterations can be made dramatically faster by focusing new compilations on particular design partitions and merging results with the results of previous compilations of other partitions. You can perform optimization on individual blocks and then integrate them into a final design and optimize the design at the top-level.

The first step in an incremental design flow is to make sure that different parts of your design do not affect each other. You must ensure that you have separate netlists for each partition in your design. If the whole design is in one netlist file, changes in one partition affect other partitions because of possible node name changes when you resynthesize the design.

You can create different implementations for each partition in your Precision Synthesis project, which allows you to switch between partitions without leaving the current project file. You can also create a separate project for each partition if you require separate projects for a team-based design flow. Alternatively, you can use the incremental synthesis capability in the Precision RTL Plus software.

- For more information about creating partitions and using incremental compilation in the Quartus II software, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.

Creating a Design with Precision RTL Plus Incremental Synthesis

The Precision RTL Plus incremental synthesis flow for Quartus II incremental compilation uses a partition-based approach to achieve faster design cycle time.

Using the incremental synthesis feature, you can create different netlist files for different partitions of a design hierarchy within one partition implementation, which makes each partition independent of the others in an incremental compilation flow. Only the portions of a design that have been updated must be recompiled during design iterations. You can make changes and resynthesize one partition in a design to create a new netlist without affecting the synthesis results or fitting of other partitions.

The following steps show a general flow for partition-based incremental synthesis with Quartus II incremental compilation.

1. Create Verilog HDL or VHDL design files.
2. Determine which hierarchical blocks you want to treat as separate partitions in your design, and designate the partitions with the `incr_partition` attribute. For the syntax to create partitions, refer to “[Creating Partitions with the incr_partition Attribute](#)” on page 15-25.
3. Create a project in the Precision RTL Plus Synthesis software and add the HDL design files to the project.
4. Enable incremental synthesis in the Precision RTL Plus Synthesis software using one of these methods:
 - On the Tools menu, click **Set Options**. On the **Optimization** page, turn on **Enable Incremental Synthesis**.
 - Run the following command in the Transcript Window:

```
setup_design -enable_incr_synth ←
```
5. Run the basic Precision Synthesis flow of compilation, synthesis, and place-and-route on your design. In subsequent runs, the Precision RTL Plus Synthesis software processes only the parts of the design that have changed, resulting in a shorter iteration than the initial run. The performance of the unchanged partitions is preserved.

The Precision RTL Plus Synthesis software sets the netlist types of the unchanged partitions to **Post-Fit** and the changed partitions to **Post-Synthesis**. You can change the netlist type during timing closure in the Quartus II software to obtain the best QoR.

6. Import the EDIF or VQM netlist for each partition and the top-level `.tcl` file into the Quartus II software, and set up the Quartus II project to use incremental compilation.
7. Compile your Quartus II project.
8. If you want, you can change the Quartus II incremental compilation netlist type for a partition with the **Design Partitions Window**. You can change the **Netlist Type** to one of the following options:
 - To preserve the previous post-fit placement results, change the **Netlist Type** of the partition to **Post-Fit**.
 - To preserve the previous routing results, set the **Fitter Preservation Level** of the partition to **Placement and Routing**.

Creating Partitions with the `incr_partition` Attribute

Partitions are set using the HDL `incr_partition` attribute. The Precision Synthesis software creates or deletes partitions by reading this attribute during compilation iterations. The attribute can be attached to either the design unit definition or an instance. [Example 15-22](#) and [Example 15-23](#) show how to use the attribute to create partitions.

To delete partitions, you can remove the attribute or set the attribute value to false.



The Precision Synthesis software ignores partitions set in a black box.

Example 15-22. Using `incr_partition` Attribute to Create a Partition in Verilog HDL

Design unit partition:

```
module my_block(  
    input clk;  
    output reg [31:0] data_out) /* synthesis incr_partition */ ;
```

Instance partition:

```
my_block my_block_inst(.clk(clk), .data_out(data_out));  
// synthesis attribute my_block_inst incr_partition true
```

Example 15-23. Using `incr_partition` Attribute to a Create Partition in VHDL

Design unit partition:

```
entity my_block is  
    port(  
        clk : in std_logic;  
        data_out : out std_logic_vector(31 downto 0)  
    );  
    attribute incr_partition : boolean;  
    attribute incr_partition of my_block : entity is true;  
end entity my_block;
```

Instance partition:

```
component my_block is  
    port(  
        clk : in std_logic;  
        data_out : out std_logic_vector(31 downto 0)  
    );  
end component;  
  
attribute incr_partition : boolean;  
attribute incr_partition of my_block_inst : label is true;  
  
my_block_inst my_block  
    port map(clk, data_out);
```

Creating Multiple Mapped Netlist Files With Separate Precision Projects or Implementations

This section describes how to manually generate multiple netlist files, which can be VQM or EDIF files, for incremental compilation using black boxes and separate Precision projects or implementations for each design partition. This manual flow is supported in versions of the Precision software that do not include the incremental synthesis feature. You might also use this feature if you perform synthesis in a team-based environment without a top-level synthesis project that includes all of the lower-level design blocks.

In the Precision Synthesis software, create a separate implementation, or a separate project, for each lower-level module and for the top-level design that you want to maintain as a separate netlist file. Implement black box instantiations of lower-level modules in your top-level implementation or project.

 For more information about managing implementations and projects, refer to the *Precision RTL Synthesis User's Manual*. To access this manual, in the Precision Synthesis software, click **Help** and select **Open Manuals Bookcase**.

When synthesizing the implementations for lower-level modules, perform these steps in the Precision Synthesis software:

1. On the Tools menu, turn off **Add IO Pads** on the **Optimization** page under **Set Options**.

 You must turn off the **Add IO Pads** option while synthesizing the lower-level modules individually. Enable the **Add IO Pads** option only while synthesizing the top-level module.

2. Read the HDL files for the modules.

 Modules can include black box instantiations of lower-level modules that are also maintained as separate netlist files.

3. Add constraints for all partitions in the design.

When synthesizing the top-level design implementation, perform these steps:

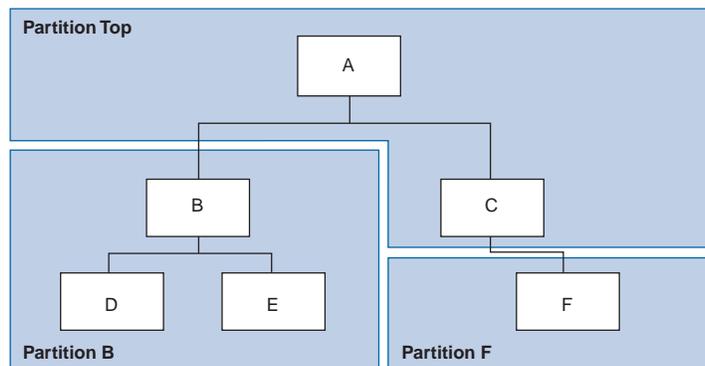
1. Read the HDL files for top-level designs.
2. On the Tools menu, click **Set Options**. On the **Optimization** page, turn on **Add IO Pads**.
3. Create black boxes for lower-level modules in the top-level design.
4. Add constraints.

 In a standard Quartus II incremental compilation flow, Precision Synthesis software constraints made on lower-level modules are not passed to the Quartus II software. Ensure that appropriate constraints are made in the top-level Precision Synthesis project, or in the Quartus II project.

Creating Black Boxes to Create EDIF Netlists

This section describes how to create black boxes to create separate EDIF netlists. Figure 15-2 shows an example of a design hierarchy separated into various partitions.

Figure 15-2. Partitions in a Hierarchical Design



In Figure 15-2, the top-level partition contains the top-level block in the design (block A) and the logic that is not defined as part of another partition. In this example, the partition for top-level block A also includes the logic in the sub-block C. Because block F is contained in its own partition, it is not treated as part of the top-level partition A. Another separate partition, B, contains the logic in blocks B, D, and E. In a team-based design, different engineers may work on the logic in different partitions. One netlist is created for the top-level module A and its submodule C, another netlist is created for module B and its submodules D and E, while a third netlist is created for module F. To create multiple EDIF netlist files for this design, follow these steps:

1. Generate an `.edf` file for module B. Use `B.v/.vhd`, `D.v/.vhd`, and `E.v/.vhd` as the source files.
2. Generate an `.edf` file for module F. Use `F.v/.vhd` as the source file.
3. Generate a top-level `.edf` file for module A. Use `A.v/.vhd` and `C.v/.vhd` as the source files. Ensure that you create black boxes for modules B and F, which were optimized separately in the previous steps.

The goal is to individually synthesize and generate an `.edf` netlist file for each lower-level module and then instantiate these modules as black boxes in the top-level file. You can then synthesize the top-level file to generate the `.edf` netlist file for the top-level design. Finally, both the lower-level and top-level `.edf` netlist files are provided to your Quartus II project.



When you make design or synthesis optimization changes to part of your design, resynthesize only the changed partition to generate the new `.edf` netlist file. Do not resynthesize the implementations or projects for the unchanged partitions.

Creating Black Boxes in Verilog HDL

Any design block that is not defined in the project or included in the list of files to be read for a project is treated as a black box by the software. In Verilog HDL, you must provide an empty module declaration for any module that is treated as a black box.

A black box for the top-level file **A.v** is shown in the following example. Provide an empty module declaration for any lower-level files, which also contain a black box for any module beneath the current level of hierarchy.

Example 15–24. Verilog HDL Black Box for Top-Level File A.v

```
module A (data_in, clk, e, ld, data_out);
    input data_in, clk, e, ld;
    output [15:0] data_out;
    wire [15:0] cnt_out;
    B U1 (.data_in (data_in), .clk(clk), .ld (ld), .data_out(cnt_out));
    F U2 (.d(cnt_out), .clk(clk), .e(e), .q(data_out));
    // Any other code in A.v goes here.
endmodule
// Empty Module Declarations of Sub-Blocks B and F follow here.
// These module declarations (including ports) are required for black
// boxes.
module B (data_in, clk, ld, data_out);
    input data_in, clk, ld;
    output [15:0] data_out;
endmodule
module F (d, clk, e, q);
    input [15:0] d;
    input clk, e;
    output [15:0] q;
endmodule
```

Creating Black Boxes in VHDL

Any design block that is not defined in the project or included in the list of files to be read for a project is treated as a black box by the software. In VHDL, you must provide a component declaration for the black box.

A black box for the top-level file **A.vhd** is shown in [Example 15–25](#). Provide a component declaration for any lower-level files that also contain a black box or for any block beneath the current level of hierarchy.

Example 15-25. VHDL Black Box for Top-Level File A.vhd

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;
ENTITY A IS
    PORT ( data_in : IN INTEGER RANGE 0 TO 15;
          clk, e, ld : IN STD_LOGIC;
          data_out : OUT INTEGER RANGE 0 TO 15);
END A;
ARCHITECTURE a_arch OF A IS
    COMPONENT B PORT(
        data_in : IN INTEGER RANGE 0 TO 15;
        clk, ld : IN STD_LOGIC;
        d_out : OUT INTEGER RANGE 0 TO 15);
    END COMPONENT;
    COMPONENT F PORT(
        d : IN INTEGER RANGE 0 TO 15;
        clk, e: IN STD_LOGIC;
        q : OUT INTEGER RANGE 0 TO 15);
    END COMPONENT;
    -- Other component declarations in A.vhd go here
    signal cnt_out : INTEGER RANGE 0 TO 15;
BEGIN
    U1 : B
        PORT MAP (
            data_in => data_in,
            clk => clk,
            ld => ld,
            d_out => cnt_out);
    U2 : F
        PORT MAP (
            d => cnt_out,
            clk => clk,
            e => e,
            q => data_out);
    -- Any other code in A.vhd goes here
END a_arch;

```

After you complete the steps outlined in this section, you have different EDIF netlist files for each partition of the design. These files are ready for use with incremental compilation in the Quartus II software.

Creating Quartus II Projects for Multiple EDIF Files

The Precision Synthesis software creates a `.tcl` file for each implementation, and provides the Quartus II software with the appropriate constraints and information to set up a project. When using incremental synthesis, the Precision RTL Plus Synthesis software creates only a single `.tcl` file, `<project name>_incr_partitions.tcl`, to pass the partition information to the Quartus II software. For details about using this Tcl script to set up your Quartus II project and to pass your top-level constraints, refer to [“Running the Quartus II Software Manually Using the Precision Synthesis-Generated Tcl Script”](#) on page 15-11.

Depending on your design methodology, you can create one Quartus II project for all EDIF netlists, or a separate Quartus II project for each EDIF netlist. In the standard incremental compilation design flow, you create design partition assignments for each partition in the design within a single Quartus II project. This methodology provides the best QoR and performance preservation during incremental changes to your design. You might require a bottom-up design flow if each partition must be optimized separately, such as for third-party IP delivery.

To follow this design flow in the Quartus II software, create separate Quartus II projects and export each design partition and incorporate it into a top-level design using the incremental compilation features to maintain placement results.

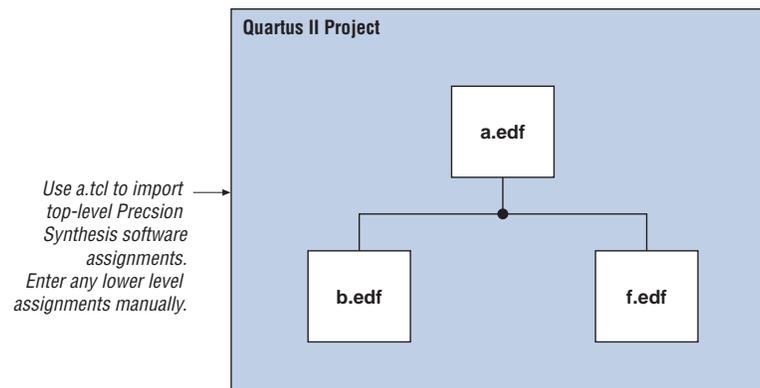
The following sections describe how to create the Quartus II projects for these two design flows.

Creating a Single Quartus II Project for a Standard Incremental Compilation Flow

Use the `<top-level project>.tcl` file generated for the top-level partition to create your Quartus II project and import all the netlists into this one Quartus II project for an incremental compilation flow. You can optimize all partitions within the single Quartus II project and take advantage of the performance preservation and compilation time reduction that incremental compilation provides. Figure 15-3 shows the design flow for the example design in Figure 15-2 on page 15-28.

All the constraints from the top-level implementation are passed to the Quartus II software in the top-level `.tcl` file, but any constraints made only in the lower-level implementations within the Precision Synthesis software are not forward-annotated. Enter these constraints manually in your Quartus II project.

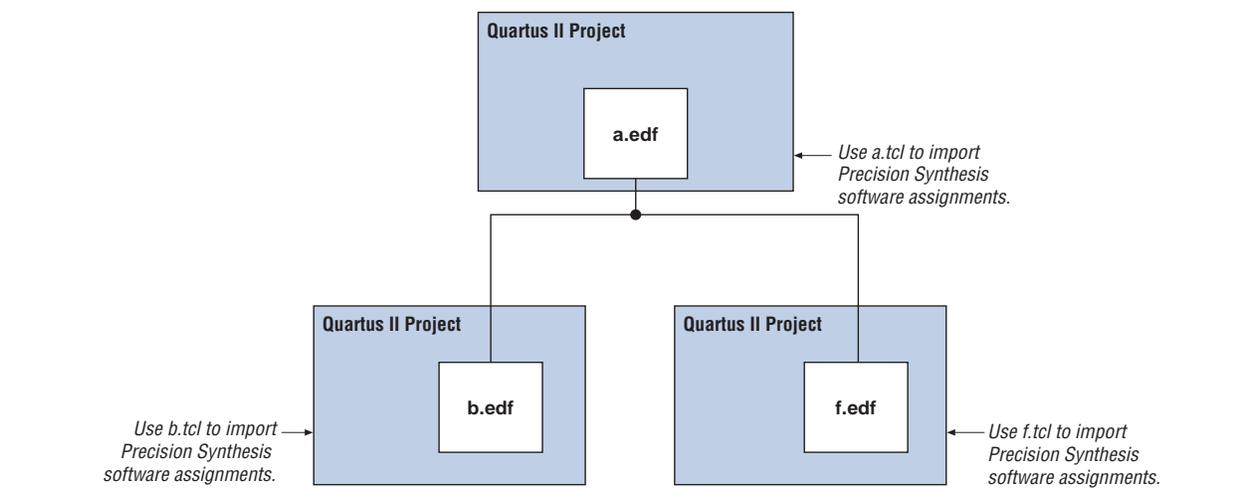
Figure 15-3. Design Flow Using Multiple EDIF Files with One Quartus II Project



Creating Multiple Quartus II Projects for a Bottom-Up Flow

Use the .tcl files generated by the Precision Synthesis software for each Precision Synthesis software implementation or project to generate multiple Quartus II projects, one for each partition in the design. Each designer in the project can optimize their block separately in the Quartus II software and export the placement of their blocks using incremental compilation. Designers should create a LogicLock region to provide a floorplan location assignment for each block; the top-level designer should then import all the blocks and assignments into the top-level project. Figure 15-4 shows the design flow for the example design in Figure 15-2 on page 15-28.

Figure 15-4. Design Flow: Using Multiple EDIF Files with Multiple Quartus II Projects



Hierarchy and Design Considerations

To ensure the proper functioning of the synthesis flow, you can create separate partitions only for modules, entities, or existing netlist files. In addition, each module or entity must have its own design file. If two different modules are in the same design file, but are defined as being part of different partitions, incremental synthesis cannot be maintained because both regions must be recompiled when you change one of the modules.

Altera recommends that you register all inputs and outputs of each partition. This makes logic synchronous and avoids any delay penalty on signals that cross partition boundaries.

If you use boundary tri-states in a lower-level block, the Precision Synthesis software pushes the tri-states through the hierarchy to the top-level to make use of the tri-state drivers on output pins of Altera devices. Because pushing tri-states requires optimizing through hierarchies, lower-level tri-states are not supported with a block-based compilation methodology. You should use tri-state drivers only at the external output pins of the device and in the top-level block in the hierarchy.



For more tips on design partitioning, refer to the *Best Practices for Incremental Compilation Partitions and Floorplan Assignments* chapter in volume 1 of the *Quartus II Handbook*.

Conclusion

The Mentor Graphics Precision Synthesis software and Quartus II design flow allow you to control how to prepare your design files for the Quartus II place-and-route process, which allows you to improve performance and optimizes your design for use with Altera devices. Several of the methodologies outlined in this chapter can help you optimize your design to achieve performance goals and decrease design time.

Document Revision History

Table 15-5 shows the revision history for this document.

Table 15-5. Document Revision History

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Changed to new document template. ■ Removed Classic Timing Analyzer support. ■ Added support for .vqm netlist files. ■ Edited the “Creating Quartus II Projects for Multiple EDIF Files” on page 15–30 section for changes with the incremental compilation flow. ■ Editorial changes.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Minor updates for the Quartus II software version 10.0 release
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Minor updates for the Quartus II software version 9.1 release
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Updated list of supported devices for the Quartus II software version 9.0 release ■ Chapter 11 was previously Chapter 10 in software version 8.1
November 2008	8.1.0	<ul style="list-style-type: none"> ■ Changed to 8-1/2 x 11 page size ■ Title changed to <i>Mentor Graphics Precision Synthesis Support</i> ■ Updated list of supported devices ■ Added information about the Precision RTL Plus incremental synthesis flow ■ Updated Figure 10-1 to include SystemVerilog ■ Updated “Guidelines for Altera Megafunctions and Architecture-Specific Features” on page 10–19 ■ Updated “Incremental Compilation and Block-Based Design” on page 10–28 ■ Added section “Creating Partitions with the incr_partition Attribute” on page 10–29
May 2008	8.0.0	<ul style="list-style-type: none"> ■ Removed Mercury from the list of supported devices ■ Changed Precision version to 2007a update 3 ■ Added note for Stratix IV support ■ Renamed “Creating a Project and Compiling the Design” section to “Creating and Compiling a Project in the Precision RTL Synthesis Software” ■ Added information about constraints in the Tcl file ■ Updated document based on the Quartus II software version 8.0

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this chapter.

This chapter documents key design methodologies and techniques for Altera® devices using the LeonardoSpectrum and Quartus II design flow.

This chapter includes the following sections:

- “Altera Device Family Support”
- “Design Flow” on page 16–2
- “LeonardoSpectrum Optimization Strategies” on page 16–4
- “Timing Analysis with the LeonardoSpectrum Software” on page 16–6
- “Exporting Designs Using NativeLink Integration” on page 16–7
- “Guidelines for Altera Megafunctions and LPM Functions” on page 16–8
- “Block-Based Design with the Quartus II Software” on page 16–16

 Altera recommends using the advanced Mentor Graphics Precision RTL Synthesis software for new designs in new device families. For more information about Precision RTL Synthesis, refer to the *Mentor Graphics Precision Synthesis Support* chapter in volume 1 of the *Quartus II Handbook*.

 This chapter assumes that you have set up, licensed, and are familiar with the LeonardoSpectrum software.

 To obtain and license the LeonardoSpectrum software, refer to the Mentor Graphics website at www.mentor.com. For information about installing the LeonardoSpectrum software and setting up your working environment, refer to the *LeonardoSpectrum Installation Guide* and the *LeonardoSpectrum User's Manual*.

Altera Device Family Support

The LeonardoSpectrum software is a mature synthesis tool supporting legacy devices and many current devices; however, newly-released devices may not be supported.

 Contact Mentor Graphics for more information about support for newly-released devices.

The LeonardoSpectrum software also supports the FLEX 8000 and MAX 9000 legacy devices that are supported only in the Altera MAX+PLUS® II software, as well as ACEX® 1K, APEX™ II, APEX 20K, APEX 20KC, APEX 20KE, FLEX® 10K, and FLEX 6000 legacy devices that are supported by the Quartus II software version 9.0 and earlier.

Design Flow

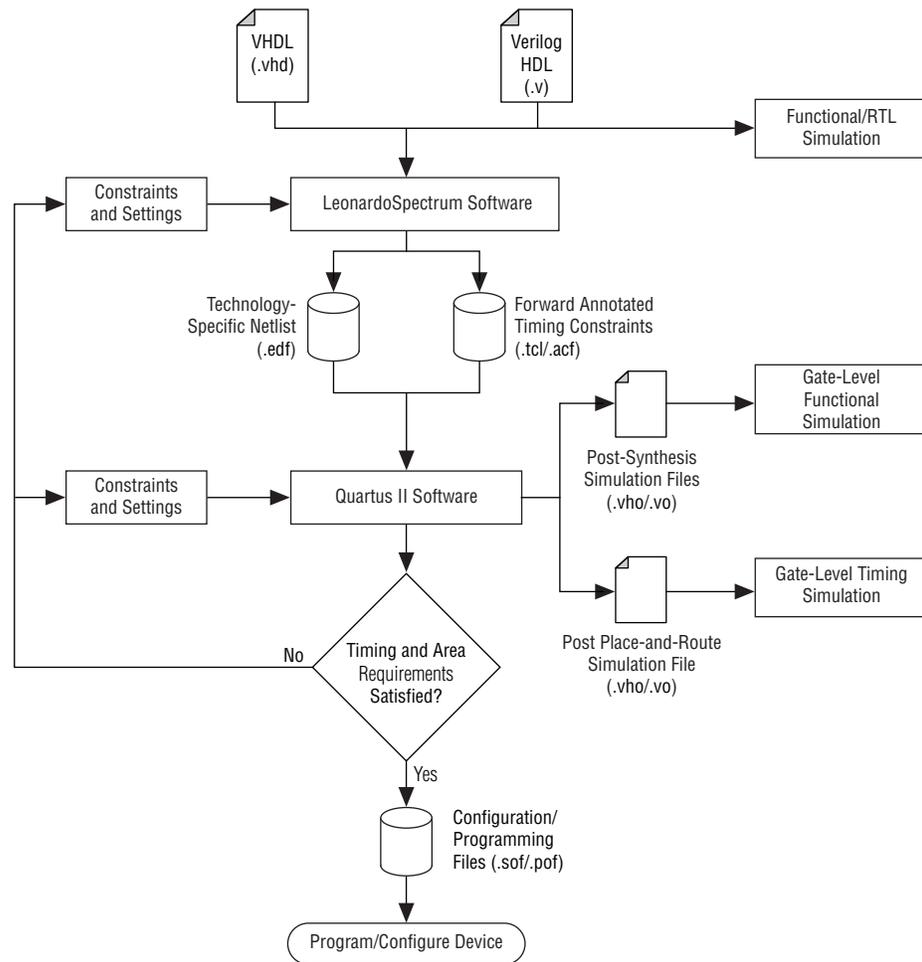
The following steps describe a basic Quartus II software design flow using the LeonardoSpectrum software:

1. Create Verilog HDL or VHDL design files.
2. Import the Verilog HDL or VHDL design files into the LeonardoSpectrum software for synthesis.
3. Select a target device and add timing constraints and compiler directives to help optimize the design during synthesis.
4. Synthesize the project in the LeonardoSpectrum software.
5. Create a Quartus II project and import the technology-specific EDIF Input File (.edf) netlist and the Tcl Script File (.tcl) generated by the LeonardoSpectrum software into the Quartus II software for placement and routing and performance evaluation.
6. After obtaining place-and-route results that meet your requirements, configure or program the Altera device.

Figure 16-1 on page 16-3 shows the recommended design flow using the LeonardoSpectrum and Quartus II software.

If your area and timing requirements are met, use the programming files generated by the Quartus II software to program or configure the Altera device. If the area or timing requirements are not met, change the constraints in the LeonardoSpectrum software and rerun synthesis. Repeat the process until the area and timing requirements are met. You can also use other Quartus II software options and techniques to meet the area and timing requirements.

Figure 16–1. Recommended Design Flow Using LeonardoSpectrum and Quartus II Software



The LeonardoSpectrum software supports both VHDL and Verilog HDL source files. With the appropriate license, the software also supports mixed synthesis, allowing a combination of VHDL and Verilog HDL source files. During synthesis, the LeonardoSpectrum software produces several intermediate and output files, which are listed and described in [Table 16–1](#).

Table 16–1. LeonardoSpectrum Intermediate and Output Files

File Extension(s)	File Description
.xdb	Technology-independent register transfer level (RTL) netlist file that can only be read by the LeonardoSpectrum software.
.edf	Technology-specific output netlist file in electronic design interchange format (EDIF).
.tcl	Forward-annotated constraints file containing constraints and assignments. A .tcl file for the Quartus II software is created for all devices. The .tcl file contains the appropriate Tcl commands to create and set up a Quartus II project and pass placement constraints.
.acf	Assignment and Configurations file for backward compatibility with the MAX+PLUS II software. For devices supported by the MAX+PLUS II software, the MAX+PLUS II assignments are imported from the MAX+PLUS II .acf file.

Altera recommends that you do not use project directory names that include spaces. Some file operations in the LeonardoSpectrum software do not work correctly if the path name contains spaces.

Specify timing constraints and compiler directives for the design in the LeonardoSpectrum software, or in a constraint file (.ctr). These constraints are forward-annotated in the .tcl file for use with the Classic Timing Analyzer in the Quartus II software version 10.0 and earlier.



The LeonardoSpectrum software does not generate a Synopsys Design Constraint (SDC) format file for the TimeQuest Timing Analyzer. If you are using a current version of the Quartus II software, you must convert your timing constraints to SDC format. Altera recommends using the advanced Precision RTL Synthesis software for new designs in new device families, instead of using the LeonardoSpectrum software.

The LeonardoInsight™ Schematic Viewer is an add-on graphical tool for schematic views of the technology-independent RTL netlist (.xdb) and the technology-specific gate-level results. You can use the Schematic Viewer to visually analyze and debug the design. It also supports cross-probing between the RTL and gate-level schematics, the design browser, and the source code in the HDLInventor™ text editor.

LeonardoSpectrum Optimization Strategies

You can configure most general settings in the **Quick Setup** tab in the LeonardoSpectrum user interface. FlowTabs provide additional options, some including multiple PowerTabs at the bottom of the screen, with more options. Advanced optimization options in the LeonardoSpectrum software include timing-driven synthesis, encoding style, resource sharing, and mapping I/O registers.

Timing-Driven Synthesis

The LeonardoSpectrum software supports timing-driven synthesis through user-assigned timing constraints that optimize the performance of the design. Constraints, such as clock frequency, can be specified globally or for individual clock signals. The following sections describe how to set the various types of timing constraints in the LeonardoSpectrum software.

The timing constraints described in “**Global PowerTab**” are set in the **Constraints** FlowTab. In this FlowTab, there are PowerTabs at the bottom, such as **Global** and **Clock**, for setting various constraints.

Global PowerTab

The **Global** PowerTab is the default PowerTab in the **Constraints** FlowTab where you can specify the global clock frequency. The **Clock Frequency** setting on the **Quick Setup** tab is equivalent to the **Registers to Registers** delay setting. You can also specify the **Input Ports to Registers**, **Registers to Output Ports**, and **Inputs to Outputs** delays that correspond to global t_{SU} , t_{CO} , and t_{PD} requirements, respectively, in the Quartus II software. The timing diagram on the **Global** PowerTab reflects the settings you have made.

Clock PowerTab

You can set various constraints for each clock in your design. First, select the clock name in the Clock(s) window. The clock names appear after the design is read from the **Input** FlowTab. Configure settings for that particular clock and click **Apply**. You can also specify the **Duty Cycle**, which is set to 5-% by default. The timing diagram shows these settings.

If a clock has an **Offset** from the main clock, which is considered to be time “0”, this constraint corresponds to the `OFFSET_FROM_BASE_CLOCK` setting in the Quartus II software.

You can specify the pin number for the clock input pin in the **Pin Location** box. This pin number is passed to the Quartus II software for place-and-route, but does not affect synthesis in the LeonardoSpectrum software.

Input and Output PowerTabs

Configure settings for individual input or output pins in the **Input** and **Output** tabs. First, select a name in the Input Ports or Output Ports window. The names appear after the design is read from the **Input** FlowTab. Then, make the following settings for that pin as described below.

The **Arrival Time** indicates that the input signal arrives a specified time after the rising clock edge (time “0”). This setting constrains the path from the pin to the first register by including the arrival time in the total delay, and corresponds to the `EXTERNAL_INPUT_DELAY` assignment in the Quartus II software.

The **Required Time** indicates the maximum delay after time “0” that the output signal should arrive at the output pin. This setting directly constrains the register to output delay, and corresponds to the `EXTERNAL_OUTPUT_DELAY` assignment in the Quartus II software.

Specify the pin number for the I/O pin in the **Pin Location** box. This pin number is passed to the Quartus II software for place-and-route, but does not affect synthesis in the LeonardoSpectrum software.

Other Constraints

The following sections describe other constraints that can be set with the LeonardoSpectrum user interface:

- [“Encoding Style”](#)
- [“Resource Sharing”](#) on page 16-6
- [“Mapping I/O Registers”](#) on page 16-6

Encoding Style

The LeonardoSpectrum software encodes state machines during the synthesis process. To improve performance when coding state machines, separate state machine logic from all arithmetic functions and data paths. When encoded, a design cannot be re-encoded later in the optimization process. You must follow a particular VHDL or Verilog HDL coding style for the LeonardoSpectrum software to identify the state machine.

Table 16-2 describes the state machine encoding styles supported by the LeonardoSpectrum software.

Table 16-2. State Machine Encoding Styles in the LeonardoSpectrum Software

Style	Description
Binary	Generates state machines with the fewest possible flipflops. Binary state machines are useful for area-critical designs when timing is not the primary concern.
Gray	Generates state machines where only one flipflop changes during each transition. Gray-encoded state machines tend to be free of glitches.
One-hot	Generates state machines containing one flipflop for each state. One-hot state machines provide the best performance and shortest clock-to-output delays. However, one-hot implementations are usually larger than binary implementations.
Random	Generates state machines using random state machine encoding. Use random state machine encoding only when no other implementation achieves the desired results.
Auto (Default)	Implements binary or one-hot encoding, depending on the size of enumerated types in the state machine.

The **Encoding Style** is created in the **Input FlowTab**. The setting instructs the software to use a particular state machine encoding style for all state machines. The default **Auto** encoding style implements binary or one-hot encoding, depending on the size of enumerated types in the state machine.

 To ensure proper recognition and improve performance when coding state machines, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook* for design guidelines.

Resource Sharing

You should turn on **Resource Sharing** in the **Input FlowTab** to allow optimization, which reduces device resources.

Mapping I/O Registers

The **Map I/O Registers** option in the **Technology FlowTab**, is available for Altera FPGAs containing I/O cells (IOC) or I/O elements (IOE). If this option is turned on, input or output registers are moved into the device's I/O cells for faster setup or clock-to-output times.

Timing Analysis with the LeonardoSpectrum Software

The LeonardoSpectrum software reports successful synthesis with an information message in the Transcript or Information window. Estimated device usage and timing results are reported in the Device Utilization section of this window. [Figure 16-2 on page 16-7](#) shows an example of a LeonardoSpectrum compilation report.

Figure 16–2. LeonardoSpectrum Compilation Report

```

*****
Device Utilization for EP20K200EQC208
*****
Resource          Used    Avail    Utilization
-----
IOs                22     136     16.18%
LCs               114    8320     1.37%
Memory Bits       0     106496    0.00%
-----

                          Clock Frequency Report

Clock              : Frequency
-----
clk                : 52.2 MHz
clk2               : 149.5 MHz

                          Critical Path Report
    
```

The LeonardoSpectrum software estimates the timing results based on timing models. The software does not have the design’s place-and-route information in the Quartus II software, so it cannot report accurate routing delays. Additionally, if the design includes any black boxed Altera-specific functions, the LeonardoSpectrum software does not report timing information for these functions.

Exporting Designs Using NativeLink Integration

You can use NativeLink® integration to integrate the LeonardoSpectrum and the Quartus II software with a single GUI for both the synthesis and place-and-route operations. You can run the Quartus II software from within the LeonardoSpectrum software GUI with NativeLink integration or you can run the LeonardoSpectrum software from within the Quartus II software GUI for device families supported in the Quartus II software.

Generating Netlist Files

The LeonardoSpectrum software generates an **.edf** netlist file readable as an input file in the Quartus II software for place-and-route. Select the **.edf** file option name in the **Output FlowTab**. The **.edf** file is also generated if the **Auto** option is turned on in the **Output FlowTab**.

Including Design Files for Black Boxed Modules

Be sure to include the MegaWizard™ Plug-In Manager-generated custom megafunction variation design file in the Quartus II project directory, or add it to the list of project files for place-and-route, if the design has black boxed megafunctions.

Passing Constraints with Scripts

The LeonardoSpectrum software can create a Tcl script (**.tcl**) file called *<project name>.tcl*. This file contains commands to create a Quartus II project along with constraints and other assignments. To output a **.tcl** file, turn on **Write Vendor Constraint Files** in the **Output** FlowTab.

To create and compile a Quartus II project using the **.tcl** file generated from the LeonardoSpectrum software, perform the following steps in the Quartus II software:

1. Place the **.edf** and **.tcl** files in the same directory.
2. On the View menu, point to **Utility Windows**, and click **Tcl Console** to open the Quartus II Tcl Console.
3. At the Tcl prompt, type `source <path>/<project name>.tcl`.
4. On the File menu, click **Open Project** to open the new project.
5. On the Processing menu, click **Start Compilation**.



The LeonardoSpectrum software does not generate a Synopsys Design Constraint (SDC) format file for the TimeQuest Timing Analyzer. If you are using a current version of the Quartus II software, you must convert your timing constraints to SDC format. Altera recommends using the advanced Precision RTL Synthesis software for new designs in new device families, instead of using the LeonardoSpectrum software.

Integration with the Quartus II Software

You can launch the Quartus II software from within the LeonardoSpectrum software with the options in the **Place And Route** section of the **Quick Setup** tab. Turn on the **Run Integrated Place and Route** option to start the compilation in the Quartus II software to show the fitting and performance results. You can also run place-and-route in the Quartus II software by turning on **Run Quartus** on the **Physical** FlowTab, and then clicking **Run PR**.

To use integrated place-and-route, on the Options menu, point to **Place and Route Path** and click **Tools**. Specify the location of the Quartus II software executable file (*<Quartus II software installation directory>/bin*).

Guidelines for Altera Megafunctions and LPM Functions

Altera provides parameterizable megafunctions ranging from simple arithmetic units, such as adders and counters, to advanced phase-locked loop (PLL) blocks, multipliers, and memory structures. These functions are performance-optimized for Altera devices. Megafunctions include the library of parameterized modules (LPM), device-specific megafunctions such as PLLs, LVDS, and digital signal processing (DSP) blocks, intellectual property (IP) available as Altera MegaCore[®] functions, and IP available through the Altera Megafunction Partners Program (AMPPsm).



Some IP cores require synthesis in the LeonardoSpectrum software. Refer to the user guide for the specific IP.

You can infer or instantiate megafunctions in the LeonardoSpectrum software. The LeonardoSpectrum software supports inferring some Altera megafunctions, such as multipliers, DSP functions, and RAM and ROM blocks. The LeonardoSpectrum software supports all Altera megafunctions through instantiation.

Instantiating Altera Megafunctions

There are two methods of instantiating Altera megafunctions in the LeonardoSpectrum software. The more common method, which maintains target technology awareness, is to set up and parameterize a megafunction variation in the MegaWizard Plug-In Manager in the Quartus II software. The megafunction wizard creates a wrapper file that instantiates the megafunction. The less common method is to directly instantiate the megafunction in the Verilog HDL or VHDL code. The advantage of using the megafunction wizard method over the instantiation method is that the megafunction wizard properly sets all the parameters. Also you do not need the library support, which is required in the instantiation method. This is referred to as black box methodology.



Altera recommends using the MegaWizard Plug-In Manager to ensure that the ports and parameters are set correctly.



When directly instantiating megafunctions, refer to the respective megafunction user guide on the [Altera IP and Megafunction](#) page.

Inferring Altera Memory Elements

The LeonardoSpectrum software can infer memory blocks from Verilog HDL or VHDL code. When the LeonardoSpectrum software detects a RAM or ROM from the style of the RTL code at a technology-independent level, the software maps the element to a generic module in the RTL database. During the technology-mapping phase of synthesis, the LeonardoSpectrum software maps the generic module to the most optimal primitive memory cells, or Altera megafunction, for the target Altera technology.



For more information about inferring RAM and ROM megafunctions, including examples of VHDL and Verilog HDL code, see the [Recommended HDL Coding Styles](#) chapter in volume 1 of the *Quartus II Handbook*.

Inferring RAM

The LeonardoSpectrum software supports RAM inference for various device families. The following are the restrictions for the LeonardoSpectrum software to successfully infer RAM in a design:

- The write process must be synchronous.
- The read process can be asynchronous or synchronous, depending on the target Altera architecture.
- Resets on the memory are not supported.

The Stratix series and Cyclone series devices support the RAM primitive `altsyncram` with a minimum RAM size of two bits, and a minimum RAM address width of one bit.

To disable RAM inference, set the `extract_ram` and `infer_ram` variables to `false`.

To enter the value `false` when performing synthesis in the user interface with the **Advanced FlowTabs**, on the Tools menu, click **Variable Editor**, or add the `set extract_ram false` and `set infer_ram false` commands to your synthesis script.

Inferring ROM

You can implement ROM behavior in HDL source code with CASE statements or specify the ROM as a table. The LeonardoSpectrum software infers both synchronous and asynchronous ROM, depending on the target Altera device. For example, memory for Stratix series devices must be synchronous to be inferred.

To disable ROM inference, set the `extract_rom` variable to `false`. To enter the value `false` when performing synthesis in the user interface with the **Advanced FlowTabs**, on the Tools menu, click **Variable Editor**, or add the `set extract_rom false` commands to your synthesis script.

Inferring Multipliers and DSP Functions

Some Altera devices include dedicated DSP blocks optimized for DSP applications. The following Altera megafunctions are used with DSP block modes:

- `LPM_MULT`
- `ALTMULT_ACCUM`
- `ALTMULT_ADD`

You can instantiate these megafunctions in the design or direct the LeonardoSpectrum software to infer the appropriate megafunction by recognizing a multiplier, multiplier-accumulator (MAC), or multiplier-adder in the design. The Quartus II software maps the functions to the DSP blocks in the device during place-and-route.



For more information about inferring multipliers and DSP functions, including examples of VHDL and Verilog HDL code, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*.

Simple Multipliers

The `LPM_MULT` megafunction implements the DSP block in the simple multiplier mode. The following functionality is supported in this mode:

- The DSP block includes registers for the input and output stages, and an intermediate pipeline stage.
- Signed and unsigned arithmetic is supported.

Multiplier Accumulators

The `ALTMULT_ACCUM` megafunction implements the DSP block in the multiplier-accumulator mode. The following functionality is supported in this mode:

- The DSP block includes registers for the input and output stages, and an intermediate pipeline stage.
- The output registers are required for the accumulator.
- The input and pipeline registers are optional.
- Signed and unsigned arithmetic is supported.

 If the design requires input registers to be used as shift registers, use the black box method to instantiate the ALTMULT_ACCUM megafunction.

Multiplier Adders

The LeonardoSpectrum software can infer multiplier adders and map them to either the two-multiplier adder mode or the four-multiplier adder mode of the DSP blocks. The LeonardoSpectrum software maps the HDL code to the correct ALTMULT_ADD function.

The following functionality is supported in these modes:

- The DSP block includes registers for the input and output stages, and an intermediate pipeline stage.
- Signed and unsigned arithmetic is supported, but support for the Verilog HDL signed construct is limited.

Controlling DSP Block Inference

Device features, such as dedicated DSP blocks, multipliers, multiply-accumulators, and multiply-adders can be implemented in DSP blocks or in logic. You can control this implementation through attribute settings in the LeonardoSpectrum software.

As described in Table 16-3, attribute settings in the LeonardoSpectrum software control the implementation of the multipliers in DSP blocks or logic at the signal block (or module), and project level.

Table 16-3. Attribute Settings for DSP Blocks in the LeonardoSpectrum Software (Note 1)

Level	Attribute Name	Value	Description
Global	extract_mac (2)	TRUE	All multipliers in the project are mapped to DSP blocks.
		FALSE	All multipliers in the project are mapped to logic.
Module	extract_mac (3)	TRUE	Multipliers inside the specified module are mapped to DSP blocks.
		FALSE	Multipliers inside the specified module are mapped to logic.
Signal	dedicated_mult	ON	LPM is inferred and multipliers are implemented in DSP block.
		OFF	LPM is inferred, but multipliers are implemented in logic by the Quartus II software.
		LCELL	LPM is not inferred, and multipliers are implemented in logic by the LeonardoSpectrum software.
		AUTO	LPM is inferred, but the Quartus II software automatically maps the multipliers to either logic or DSP blocks based on the Quartus II software place-and-route.

Notes to Table 16-3:

- (1) The extract_mac attribute takes precedence over the dedicated_mult attribute.
- (2) For devices with DSP blocks, the extract_mac attribute is set to “true” by default for the entire project.
- (3) For devices with DSP blocks, the extract_mac attribute is set to “true” by default for all modules.

Global Attribute

You can set the extract_mac global attribute to control the implementation of multipliers in DSP blocks for the entire project. You can set this attribute with the following script command:

```
set extract_mac <value>
```

Module Level Attributes

You can control the implementation of multipliers inside a module or component by setting the `extract_mac` attribute in the Verilog HDL source code. Setting this attribute for a module affects only the multipliers inside that module. Use the following command:

```
//synthesis attribute <module name> extract_mac <value>
```

The Verilog HDL and VHDL code samples in [Example 16-1](#) and [Example 16-2](#) show how to use the `extract_mac` attribute.

Example 16-1. Using Module Level Attributes in Verilog HDL Code

```
module mult_add ( dataa, datab, datac, datad, result);  
//synthesis attribute mult_add extract_mac FALSE  
// Port Declaration  
input [15:0] dataa;  
input [15:0] datab;  
input [15:0] datac;  
input [15:0] datad;  
  
output [32:0] result;  
  
// Wire Declaration  
wire [31:0] mult0_result;  
wire [31:0] mult1_result;  
  
// Implementation  
// Each of these can go into one of the 4 mults in a  
// DSP block  
assign mult0_result = dataa * `signed datab;  
//synthesis attribute mult0_result preserve_signal TRUE  
  
assign mult1_result = datac * datad;  
  
// This adder can go into the one-level adder in a DSP  
// block  
assign result = (mult0_result + mult1_result);  
  
endmodule
```

Example 16-2. Using Module Level Attributes in VHDL Code

```

library ieee ;
USE ieee.std_logic_1164.all;

USE ieee.std_logic_arith.all;

entity mult_acc is
    generic (size : integer := 4) ;
    port (
        a: in std_logic_vector (size-1 downto 0) ;
        b: in std_logic_vector (size-1 downto 0) ;
        clk : in std_logic;
        accum_out: inout std_logic_vector (2*size downto 0)
    ) ;
    attribute extract_mac : boolean;
    attribute extract_mac of mult_acc : entity is FALSE;
end mult_acc;

architecture synthesis of mult_acc is
    signal a_int, b_int : signed (size-1 downto 0);
    signal pdt_int : signed (2*size-1 downto 0);
    signal adder_out : signed (2*size downto 0);

begin
    a_int <= signed (a);
    b_int <= signed (b);
    pdt_int <= a_int * b_int;
    adder_out <= pdt_int + signed(accum_out);
    process (clk)
    begin
        if (clk'event and clk = '1') then
            accum_out <= std_logic_vector (adder_out);
        end if;
    end process;
end synthesis ;
    
```

Signal Level Attributes

You can control the implementation of individual LPM_MULT multipliers with the dedicated_mult attribute, as shown in the following example:

```
//synthesis attribute <signal_name> dedicated_mult <value>
```



The dedicated_mult attribute is only applicable to signals or wires; it is not applicable to registers.

Table 16-4 describes the supported values for the dedicated_mult attribute.

Table 16-4. Values for the dedicated_mult Attribute (Part 1 of 2)

Value	Description
ON	LPM is inferred and multipliers are implemented in the DSP block.
OFF	LPM is inferred and multipliers are synthesized, implemented in logic, and optimized by the Quartus II software. (1)
LCELL	LPM is not inferred and multipliers are synthesized, implemented in logic, and optimized by the LeonardoSpectrum software. (1)

Table 16-4. Values for the dedicated_mult Attribute (Part 2 of 2)

Value	Description
AUTO	LPM is inferred, but the Quartus II software maps the multipliers automatically to either the DSP block or logic based on resource availability.

Note to Table 16-4:

- (1) Although both dedicated_mult=OFF and dedicated_mult=LCELLS result in logic implementations, the optimized results in these two cases may differ.



Some signals for which the dedicated_mult attribute is set may be removed during synthesis by the LeonardoSpectrum software due to design optimization. In such cases, if you want to force the implementation, you can preserve the signal from being removed during synthesis by setting the preserve_signal attribute to true.

The extract_mac attribute must be set to false for the module or project level when using the dedicated_mult attribute.

Example 16-3 and Example 16-4 are samples of Verilog HDL and VHDL codes, respectively, using the dedicated_mult attribute.

Example 16-3. Signal Attributes for Controlling DSP Block Inference in Verilog HDL Code

```

module mult (AX, AY, BX, BY, m, n, o, p);
input [7:0] AX, AY, BX, BY;
output [15:0] m, n, o, p;
wire [15:0] m_i = AX * AY; // synthesis attribute m_i dedicated_mult ON
// synthesis attribute m_i preserve_signal TRUE
//Note that the preserve_signal attribute prevents
// signal m_i from being removed during synthesis
wire [15:0] n_i = BX * BY; // synthesis attribute n_i dedicated_mult OFF
wire [15:0] o_i = AX * BY; // synthesis attribute o_i dedicated_mult AUTO
wire [15:0] p_i = BX * AY; // synthesis attribute p_i dedicated_mult
LCELL
// since n_i , o_i , p_i signals are not preserved,
// they may be removed during synthesis based on the design
assign m = m_i;
assign n = n_i;
assign o = o_i;
assign p = p_i;
endmodule

```

Example 16-4. Signal Attributes for Controlling DSP Block Inference in VHDL Code

```
library ieee ;
USE ieee.std_logic_1164.all;
USE ieee.std_logic_arith.all;
USE ieee.std_logic_unsigned.all;
USE ieee.std_logic_signed.all;
ENTITY mult is

PORT( AX,AY,BX,BY: IN
std_logic_vector (17 DOWNT0 0);
m,n,o,p: OUT
std_logic_vector (35 DOWNT0 0));
attribute dedicated_mult: string;
attribute preserve_signal : boolean
END mult;
ARCHITECTURE struct of mult is

signal m_i, n_i, o_i, p_i : unsigned (35 downto 0);
attribute dedicated_mult of m_i:signal is "ON";
attribute dedicated_mult of n_i:signal is "OFF";
attribute dedicated_mult of o_i:signal is "AUTO";
attribute dedicated_mult of p_i:signal is "LCELL";

begin

m_i <= unsigned (AX) * unsigned (AY);
n_i <= unsigned (BX) * unsigned (BY);
o_i <= unsigned (AX) * unsigned (BY);
p_i <= unsigned (BX) * unsigned (AY);

m <= std_logic_vector(m_i);
n <= std_logic_vector(n_i);
o <= std_logic_vector(o_i);
p <= std_logic_vector(p_i);
end struct;
```

Guidelines for Using DSP Blocks

In addition to the guidelines mentioned earlier in this section, use the following guidelines when designing with DSP blocks in the LeonardoSpectrum software:

- To access all the control signals for the DSP block, such as sign A, sign B, and dynamic addnsub, use the black box technique.
- While performing signed operations, ensure that the specified data width of the output port matches the data width of the expected result. Otherwise, the sign bit might be lost or data might be incorrect because the sign is not extended. For example, if the data widths of input A and B are width_a and width_b, respectively, the maximum data width of the result can be (width_a + width_b +2) for the four-multipliers adder mode. Thus, the data width of the output port should be less than or equal to (width_a + width_b +2).
- While using the accumulator, the data width of the output port should be equal to or greater than (width_a + width_b). The maximum width of the accumulator can be (width_a + width_b + 16). Accumulators wider than this are implemented in logic.

- If the design uses more multipliers than are available in a particular device, the Quartus II software may issue a no fit error. In such cases, use the attribute settings in the LeonardoSpectrum software to control the mapping of multipliers in your design to DSP blocks or logic.

Block-Based Design with the Quartus II Software

The incremental compilation design flow with LogicLock™ constraints enables users to design, optimize, and lock down a design one section at a time. You can independently create and implement each logic module into a hierarchical or team-based design. With this method, you can preserve the performance of each module during system integration and have more control over the placement of your design. To maximize the benefits of the incremental compilation in the Quartus II software, you can partition a new design into a hierarchy of netlist files during synthesis in the LeonardoSpectrum software.

You can create different netlist files with the LeonardoSpectrum software for different sections of a design hierarchy. Having different netlist files means that each section is independent of the others. When synthesizing the entire project, only portions of a design that have been updated must be resynthesized when you compile the design. You can make changes, optimize, and resynthesize your section of a design without affecting other sections.

-  For more information about incremental compilation, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*. For more information about the LogicLock feature, refer to the *Analyzing and Optimizing the Design Floorplan* chapter in volume 2 of the *Quartus II Handbook*.

Hierarchy and Design Considerations

You must plan your design's structure and partitioning carefully to use incremental compilation and LogicLock features effectively. Optimal hierarchical design practices include partitioning the blocks at functional boundaries, registering the boundaries of each block, minimizing the I/O between each block, separating timing-critical blocks, and keeping the critical path within one hierarchical block.

-  For more recommendations for hierarchical design partitioning, refer to the *Design Recommendations for Altera Devices and the Quartus II Design Assistant* chapter in volume 1 of the *Quartus II Handbook*.

To ensure the proper functioning of the synthesis tool, you can apply the LogicLock option in the LeonardoSpectrum software only to modules, entities, or netlist files. In addition, each module or entity should have its own design file. It is difficult to maintain incremental synthesis if two different modules are in the same design file, but are defined as being part of different regions, because both regions must be recompiled when you change one of the modules or entities.

If you use boundary tri-states in a lower-level block, the LeonardoSpectrum software pushes, or bubbles, the tri-states through the hierarchy to the top-level to take advantage of the tri-state drivers on the output pins of the Altera device. Because bubbling tri-states require optimizing through hierarchies, lower-level tri-states are not supported with a block-level design methodology. You should use tri-state drivers only at the external output pins of the device and in the top-level block in the hierarchy.

If the hierarchy is flattened during synthesis, logic is optimized across boundaries, preventing you from making LogicLock assignments to the flattened blocks. Altera recommends preserving the hierarchy when compiling the design. In the **Optimize** command of your script, use the **Hierarchy Preserve** command, or in the user interface, specify **Preserve** in the **Hierarchy** section of the **Optimize** FlowTab.

If you are compiling your design with a script, you can use an alternative method for preventing optimization across boundaries. In this case, use the **Auto** hierarchy setting and set the `auto_dissolve` attribute to `false` on the instances or views that you want to preserve (that is, the modules with LogicLock assignments) using the following syntax:

```
set_attribute -name auto_dissolve -value false \  
    .work.<block1>.INTERFACE
```

This alternative method flattens your design according to the `auto_dissolve` limits, but does not optimize across boundaries where you apply the attribute.



For more details about LeonardoSpectrum attributes and hierarchy levels, refer to the LeonardoSpectrum documentation in the Help menu.

Creating a Design with Multiple .edf Files

The first stage of a hierarchical design flow is to generate multiple **.edf** files, which allows you to take advantage of the incremental compilation flow in the Quartus II software. If the whole design is in one **.edf** file, changes in one block affect other blocks because of possible node name changes. You can generate multiple **.edf** files either by using the LogicLock option in the LeonardoSpectrum software, or by manually using a black box methodology on each block that you want to include in a LogicLock region.

After you create multiple **.edf** files with one of these methods, you must create the appropriate Quartus II project(s) to place-and-route the design.

Generating Multiple .edf Files Using the LogicLock Option

This section describes how to generate multiple **.edf** files using the LogicLock option in the LeonardoSpectrum software.

When synthesizing a top-level design that includes LogicLock regions, perform the following general steps:

1. Read in the Verilog HDL or VHDL source files.
2. Add LogicLock constraints.
3. Optimize and write output netlist files, or click **Run Flow**.

To set the correct constraints and compile the design, perform the following steps in the LeonardoSpectrum software:

1. On the Tools menu, switch to the **Advanced** FlowTab instead of the **Quick Setup** tab.
2. Set the target technology and speed grade for the device on the **Technology FlowTab**.
3. Open the input source files on the **Input** FlowTab.
4. Click **Read** on the **Input** FlowTab to read the source files, but not begin optimization.
5. Click the **Module** PowerTab located at the bottom of the **Constraints** FlowTab.
6. Select a module to be placed in a LogicLock region in the **Modules** section.
7. Turn on **LogicLock**.
8. Type the desired LogicLock region name under **LogicLock**.
9. Click **Apply**.
10. Repeat steps 6-9 for any other modules that you want to place in LogicLock regions.



In some cases, you are prompted to save your LogicLock and other non-global constraints in a Constraints File (.ctr) when you click anywhere off the **Constraints** FlowTab. The default name is *<project name>.ctr*. This file is added to your **Input** file list, and must be manually included later if you recreate the project.

The command written into the LeonardoSpectrum Information or Transcript Window is the Tcl command that is written into the .ctr file. The format of the “path” for the module specified in the command should be work.*<module>.INTERFACE*. To ensure that you do not see an optimized version of the module, do not perform a **Run Flow** on the **Quick Setup** tab prior to setting LogicLock constraints. Always use the **Read** command, as described in step 4.

11. Make any other settings as required on the **Constraints** FlowTab.
12. Select **Preserve** in the **Hierarchy** section of the **Optimize** FlowTab to ensure that the hierarchy names are not flattened during optimization.
13. Make any other settings as required on the **Optimize** FlowTab.
14. Run your synthesis flow with each FlowTab, or click **Run Flow**.

Synthesis creates an .edf file for each module that has a LogicLock assignment in the **Constraints** FlowTab. You can now use these files with the incremental compilation flow in the Quartus II software.



You might occasionally see multiple .edf files and LogicLock commands for the same module. An “unfolded” version of a module is created when you instantiate a module more than once and the boundary conditions of the instances are different. For example, if you apply a constant to one instance of the block, it might be optimized to eliminate unneeded logic. In this case, the LeonardoSpectrum software must create a

separate module for each instantiation (unfolding). If this unfolding occurs, you see more than one .edf file, and each .edf file has a LogicLock assignment to the same LogicLock region. When you import the .edf files to the Quartus II software, the .edf files created from the module are placed in different LogicLock regions. Any optimizations performed in the Quartus II software using the LogicLock methodology must be performed separately for each .edf file.

Creating a Quartus II Project for Multiple .edf Files Including LogicLock Regions

The LeonardoSpectrum software creates .tcl files that provide the Quartus II software with the appropriate LogicLock assignments, creating a region for each .edf file along with the information to set up a Quartus II project.

The .tcl file contains the commands shown in [Example 16-5](#) for each LogicLock region. This example is for module taps where the name taps_region is typed as the LogicLock region name in the **Constraints** FlowTab in the LeonardoSpectrum software.

Example 16-5. Tcl File for Module Taps with taps_region as LogicLock Region Name

```
project add_assignment {taps} {taps_region} {} {}  
  {LL_AUTO_SIZE} {ON}  
project add_assignment {taps} {taps_region} {} {}  
  {LL_STATE} {FLOATING}  
project add_assignment {taps} {taps_region} {} {}  
  {LL_MEMBER_OF} {taps_region}
```

These commands create a LogicLock region with auto-size and floating-origin properties. This flexible LogicLock region allows the Quartus II Compiler to select the size and location of the region.



For more information about Tcl commands, refer to the *Tcl Scripting* chapter in volume 2 of the *Quartus II Handbook*.

You can use the following methods to import the .edf file and corresponding .tcl file into the Quartus II software:

- Use the .tcl file that is created for each .edf file by the LeonardoSpectrum software. This method allows you to generate multiple Quartus II projects, one for each block in the design. Each designer in the project can optimize their block separately in the Quartus II software and preserve their results. Altera recommends this method for bottom-up incremental and hierarchical design methodologies because it allows each block in the design to be treated separately. Each block can be brought into one top-level project with the import function.

or

- Use the <top-level project>.tcl file that contains the assignments for all blocks in the project. This method allows the top-level designer to import all the blocks into one Quartus II project. You can optimize all modules in the project at once in a top-down design flow. If additional optimization is required for individual blocks, each designer can use their .edf file to create a separate project at that time. You must then add new assignments to the top-level project using the import function.

In both methods, use the following steps to create the Quartus II project, import the appropriate LogicLock assignments, and compile the design:

1. Place the **.edf** and **.tcl** files in the same directory.
2. On the View menu, point to **Utility Windows** and click **Tcl Console** to open the **Quartus II Tcl Console**.
3. At the Tcl prompt, type `source <path>/<project name>.tcl` ↵.
4. To open the newly completed project, on the File menu, click **Open Project**. Browse to and select the project name, and click **Open**.



For more information about importing a design using incremental compilation, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*. For more information about importing LogicLock assignments, see the *Analyzing and Optimizing the Design Floorplan* chapter in volume 2 of the *Quartus II Handbook*.

Generating Multiple .edf Files Using Black Boxes

This section describes how to manually generate multiple **.edf** files using the black box technique. The manual flow, which was supported in older versions of the LeonardoSpectrum software, is discussed here because some designers want more control over the project for each submodule.

To create multiple **.edf** files in the LeonardoSpectrum software, create a separate project for each module and top-level design that you want to maintain as a separate **.edf** file. Implement black box instantiations of lower-level modules in your top-level project.

When synthesizing the projects for the lower-level modules and the top-level design, use the following general guidelines:

For lower-level modules:

- Turn off **Map IO Registers** for the target technology on the **Technology** FlowTab.
- Read the HDL files for the modules. Modules may include black box instantiations of lower-level modules that are also maintained as separate **.edf** files.
- Add constraints.
- Turn off **Add I/O Pads** on the **Optimize** FlowTab.

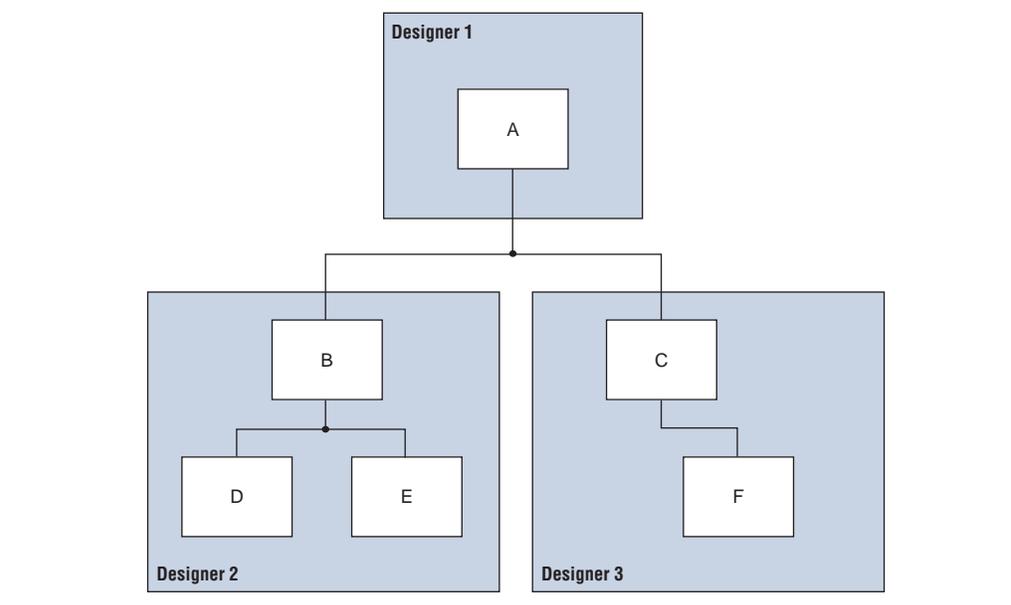
For the top-level design:

- Turn on **Map IO Registers** if you want to implement input and/or output registers in the IOEs for the target technology on the **Technology** FlowTab.
- Read the HDL files for the top-level design.
- Black box lower-level modules in the top-level design.
- Add constraints (clock settings should be made at this time).

The following sections describe examples of black box modules in a block-based and team-based design flow.

In [Figure 16-3](#), the top-level module A is assigned to one engineer (designer 1), while two engineers work on the lower levels of the design. Designer 2 works on module B and its submodules D and E, while designer 3 works on module C and its submodule F.

Figure 16-3. Block-Based and Team-Based Design Example



One netlist is created for the top-level module A, another netlist is created for module B and its submodules D and E, and another netlist is created for module C and its submodule F. To create multiple `.edf` files, perform the following steps:

1. Generate an `.edf` file for module C. Use `C.v` and `F.v` as the source files.
2. Generate an `.edf` file for module B. Use `B.v`, `D.v`, and `E.v` as the source files.
3. Generate a top-level `.edf` file `A.v` for module A. Ensure that your black box modules B and C were optimized separately in steps 1 and 2.

Black Box Methodology in Verilog HDL

Any design block that is not defined in the project, or included in the list of files to be read for a project, is treated as a black box by the software. In Verilog HDL, you must also provide an empty module declaration for the module that you plan to treat as a black box.

Example 16-6 shows an example of the **A.v** top-level file. If any of your lower-level files also contain a black-boxed lower-level file in the next level of hierarchy, follow the same procedure.

Example 16-6. Verilog HDL Top-Level File Black Boxing Example

```
module A (data_in,clk,e,ld,data_out);
    input data_in, clk, e, ld;
    output [15:0] data_out;

    reg [15:0] cnt_out;
    reg [15:0] reg_a_out;

    B U1 ( .data_in (data_in),.clk (clk), .e(e), .ld (ld),
        .data_out(cnt_out) );

    C U2 ( .d(cnt_out), .clk (clk), .e(e), .q (reg_out));
    // Any other code in A.v goes here.

endmodule

// Empty Module Declarations of Sub-Blocks B and C follow here.
// These module declarations (including ports) are required for
blackboxing.

module B (data_in,e,ld,data_out );
    input data_in, clk, e, ld;
    output [15:0] data_out;
endmodule

module C (d,clk,e,q );
    input d, clk, e;
    output [15:0] q;
endmodule
```



Previous versions of the LeonardoSpectrum software require the attribute statement `//exemplar attribute U1 NOOPT TRUE`, which instructs the software to treat the instance `U1` as a black box. This attribute is no longer required, although it is still supported in the software.

Black Boxing in VHDL

Any design block that is not defined in the project, or included in the list of files to be read for a project, is treated as a black box by the LeonardoSpectrum software. In VHDL, a component declaration is required for the black box.

Example 16-7 shows an example of the **A.vhd** top-level file. If any of your lower-level files also contain a black-boxed lower-level file in the next level of hierarchy, follow the same procedure.

Example 16-7. VHDL Top-Level File Black Boxing Example

```

LIBRARY ieee;
USE ieee.std_logic_1164.all;
ENTITY A IS
PORT ( data_in : IN INTEGER RANGE 0 TO 15;
      clk : IN STD_LOGIC;
      e : IN STD_LOGIC;
      ld : IN STD_LOGIC;
      data_out : OUT INTEGER RANGE 0 TO 15
);
END A;

ARCHITECTURE a_arch OF A IS

COMPONENT B PORT(
  data_in : IN INTEGER RANGE 0 TO 15;
  clk : IN STD_LOGIC;
  e : IN STD_LOGIC;
  ld : IN STD_LOGIC;
  data_out : OUT INTEGER RANGE 0 TO 15
);
END COMPONENT;

COMPONENT C PORT(
  d : IN INTEGER RANGE 0 TO 15;
  clk : IN STD_LOGIC;
  e : IN STD_LOGIC;
  q : OUT INTEGER RANGE 0 TO 15
);
END COMPONENT;

-- Other component declarations in A.vhd go here

signal cnt_out : INTEGER RANGE 0 TO 15;
signal reg_a_out : INTEGER RANGE 0 TO 15;
BEGIN
CNT : C
PORT MAP (
  data_in => data_in,
  clk => clk,
  e => e,
  ld => ld,
  data_out => cnt_out
);

REG_A : D
PORT MAP (
  d => cnt_out,
  clk => clk,
  e => e,
  q => reg_a_out
);

-- Any other code in A.vhd goes here

END a_arch;

```

 Previous versions of the LeonardoSpectrum software require the attribute statement `noopt of C: component is TRUE`, which instructs the software to treat the component C as a black box. This attribute is no longer required, although it is still supported in the software.

After you have completed the steps outlined in this section, you have a different `.edf` netlist file for each block of code. You can now use these files for the incremental compilation flow in the Quartus II software.

Creating a Quartus II Project for Multiple `.edf` Files

The LeonardoSpectrum software creates a `.tcl` file for each `.edf` file, which provides the Quartus II software with the information to set up a project.

There are two different methods for importing each `.edf` file and corresponding `.tcl` file into the Quartus II software:

- Use the `.tcl` file that is created for each `.edf` file by the LeonardoSpectrum software. This method generates multiple Quartus II projects, one for each block in the design. Each designer in the project can optimize their block separately in the Quartus II software and preserve their results. Designers should create a LogicLock region for each block; the top-level designer should then import all the blocks and assignments into the top-level project. Altera recommends this method for bottom-up incremental and hierarchical design methodology because it allows each block in the design to be treated separately; each block can be imported into one top-level project.

or

- Use the `<top-level project>.tcl` file that contains the information to set up the top-level project. This method allows the top-level designer to create LogicLock regions for each block and bring all the blocks into one Quartus II project. Designers can optimize all modules in the project at once in a top-down design flow. If additional optimization is required for individual blocks, each designer creates a separate Quartus II project with each `.edf` file. New assignments must then be added to the top-level project manually or through the import function.

 For more information about importing designs using incremental compilation, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*. For more information about importing LogicLock regions, refer to the *Analyzing and Optimizing the Design Floorplan* chapter in volume 2 of the *Quartus II Handbook*.

With both methods, use the following steps to create the Quartus II project and compile the design:

1. Place the `.edf` and `.tcl` files in the same directory.
2. On the View menu, point to **Utility Windows** and click **Tcl Console**. The Quartus II Tcl Console appears.
3. At a Tcl prompt, type `source <path>/<project name>.tcl` .
4. On the File menu, click **Open Project**. In the New Project window, browse to and select the project name. Click **Open**.

5. To create LogicLock assignments, on the Assignments menu, click **LogicLock Regions Window**. Use the LogicLock Regions window to create LogicLock regions.
6. On the Processing menu, click **Start Compilation**.

Incremental Synthesis Flow

If you make changes to one or more submodules, you can manually create new projects in the LeonardoSpectrum software to generate a new **.edf** file when there are changes to the source files. Alternatively, you can use incremental synthesis to generate a new netlist for the changed submodule(s). To perform incremental synthesis in the LeonardoSpectrum software, use the script described in this section to reoptimize and generate a new **.edf** file for only the affected modules using the LeonardoSpectrum top-level project. This method applies only when you are using the **LogicLock** option in the LeonardoSpectrum software.

Modifications Required for the **LogicLock_Incremental.tcl** Script File

There are three sets of entries in the file that must be modified before beginning incremental synthesis. The variables in the **.tcl** file are surrounded by angle brackets (< >).

1. Add the list of source files that are included in the project. You can enter the full path to the file or just the file name if the files are located in the working directory.
2. Indicate which modules in the design have changed. These modules are the **.edf** files that are regenerated by the LeonardoSpectrum software and contain a LogicLock assignment in the original compilation.



Determine the LeonardoSpectrum software path for each module by looking at the **.ctr** file that contains the LogicLock assignments from the original project. Each LogicLock assignment is applied to a particular module in the design.

3. Enter the target device family using the appropriate device keyword. The device keyword is displayed on the Transcript or Information window when you select a target Technology and click **Load Library** or **Apply** on the **Technology** FlowTab.

Example 16-8 shows the `LogicLock_Incremental.tcl` file for the incremental synthesis flow. You must modify the `.tcl` file before you can use it for your project.

Example 16-8. LogicLock_Interface.tcl Script File for Incremental Synthesis

```
#####
#### LogicLock Incremental Synthesis Flow ####
#####

## You must indicate which modules have changed (based on the source files
## that have changed) and provide the complete path to each module

## You must also specify the list of design files and the target Altera
## technology being used

# Read the design source files.
read <list of design files separated by spaces (such as block1.v block2.v)>

# Get the list of modified modules in bottom-up "depth first search" order
# where the lower-level blocks are listed first (these should be modules
# that had LogicLock assignments and separate EDIF netlist files in the
# first pass and had their source code modified)

set list_of_modified_modules {.work.<block2>.INTERFACE .work.<block1>.INTERFACE}

foreach module $list_of_modified_modules {
    set err_rc [regexp {\.(.*)\.(.*)\.(.*)} $module unused lib module_name arch]
    present_design $module

    # Run optimization, preserving hierarchy. You must specify a technology.
    optimize -ta <technology> -hierarchy preserve

    # Ensure that the lower-level module is not optimized again when
    # optimizing higher-level modules.
    dont_touch $module
}

foreach module $list_of_modified_modules {
    set err_rc [regexp {\.(.*)\.(.*)\.(.*)} $module unused lib module_name arch]
    present_design $module
    undont_touch $module
    auto_write $module_name.edf
    # Ensure that the lower-level module is not written out in the EDIF file
    # of the higher-level module.
    noopt $module
}

```

Running the Tcl Script File in LeonardoSpectrum

After you modify the Tcl script, as described in “[Modifications Required for the LogicLock_Incremental.tcl Script File](#)” on page 16-25, you can compile your design using the script.

You can run the script in batch mode at the command-line prompt using the following command:

```
spectrum -file <Tcl_file> ←
```

To run the script from the GUI, on the File menu, click **Run Script**, then browse to your `.tcl` file and click **Open**.

The LogicLock incremental design flow uses module-based design to help you preserve performance of modules and have control over placement. By tagging the modules that require separate **.edf** files, you can make multiple **.edf** files for use with the Quartus II software from a single LeonardoSpectrum software project.

Conclusion

Taking advantage of the Mentor Graphics LeonardoSpectrum software and the Quartus II design flow allows you to control how your design files are prepared for the Quartus II place-and-route process, as well as to improve performance and optimize a design for use with Altera devices. The methodologies outlined in this chapter can help optimize a design to achieve performance goals and save design time with the LeonardoSpectrum software. For the best results with new designs in new device families, Altera recommends migrating to the advanced Mentor Graphics Precision RTL Synthesis software.

Document Revision History

Table 16-5 shows the revision history for this chapter.

Table 16-5. Document Revision History

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Changed to new document template. ■ Removed support for the Classic Timing Analyzer. ■ Editorial changes.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Updated supported devices. ■ Removed Referenced Documents section. ■ Minor updates for the version 10.0 release.
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Minor updates for the Quartus II software version 9.1 release. ■ Removed Table 12-3, Inferring RAM Summary.
March 2009	9.0.0	<ul style="list-style-type: none"> ■ No change to content. ■ Chapter 12 was previously Chapter 11 in software release 8.1.
November 2008	8.1.0	<ul style="list-style-type: none"> ■ Changed to 8-1/2" x 11" page size. ■ Updated Table 12-3.
May 2008	8.0.0	Updated date and part number and added hypertext links.

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this chapter.

This chapter describes how you can use the Quartus® II netlist viewers to analyze and debug your designs. As FPGA designs grow in size and complexity, the ability to analyze, debug, optimize, and constrain your design is critical. Often, with today's advanced designs, several design engineers are involved in coding and synthesizing different design blocks, making it difficult to analyze and debug the design. The Quartus II RTL Viewer, State Machine Viewer, and Technology Map Viewer provide powerful ways to view your initial and fully mapped synthesis results during the debugging, optimization, and constraint entry processes.

This chapter contains the following sections:

- “Introduction to the User Interface” on page 17–6
- “Navigating the Schematic View” on page 17–22
- “Filtering in the Schematic View” on page 17–33
- “Probing to a Source Design File and Other Quartus II Windows” on page 17–39
- “Probing to the Netlist Viewers from Other Quartus II Windows” on page 17–40
- “Viewing a Timing Path” on page 17–41
- “Other Features in the Schematic Viewer” on page 17–42

The first section in this chapter, “When to Use the Netlist Viewers: Analyzing Design Problems”, describes examples of using the netlist viewers to analyze your design at various stages of the design cycle. The sections following this section provide an introduction to the Quartus II design flow using the netlist viewers, an overview of each viewer, and an explanation of the user interface. These sections describe the following tasks:

- How to navigate and filter schematics
- How to probe to and from other windows in the Quartus II software
- How to view a timing path from the Timing Analyzer report

When to Use the Netlist Viewers: Analyzing Design Problems

You can use the netlist viewers to analyze and debug your design. This section provides simple examples of how to use the RTL Viewer, State Machine Viewer, and Technology Map Viewer to analyze problems encountered in the design process.

The following sections contain information about how the netlist viewers display your design:

- “Quartus II Design Flow with the Netlist Viewers” on page 17–2
- “RTL Viewer Overview” on page 17–4
- “State Machine Viewer Overview” on page 17–5

- [“Technology Map Viewer Overview” on page 17-5](#)

Using the RTL Viewer is a good way to view your initial synthesis results to determine whether you have created the necessary logic, and that the logic and connections have been interpreted correctly by the software. You can use the RTL Viewer and State Machine Viewer to check your design visually before simulation or other verification processes. Catching design errors at this early stage of the design process can save you valuable time.

If you see unexpected behavior during verification, use the RTL Viewer to trace through the netlist and ensure that the connections and logic in your design are as expected. You can also view state machine transitions and transition equations with the State Machine Viewer. Viewing your design helps you find and analyze the source of design problems. If your design looks correct in the RTL Viewer, you know to focus your analysis on later stages of the design process and investigate potential timing violations or issues in the verification flow itself.

You can use the Technology Map Viewer to look at the results at the end of synthesis and technology mapping by running the netlist viewer after performing Analysis and Synthesis. If you have compiled your design through the Fitter stage, you can view your post-mapping netlist in the Technology Map Viewer (Post-Mapping) and your post-fitting netlist in the Technology Map Viewer. If you perform only Analysis and Synthesis, both the netlist viewers display the same post-mapping netlist.

In addition, you can use the RTL Viewer or Technology Map Viewer to locate the source of a particular signal, which can help you debug your design. Use the navigation techniques described in this chapter to search easily through your design. You can trace back from a point of interest to find the source of the signal and ensure the connections are as expected.

The Technology Map Viewer can help you locate post-synthesis nodes in your netlist and make assignments when optimizing your design. This functionality is useful, for example, when making a multicycle clock timing assignment between two registers in your design. Start at an I/O port and trace forward or backward through the design and through levels of hierarchy to find nodes that interest you, or locate a specific register by visually inspecting the schematic.

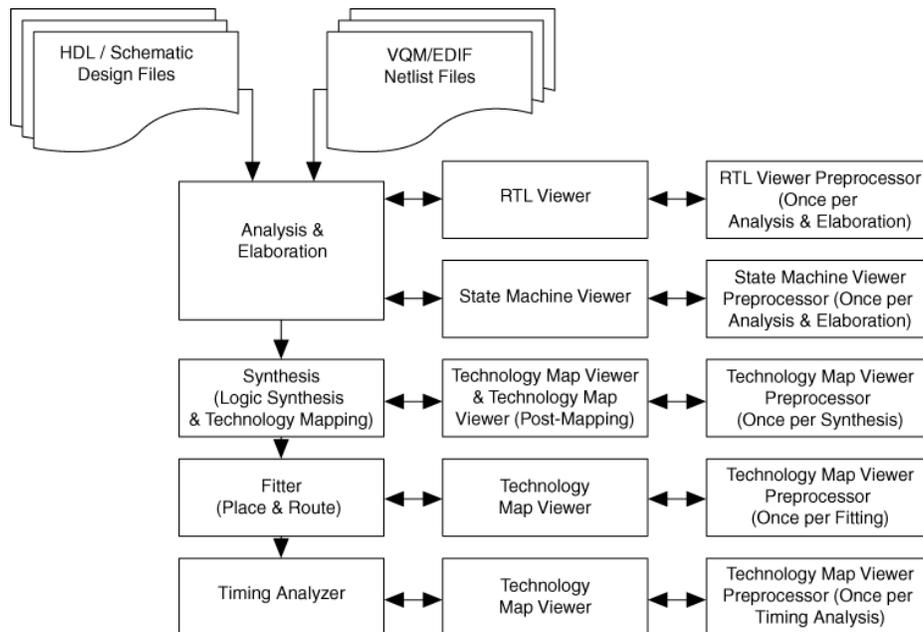
You can use the RTL Viewer, State Machine Viewer, and Technology Map Viewer in many other ways throughout the design, debugging, and optimization stages. Viewing the design netlist is a powerful way to analyze design problems. This chapter shows you how to use the various features of the netlist viewers to increase your productivity when analyzing a design.

Quartus II Design Flow with the Netlist Viewers

When you first open one of the netlist viewers after compiling the design, a preprocessor stage runs automatically before the netlist viewer opens. If you close the netlist viewer and open it again later without recompiling the design, the netlist viewer opens immediately without performing the preprocessing stage.

Figure 17-1 shows how the netlist viewers fit into the basic Quartus II design flow.

Figure 17-1. Quartus II Design Flow Including the RTL Viewer and Technology Map Viewer



To use a netlist viewer, and before the netlist viewer can run the preprocessor and open the design, compile your design with the following minimum compilation:

- To open the RTL Viewer or State Machine Viewer, first perform Analysis and Elaboration.
- To open the Technology Map Viewer or the Technology Map Viewer (Post-Mapping), first perform Analysis and Synthesis.

 If you open one of the netlist viewers without first compiling the design with the appropriate minimum compilation stage, the netlist viewer does not appear. Instead, the Quartus II software issues an error message instructing you to run the necessary compilation stage and restart the netlist viewer.

The netlist viewers display the results of the last successful compilation. Therefore, if you make a design change that causes an error during Analysis and Elaboration, you cannot view the netlist for the new design files, but you can still see the results from the last successfully compiled version of the design files. If you receive an error during compilation and you have not yet successfully run the appropriate compilation stage for your project, the netlist viewer cannot be displayed; in this case, the Quartus II software issues an error message when you try to open the netlist viewer.

 If the netlist viewer is open when you start a new compilation, the netlist viewer closes automatically. You must open the netlist viewer again to view the new design netlist after compilation completes successfully.

RTL Viewer Overview

The Quartus II RTL Viewer allows you to view a register transfer level (RTL) graphical representation of your Quartus II integrated synthesis results or your third-party netlist file in the Quartus II software.

You can view results after Analysis and Elaboration when your design uses any supported Quartus II design entry method, including Verilog HDL Design Files (.v), SystemVerilog Design Files (.sv), VHDL Design Files (.vhd), AHDL Text Design Files (.tdf), schematic Block Design Files (.bdf), or schematic Graphic Design Files (.gdf) imported from the MAX+PLUS® II software. You can also view the hierarchy of atom primitives (such as device logic cells and I/O ports) when your design uses a synthesis tool to generate a Verilog Quartus Mapping File (.vqm) or Electronic Design Interchange Format (.edf) netlist file. For a flow diagram, refer to [Figure 17-1](#).

The Quartus II RTL Viewer displays a schematic view of the design netlist after Analysis and Elaboration or netlist extraction is performed by the Quartus II software, but before technology mapping and any synthesis or fitter optimization algorithms occur. This view is not the final design structure because optimizations have not yet occurred. This view most closely represents your original source design. If you synthesized your design with the Quartus II integrated synthesis, this view shows how the Quartus II software interpreted your design files. If you are using a third-party synthesis tool, this view shows the netlist written by your synthesis tool.

When displaying your design, the RTL Viewer optimizes the netlist to maximize readability in the following ways:

- Logic with no fan-out (its outputs are unconnected) and logic with no fan-in (its inputs are unconnected) are removed from the display.
- Default connections such as V_{CC} and GND are not shown.
- Pins, nets, wires, module ports, and certain logic are grouped into buses where appropriate.
- Constant bus connections are grouped.
- Values are displayed in hexadecimal format.
- NOT gates are converted to bubble inversion symbols in the schematic.
- Chains of equivalent combinational gates are merged into a single gate. For example, a 2-input AND gate feeding a 2-input AND gate is converted to a single 3-input AND gate.
- State machine logic is converted into a state diagram, state transition table, and state encoding table, which are displayed in the State Machine Viewer.

To run the RTL Viewer for a Quartus II project, first analyze the design to generate an RTL netlist. To analyze the design and generate an RTL netlist, on the Processing menu, point to **Start** and click **Start Analysis & Elaboration**. You can also perform a full compilation on any process that includes the initial Analysis and Elaboration stage of the Quartus II compilation flow.

To run the RTL Viewer, on the Tools menu, point to **Netlist Viewers** and click **RTL Viewer**.

You can set the RTL Viewer preprocessing to run during a full compilation, which allows you to launch the RTL Viewer after Analysis and Synthesis has completed, but while the Fitter is still running. In this case, you do not have to wait for the Fitter to finish before viewing the schematic. This technique is useful for a large design that requires a substantial amount of time in the place-and-route stage.

To set the RTL Viewer preprocessing to run during compilation, on the Assignments menu, click **Settings**. In the **Category** list, select **Compilation Process Settings** and turn on **Run RTL Viewer preprocessing during compilation**. By default, this option is turned off.

State Machine Viewer Overview

The State Machine Viewer presents a high-level view of finite state machines in your design. The State Machine Viewer provides a graphical representation of the states and their related transitions, as well as a state transition table that displays the condition equation for each of the state transitions, and encoding information for each state.

To run the State Machine Viewer, on the Tools menu, point to **Netlist Viewers** and click **State Machine Viewer**. To open the State Machine Viewer for a particular state machine, double-click the state machine instance in the RTL Viewer or right-click the state machine instance and click **Hierarchy Down**.

Technology Map Viewer Overview

The Quartus II Technology Map Viewer provides a technology-specific, graphical representation of your design after Analysis and Synthesis or after the Fitter has mapped your design into the target device. The Technology Map Viewer shows the hierarchy of atom primitives (such as device logic cells and I/O ports) in your design. For supported families, you can also view internal registers and look-up tables (LUTs) inside logic cells (LCELLs) and registers in I/O atom primitives. For more information, refer to [“Viewing Contents of Atom Primitives”](#) on page 17–23.



Where possible, the port names of each hierarchy are maintained throughout synthesis; however, port names might change or be removed from the design. For example, if a port is unconnected or driven by GND or V_{CC} , it is removed during synthesis. When a port name changes, the port is assigned a related user logic name in the design or a generic port name such as IN1 or OUT1.

You can view your Quartus II technology-mapped results after synthesis, fitting, or timing analysis. To run the Technology Map Viewer for a Quartus II project, on the Processing menu, point to **Start** and click **Start Analysis & Synthesis** to synthesize and map the design to the target technology. At this stage, the Technology Map Viewer shows the same post-mapping netlist as the Technology Map Viewer (Post-Mapping). You can also perform a full compilation, or any process that includes the synthesis stage in the compilation flow.

If you have completed the Fitter stage, the Technology Map Viewer shows the changes made to your netlist by the Fitter, such as physical synthesis optimizations, while the Technology Map Viewer (Post-Mapping) shows the post-mapping netlist. If you have completed the Timing Analysis stage, you can locate timing paths from the Timing Analyzer report in the Technology Map Viewer (for more information, refer to “Viewing a Timing Path” on page 17-41). For a flow diagram, refer to Figure 17-1 on page 17-3.

To run the Technology Map Viewer, on the Tools menu, point to **Netlist Viewers** and click **Technology Map Viewer**, or select **Technology Map Viewer** from the Applications toolbar.

To run the Technology Map Viewer (Post-Mapping), on the Tools menu, point to **Netlist Viewers** and click **Technology Map Viewer (Post-Mapping)**.

Introduction to the User Interface

The RTL Viewer and Technology Map Viewer each consist of three main parts:

- The schematic view
- The **Netlist Navigator** pane
- The **Find** pane

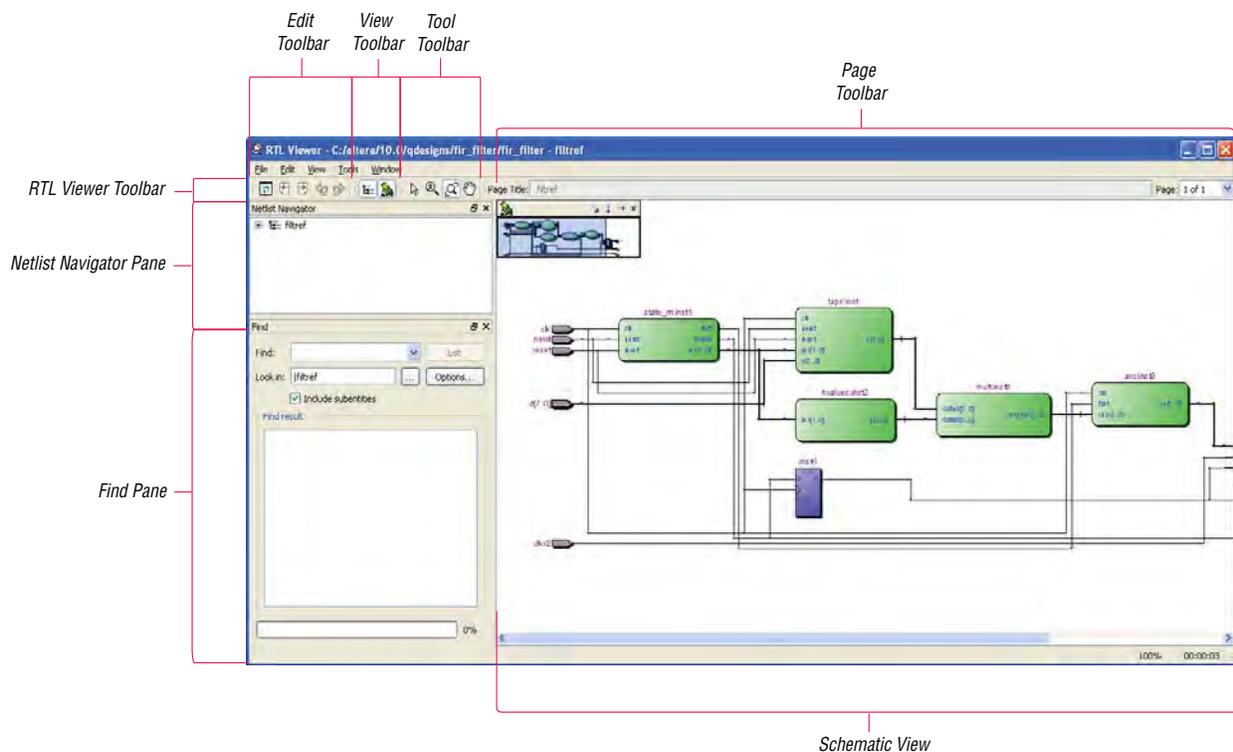
The schematic view displays a graphical representation of the internal structure of your design; the **Netlist Navigator** pane displays a representation of the project hierarchy; and the **Find** pane allows you to find and locate specific design elements in the schematic view.

Figure 17-2 shows the RTL Viewer and indicates these three parts, along with other elements of the user interface. Both the netlist viewers also contain a toolbar that provides tools to use in the schematic view.

You can have only one RTL Viewer, one Technology Map Viewer, one Technology Map Viewer (Post-Mapping), and one State Machine Viewer window open at the same time, although each window can show multiple pages. For example, you cannot have two RTL Viewer windows open at the same time. The netlist viewer window has characteristics similar to other “child” windows in the Quartus II software; it can be resized and moved, minimized or maximized, tiled or cascaded, and moved in front of or behind other windows.

Figure 17-2 shows the schematic view and the Netlist Navigator pane of the RTL Viewer.

Figure 17-2. RTL Viewer



Schematic View

The schematic view is shown on the right side of the RTL Viewer and Technology Map Viewer. It contains a schematic representing the design logic in the netlist. This view is the main screen for viewing your gate-level netlist in the RTL Viewer and your technology-mapped netlist in the Technology Map Viewer.

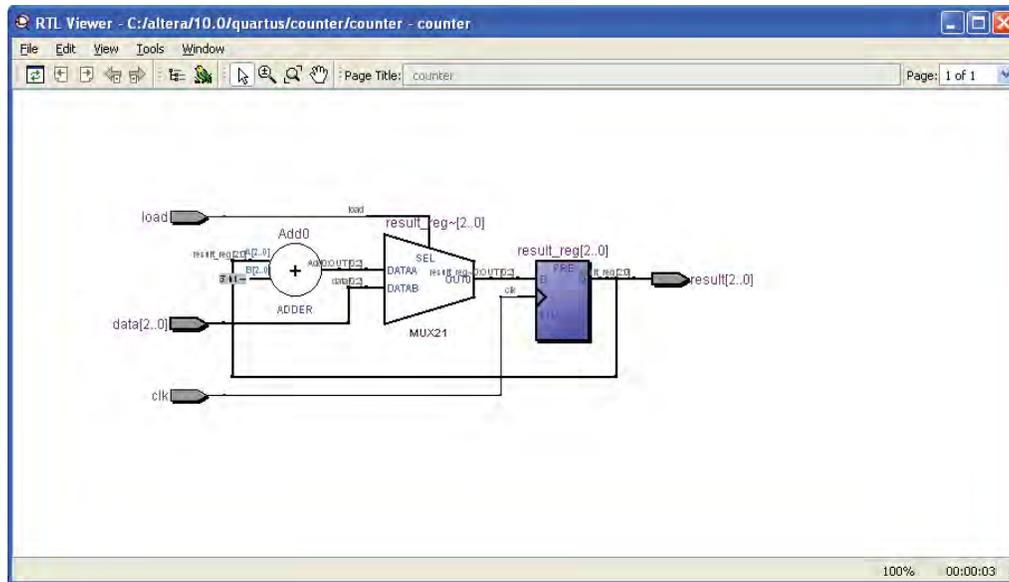
Schematic Symbols

The symbols for nodes in the schematic represent elements of your design netlist. These elements include input and output ports, registers, logic gates, Altera® primitives, high-level operators, and hierarchical instances.

Figure 17-3 shows an example of an RTL Viewer schematic for a 3-bit synchronous loadable counter. Example 17-1 shows the Verilog HDL code that produced this schematic. This example includes multiplexers and a group of registers (Table 17-1) in a bus along with an ADDER operator (Table 17-3 on page 17-13) inferred by the counting function in the HDL code.

The schematic in [Figure 17-3](#) displays wire connections between nodes with a thin black line and bus connections with a thick black line.

Figure 17-3. Example Schematic Diagram in the RTL Viewer



Example 17-1. Code Sample for Counter Schematic Shown in [Figure 17-3](#)

```

module counter (input [2:0] data, input clk, input load, output [2:0]
result);
    reg [2:0] result_reg;
    always @ (posedge clk)
        if (load)
            result_reg <= data;
        else
            result_reg <= result_reg + 1;
    assign result = result_reg;
endmodule

```

Figure 17-4 shows a portion of the corresponding Technology Map Viewer schematic with a compiled design that targets a Stratix® device. In this schematic, you can see the LCELL (logic cell) device-specific primitives that represent the counter function, labeled with their post-synthesis node names. The REGOUT port represents the output of the register in the LCELL; the COMBOUT port represents the output of the combinational logic in the LUT of the LCELL. The hexadecimal number in parentheses below each LCELL primitive represents the LUT mask, which is a hexadecimal representation of the logic function of the LCELL.

Figure 17-4. Example Schematic Diagram in the Technology Map Viewer

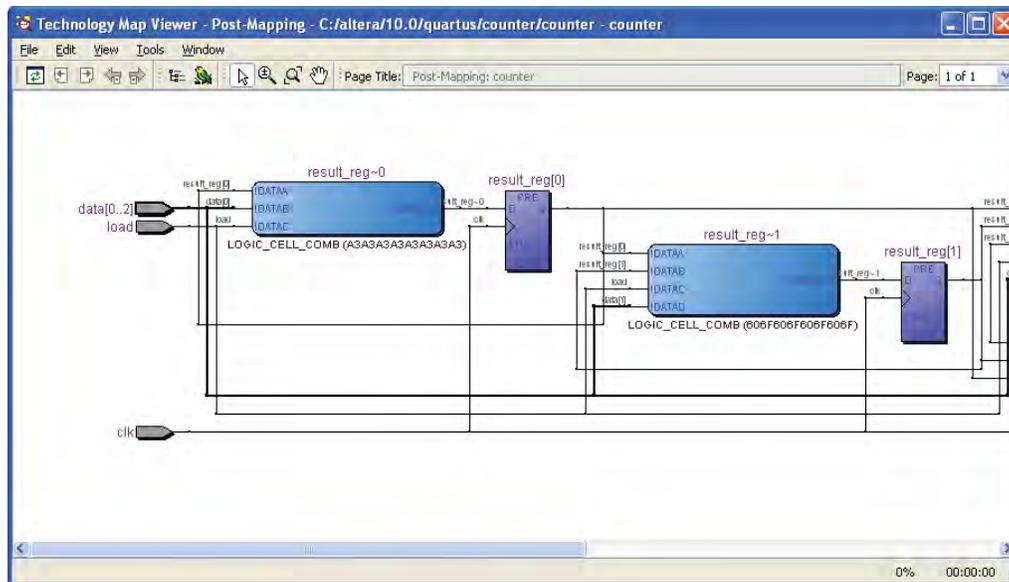


Table 17-1 lists and describes the primitives and basic symbols that you can display in the schematic view of the RTL Viewer and Technology Map Viewer. Table 17-3 on page 17-13 lists and describes the additional higher-level operator symbols used in the RTL Viewer schematic view.

 The logic gates and operator primitives appear only in the RTL Viewer. Logic in the Technology Map Viewer is represented by atom primitives, such as registers and LCELLs.

Table 17-1. Symbols in the Schematic View (Part 1 of 3)

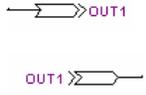
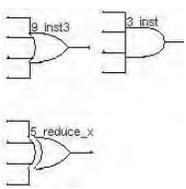
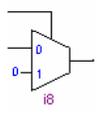
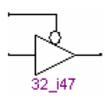
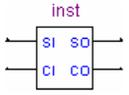
Symbol	Description
<p>I/O Ports</p> 	<p>An input, output, or bidirectional port in the current level of hierarchy. A device input, output, or bidirectional pin when viewing the top-level hierarchy. The symbol can also represent a bus. Only one wire is shown connected to the bidirectional symbol, representing the input and output paths.</p> <p>Input symbols appear on the left-most side of the schematic. Output and bidirectional symbols appear on the right-most side of the schematic.</p>
<p>I/O Connectors</p> 	<p>An input or output connector, representing a net that comes from another page of the same hierarchy (refer to “Partitioning the Schematic into Pages” on page 17-30). To go to the page that contains the source or the destination, right-click on the net and choose the page from the menu (refer to “Following Nets Across Schematic Pages” on page 17-31).</p>
<p>Hierarchy Port Connector</p> 	<p>A connector representing a port relationship between two different hierarchies. A connector indicates that a path passes through a port connector in a different level of hierarchy.</p>
<p>OR, AND, XOR Gates</p> 	<p>An OR, AND, or XOR gate primitive (the number of ports can vary). A small circle (bubble symbol) on an input or output port indicates the port is inverted.</p>
<p>MULTIPLEXER</p> 	<p>A multiplexer primitive with a selector port that selects between port 0 and port 1. A multiplexer with more than two inputs is displayed as an operator (refer to “Operator Symbols in the RTL Viewer Schematic View” on page 17-13).</p>
<p>BUFFER</p> 	<p>A buffer primitive. The figure shows the tri-state buffer, with an inverted output enable port. Other buffers without an enable port include LCELL, SOFT, CARRY, and GLOBAL. The NOT gate and EXP expander buffers use this symbol without an enable port and with an inverted output port.</p>
<p>CARRY_SUM</p> 	<p>A CARRY_SUM buffer primitive with the following ports:</p> <ul style="list-style-type: none"> ■ SI – SUM IN ■ SO – SUM OUT ■ CI – CARRY IN ■ CO – CARRY OUT

Table 17-1. Symbols in the Schematic View (Part 2 of 3)

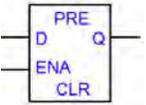
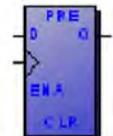
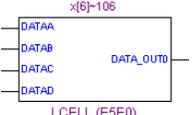
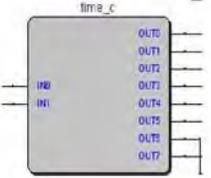
Symbol	Description
<p>LATCH</p> 	<p>A latch primitive with the following ports:</p> <ul style="list-style-type: none"> ■ D – data input ■ ENA – latch enable input ■ Q – data output ■ PRE – preset ■ CLR – clear
<p>DFFE/DFFEAS/DFFEAS</p> 	<p>A DFFE (data flipflop with clock enable) primitive, with the same ports as a latch and a clock trigger. The other flipflop primitives are similar:</p> <ul style="list-style-type: none"> ■ DFFEAS (data flipflop with enable and asynchronous load) primitive with additional <code>ALOAD</code> asynchronous load and <code>ADATA</code> data signals ■ DFFEAS (data flipflop with enable and synchronous and asynchronous load), which has <code>ASDATA</code> as the secondary data port
<p>Atom Primitive</p> 	<p>An atom primitive. The symbol displays the atom name, the port names, and the atom type. The blue shading indicates an atom primitive for which you can view the internal details. For more information, refer to “Viewing Contents of Atom Primitives” on page 17-23.</p>
<p>Other Primitive</p> 	<p>Any primitive that does not fall into the previous categories. Primitives are low-level nodes that cannot be expanded to any lower hierarchy. The symbol displays the port names, the primitive or operator type, and its name.</p> <p>The figure shows an LCELL WYSIWYG primitive, with <code>DATAA</code> to <code>DATAD</code> and <code>COMBOUT</code> port connections. This type of LCELL primitive is found in the Technology Map Viewer for technology-specific atom primitives when the contents of the atom primitive cannot be viewed. The RTL Viewer contains similar primitives if the source design is a VQM or EDIF netlist.</p>
<p>Instance</p> 	<p>An instance in the design that does not correspond to a primitive or operator (generally a user-defined hierarchy block). The symbol displays the port name and the instance name.</p> <p>For more information about opening the schematic for the lower-level hierarchy, refer to “Traversing and Viewing the Design Hierarchy” on page 17-22.</p>
<p>Encrypted Instance</p> 	<p>A user-defined encrypted instance in the design. The symbol displays the instance name. You cannot open the schematic for the lower-level hierarchy, because the source design is encrypted.</p>
<p>State Machine Instance</p> 	<p>A finite state machine instance in the design. For more information, refer to “State Machine Viewer” on page 17-16.</p>

Table 17-1. Symbols in the Schematic View (Part 3 of 3)

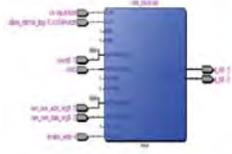
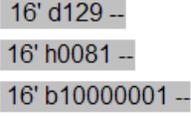
Symbol	Description
	<p>A synchronous memory instance with registered inputs and optionally registered outputs. The symbol shows the device family and the type of memory block. This figure shows a true dual-port memory block in a Stratix M-RAM block.</p>
	<p>A combinational logic cloud in the design. For more information, refer to “Grouping Combinational Logic into Logic Clouds” on page 17-26.</p>
	<p>A constant signal value that is highlighted in gray and displayed in hexadecimal format by default throughout the schematic. To change the format, refer to “Changing the Constant Signal Value Formatting” on page 17-28.</p>

Table 17-2 lists and describes the symbol used only in the State Machine Viewer.

Table 17-2. Symbol Available Only in the State Machine Viewer

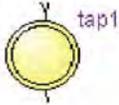
Symbol	Description
	<p>The node representing a state in a finite state machine. State transitions are indicated with arcs between state nodes. The double circle border indicates the state connects to logic outside the state machine, and a single circle border indicates the state node does not feed outside logic.</p>

Table 17-3 lists and describes the additional higher level operator symbols used in the RTL Viewer schematic view.

Table 17-3. Operator Symbols in the RTL Viewer Schematic View (Part 1 of 2)

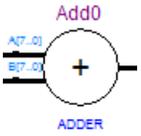
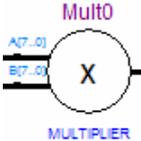
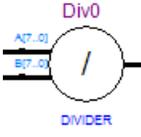
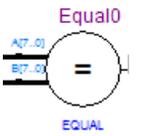
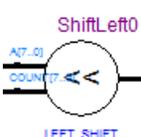
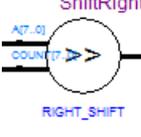
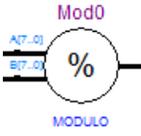
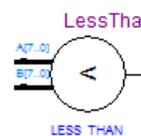
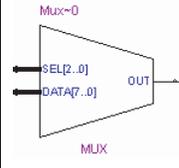
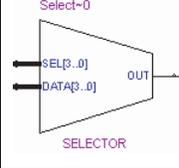
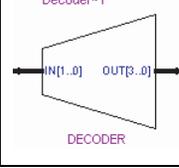
Symbol	Description
	<p>An adder operator: $OUT = A + B$</p>
	<p>A multiplier operator: $OUT = A \times B$</p>
	<p>A divider operator: $OUT = A / B$</p>
	<p>Equals</p>
	<p>A left shift operator: $OUT = (A \ll COUNT)$</p>
	<p>A right shift operator: $OUT = (A \gg COUNT)$</p>
	<p>A modulo operator: $OUT = (A \% B)$</p>
	<p>A less than comparator: $OUT = (A < B : A > B)$</p>

Table 17-3. Operator Symbols in the RTL Viewer Schematic View (Part 2 of 2)

Symbol	Description
	<p>A multiplexer:</p> $\text{OUT} = \text{DATA} [\text{SEL}]$ <p>The data range size is $2^{\text{sel range size}}$</p>
	<p>A selector:</p> <p>A multiplexer with one-hot select input and more than two input signals</p>
	<p>A binary number decoder:</p> $\text{OUT} = (\text{binary_number} (\text{IN}) == x)$ <p>for $x = 0$ to $x = 2^{(n+1)} - 1$</p>

Selecting an Item in the Schematic View

To select an item in the schematic view, ensure that the Selection Tool is enabled in the netlist viewer toolbar (this tool is enabled by default). Click an item in the schematic view to highlight it in red.

Select multiple items by pressing the Shift or Ctrl key while selecting with your mouse. You can also select all nodes in a region by selecting a rectangular box area with your mouse cursor when the Selection Tool is enabled. To select nodes in a box, left-click at one corner of the area you want to select and drag the mouse to the diagonally opposite corner. By default, this highlights and selects all nodes in the selected area (instances, primitives, and pins), but not the nets. The **Viewer Options** dialog box provides an option to select nets. To include nets, right-click in the schematic and click **Viewer Options**. Under **Net Selection**, turn on the **Select entire net when segment is selected** option.

If you enable the **Enable auto hierarchy expansion** option, items selected in the schematic view are automatically selected in the **Netlist Navigator** pane (for more information about how to enable the auto hierarchy expansion option, refer to [“Netlist Navigator Pane” on page 17-15](#)). The folder then expands automatically if it is required to show the selected entry; however, the folder does not collapse automatically when you are not using or you have deselected the entries.

When you select a hierarchy box, node, or port in the schematic view, the item is highlighted in red but none of the connecting nets are highlighted. When you select a net (wire or bus) in the schematic view, all connected nets are highlighted in red. The selected nets are highlighted across all hierarchy levels and pages. Net selection can be useful when navigating a netlist because you see the net highlighted when you traverse between hierarchy levels or pages.

In some cases, when you select a net that connects to nets in other levels of the hierarchy, these connected nets are also highlighted in the current hierarchy. If you prefer that these nets not be highlighted, use the **Viewer Options** dialog box option to highlight a net only if the net is in the current hierarchy. Right-click in the schematic and click **Viewer Options**. In the **Net Selection** section, turn on the **Limit selections to current hierarchy** option.

Moving and Panning in the Schematic View

When the schematic view page is larger than the portion currently displayed, you can use the scroll bars at the bottom and right side of the schematic view to see other areas of the page.

You can also use the Hand Tool to “grab” the schematic page and drag it in any direction. Enable the Hand Tool with the toolbar button. Click and drag to move around the schematic view without using the scroll bars.

In addition to the scroll bars and Hand Tool, you can use the middle-mouse or wheel button to move and pan in the schematic view. Click the middle-mouse or wheel button once to enable the feature. Move the mouse or scroll the wheel to move around the schematic view. Click the middle-mouse or wheel button again to turn the feature off.

Netlist Navigator Pane

The **Netlist Navigator** pane displays the entire netlist in a tree format based on the hierarchical levels of the design. In each level, similar elements are grouped into subcategories. You can use the **Netlist Navigator** pane to traverse through the design hierarchy to view the logic schematic for each level. You can also select an element in the netlist navigator to highlight in the schematic view.



Nodes inside atom primitives are not listed in the **Netlist Navigator** pane.

For each module in the design hierarchy, the **Netlist Navigator** pane displays the applicable elements listed in [Table 17-4](#). Click the “+” icon to expand an element.

Table 17-4. Netlist Navigator Pane Elements (Part 1 of 2)

Elements	Description
Instances	Modules or instances in the design that can be expanded to lower hierarchy levels.
State Machines	State machine instances in the design that can be viewed in the State Machine Viewer.
Primitives	<p>Low-level nodes that cannot be expanded to any lower hierarchy level. These primitives include:</p> <ul style="list-style-type: none"> ■ Registers and gates that you can view in the RTL Viewer when using Quartus II integrated synthesis ■ Logic cell atoms in the Technology Map Viewer or in the RTL Viewer when using a VQM or EDIF from third-party synthesis software <p>In the Technology Map Viewer, you can view the internal implementation of certain atom primitives, but you cannot traverse into a lower-level of hierarchy.</p>

Table 17-4. Netlist Navigator Pane Elements (Part 2 of 2)

Elements	Description
Pins	<p>The I/O ports in the current level of hierarchy.</p> <ul style="list-style-type: none"> ■ Pins are device I/O pins when viewing the top hierarchy level and are I/O ports of the design when viewing the lower-levels. ■ When a pin represents a bus or an array of pins, expand the pin entry in the list view to see individual pin names.
Nets	<p>Nets or wires connecting the nodes. When a net represents a bus or array of nets, expand the net entry in the tree to see individual net names.</p>
Logic Clouds	<p>A group of related combinational logics of a particular source. You can automatically or manually group combinational logics or ungroup logic clouds in your design.</p>

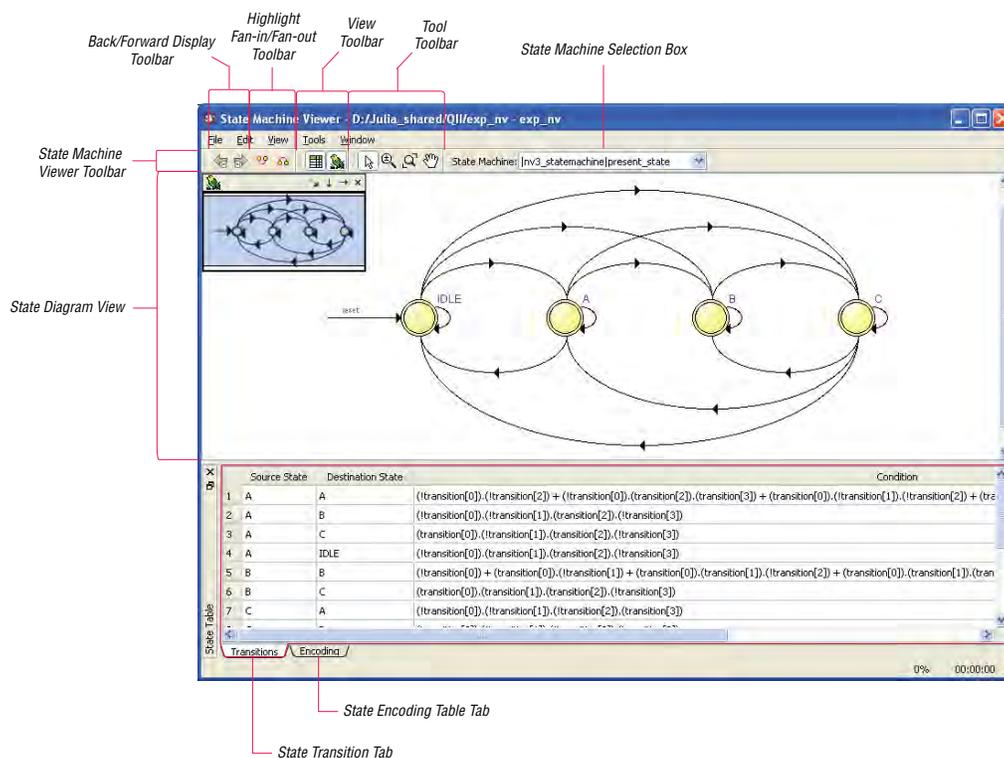
State Machine Viewer

The State Machine Viewer displays a graphical representation of the state machines in your design. You can open the State Machine Viewer in any of the following ways:

- On the Tools menu, point to **Netlist Viewers** and click **State Machine Viewer**.
- Double-click a state machine instance in the RTL Viewer
- Right-click a state machine instance in the RTL Viewer and click **Hierarchy Down**.
- Select a state machine instance in the RTL Viewer, and on the Project menu, point to **Hierarchy** and click **Down**.

Figure 17-5 shows an example of the State Machine Viewer for a simple state machine.

Figure 17-5. The State Machine Viewer



State Diagram View

The state diagram view appears at the top of the State Machine Viewer. It contains a diagram of the states and state transitions.

The nodes that represent each state are arranged horizontally in the state diagram view with the initial state (the state node that receives the reset signal) in the left-most position. Nodes that connect to logic outside of the state machine instance are represented by a double circle. The state transition is represented by an arc with an arrow pointing in the direction of the transition.

When you select a node in the state diagram view, and turn on the **Highlight Fan-in** or **Highlight Fan-out** command from the View menu or the State Machine Viewer toolbar, the respective fan-in or fan-out transitions from the node are highlighted in red.



An encrypted block with a state machine displays encoding information in the state encoding table, but does not display a state transition diagram or table.

State Transition Table

The state transition table on the **Transitions** tab at the bottom of the State Machine Viewer displays the condition equation for each state transition. Each row in the table represents a transition (each arc in the state diagram view). The table has the following columns:

- **Source State**—the name of the source state for the transition
- **Destination State**—the name of the destination state for the transition
- **Condition**—the condition equation that causes the transition from source state to destination state

To see all of the transitions to and from each state name, click the appropriate column heading to sort on that column.

The text in each column is left-aligned by default; to change the alignment and to make it easier to see the relevant part of the text, right-click the column and click **Align Right**. To revert to left alignment, click **Align Left**.

Click in any cell in the table to select it. To select all cells, right-click in the cell and click **Select All**; or, on the Edit menu, click **Select All**. To copy selected cells to the clipboard, right-click the cells and click **Copy Table**; or, on the Edit menu, point to **Copy** and click **Copy Table**. You can paste the table into any text editor as tab-separated columns.

State Encoding Table

The state encoding table on the **Encoding** tab at the bottom of the State Machine Viewer displays encoding information for each state transition.

To view state encoding information in the State Machine Viewer, you must synthesize your design with the **Start Analysis & Synthesis** command. If you have only elaborated your design with the **Start Analysis & Elaboration** command, the encoding information is not displayed.

Selecting an Item in the State Machine Viewer

You can select and highlight each state node and transition in the State Machine Viewer. To select a state transition, click the arc that represents the transition.

When you select a node or transition arc in the state diagram view, the matching state node or equation conditions in the state transition table are highlighted; conversely, when you select a state node or equation condition in the state transition table, the corresponding state node or transition arc is highlighted in the state diagram view.

Switching Between State Machines

A design may contain multiple state machines. To choose which state machine to view, use the **State Machine** selection box located at the top of the State Machine Viewer. Click in the drop-down box and select the necessary state machine.

Global Options

The **Options** dialog box allows you to customize the following settings:

- Display Settings

- Colors
- Fonts
- Tracing
- Customize View
- Shortcut Commands

Display Settings

If you want to customize the display settings for your preferred viewing, you can direct the RTL Viewer and Technology Map Viewer to adjust the settings in the schematic view. To adjust the display settings, on the Tools menu, click **Options**. Under **Display Settings**, you can select the options to customize your display.

You can divide the schematic representation of a large design into multiple pages. Dividing the display of the schematic into multiple pages does not affect your design, and only controls the number of elements per page. Under **Display Settings**, the **Nodes per page** option allows you to specify the number of nodes displayed per page. The default value is 50 nodes; however, you can view from one to 1,000 nodes per page. The **Ports per page** option allows you to specify the number of ports (or pins) displayed per page. The default value is 1,000 ports (or pins); the range is 1 to 2,000 ports (or pins). The netlist viewers partition your design into a new page if either the number of nodes or the number of ports exceeds the limit you specified. Occasionally, the number of ports displayed on the page might exceed the limit you specified, depending on the configuration of nodes on the page. If you turned on the **Display boundary around hierarchy levels** option and the total number of nodes or ports in the hierarchy exceeds the value you specified for the **Nodes per page** or **Ports per page** options, the netlist viewers display the boundary as a hierarchy port connector (refer to [Table 17-1 on page 17-10](#)).

To display net names in your schematic, turn on the **Show Net Name** option. If you turn on this option, the schematic view refreshes automatically to display the net names. To show node names in the schematic view, turn on the **Show node name** option. To change the value formatting, select the necessary format in the **Constant signal format** list.

To view highlighting around the design element in range of the mouse pointer, turn on the **Enable rollover** option. To more easily navigate through your design hierarchy, you can direct the netlist viewers to automatically expand the hierarchy list in **Netlist Navigator** pane and highlight the design element you selected in the schematic view. This automatic expansion and selection of corresponding design elements is useful when you have a complex schematic that spans multiple display pages in the netlist viewers. To use the automatic expansion and selection feature, turn on the **Enable auto hierarchy expansion** option.

To enable the radial menu, turn on the **Enable radial menu** option. The radial menu is an octagonal menu with eight commands from which you can choose. This menu provides a quick way to perform any of the commands with a single click whenever you are in the schematic view.

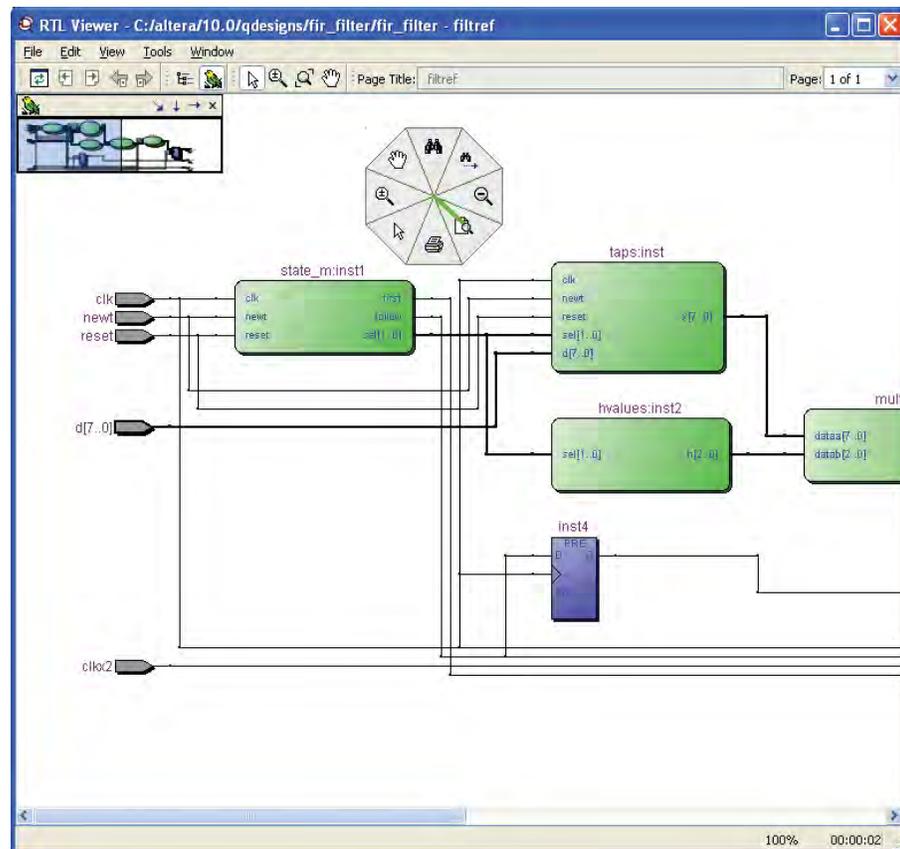
To open the radial menu, right-click and hold anywhere in the schematic view and wait for the menu to appear. By default, the menu appears after 0.2 seconds. The radial menu appears with the mouse pointer always at the center point. The center point of the menu is a non-trigger boundary in which no command is started.

To invoke the necessary command, hold down the right mouse button, drag the mouse onto the command, and then press the left mouse button. If you decide not to trigger any command after the radial menu appears, press the Esc key or drag the pointer back into the center point and release the mouse button.

To change the delay time before the radial menu appears, select the necessary interval time in the drop-down list for **Delay showing radial menu for**. The default delay is 0.2 seconds.

Figure 17-6 shows the radial menu.

Figure 17-6. Radial Menu



Tracing

If you want to filter information from the schematic view of your design to isolate specific design elements for further inspection, or if you want to expand specific design elements, you can direct the RTL Viewer and the Technology Map Viewer to adjust the elements the viewers show in the schematic view. To adjust the filtering and expansion settings in the RTL Viewer and Technology Map Viewer, on the Tools menu, click **Options**. Under **Tracing**, you can select the options to control filtering and expansion settings.

For all filtering commands, the netlist viewers stop tracing through the netlist when they reach one of the following objects:

- A pin

- A specified number of filtering levels, counting from the selected node or port
- A register

To specify the number of filtering levels, set the **Number of filtering levels** option to specify the number of levels to expand. You can specify a value from one to 100.

To enable the **Stop filtering at register** option, turn on the **Stop filtering at register** option. You can filter across hierarchies when you turn on the **Filter across hierarchy** option.

By default, the filtered schematic shows all possible connections between the nodes shown in the schematic. To remove the connections that are not directly part of the path that was traced to generate a filtered netlist, turn off the **Show all connections between nodes** option.

To set the amount of logic you want to expand, set the **Number of expansion levels** option to specify the number of levels to expand. You can specify a range from one to 100 levels. You can also set the **Stop expanding at register** option to specify whether netlist expansion should stop when a register is reached.

Customize View

If you want to customize the schematic display for better viewing and to speed up your debugging process, you can direct the RTL Viewer and the Technology Map Viewer to remove fan-out free nodes, show simplify logic, group or ungroup related nodes, and group combinational logic into a logic cloud. To adjust the options that control the schematic display in the RTL Viewer and the Technology Map Viewer, on the Tools menu, click **Options**. Under **Customize View**, you can select the options to customize your view. These options are also available in the **Customize View** tab of the **RTL/Technology Map Viewer Options** dialog box. To open the dialog box, right-click in the schematic and click **Viewer Options**.



When you change settings, the list of previously viewed pages is cleared. The settings are revision-specific, so different revisions can have different settings.

To remove fan-out free registers from your schematic display, turn on the **Remove registers without fan-out** option. To remove all single-input nodes and merge a chain of equivalent combinational gates that have direct connections (without inversion in between) into a single multiple-input gate, turn on the **Show simplified logic** option.

To group all related nodes into a single node, turn on the **Group all related nodes** option. You can manually group or ungroup any nodes by right-clicking the selected nodes in the schematic and selecting **Group Related Nodes** or **Ungroup Selected Nodes**. To group combinational logic into logic clouds, turn on the **Group combinational logic into logic cloud** option.



Customize Logic and Customize Group options are available for the RTL Viewer, whereas only the Customize Group option is available for the Technology Map Viewer.

Shortcut Commands

You can choose eight commands to appear on the radial menu from a list of 18 available commands. To customize the command list on the menu, first launch the RTL Viewer or the Technology Map Viewer. On the Tools menu, select **Options**. On the **Shortcut Commands** category list, drag and drop the icon under **Shortcut buttons** into any region under **Shortcut commands popup**. You can click the icon under **Shortcut buttons** to see its description.

Navigating the Schematic View

This section describes methods to navigate through the pages and hierarchy levels in the schematic view of the RTL Viewer and the Technology Map Viewer.

Traversing and Viewing the Design Hierarchy

You can open different hierarchy levels in the schematic view from the **Netlist Navigator** pane (refer to “[Netlist Navigator Pane](#)” on page 17-15), or the **Hierarchy Up** and **Hierarchy Down** commands (Shortcut menu) in the schematic view.

Use the **Hierarchy Down** command to go down in an instance’s hierarchy, and open a lower-level schematic showing the internal logic of the instance. Use the **Hierarchy Up** command to go up in hierarchy or collapse a lower-level hierarchy, and open the parent higher level hierarchy. When the Selection Tool is selected, the appropriate option is available when your mouse pointer is located over an area of the schematic view that has a corresponding lower or higher level hierarchy.

The mouse pointer changes as it moves over different areas of the schematic to indicate whether you can move up, down, or both up and down in the hierarchy ([Figure 17-7](#)). To open the next hierarchy level, right-click in that area of the schematic and click **Hierarchy Down** or **Hierarchy Up**, as appropriate, or double-click in that area of the schematic.

Figure 17-7. Mouse Pointers Indicate How to Traverse Hierarchy



Flattening the Design Hierarchy

You can flatten the design hierarchy to view the design without hierarchical boundaries. To flatten the hierarchy from the current level and all lower-level hierarchies of the current design hierarchy, right-click in the schematic and click **Flatten Netlist**. To flatten the entire design, choose **Flatten Netlist** from the top-level schematic of the design.

Viewing the Contents of a Design Hierarchy in the Current Schematic

You can use the **Display Content** and **Hide Content** (Shortcut menu) commands to show or hide a lower hierarchy level for a specific instance in the schematic for the current hierarchy level.

To display the lower hierarchy netlist of an instance on the same schematic as the remaining logic in the currently viewed netlist, right-click the selected instance and click **Display Content**.

To hide all of the lower hierarchy logic of a hierarchy box into a closed instance, right-click the selected instance and click **Hide Content**.

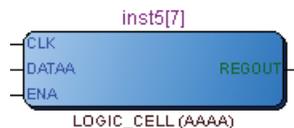
Viewing Contents of Atom Primitives

In the Technology Map Viewer, you can view the contents of certain device atom primitives to see their underlying implementation details. For logic cell (LCELL) atoms in Arria® GX, Cyclone® series, MAX® II, and Stratix series of devices, you can view LUTs, registers, and logic gates. For I/O atoms in the Arria GX, Cyclone series, HardCopy® IV, and Stratix series of devices, you can view registers and logic gates.

In addition, you can view the implementation of RAM and DSP blocks in certain devices in the RTL Viewer or Technology Map Viewer. You can view the implementation of RAM blocks in the Arria GX, Cyclone series, and Stratix series of devices. You can view the implementation of DSP blocks only in Arria GX and Stratix series of devices.

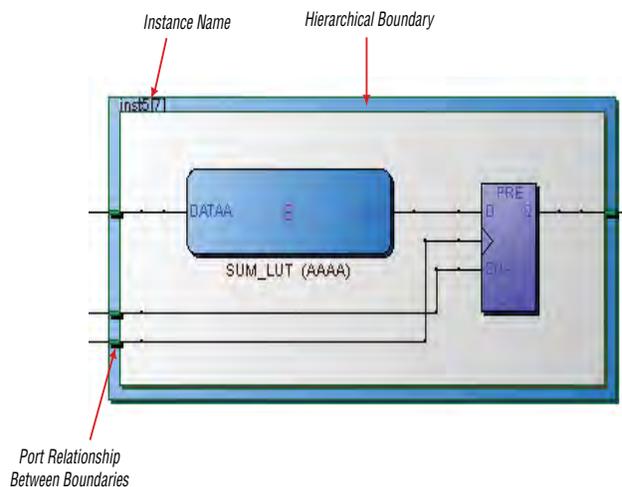
If you can view the contents of an atom instance, the internal contents are shown in blue in the schematic view (Figure 17-8).

Figure 17-8. Instance That Can Be Expanded to View Internal Contents



To view the contents of one or more atom primitive instances, select the necessary atom instances. Right-click a selected instance and click **Display Content**. You can also double-click the necessary atom instance to view the contents. Figure 17-9 shows an expanded version of the instance in Figure 17-8.

Figure 17-9. Internal Contents of the Atom Instance in Figure 17-8.



To hide the contents (and revert to the compact format), select and right-click the atom instance or instances, and click **Hide Content**.



In the schematic view, the internal details in an atom instance cannot be selected as individual nodes. Any mouse action on any of the internal details is treated as a mouse action on the atom instance.

Viewing the Properties of Instances and Primitives

You can view the properties of an instance or primitive using the **Properties** dialog box. To view the properties of an instance or primitive in the RTL Viewer or Technology Map Viewer, right-click the node and click **Properties**.

The **Properties** dialog box contains the following information about the selected node:

- The parameter values of an instance.
- The active level of the port (for example, active high or active low). An active low port is denoted with an exclamation mark “!”.
- The port’s constant value (for example, V_{CC} or GND). Table 17-5 describes the possible value of a port.

Table 17-5. Possible Port Values

Value	Description
V_{CC}	The port is not connected and has V_{CC} value (tied to V_{CC})
GND	The port is not connected and has GND value (tied to GND)
--	The port is connected and has value (other than V_{CC} or GND)
Unconnected	The port is not connected and has no value (hanging)

In the LUT of a logic cell (LCELL), the **Properties** dialog box contains the following additional information:

- The schematic of the LCELL
- The Truth Table representation of the LCELL
- The Karnaugh map representation of the LCELL

Viewing LUT Representations in the Technology Map Viewer

You can view different representations of a LUT by right-clicking the selected LUT and clicking **Properties**. This feature is supported for the Arria GX, Cyclone series, MAX II, and Stratix series of devices only. You can view the LUT representations in the following three tabs in the **Properties** dialog box:

- The **Schematic** tab (Figure 17-10) shows the equivalent gate representations of the LUT.
- The **Truth Table** tab (Figure 17-11) shows the truth table representations.
- The **Karnaugh Map** tab (Figure 17-12) shows the Karnaugh map representations of the LUT. The Karnaugh map supports up to 6 input LUTs.

For more information about the **Ports** tab, refer to “Viewing the Properties of Instances and Primitives” on page 17-24.

Figure 17-10. Schematic Tab

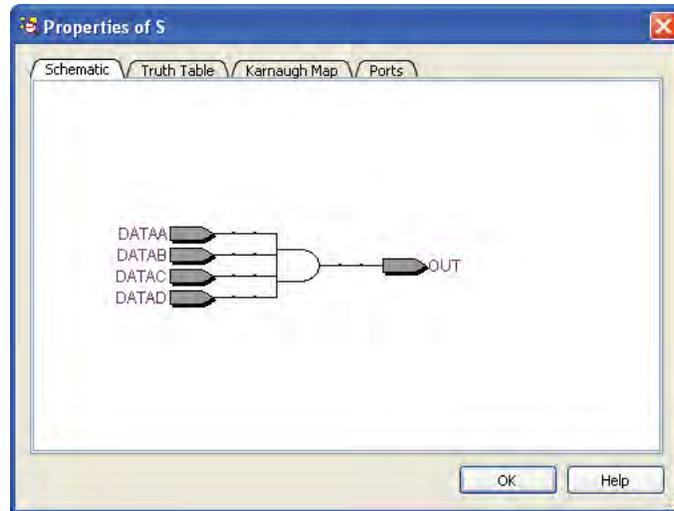
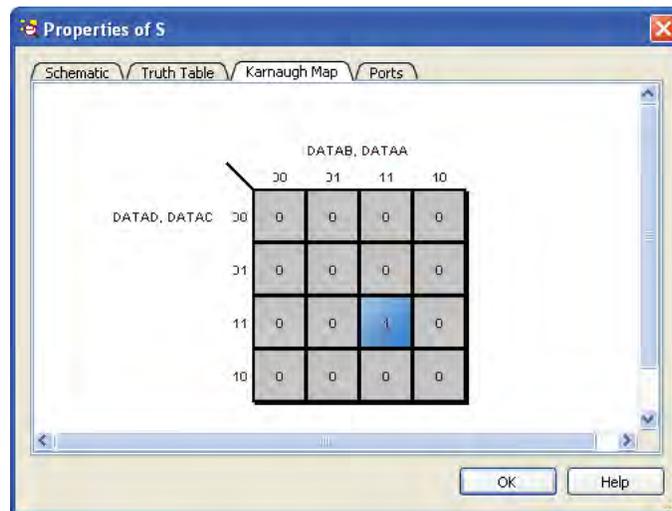


Figure 17-11. Truth Table Tab

The screenshot shows the 'Properties of S' dialog box with the 'Truth Table' tab selected. The truth table has 12 rows and 6 columns. The columns are labeled DATAD, DATAC, DATAB, DATAA, and OUT. The rows represent all possible combinations of the four input variables.

	DATAD	DATAC	DATAB	DATAA	OUT
1	0	0	0	0	0
2	0	0	0	1	0
3	0	0	1	0	0
4	0	0	1	1	0
5	0	1	0	0	0
6	0	1	0	1	0
7	0	1	1	0	0
8	0	1	1	1	0
9	1	0	0	0	0
10	1	0	0	1	0
11	1	0	1	0	0
12	1	0	1	1	0

Figure 17-12. Karnaugh Map Tab



Grouping Combinational Logic into Logic Clouds

The following sections describe how to group combinational logic into logic clouds.

 For the definition of a logic cloud, refer to [Table 17-1 on page 17-10](#).

Logic Clouds in the RTL Viewer

You can automatically group all combinational logic nodes in your design into logic clouds. For more information about how to group combinational logic clouds, refer to [“Customize View” on page 17-21](#). [Figure 17-13](#) and [Figure 17-14](#) show the schematic before and after the combinational logic grouping operation in the RTL Viewer.

Logic Clouds in the Technology Map Viewer

In the Technology Map Viewer, the **Group combinational logic into logic clouds** option is supported for Cyclone II, HardCopy, and Stratix II devices only. For more information about how to group combinational logic clouds, refer to “[Customize View](#)” on page 17-21.

Figure 17-13. Schematic Before Combinational Logic Grouping

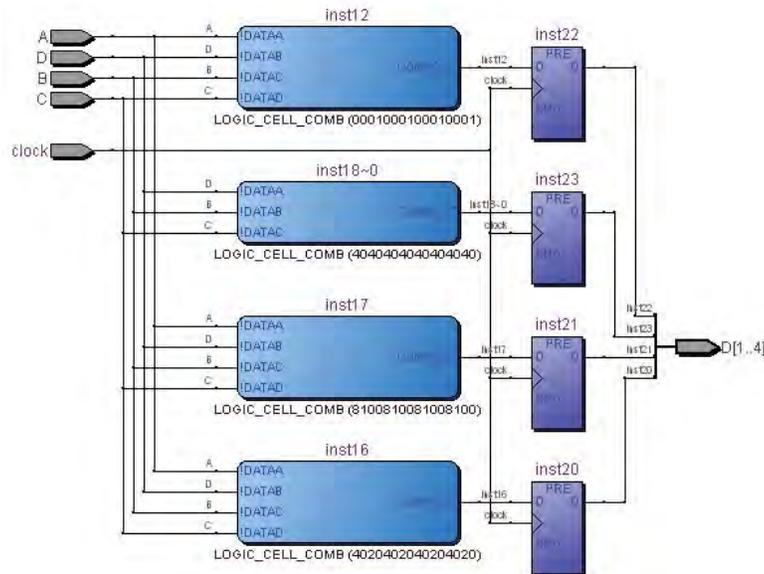
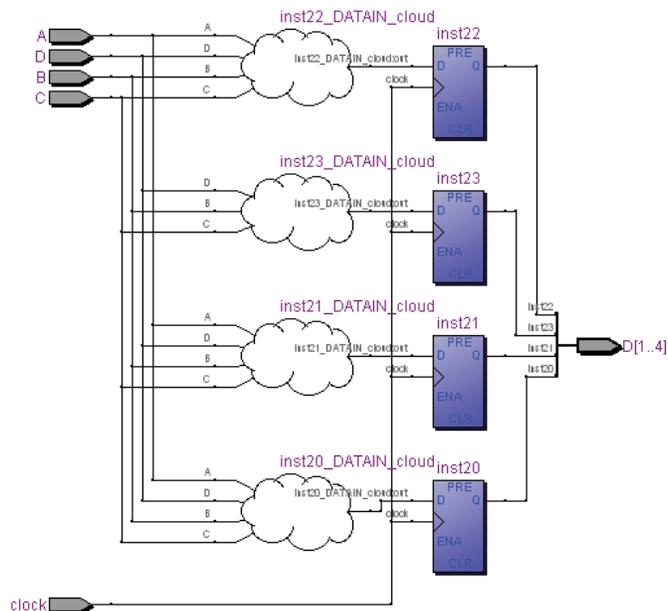


Figure 17-14. Schematic After Combinational Logic Grouping



Grouping and Ungrouping Logic Clouds

To group logic nodes into a logic cloud manually, right-click the selected node or input port and click **Group source logic into logic cloud**. To ungroup a logic cloud manually, right-click the selected logic cloud and click **Ungroup source logic from logic cloud**. You can also ungroup a logic cloud manually by double-clicking the selected logic cloud. These options are not available if the nodes cannot be grouped.

Changing the Constant Signal Value Formatting

The constant signal value is highlighted in gray in the schematic view. By default, the value is displayed in hexadecimal format, but you can also choose binary or decimal format. For more information about changing the value formatting, refer to [“Display Settings” on page 17-19](#).

Changing the format affects all constant signal values throughout the schematic. For descriptions of constant signal values in the schematic, refer to [Table 17-3 on page 17-13](#).

Zooming and Magnification

You can control the magnification of your schematic on the View menu, with the Zoom Tool in the toolbar, or the Ctrl key and mouse wheel button, as described in this section.

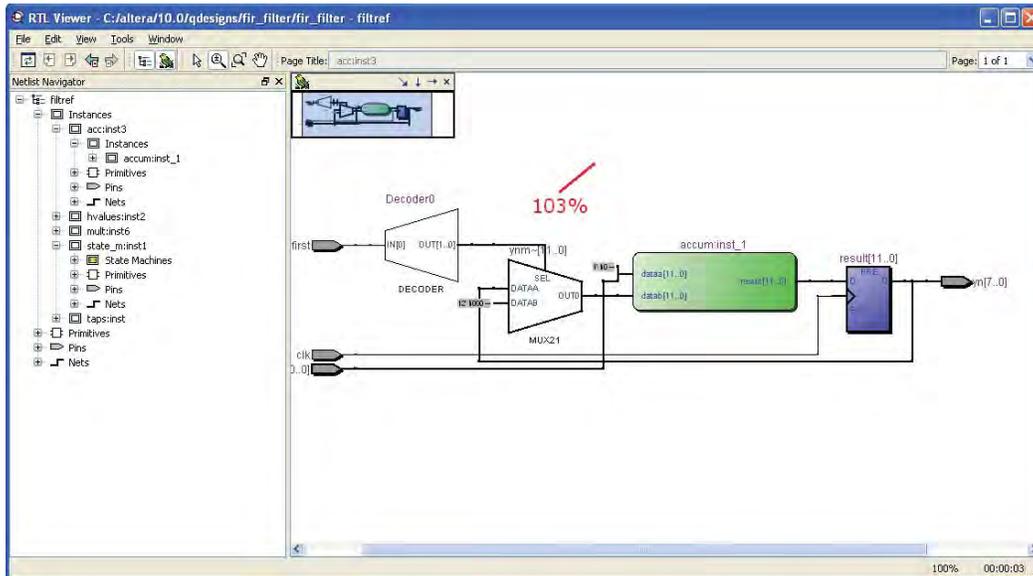
By default, the netlist viewer displays most pages sized to fit in the window. If the schematic page is very large, the schematic is displayed at the minimum zoom level, and the view is centered on the first node. Click **Zoom In** to view the image at a larger size, and click **Zoom Out** to view the image (when the entire image is not displayed) at a smaller size. The **Zoom** command allows you to specify a magnification percentage (100% is considered the normal size for the schematic symbols).

The **Fit Selection in Window** command zooms in on the selected nodes in a schematic to fit in the window. Use the Selection Tool to select one or more nodes (instances, primitives, pins, and nets), then click **Fit Selection in Window** to enlarge the area covered by the selection. This feature is helpful when you want to see a particular element in a large schematic. After you select a node, you can easily zoom in to view the particular node. You can temporarily enlarge a portion of the schematic with the magnifying glass tool in the toolbar.

You can also use the Zoom Tool on the netlist viewer toolbar to control magnification in the schematic view. When you select the Zoom Tool in the toolbar, clicking in the schematic zooms in and centers the view on the location you clicked. Right-click in the schematic to zoom out and center the view on the location you clicked. When you select the Zoom Tool, you can also zoom in to a certain portion of the schematic by selecting a rectangular box area with your mouse cursor. The schematic is enlarged to show the selected area.

Alternatively, you can specify the magnification percentage by right-clicking the necessary area and dragging the mouse to the right to zoom in or to the left to zoom out with the Zoom Tool. You will see a red line with the zoom percentage above it. The zoom percentage is proportional to the length of the green line (Figure 17-15). Release the mouse button at the necessary zoom percentage.

Figure 17-15. Dragging the Mouse Pointer to Change Zoom Percentage



By default, the netlist viewers maintain the zoom level when filtering on the schematic (refer to “Filtering in the Schematic View” on page 17-33). To change the behavior so that the zoom level is always reset to “Fit in Window,” on the Tools menu, click **Options**. In the **Category** list, select **Netlist Viewers**, and turn off **Maintain zoom level**.

Schematic Debugging and Tracing Using the Bird’s Eye View

Viewing the entire schematic can be useful when debugging and tracing through a large netlist. The Quartus II software allows you to quickly navigate to a specific section of the schematic using the Bird’s Eye View feature, which is available in the RTL Viewer and Technology Map Viewer.

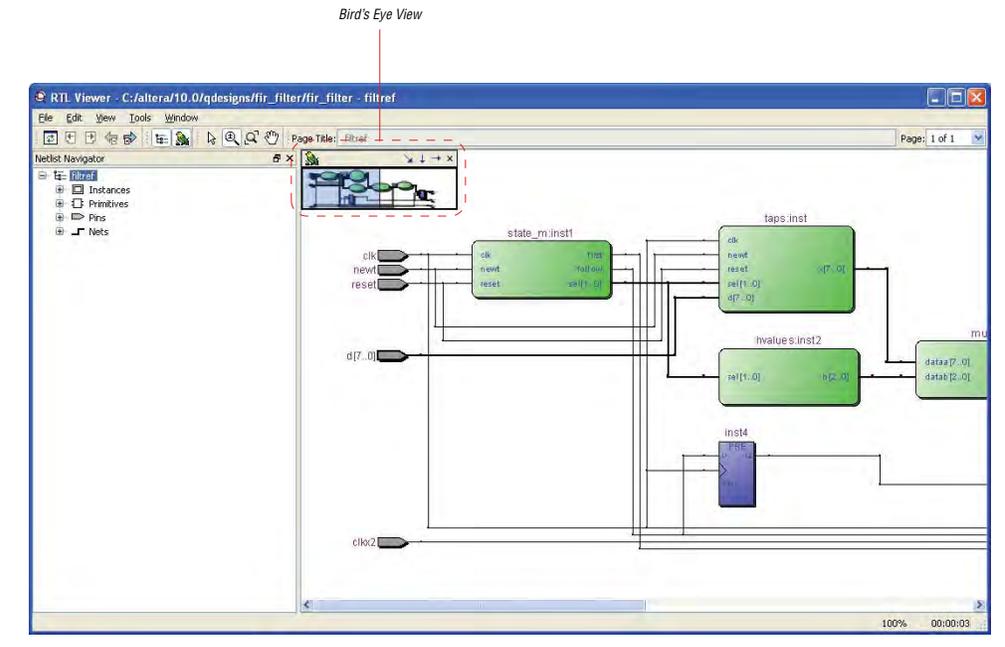
The Bird’s Eye View shows the current area of interest. Select the necessary area by clicking and dragging the indicator or right-clicking to form a rectangular box around the necessary area. You can also click and drag the rectangular box to move around the schematic. To open the Bird’s Eye View, on the View menu, click **Bird’s Eye View**, or click on the **Bird’s Eye View** icon in the toolbar (Figure 17-16).

Figure 17-16. Bird’s Eye View Icon



Figure 17-17 shows the Bird's Eye View window in the schematic diagram.

Figure 17-17. Bird's Eye View Window



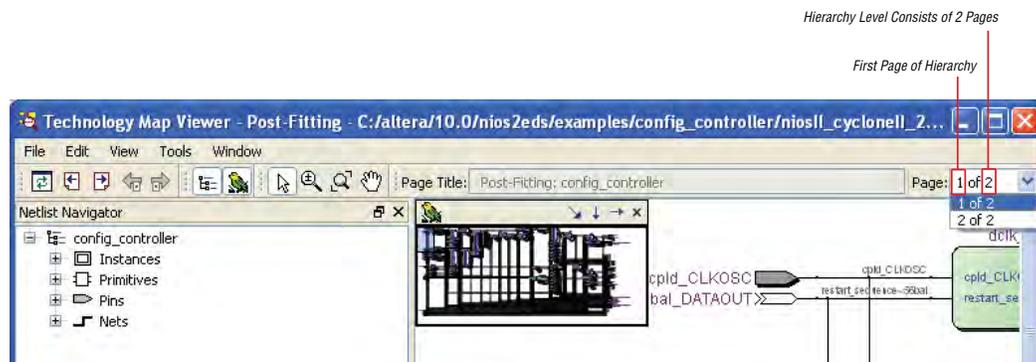
Partitioning the Schematic into Pages

For large design hierarchies, the RTL Viewer and Technology Map Viewer partition your netlist into multiple pages in the schematic view. For more information about controlling how much of the design is visible on each page, refer to [“Display Settings” on page 17-19](#).

When a hierarchy level is partitioned into multiple pages, the title bar for the schematic window indicates which page is displayed and how many total pages exist for this level of hierarchy (shown in the format:

Page <current page number> of <total number of pages>), as shown in [Figure 17-18](#).

Figure 17-18. RTL Viewer Title Bars Indicating Page Number Information



When you change the number of nodes or ports per page, the change applies only to new pages that are shown or opened in the netlist viewer. To refresh the current page so that it displays the changed number of nodes or ports, click the **Refresh** button on the toolbar.

Moving Between Schematic Pages

To move to another schematic page, on the View menu, click **Previous Page** or **Next Page**, or click the **Previous Page** icon or the **Next Page** icon on the netlist viewer toolbar.

To go to a particular page of the schematic, on the Edit menu, click **Go To**, or right-click in the schematic view and click **Go To**. In the **Page** list, select the necessary page number. You can also go to a particular page by selecting the necessary page number from the pull-down list on the top right of the netlist viewer.

Moving Back and Forward Through Schematic Pages

To return to the previous view after changing the page view, click **Back** on the View menu, or click the **Back** icon on the netlist viewer toolbar. To go to the next view, click **Forward** on the View menu, or click the **Forward** icon on the netlist viewer toolbar.

 You can go forward only if you have not made any changes to the view since going back. Use the **Back** and **Forward** commands to switch between page views. These commands do not undo an action, such as selecting a node.

Following Nets Across Schematic Pages

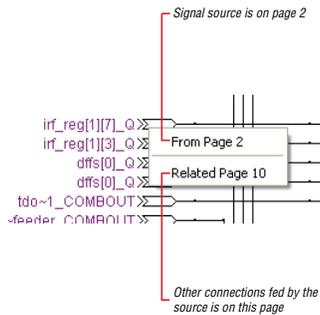
Input and output connectors indicate nodes that connect across pages of the same hierarchy. Right-click a connector to display a menu of commands that trace the net through the pages of the hierarchy.

 After you right-click to follow a connector port, the netlist viewer opens a new page, which centers the view on the particular source or destination net using the same zoom factor as the previous page. To trace a specific net to the new page of the hierarchy, Altera recommends that you first select the necessary net, which highlights it in red, before you right-click to traverse pages.

Input Connectors

Figure 17-19 shows an example of the menu that appears when you right-click an input connector. The **From** command opens the page containing the source of the signal. The **Related** commands, if applicable, open the specified page containing another connection fed by the same source.

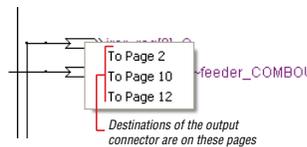
Figure 17-19. Input Connector Shortcut Menu



Output Connectors

Figure 17-20 shows an example of the menu that appears when you right-click an output connector. The **To** command opens the specified page that contains a destination of the signal.

Figure 17-20. Output Connector Shortcut Menu



Go to Net Driver

To locate the source of a particular net in the schematic view, right-click the net, point to **Go to Net Driver** and click **Current page**, **Current hierarchy**, or **Across hierarchies**. Table 17-6 lists the **Go to Net Driver** commands.

Table 17-6. Go to Net Driver Commands

Command	Action
Current page	Locates the source or driver on the current page of the schematic only.
Current hierarchy	Locates the source in the current level of hierarchy, even if the source is located on another page of the netlist schematic.
Across hierarchies	Locates the source across hierarchies until the software reaches the source at the top hierarchy level.

The schematic view opens the correct page of the schematic, if required, and adjusts the centering of the page so that you can see the net source. The schematic shows the default page for the net driver. The view is unfiltered, so no filtering results are kept.

Filtering in the Schematic View

Filtering allows you to filter out nodes and nets in your netlist to view only the logic elements of interest to you.

You can filter your netlist by selecting hierarchy boxes, nodes, ports of a node, nets, or states in a state machine that are part of the path you want to see. The following filter commands are available:

- **Sources**—Displays the sources of the selection
- **Destinations**—Displays the destinations of the selection
- **Sources & Destinations**—Displays the sources and destinations of the selection
- **Selected Nodes and Nets**—Displays only the selected nodes and nets with the connections between them
- **Between Selected Nodes**—Displays nodes and connections in the path between the selected nodes
- **Bus Index**—Displays the sources or destinations for one or more indices of an output or input bus port

To filter your netlist, select a hierarchy box, node, port, net, or state node, right-click in the window, point to **Filter** and click the appropriate filter command. The netlist viewer generates a new page showing the netlist that remains after filtering.

When filtering in a state diagram in the State Machine Viewer, sources and destinations refer to the previous and next transition states or paths between transition states in the state diagram. The transition table and encoding table also reflect the filtering.

You can go back to the netlist page before it was filtered using the **Back** command, as described in “[Moving Back and Forward Through Schematic Pages](#)” on page 17-31.

 When viewing a filtered netlist, clicking an item in the **Netlist Navigator** pane causes the schematic view to display an unfiltered view of the appropriate hierarchy level. You cannot use the **Netlist Navigator** pane to select items or navigate in a filtered netlist.

Filter Sources Command

To filter out all but the source of the selected item, right-click the item, point to **Filter**, and click **Sources**. The selected object type determines what is displayed, as listed in [Table 17-7](#) and shown in [Figure 17-21](#).

Table 17-7. Selected Objects Determine Filter Sources Display

Selected Object	Result Shown in Filtered Page
Node or hierarchy box	Shows all the sources of the node's input ports. For an example, refer to Figure 17-21 .
Net	Shows the sources that feed the net.
Input port of a node	Shows only the input source nodes that feed this port.
Output port of a node	Shows only the selected node.
State node in a state machine	Shows the states that feed the selected state (previous transition states).

Filter Destinations Command

To filter out all but the destinations of the selected node or port as outlined in Table 17-8 and shown in Figure 17-21, right-click the node or port, point to **Filter**, and click **Destinations**.

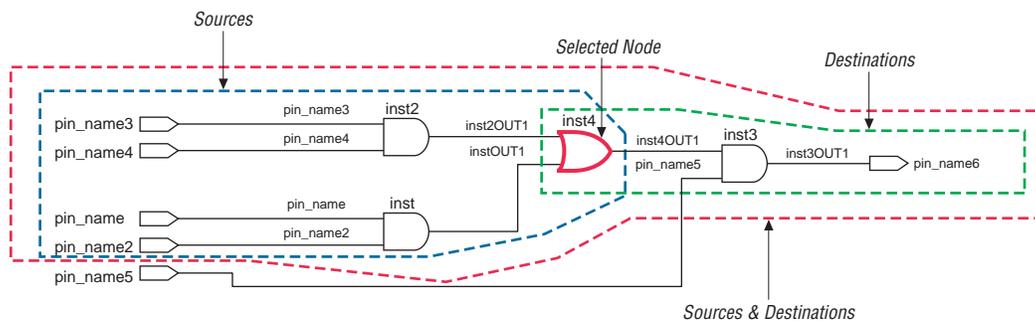
Table 17-8. Selected Objects Determine Filter Destinations Display

Selected Object	Result Shown in Filtered Page
Node or hierarchy box	Shows all the destinations of the node's output ports. For an example, refer to Figure 17-21.
Net	Shows the destinations fed by the net.
Input port of a node	Shows only the selected node.
Output port of a node	Shows only the fan-out destination nodes fed by this port.
State node in a state machine	Shows the states that are fed by the selected states (next transition states).

Filter Sources and Destinations Command

The **Sources & Destinations** command is a combination of the **Sources** and **Destinations** filtering commands, in which the filtered page shows the sources and the destinations of the selected item. To select this option, right-click the necessary object, point to **Filter**, and click **Sources & Destinations**. Refer to the example in Figure 17-21.

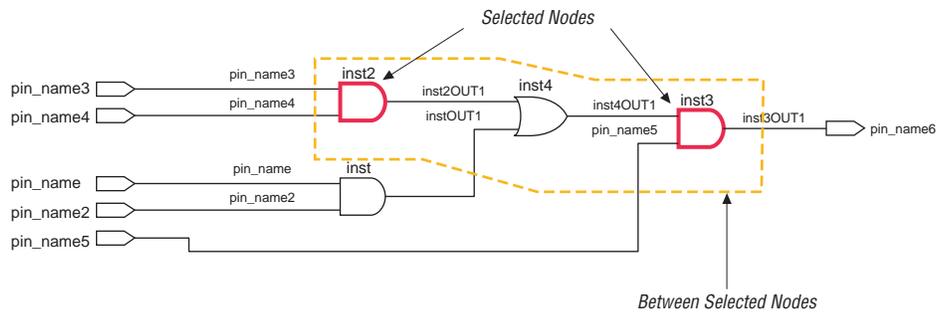
Figure 17-21. Sources, Destinations, and Sources and Destinations Filtering for inst4



Filter Between Selected Nodes Command

To show the nodes in the path between two or more selected nodes or hierarchy boxes, right-click the necessary object, point to **Filter**, and click **Between Selected Nodes**. For this option, selecting a port of a node is the same as selecting the node. For an example, refer to [Figure 17-22](#).

Figure 17-22. Between Selected Nodes Filtering Between inst2 and inst3



Filter Selected Nodes and Nets Command

To create a filtered page that shows only the selected nodes, nets, or both, and, if applicable, the connections between the selected nodes, nets, or both, right-click the necessary object, point to **Filter**, and click **Selected Nodes & Nets**. Figure 17-23 shows a schematic with several nodes selected.

Figure 17-23. Using Selected Nodes and Nets to Select Nodes

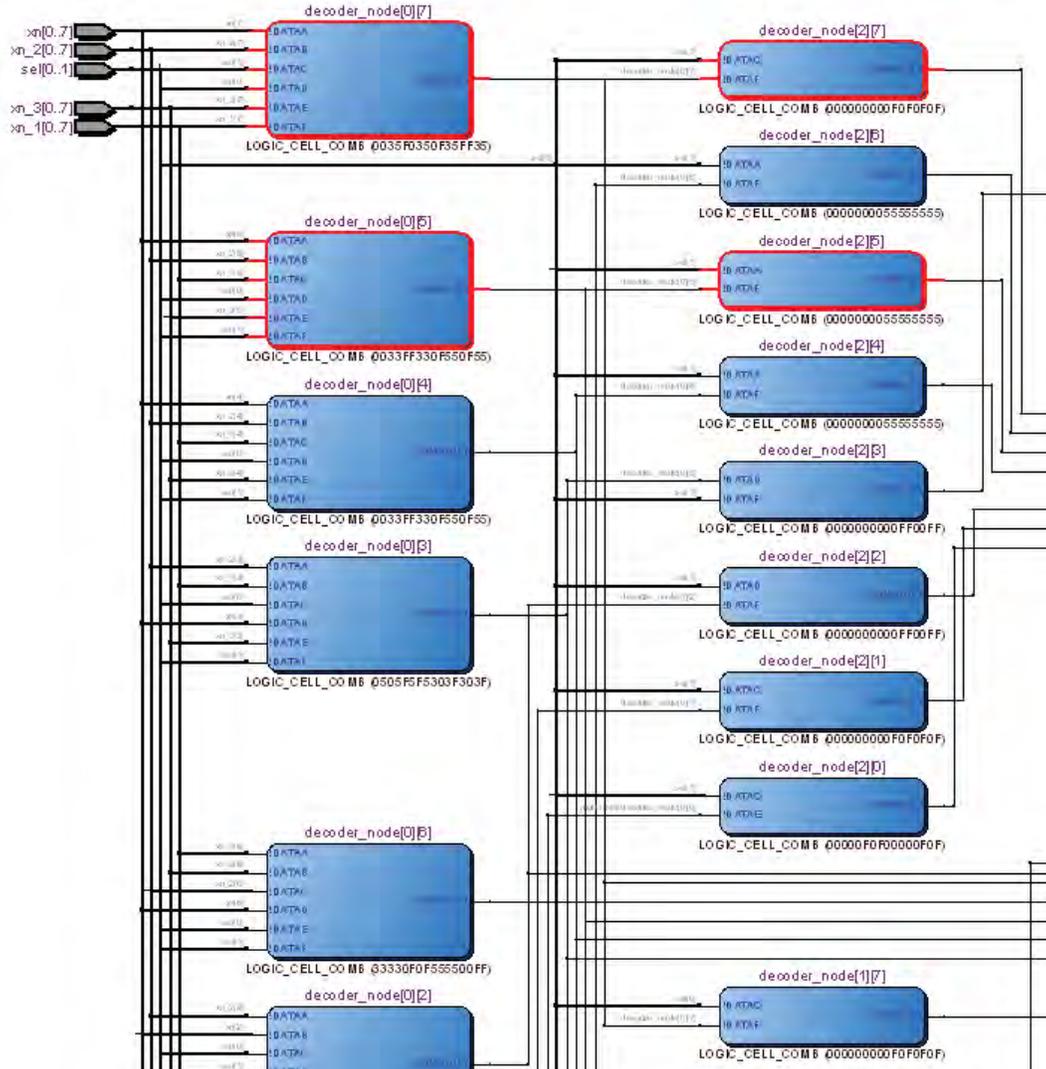
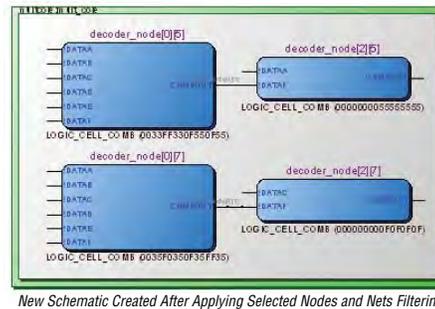


Figure 17-24 shows the schematic after filtering. If you select a net, the filtered page shows the immediate sources and destinations of the selected net.

Figure 17-24. Selected Nodes and Nets Filtering on Figure 17-23 Schematic



Filter Bus Index Command

To show the path related to a specific index of a bus input or output port in the RTL Viewer, right-click the port, point to **Filter**, and click **Bus Index**. The **Select Bus Index** dialog box allows you to select the indices of interest.

Filter Command Processing

The options to control filtering are available in the **Tracing** tab of the **RTL/Technology Map Viewer Options** dialog box. Right-click in the schematic view and click **Viewer Options** to open the dialog box. For more information about filter command processing, refer to “[Tracing](#)” on page 17-20.

Filtering Across Hierarchies

The filtering commands display nodes in all hierarchies by default. When the filtered path passes through levels of hierarchy on the same schematic page, green hierarchy boxes group the logic and show the hierarchy boundaries. A green rectangular symbol that appears on the border represents the port relationship between two different hierarchies (Figure 17-25 and Figure 17-26).

For more information about how to control filtering, refer to “[Tracing](#)” on page 17-20.

For more information about how to disable the box hierarchy display, refer to “[Display Settings](#)” on page 17-19.



Netlists of the same hierarchy displayed over more than one page are not grouped with a box. Filtering and expanding on a blue atom primitive does not trace the underlying netlist, even when **Filter across hierarchy** is enabled.

Figure 17-25 and Figure 17-26 show examples of filtering across hierarchical boundaries. Figure 17-25 shows a smaller example of an input port of the taps instance after the **Sources** filter is applied, in which the input port of the lower-level hierarchical block connects directly to an input pin of the design. The name of the instance appears in the green border and as a tooltip when you move your mouse pointer over the instance.

Figure 17-25. Filtering Across Hierarchical Boundaries

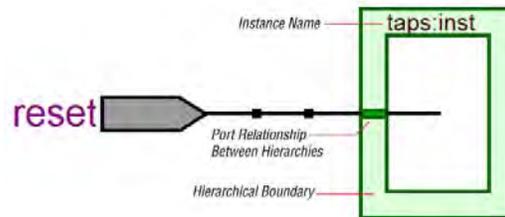
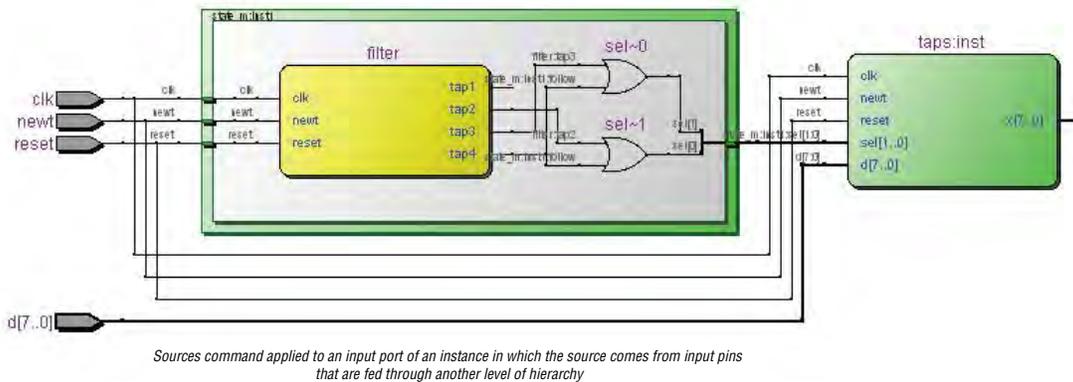


Figure 17-26 shows a larger example of an input port of an instance after the **Sources** filter is applied, in which the source comes from input pins that are fed through another level of hierarchy.

Figure 17-26. Filtering Across Hierarchical Boundaries



Expanding a Filtered Netlist

After a netlist is filtered, some ports might not have connections displayed because their connections are not part of the main path through the netlist. Two expansion features, immediate expansion and the **Expand** command, allow you to add the fan-in or fan-out signals of these ports to the schematic display of a filtered netlist.

You can immediately expand any port whose connections are not displayed. When you double-click that port in the filtered schematic, one level of logic is expanded.

To expand more than one level of logic, right-click the port and click the **Expand** command. This command expands logic from the selected port by the amount specified in **Viewer Options**. To set these options, right-click in the schematic view and click **Viewer Options**. In the **Expansion** section, set the **Number of expansion levels** option to specify the number of levels to expand (the default value is 3 and the range is 1 to 15 levels). You can also set the **Stop expanding at register** option (which is turned on by default) to specify whether netlist expansion should stop when a register is reached. These options are also available under **Tracing** in the **Options** dialog box (refer to “[Tracing](#)” on page 17-20).

You can select multiple nodes to expand when you use the **Expand** command. If you select ports that are located on multiple schematic pages, only the ports on the currently viewed page appear in the expanded schematic.

In the State Machine Viewer, the **Expand** command has the following three options:

- **Sources**—Displays the states that feed the selected states (previous transition states)
- **Destinations**—Displays the states that are fed by the selected states (next transition states)
- **Sources & Destinations**—Displays the previous and next transition states

The state transition table and state encoding table also reflect the changes to the filter.

The expansion feature works across hierarchical boundaries if the filtered page containing the port you want to expand was generated with the **Filter across hierarchy** option turned on (for details about this option, refer to “[Filtering in the Schematic View](#)” on page 17-33). When viewing timing paths in the Technology Map Viewer, the **Expand** command always works across hierarchical boundaries because filtering across hierarchy is always turned on for these schematics (for details about these schematics, refer to “[Viewing a Timing Path](#)” on page 17-41).

Reducing a Filtered Netlist

In some cases, removing logic from a filtered schematic or state diagram makes the schematic view easier to read and minimizes distracting logic in the schematic that you do not need to view.

To reduce elements in the filtered schematic or state diagram view, right-click the node or nodes you want to remove and click **Reduce**.

Probing to a Source Design File and Other Quartus II Windows

The RTL Viewer, Technology Map Viewer, and State Machine Viewer allows you to cross-probe to the source design file and to various other windows in the Quartus II software. You can select one or more hierarchy boxes, nodes, nets, state nodes, or state transition arcs that interest you in the netlist viewer and locate the corresponding items in another applicable Quartus II software window. You can then view and make changes or assignments in the appropriate editor or floorplan.

To locate an item from the netlist viewer in another window, right-click the items of interest in the schematic or state diagram, point to **Locate**, and click the appropriate command. The following commands are available:

- **Locate in Assignment Editor**
- **Locate in Pin Planner**
- **Locate in Timing Closure Floorplan**
- **Locate in Chip Planner**
- **Locate in Resource Property Editor**
- **Locate in RTL Viewer**
- **Locate in Technology Map Viewer**
- **Locate in Design File**

The options available for locating an item depend on the type of node and whether it exists after placement and routing. If a command is enabled in the menu, it is available for the selected node. You can use the **Locate in Assignment Editor** command for all nodes, but assignments might be ignored during placement and routing if they are applied to nodes that do not exist after synthesis.

The netlist viewer automatically opens another window for the appropriate editor or floorplan and highlights the selected node or net in the newly opened window. You can switch back to the netlist viewer by selecting it in the Window menu or by closing, minimizing, or moving the new window.



When probing to a logic cloud in the RTL Viewer, a message box appears, prompting you to ungroup the logic cloud or allow it to remain grouped.

Moving Selected Nodes to Other Quartus II Windows

You can drag selected nodes from the netlist viewers to the Text Editor, Block Editor, Pin Planner, SignalTap® II Embedded Logic Analyzer, and Waveform Editor windows in the Quartus II software. Whenever you see the drag-and-drop pointer on the selected node in the netlist viewers, it means that the node can be dragged to other child windows in the Quartus II software.

To tap a node from the schematic in the Technology Map Viewer to an open SignalTap II Embedded Logic Analyzer window or to a new SignalTap II file (.stp), right-click the selected node in the schematic diagram or in the netlist navigator, and then click **Add Node to SignalTap II Logic Analyzer**. If the node cannot be tapped, the option is unavailable.

Probing to the Netlist Viewers from Other Quartus II Windows

You can cross-probe to the RTL Viewer and Technology Map Viewer from other windows in the Quartus II software. You can select one or more nodes or nets in another window and locate them in one of the netlist viewers.

You can locate nodes between the RTL Viewer, State Machine Viewer, and Technology Map Viewer, and you can locate nodes in the RTL Viewer and Technology Map Viewer from the following Quartus II software windows:

- Project Navigator
- Timing Closure Floorplan

- Chip Planner
- Resource Property Editor
- Node Finder
- Assignment Editor
- Messages Window
- Compilation Report
- TimeQuest Timing Analyzer (supports the Technology Map Viewer only)

To locate elements in the netlist viewer from another Quartus II window, select the node or nodes in the appropriate window; for example, select an entity in the **Entity** list on the **Hierarchy** tab in the Project Navigator, or select nodes in the Timing Closure Floorplan, or select node names in the **From** or **To** column in the Assignment Editor. Next, right-click the selected object, point to **Locate**, and click **Locate in RTL Viewer** or **Locate in Technology Map Viewer**. After you click this command, the netlist viewer opens, or is brought to the foreground if the netlist viewer is already open.



The first time the window opens after a compilation, the preprocessor stage runs before the netlist viewer opens.

The netlist viewer shows the selected nodes and, if applicable, the connections between the nodes. The display is similar to what you see if you right-click the object, point to **Filter**, and click **Selected Nodes & Nets using Filter Across Hierarchy**. If the nodes cannot be found in the netlist viewer, a message box displays the message: **Can't find requested location**.

Viewing a Timing Path

To see a visual representation of a timing path, cross-probe from a report panel in the TimeQuest analyzer.

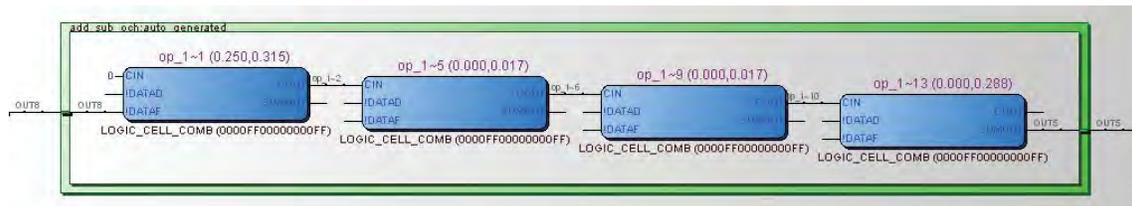
To take advantage of this feature, you must complete a full compilation of your design, including the timing analyzer stage. To see the timing results for your design, on the Processing menu, click **Compilation Report**. On the left side of the Compilation Report, select **Timing Analyzer** or **TimeQuest Timing Analyzer**. When you select a detailed report, the timing information is listed in a table format on the right side of the Compilation Report; each row of the table represents a timing path in the design. You can also view timing paths in TimeQuest analyzer report panels. To view a particular timing path in the Technology Map Viewer or RTL Viewer, right-click the appropriate row in the table, point to **Locate**, and click **Locate in Technology Map Viewer** or **Locate in RTL Viewer**.

To locate a path, on the **Tasks** pane, in the **Custom Reports** folder, double-click **Report Timing**. In the **Report Timing** dialog box, make any necessary settings, and then click the **Report Timing** button. After the TimeQuest analyzer generates the report, right-click on the node in the table and select **Locate Path**. In the Technology Map Viewer, the schematic page displays the nodes along the timing path with a summary of the total delay.

When you locate the timing path from the TimeQuest analyzer to the Technology Map Viewer, the interconnect and cell delay associated with each node is displayed on top of the schematic symbols. The total slack of the selected timing path is displayed in the Page Title section of the schematic. If the nodes are grouped in a logic cloud, the delay information displayed with the logic cloud is the total sum delay of the grouped nodes. The delay information for each node in the logic cloud is displayed in a tooltip. Move the mouse pointer over the logic cloud to see the tooltip. For more information about tooltips, refer to “Tooltips” on page 17-42.

Figure 17-27 shows a portion of a timing path represented in the Technology Map Viewer.

Figure 17-27. Timing Path Schematic in the Technology Map Viewer



In the RTL Viewer, the schematic page displays the nodes in the paths between the source and destination registers with a summary of the total delay.

The RTL Viewer netlist is based on an initial stage of synthesis, so the post-fitting nodes might not exist in the RTL Viewer netlist. Therefore, the internal delay numbers are not displayed in the RTL Viewer as they are in the Technology Map Viewer, and the timing path might not be displayed exactly as it appears in the timing analysis report. If multiple paths exist between the source and destination registers, the RTL Viewer might display more than just the timing path. There are also some cases in which the path cannot be displayed, such as paths through state machines, encrypted intellectual property (IP), or registers that are created during the fitting process. In cases where the timing path displayed in the RTL Viewer might not be the correct path, the compiler issues messages.

Other Features in the Schematic Viewer

This section describes other features in the schematic view that enhance usability and help you analyze your design.

Tooltips

A tooltip is displayed whenever the mouse pointer is held over an element in the schematic. The tooltip contains useful information about a node, net, logic cloud, input port, or output port. Table 17-9 lists the information contained in the tooltip for each type of node.

The tooltip information for an instance (the first row in Table 17-9) includes a list of the primitives found in that level of hierarchy and the number of each primitive contained in the current instance. The number includes all hierarchical blocks below the current instance in the hierarchy. This information lets you estimate the size and complexity of a hierarchical block without navigating into the block.

The tooltip information for atom primitives in the Technology Map Viewer (the second row in Table 17-9) shows the equation for the design atom. The equations are an expanded version of the equations you can view in the Equations window in the Timing Closure Floorplan. Advanced users can use these equations to analyze the design implementation in detail.

- For more information about understanding equations, refer to the Quartus II Help.

To copy tooltips into the clipboard for use in other applications, right-click the necessary node or netlist and click **Copy Tooltip**.

To turn off tooltips or change the duration of time that a tooltip appears in the view, in the main Quartus II window, select **Options**. In the **Tooltip Settings** pane, turn on the **Enable tooltips** option.

The **Show names in tooltip for** option specifies the number of seconds to display the names of assigned nodes and pins in a tooltip when the pointer is over the assigned nodes and pins. Selecting **Unlimited** displays the tooltip as long as the pointer remains over the node or pin. Selecting **0** turns off tooltips. The default value is 5 seconds.

The **Delay showing tooltip for** option specifies the number of seconds you must hold the mouse pointer over assigned nodes and pins before the tooltip displays the names of the assigned nodes and pins. Selecting **0** displays the tooltip immediately when the pointer is over an assigned node or pin. Selecting **Unlimited** prevents the display of tooltips. The default value is 1 second.

Table 17-9. Tooltip Information (Part 1 of 2)

Tooltip Format	Description	Example Tooltips
Instance	Format: <instance name>, <instance type> <primitive type>, <number of primitives>... <primitive type>, <number of primitives>	taps:inst_1 INST DFF_32 OPERATOR(SELECTOR) 8 OPERATOR(DECODER) 1
Atom Primitive	Format: <instance name>, <primitive name> (<LUT Mask Value>) {(r c <Register or Combinational equation>)} ... An r (as in the first example) represents the equation for a register, and a c (as in the second example) represents the equation for combinational logic.	inst5[3].LCELL (0000) <r> inst5[3] = DFFEAS((GND), GLOBAL(CLK), VCC, , ENA, SYNCH_DATA, , VCC) CLK = clkx2 ENA = inst4 SYNCH_DATA = result[7] acc:inst3lynm[2]^133.LCELL (00F0) <c> ynm[2]^133 = DATAC & IDATAD DATAC = result[2] DATAD = filter.tap1
Primitive	Format:<primitive name>, <primitive type>	clocks:inst7 Mux^1, OPER (MUX) md_me:inst18 data[3..3].DFFE
Pin	Format: <pin name>, <pin type>	pc_clock, INPUT Test_probe, OUTPUT
Connector	Format: <connector name>	inst4_CLK
Net	Format: <net name>, fan-out = <number of fan-out signals>	state_m:inst1.decoder_node[2][0], fan-out = 1

Table 17-9. Tooltip Information (Part 2 of 2)

Tooltip Format	Description	Example Tooltips
Output Port	Format: fan-out = <number of fan-out signals>	[fan-out = 9]
Input Port	<p>The information displayed depends on the type of source net. The examples of the tooltips shown represent the following types of source nets:</p> <p>(1) Single net</p> <p>(2) Individual nets, part of the same bus net</p> <p>(3) Combination of different bus nets</p> <p>(4) Constant inputs</p> <p>(5) Combination of single net and constant input</p> <p>(6) Bus net</p> <p>Source from—refers to the source net name that connects to the input port.</p> <p>Destination Index—refers to the bit(s) at the destination input port to which the source net is connected (not applicable for single nets).</p>	<p>[Source from: reset:reset_inst] (1)</p> <p>< Destination Index > Source from: (2) < [11] > sample~0:OUT1 < [10] > sample~1:OUT1 < [9] > sample~2:OUT1 < [8] > sample~3:OUT1 < [7] > sample~4:OUT1 < [6] > sample~5:OUT1 < [5] > sample~6:OUT1 < [4] > sample~7:OUT1 < [3] > sample~8:OUT1 < [2] > sample~9:OUT1 < [1] > sample~10:OUT1 < [0] > sample~11:OUT1</p> <p>< Destination Index > Source from: (3) < [7..6] > node2:OUT1 < [5] > ct[3]:OUT1 < [4] > node2:OUT1 < [3..2] > ct[3]:OUT1 < [1] > node2:OUT1 < [0] > ct[3]:OUT1</p> <p>< Destination Index > Source from: (4) < [11..0] > 12' h000</p> <p>< Destination Index > Source from: (5) < [2..1] > 2' h1 < [0] > always7~2:OUT1</p> <p>< Destination Index > Source from: (6) < [15..0] > md_me:inst18:dout[15:0]</p>
State Machine Node	Format: <node name>	[state_m:inst1 filter.tap1]
State Machine Transition Arc	<p>This information is displayed when you hold your mouse over the arrow on the arc representing the transition between two states.</p> <p>Format: (<equation for transition between states>)</p>	[!newt]

Finding Design Elements in the Netlist Viewers

You can narrow the range of the search process by setting the following options in the **Find** pane:

- Click **Browse** in the **Find** pane to specify the hierarchy level of the search. In the **Select Hierarchy Level** dialog box, select the particular instance you want to search.
- Turn on the **Include subtentities** option to include child hierarchies of the parent instance during the search.
- Click **Options** to open the **Find Options** dialog box. Turn on **Instances**, **Nodes**, **Pins**, or any combination of the three to further refine the parameters of the search.

When you click the **List** button, a progress bar appears below the **Find** box.

All results that match the criteria you set are listed in a table. When you double-click an item in the table, the related node is highlighted in red in the schematic view.

Exporting and Copying a Schematic Image

You can export the schematic view of the RTL Viewer or Technology Map Viewer into various image formats. This allows you to include the schematic in project documentation or share it with other project members. The currently supported formats are JPEG File Interchange Format (.jpg), Portable Network Graphics (.png), Graphics Interchange Format (.gif), and Windows Bitmap (.bmp). To export the schematic view, on the File menu, click **Export**. In the **Export** dialog box, type a file name and location and select the necessary file type. The default file name is based on the current instance name; the default file type is .jpg. However, for pages that use filtering, expanding, or reducing operations, the default name is **Filter**<number of export operation>.jpg.



Nodes grouped as logic clouds are not shown in the exported or copied schematic image; the logic clouds are shown instead.

You can copy the entire image or a portion of the image. To copy the entire image, right-click on the schematic, point to **Copy**, and then click **Full Image**. To copy a portion of the image, right-click on the schematic, point to **Copy**, and then click **Partial Image**. The cursor changes to a “+” sign to indicate that you can draw a box shape. Drag the mouse pointer around the portion of the schematic you want to copy. When you release the mouse button, the partial image is copied to the clipboard.



Occasionally, due to the design size and objects selected, an image is too large to copy to the clipboard. In this case, the Quartus II software displays an error message.

To export or copy a schematic that is too large to copy in one piece, split the design into multiple pages to export or to copy smaller portions of the design. For more information about controlling how much of your design is shown on each schematic page, refer to [“Partitioning the Schematic into Pages” on page 17–30](#). As an alternative, use the Partial Image feature to copy a portion of the image.

Printing

To print your schematic page, on the File menu, click **Print**. You can print each schematic page onto one page, or you can print selected parts of your schematic onto one page with the **Selection** option. To control how much of your design is shown on each schematic page, refer to [“Partitioning the Schematic into Pages” on page 17–30](#).

You cannot print the **Netlist Navigator** pane in the RTL Viewer and Technology Map Viewer and the table view of the State Machine Viewer. You can use the State Machine Viewer **Copy** command to copy the table to a text editor and print from the text editor.

Conclusion

The Quartus II RTL Viewer, State Machine Viewer, and Technology Map Viewer allow you to explore and analyze your initial synthesis netlist, post-synthesis netlist, or post-fitting and physical synthesis netlist. The netlist viewers provide a number of features in the **Netlist Navigator** pane and schematic view to help you quickly trace through your netlist and find specific hierarchies or nodes of interest. These capabilities can help you debug, optimize, and constrain your design more efficiently to increase your productivity.

Document Revision History

Table 17-10 shows the revision history for this chapter.

Table 17-10. Document Revision History

Date	Document Version	Changes
December 2010	10.0.1	Changed to new document template.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Updated screenshots ■ Updated chapter for the Quartus II software version 10.0, including major user interface changes
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Updated devices ■ Minor text edits
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Chapter 13 was formerly Chapter 12 in version 8.1.0 ■ Updated Figure 13-2, Figure 13-3, Figure 13-4, Figure 13-14, and Figure 13-30 ■ Added “Enable or Disable the Auto Hierarchy List” on page 13-15 ■ Updated “Find Command” on page 13-44
November 2008	8.1.0	Changed page size to 8.5” × 11”
May 2008	8.0.0	<ul style="list-style-type: none"> ■ Added Arria GX support ■ Updated operator symbols ■ Updated information about the radial menu feature ■ Updated zooming feature ■ Updated information about probing from schematic to SignalTap II Analyzer ■ Updated constant signal information ■ Added .png and .gif to the list of supported image file formats ■ Updated several figures and tables ■ Added new sections “Enabling and Disabling the Radial Menu”, “Changing the Time Interval”, “Changing the Constant Signal Value Formatting”, “Logic Clouds in the RTL Viewer”, “Logic Clouds in the Technology Map Viewer”, “Manually Group and Ungroup Logic Clouds”, “Customizing the Shortcut Commands” ■ Renamed several sections ■ Removed section “Customizing the Radial Menu” ■ Moved section “Grouping Combinational Logic into Logic Clouds” ■ Updated document content based on the Quartus II software version 8.0

-  For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).
-  Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter provides additional information about the document and Altera.

About this Handbook

This handbook provides comprehensive information about the Altera® Quartus® II design software, version 10.1.

How to Contact Altera

To locate the most up-to-date information about Altera products, refer to the following table.

Contact (1)	Contact Method	Address
Technical support	Website	www.altera.com/support
Technical training	Website	www.altera.com/training
	Email	custrain@altera.com
Product literature	Website	www.altera.com/literature
Non-technical support (General) (Software Licensing)	Email	nacomp@altera.com
	Email	authorization@altera.com

Note to Table:

(1) You can also contact your local Altera sales office or sales representative.

Third-Party Software Product Information

Third-party software products described in this handbook are not Altera products, are licensed by Altera from third parties, and are subject to change without notice. Updates to these third-party software products may not be concurrent with Quartus II software releases. Altera has assumed responsibility for the selection of such third-party software products and its use in the Quartus II 10.1 software release. To the extent that the software products described in this handbook are derived from third-party software, no third party warrants the software, assumes any liability regarding use of the software, or undertakes to furnish you any support or information relating to the software. EXCEPT AS EXPRESSLY SET FORTH IN THE APPLICABLE ALTERA PROGRAM LICENSE SUBSCRIPTION AGREEMENT UNDER WHICH THIS SOFTWARE WAS PROVIDED TO YOU, ALTERA AND THIRD-PARTY LICENSORS DISCLAIM ALL WARRANTIES WITH RESPECT TO THE USE OF SUCH THIRD-PARTY SOFTWARE CODE OR DOCUMENTATION IN THE SOFTWARE, INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, AND NONINFRINGEMENT. For more information, including the latest available version of specific third-party software products, refer to the documentation for the software in question.

Typographic Conventions

The following table shows the typographic conventions this document uses.

Visual Cue	Meaning
Bold Type with Initial Capital Letters	Indicate command names, dialog box titles, dialog box options, and other GUI labels. For example, Save As dialog box. For GUI elements, capitalization matches the GUI.
bold type	Indicates directory names, project names, disk drive names, file names, file name extensions, software utility names, and GUI labels. For example, \qdesigns directory, D: drive, and chiptrip.gdf file.
<i>Italic Type with Initial Capital Letters</i>	Indicate document titles. For example, <i>Stratix IV Design Guidelines</i> .
<i>italic type</i>	Indicates variables. For example, $n + 1$. Variable names are enclosed in angle brackets (< >). For example, <file name> and <project name>.pdf file.
Initial Capital Letters	Indicate keyboard keys and menu names. For example, the Delete key and the Options menu.
“Subheading Title”	Quotation marks indicate references to sections within a document and titles of Quartus II Help topics. For example, “Typographic Conventions.”
Courier type	Indicates signal, port, register, bit, block, and primitive names. For example, data1, tdi, and input. The suffix n denotes an active-low signal. For example, resetn. Indicates command line commands and anything that must be typed exactly as it appears. For example, c:\qdesigns\tutorial\chiptrip.gdf. Also indicates sections of an actual file, such as a Report File, references to parts of files (for example, the AHDL keyword SUBDESIGN), and logic function names (for example, TRI).
	An angled arrow instructs you to press the Enter key.
1., 2., 3., and a., b., c., and so on	Numbered steps indicate a list of items when the sequence of the items is important, such as the steps listed in a procedure.
	Bullets indicate a list of items when the sequence of the items is not important.
	The hand points to information that requires special attention.
	A question mark directs you to a software help system with related information.
	The feet direct you to another document or website with related information.
	A caution calls attention to a condition or possible situation that can damage or destroy the product or your work.
	A warning calls attention to a condition or possible situation that can cause you injury.
	The envelope links to the Email Subscription Management Center page of the Altera website, where you can sign up to receive update notifications for Altera documents.



Quartus II Handbook Version 10.1 Volume 2: Design Implementation and Optimization



101 Innovation Drive
San Jose, CA 95134
www.altera.com

QI15V2-10.1.0

© 2010 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, HARDCOPY, MAX, MEGACORE, NIOS, QUARTUS and STRATIX are Reg. U.S. Pat. & Tm. Off. and/or trademarks of Altera Corporation in the U.S. and other countries. All other trademarks and service marks are the property of their respective holders as described at www.altera.com/common/legal.html. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.



Chapter Revision Dates	xvii
-------------------------------------	------

Section I. Scripting and Constraint Entry

Chapter 1. Constraining Designs

Constraining Designs with the Quartus II GUI	1-1
Global Constraints	1-2
Node, Entity, and Instance-Level Constraints	1-3
Probing Between Components of the Quartus II GUI	1-4
SDC and the TimeQuest Timing Analyzer	1-4
Constraining Designs with Tcl	1-4
Quartus II Settings Files and Tcl	1-5
Timing Analysis with Synopsys Design Constraints and Tcl	1-8
A Fully Iterative Scripted Flow	1-9
Document Revision History	1-9

Chapter 2. Command-Line Scripting

Benefits of Command-Line Executables	2-1
Introductory Example	2-2
Command-Line Scripting Help	2-3
Command-Line Option Details	2-4
Option Precedence	2-5
Compilation with <code>quartus_sh --flow</code>	2-7
Text-Based Report Files	2-7
Makefile Implementation	2-9
The MegaWizard Plug-In Manager	2-11
Command-Line Support	2-12
Module and Wizard Names	2-13
Ports and Parameters	2-14
Invalid Configurations	2-15
Strategies to Determine Port and Parameter Values	2-15
Optional Files	2-15
Parameter File	2-16
Working Directory	2-17
Variation File Name	2-17
Command-Line Scripting Examples	2-17
Create a Project and Apply Constraints	2-17
Check Design File Syntax	2-18
Create a Project and Synthesize a Netlist Using Netlist Optimizations	2-19
Archive and Restore Projects	2-20
Perform I/O Assignment Analysis	2-20
Update Memory Contents Without Recompiling	2-20
Create a Compressed Configuration File	2-21
Fit a Design as Quickly as Possible	2-21
Fit a Design Using Multiple Seeds	2-22
Regenerating Megafunctions After Updating the Quartus II Software	2-23
The QFlow Script	2-23
Document Revision History	2-24

Chapter 3. Tcl Scripting

Introduction	3-1
What is Tcl?	3-2
Quartus II Tcl Packages	3-2
Loading Packages	3-3
Quartus II Tcl API Help	3-3
Command-Line Options: -t, -s, and --tcl_eval	3-6
Run a Tcl Script	3-6
Interactive Shell Mode	3-6
Evaluate as Tcl	3-6
Using the Quartus II Tcl Console Window	3-7
End-to-End Design Flows	3-7
Creating Projects and Making Assignments	3-8
HardCopy Device Design	3-8
Compiling Designs	3-9
The flow Package	3-9
Compile All Revisions	3-9
Reporting	3-9
Creating .csv Files for Excel	3-11
Timing Analysis	3-12
Automating Script Execution	3-12
Execution Example	3-13
Controlling Processing	3-14
Displaying Messages	3-14
Other Scripting Features	3-14
Natural Bus Naming	3-15
Short Option Names	3-15
Using Collection Commands	3-15
The foreach_in_collection Command	3-16
The get_collection_size Command	3-16
Using the post_message Command	3-16
Accessing Command-Line Arguments	3-17
Using the cmdline Package	3-17
Using the Quartus II Tcl Shell in Interactive Mode	3-19
Using the tclsh Shell	3-20
Tcl Scripting Basics	3-20
Hello World Example	3-20
Variables	3-20
Substitutions	3-21
Variable Value Substitution	3-21
Nested Command Substitution	3-21
Backlash Substitution	3-21
Arithmetic	3-21
Lists	3-22
Arrays	3-22
Control Structures	3-23
Procedures	3-24
File I/O	3-25
Syntax and Comments	3-25
External References	3-26
Document Revision History	3-26

Chapter 4. Managing Quartus II Projects

Managing Your Quartus II Projects	4-1
File Association	4-2
Editing Text-Based Designs with the Quartus II Text Editor	4-2
Creating Assignments	4-3
Quartus II Settings File	4-3
Preserving QSF Format	4-3
Quartus II Default Settings File	4-3
Creating Timing Assignments	4-4
Creating Revisions	4-4
Managing Project Revisions	4-4
Creating New Copies of Your Design	4-5
Archiving and Restoring Projects	4-5
Exporting and Importing Version-Compatible Database Files	4-6
Migrating to a New Version of the Quartus II Software	4-7
Saving the Database in a Version-Compatible Format	4-7
Quartus II Project Platform Migration	4-8
File Names and Hierarchies	4-8
Specifying Libraries	4-10
Quartus II Search Path Precedence Rules	4-11
Quartus II-Generated Files for Third-Party EDA Tools	4-12
Migrating Database Files Between Platforms	4-13
Working with Messages	4-13
Messages Window	4-13
Message Suppression	4-13
Managing Projects in a Team-Based Design Environment	4-15
Scripting Support	4-16
Managing Revisions	4-16
Creating Revisions	4-16
Setting the Current Revision	4-16
Getting a List of Revisions	4-17
Deleting Revisions	4-17
Archiving Projects	4-17
Restoring Archived Projects	4-18
Importing and Exporting Version-Compatible Databases	4-18
Specifying Libraries Using Scripts	4-19
Conclusion	4-20
Document Revision History	4-20

Section II. I/O and PCB Tools

Chapter 5. I/O Management

Understanding Altera Pin Terminology	5-1
Package Pins	5-2
Pads	5-3
I/O Banks	5-3
VREF Groups	5-4
I/O Planning Overview	5-5
Selecting a Device	5-7
Working with Third-Party PCB Tools	5-7
Performing Early I/O Planning with the Pin Planner	5-8
Instantiating or Importing IP Cores in the Pin Planner	5-9
Adding and Connecting Nodes	5-9

Setting Up and Creating the Top-Level Design File	5-10
Importing and Exporting Pin Assignments	5-12
Importing and Exporting Assignments with the Quartus II Software	5-12
Importing and Exporting Assignments with Third-Party PCB Tools	5-12
Creating Pin-Related Assignments	5-13
Creating Pin Assignments With the Pin Planner	5-14
Finding Compatible Pin Locations with the Pin Finder	5-15
Verifying Pin Migration Compatibility	5-15
Viewing Simultaneous Switching Noise (SSN) Results	5-17
Creating Location Assignments	5-17
Creating Exclusive I/O Group Assignments	5-18
Changing the Slew Rate and Drive Strength	5-18
Assigning Locations for Differential Pins	5-19
Creating Pin Assignments with the Chip Planner	5-20
Creating Pin Assignments with Tcl Scripts	5-20
Creating Pin Assignments in HDL Code	5-20
Synthesis Attributes	5-21
chip_pin and useioff	5-21
altera_attribute	5-22
Creating Pin Assignments with Low-Level I/O Primitives	5-22
Validating Pin Assignments	5-23
Validating Pin Assignments with the Live I/O Check Feature	5-24
Validating Pin Assignments with I/O Assignment Analysis	5-25
Running I/O Assignment Analysis without Design Files	5-26
Running I/O Assignment Analysis with Design Files	5-28
Optimizing I/O Assignment Analysis with Output Enable Group Logic Option Assignments	5-29
Validating Pin Assignments with Full Compilation	5-31
Performing I/O Timing Analysis	5-32
Enabling and Configuring Advanced I/O Timing	5-33
Defining Overall Board Trace Models	5-33
Customizing the Board Trace Model in the Pin Planner	5-34
Configuring Board Trace Models	5-34
Specifying Near-End vs Far-End Timing Analysis	5-36
Understanding Advanced I/O Timing Analysis Reports	5-36
Adjusting I/O Timing and Power with Capacitive Loading	5-37
Incorporating PCB Design Tools	5-37
Scripting Support	5-38
Running I/O Assignment Analysis	5-38
Generating a Mapped Netlist	5-38
Reserving Pins	5-38
Creating Location Assignments	5-39
Creating Exclusive I/O Group Assignments	5-39
Changing the Slew Rate and Drive Strength	5-39
Conclusion	5-40
Document Revision History	5-40

Chapter 6. Simultaneous Switching Noise (SSN) Analysis and Optimizations

Definitions	6-1
Understanding SSN	6-2
SSN Estimation Tools	6-5
SSN Analysis Overview	6-5
Performing Early Pin-Out SSN Analysis	6-6
Performing Early Pin-Out SSN Analysis with the ESE Tool	6-6

Performing Early Pin-Out SSN Analysis with the SSN Analyzer	6-7
Performing Final Pin-Out SSN Analysis	6-8
Design Factors Affecting SSN Results	6-8
Optimizing Your Design for SSN Analysis	6-8
Optimizing Pin Placements for Signal Integrity	6-9
Specifying Board Trace Model Settings	6-10
Defining PCB Layers and PCB Layer Thickness	6-11
Specifying Signal Breakout Layers	6-13
Creating I/O Assignments	6-14
Decreasing Pessimism in SSN Analysis	6-14
Excluding Pins as Aggressor Signals	6-15
Performing SSN Analysis and Viewing Results	6-15
Understanding the SSN Reports	6-15
Summary Report	6-16
Output Pins and Input Pins Reports	6-16
Unanalyzed Pins Report	6-16
Confidence Metric Details Report	6-16
Viewing SSN Analysis Results in the Pin Planner	6-16
Decreasing Processing Time for SSN Analysis	6-17
Scripting Support	6-17
Optimizing Pin Placements for Signal Integrity	6-18
Defining PCB Layers and PCB Layer Thickness	6-18
Specifying Signal Breakout Layers	6-19
Decreasing Pessimism in SSN Analysis	6-19
Performing SSN Analysis	6-19
Conclusion	6-20
Document Revision History	6-20

Chapter 7. Signal Integrity Analysis with Third-Party Tools

Introduction	7-1
I/O Model Selection: IBIS or HSPICE	7-3
FPGA to Board Signal Integrity Analysis Flow	7-4
Create I/O and Board Trace Model Assignments	7-5
Output File Generation	7-6
Customize the Output Files	7-6
Set Up and Run Simulations in Third-Party Tools	7-7
Interpret Simulation Results	7-7
Simulation with IBIS Models	7-7
Elements of an IBIS Model	7-8
Creating Accurate IBIS Models	7-8
Download IBIS Models	7-9
Generate Custom IBIS Models with the IBIS Writer	7-9
Design Simulation Using the Mentor Graphics HyperLynx Software	7-10
Configuring LineSim to Use Altera IBIS Models	7-12
Integrating Altera IBIS Models into LineSim Simulations	7-13
Running and Interpreting LineSim Simulations	7-15
Simulation with HSPICE Models	7-16
Supported Devices and Signaling	7-17
Accessing HSPICE Simulation Kits	7-17
The Double Counting Problem in HSPICE Simulations	7-17
Defining the Double Counting Problem	7-18
The Solution to Double Counting	7-19
HSPICE Writer Tool Flow	7-20
Applying I/O Assignments	7-20

Enabling HSPICE Writer	7-20
Enabling HSPICE Writer Using Assignments	7-21
Naming Conventions for HSPICE Files	7-21
Invoking HSPICE Writer	7-21
Invoking HSPICE Writer from the Command Line	7-22
Customizing Automatically Generated HSPICE Decks	7-22
Running an HSPICE Simulation	7-23
Interpreting the Results of an Output Simulation	7-23
Interpreting the Results of an Input Simulation	7-24
Viewing and Interpreting Tabular Simulation Results	7-24
Viewing Graphical Simulation Results	7-24
Making Design Adjustments Based on HSPICE Simulations	7-26
Sample Input for I/O HSPICE Simulation Deck	7-28
Header Comment	7-28
Simulation Conditions	7-29
Simulation Options	7-29
Constant Definition	7-30
Buffer Netlist	7-31
Drive Strength	7-31
I/O Buffer Instantiation	7-31
Board Trace and Termination	7-32
Stimulus Model	7-32
Simulation Analysis	7-33
Sample Output for I/O HSPICE Simulation Deck	7-33
Header Comment	7-33
Simulation Conditions	7-34
Simulation Options	7-35
Constraint Definition	7-36
I/O Buffer Netlist	7-36
Drive Strength	7-37
Slew Rate and Delay Chain	7-37
I/O Buffer Instantiation	7-38
Board and Trace Termination	7-38
Double-Counting Compensation Circuitry	7-39
Simulation Analysis	7-40
Advanced Topics	7-40
PVT Simulations	7-40
Hold Time Analysis	7-41
I/O Voltage Variations	7-41
Correlation Report	7-41
Conclusion	7-42
Document Revision History	7-42
Chapter 8. Mentor Graphics PCB Design Tools Support	
FPGA-to-PCB Design Flow	8-2
Performing Simultaneous Switching Noise (SSN) Analysis of Your FPGA	8-3
Setting Up the Quartus II Software	8-4
Generating a .pin File	8-5
Generating an .fx File	8-6
Creating a Backup .qsf	8-6
FPGA-to-Board Integration with the I/O Designer Software	8-6
I/O Designer Database Wizard	8-7
Updating Pin Assignments from the Quartus II Software	8-11
Sending Pin Assignment Changes to the Quartus II Software	8-14

Protecting Assignments in the Quartus II Software	8–15
Generating Symbols for the DxDesigner Software	8–15
Setting Up the I/O Designer Software to Work with the DxDesigner Software	8–16
Creating Symbols with the Symbol Wizard	8–17
Exporting Symbols to the DxDesigner Software	8–19
Scripting Support	8–19
FPGA-to-Board Integration with the DxDesigner Software	8–20
DxDesigner Project Settings	8–21
DxDesigner Symbol Wizard	8–23
Conclusion	8–25
Document Revision History	8–25

Chapter 9. Cadence PCB Design Tools Support

Product Comparison	9–2
FPGA-to-PCB Design Flow	9–3
Performing Simultaneous Switching Noise (SSN) Analysis of Your FPGA	9–5
Setting Up the Quartus II Software	9–5
Generating a .pin File	9–5
FPGA-to-Board Integration with the Cadence Allegro Design Entry HDL Software	9–6
Creating Symbols	9–6
Cadence Allegro PCB Librarian Part Developer Tool	9–8
Instantiating the Symbol in the Cadence Allegro Design Entry HDL Software	9–12
FPGA-to-Board Integration with Cadence Allegro Design Entry CIS Software	9–13
Creating a Cadence Allegro Design Entry CIS Project	9–13
Generating a Part	9–14
Splitting a Part	9–16
Instantiating a Symbol in a Design Entry CIS Schematic	9–18
Altera Libraries for the Cadence Allegro Design Entry CIS Software	9–18
Conclusion	9–20
Document Revision History	9–20

Chapter 10. Reviewing Printed Circuit Board Schematics with the Quartus II Software

Reviewing Quartus II Software Settings	10–1
Device and Pins Options Dialog Box Settings	10–2
Configuration Page Settings	10–2
Unused Pin Page Settings	10–2
Dual-Purpose Pins Page Settings	10–3
Voltage Page Settings	10–3
Error Detection CRC Page Settings	10–3
Voltage Page Settings	10–3
Reviewing Device Pin-Out Information in the Fitter Report	10–4
Reviewing Compilation Error and Warning Messages	10–5
Running the HardCopy Design Readiness Check	10–6
Using Additional Quartus II Software Features	10–6
Using Additional Quartus II Software Tools	10–6
Pin Planner	10–6
SSN Analyzer	10–7
Conclusion	10–7
Document Revision History	10–7

Section III. Area, Timing, Power, and Compilation Time Optimization

Chapter 11. Design Optimization Overview

Introduction	11-1
Physical Implementation	11-1
Trade-Offs and Limitations	11-1
Preserving Results and Enabling Teamwork	11-2
Reducing Area	11-2
Reducing Critical Path Delay	11-3
Reducing Power Consumption	11-3
Reducing Runtime	11-3
Using Quartus II Tools	11-4
Design Analysis	11-4
Advisors	11-4
Design Space Explorer	11-5
Conclusion	11-5
Document Revision History	11-5

Chapter 12. Reducing Compilation Time

Compilation Time Optimization Techniques	12-1
Compilation Time Advisor	12-2
Strategies to Reduce the Overall Compilation Time	12-2
Using Parallel Compilation with Multiple Processors	12-2
Using Incremental Compilation	12-4
Using the Smart Compilation Setting	12-4
Using Rapid Recompile	12-4
Reducing Synthesis Time and Synthesis Netlist Optimization Time	12-5
Settings to Reduce Synthesis Time and Synthesis Netlist Optimization Time	12-5
Use Appropriate Coding Style to Reduce Synthesis Time	12-6
Using Early Timing Estimation	12-6
Reducing Placement Time	12-7
Fitter Effort Setting	12-7
Placement Effort Multiplier Settings	12-7
Final Placement Optimization Levels	12-7
Physical Synthesis Effort Settings	12-8
Limit to One Fitting Attempt	12-8
Preserving Placement, Incremental Compilation, and LogicLock Regions	12-8
Reducing Routing Time	12-8
Identifying Routing Congestion in the Chip Planner	12-9
Placement Effort Multiplier Setting	12-9
Preserving Routing with Incremental Compilation	12-9
Reducing Static Timing Analysis Time	12-9
Setting Process Priority	12-10
Conclusion	12-10
Document Revision History	12-10

Chapter 13. Area and Timing Optimization

Optimizing Your Design	13-1
Initial Compilation: Required Settings	13-2
Device Settings	13-3
I/O Assignments	13-3
Timing Requirement Settings	13-3
Device Migration Settings	13-5

Partitions and Floorplan Assignments for Incremental Compilation	13–5
Initial Compilation: Optional Fitter Settings	13–6
Optimize Hold Timing	13–6
Limit to One Fitting Attempt	13–7
Optimize Multi-Corner Timing	13–7
Fitter Effort Setting	13–8
Design Assistant	13–9
Design Analysis	13–10
Error and Warning Messages	13–10
Ignored Timing Constraints	13–10
Resource Utilization	13–10
I/O Timing (Including t_{PD})	13–11
Register-to-Register Timing	13–12
Timing Analysis with the TimeQuest Timing Analyzer	13–12
Tips for Analyzing Failing Paths	13–14
Tips for Analyzing Failing Clock Paths that Cross Clock Domains	13–14
Global Routing Resources	13–15
Resource Utilization Optimization Techniques (LUT-Based Devices)	13–15
Using the Resource Optimization Advisor	13–16
Resolving Resource Utilization Issues Summary	13–16
I/O Pin Utilization or Placement	13–16
Use I/O Assignment Analysis	13–16
Modify Pin Assignments or Choose a Larger Package	13–16
Logic Utilization or Placement	13–17
Optimize Synthesis for Area, Not Speed	13–17
Restructure Multiplexers	13–18
Perform WYSIWYG Primitive Resynthesis with Balanced or Area Setting	13–18
Use Register Packing	13–18
Remove Fitter Constraints	13–20
Change State Machine Encoding	13–20
Flatten the Hierarchy During Synthesis	13–21
Retarget Memory Blocks	13–21
Use Physical Synthesis Options to Reduce Area	13–22
Retarget or Balance DSP Blocks	13–22
Optimize Source Code	13–23
Use a Larger Device	13–24
Routing	13–24
Set Auto Packed Registers to Sparse or Sparse Auto	13–24
Set Fitter Aggressive Routability Optimizations to Always	13–24
Increase Placement Effort Multiplier	13–24
Increase Router Effort Multiplier	13–25
Remove Fitter Constraints	13–25
Optimize Synthesis for Area, Not Speed	13–26
Optimize Source Code	13–26
Use a Larger Device	13–26
Timing Optimization Techniques (LUT-Based Devices)	13–27
Debugging Timing Failures in the TimeQuest Analyzer	13–27
Timing Optimization Advisor	13–29
I/O Timing Optimization	13–30
Improving Setup and Clock-to-Output Times Summary	13–30
Timing-Driven Compilation	13–31
Fast Input, Output, and Output Enable Registers	13–31
Programmable Delays	13–32
Use PLLs to Shift Clock Edges	13–33

Use Fast Regional Clock Networks and Regional Clocks Networks	13–33
Change How Hold Times are Optimized for MAX II Devices	13–34
Register-to-Register Timing Optimization Techniques (LUT-Based Devices)	13–34
Improving Register-to-Register Timing Summary	13–34
Physical Synthesis Optimizations	13–35
Turn Off Extra-Effort Power Optimization Settings	13–37
Optimize Synthesis for Speed, Not Area	13–37
Flatten the Hierarchy During Synthesis	13–38
Set the Synthesis Effort to High	13–38
Change State Machine Encoding	13–38
Duplicate Logic for Fan-Out Control	13–38
Prevent Shift Register Inference	13–39
Use Other Synthesis Options Available in Your Synthesis Tool	13–39
Fitter Seed	13–39
Set Maximum Router Timing Optimization Level	13–40
Optimize Source Code	13–40
LogicLock Assignments	13–41
Hierarchy Assignments	13–42
Location Assignments and Back-Annotation	13–42
Metastability Analysis and Optimization Techniques	13–43
Resource Utilization Optimization Techniques (Macrocell-Based CPLDs)	13–43
Use Dedicated Inputs for Global Control Signals	13–43
Reserve Device Resources	13–44
Pin Assignment Guidelines and Procedures	13–44
Control Signal Pin Assignments	13–45
Output Enable Pin Assignments	13–45
Estimate Fan-In When Assigning Output Pins	13–45
Outputs Using Parallel Expander Pin Assignments	13–45
Resolving Resource Utilization Problems	13–46
Resolving Macrocell Usage Issues	13–47
Resolving Routing Issues	13–47
Using LCELL Buffers to Reduce Required Resources	13–48
Timing Optimization Techniques (Macrocell-Based CPLDs)	13–49
Improving Setup Time	13–50
Improving Clock-to-Output Time	13–51
Improving Propagation Delay (t_{PD})	13–51
Improving Maximum Frequency (f_{MAX})	13–52
Optimizing Source Code—Pipelining for Complex Register Logic	13–53
Other Optimization Resources	13–53
Design Space Explorer	13–53
Other Optimization Advisors	13–53
Scripting Support	13–53
Initial Compilation Settings	13–54
Resource Utilization Optimization Techniques (LUT-Based Devices)	13–55
I/O Timing Optimization Techniques (LUT-Based Devices)	13–56
Register-to-Register Timing Optimization Techniques (LUT-Based Devices)	13–56
Duplicate Logic for Fan-Out Control	13–57
Conclusion	13–57
Document Revision History	13–57

Chapter 14. Power Optimization

Power Dissipation	14–2
Design Space Explorer	14–3
Power-Driven Compilation	14–4

Power-Driven Synthesis	14-4
Power-Driven Fitter	14-9
Area-Driven Synthesis	14-11
Gate-Level Register Retiming	14-11
Design Guidelines	14-12
Clock Power Management	14-12
LAB-Wide Clock Enable Example	14-14
Reducing Memory Power Consumption	14-14
Memory Power Reduction Example	14-16
Pipelining and Retiming	14-17
Architectural Optimization	14-18
I/O Power Guidelines	14-19
Dynamically Controlled On-Chip Terminations	14-21
Power Optimization Advisor	14-21
Power Optimization Advisor Example	14-22
Conclusion	14-24
Document Revision History	14-24

Chapter 15. Analyzing and Optimizing the Design Floorplan

Chip Planner Overview	15-1
Starting the Chip Planner	15-2
Chip Planner Toolbar	15-3
Chip Planner Tasks, Layers, and Editing Modes	15-3
Locate History Window	15-4
LogicLock Regions	15-4
Creating LogicLock Regions	15-5
Creating LogicLock Regions with the Project Navigator	15-5
Creating LogicLock Regions with the Chip Planner	15-5
Placing LogicLock Regions	15-5
Placing Device Resources into LogicLock Regions	15-6
LogicLock Regions Window	15-6
Assigning LogicLock Region Content	15-7
Hierarchical (Parent and Child) LogicLock Regions	15-8
Reserved LogicLock Region	15-8
Creating Non-Rectangular LogicLock Regions	15-8
Excluded Resources	15-9
Additional Quartus II LogicLock Design Features	15-10
Analysis and Synthesis Resource Utilization by Entity	15-10
Quartus II Revisions Feature	15-10
LogicLock Assignment Precedence	15-10
Virtual Pins	15-11
Using LogicLock Regions in the Chip Planner	15-12
Viewing Connections Between LogicLock Regions in the Chip Planner	15-12
Using LogicLock Regions with the Design Partition Planner	15-12
Design Floorplan Analysis Using the Chip Planner	15-13
Chip Planner Floorplan Views	15-13
Bird's Eye View	15-13
Properties Window	15-14
Viewing Architecture-Specific Design Information	15-14
Viewing Available Clock Networks in the Device	15-15
Viewing Critical Paths	15-16
Viewing Routing Congestion	15-17
Viewing I/O Banks	15-18
Generating Fan-In and Fan-Out Connections	15-18

Generating Immediate Fan-In and Fan-Out Connections	15–19
Highlight Routing	15–20
Show Delays	15–21
Exploring Paths in the Chip Planner	15–21
Locate Path from the Timing Analysis Report to the Chip Planner	15–21
Analyzing Connections for a Path	15–22
Viewing Assignments in the Chip Planner	15–22
Viewing Routing Channels for a Path in the Chip Planner	15–24
Delay Information Table	15–25
Viewing High-Speed and Low-Power Tiles in the Chip Planner	15–25
Scripting Support	15–26
Initializing and Uninitializing a LogicLock Region	15–27
Creating or Modifying LogicLock Regions	15–27
Obtaining LogicLock Region Properties	15–27
Assigning LogicLock Region Content	15–27
Save a Node-Level Netlist for the Entire Design into a Persistent Source File	15–28
Setting LogicLock Assignment Priority	15–28
Assigning Virtual Pins	15–28
Conclusion	15–28
Document Revision History	15–29

Chapter 16. Netlist Optimizations and Physical Synthesis

WYSIWYG Primitive Resynthesis	16–1
Performing Physical Synthesis Optimizations	16–3
Automatic Asynchronous Signal Pipelining	16–5
Physical Synthesis for Combinational Logic	16–6
Physical Synthesis for Registers—Register Duplication	16–6
Physical Synthesis for Registers—Register Retiming	16–7
Preserving Your Physical Synthesis Results	16–10
Physical Synthesis Options for Fitting	16–11
Applying Netlist Optimization Options	16–12
Scripting Support	16–12
Synthesis Netlist Optimizations	16–13
Physical Synthesis Optimizations	16–13
Incremental Compilation	16–14
Back-Annotating Assignments	16–14
Conclusion	16–14
Document Revision History	16–14

Section IV. Engineering Change Management

Chapter 17. Engineering Change Management with the Chip Planner

Engineering Change Orders	17–2
Performance Preservation	17–2
Compilation Time	17–3
Verification	17–3
Change Modification Record	17–3
ECO Design Flow	17–4
The Chip Planner Overview	17–5
Opening the Chip Planner	17–5
The Chip Planner Tasks and Layers	17–6
Performing ECOs with the Chip Planner (Floorplan View)	17–6
Creating, Deleting, and Moving Atoms	17–7

Check and Save Netlist Changes	17-7
Performing ECOs in the Resource Property Editor	17-7
Logic Elements	17-7
Logic Element Schematic View	17-8
Logic Element Properties	17-8
Modes of Operation	17-9
Sum and Carry Equations	17-9
load and sclr Signals	17-9
Register Cascade Mode	17-9
Cell Delay Table	17-9
Logic Element Connections	17-10
Delete a Logic Element	17-10
Adaptive Logic Modules	17-10
Adaptive Logic Module Schematic	17-11
Adaptive Logic Module Properties	17-11
Adaptive Logic Module Connections	17-12
FPGA I/O Elements	17-12
Stratix V I/O Elements	17-12
Arria GX, Stratix, Stratix II, and Stratix GX I/O Elements	17-14
Arria II GX, Stratix III, and Stratix IV I/O Elements	17-15
Cyclone and Cyclone II I/O Elements	17-16
Cyclone III I/O Elements	17-17
MAX II I/O Elements	17-18
FPGA RAM Blocks	17-19
FPGA DSP Blocks	17-20
Change Manager	17-21
Complex Changes in the Change Manager	17-21
Managing SignalProbe Signals	17-21
Exporting Changes	17-21
Scripting Support	17-22
Common ECO Applications	17-22
Adjust the Drive Strength of an I/O with the Chip Planner	17-22
Modify the PLL Properties With the Chip Planner	17-23
PLL Properties	17-24
Adjusting the Duty Cycle	17-25
Adjusting the Phase Shift	17-25
Adjusting the Output Clock Frequency	17-26
Adjusting the Spread Spectrum	17-26
Modify the Connectivity between Resource Atoms	17-26
Post ECO Steps	17-27
Conclusion	17-27
Document Revision History	17-28

Additional Information

How to Contact Altera	Info-1
Typographic Conventions	Info-2

The chapters in this document, Quartus II Handbook Version 10.1 Volume 2: Design Implementation and Optimization, were revised on the following dates. Where chapters or groups of chapters are available separately, part numbers are listed.

- Chapter 1. Constraining Designs
Revised: *December 2010*
Part Number: *QII52001-10.0.1*

- Chapter 2. Command-Line Scripting
Revised: *December 2010*
Part Number: *QII52002-10.1.0*

- Chapter 3. Tcl Scripting
Revised: *December 2010*
Part Number: *QII52003-10.1.0*

- Chapter 4. Managing Quartus II Projects
Revised: *December 2010*
Part Number: *QII52012-10.1.0*

- Chapter 5. I/O Management
Revised: *December 2010*
Part Number: *QII52013-10.0.1*

- Chapter 6. Simultaneous Switching Noise (SSN) Analysis and Optimizations
Revised: *December 2010*
Part Number: *QII52018-10.0.1*

- Chapter 7. Signal Integrity Analysis with Third-Party Tools
Revised: *December 2010*
Part Number: *QII53020-10.0.1*

- Chapter 8. Mentor Graphics PCB Design Tools Support
Revised: *December 2010*
Part Number: *QII52015-10.0.1*

- Chapter 9. Cadence PCB Design Tools Support
Revised: *December 2010*
Part Number: *QII52014-10.0.1*

- Chapter 10. Reviewing Printed Circuit Board Schematics with the Quartus II Software
Revised: *December 2010*
Part Number: *QII52019-10.0.1*

- Chapter 11. Design Optimization Overview
Revised: *December 2010*
Part Number: *QII52021-10.0.2*

- Chapter 12. Reducing Compilation Time

Revised: *December 2010*
Part Number: *QII52022-10.1.0*

Chapter 13. Area and Timing Optimization

Revised: *December 2010*
Part Number: *QII52005-10.1.0*

Chapter 14. Power Optimization

Revised: *December 2010*
Part Number: *QII52016-10.0.1*

Chapter 15. Analyzing and Optimizing the Design Floorplan

Revised: *December 2010*
Part Number: *QII52006-10.1.0*

Chapter 16. Netlist Optimizations and Physical Synthesis

Revised: *December 2010*
Part Number: *QII52007-10.0.1*

Chapter 17. Engineering Change Management with the Chip Planner

Revised: *December 2010*
Part Number: *QII52017-10.1.0*

As a result of the increasing complexity of today's FPGA designs and the demand for higher performance, designers must make a large number of complex timing and logic constraints to meet their performance requirements. After you create a project and design, you can use the Quartus® II software Assignment Editor and other GUI features to specify your initial design constraints, such as pin assignments, device options, logic options, and timing constraints.

This section describes how to constrain designs, how to take advantage of Quartus II modular executables, how to develop and run Tcl scripts to perform a wide range of functions, and how to manage the Quartus II projects.

This section includes the following chapters:

- [Chapter 1, Constraining Designs](#)

This chapter discusses the ways to constrain designs in the Quartus II software, including the tools available in the Quartus II software GUI, as well as Tcl scripting flows.

- [Chapter 2, Command-Line Scripting](#)

This chapter discusses Quartus II command-line executables, which provide command-line control over each step of the design flow. Each executable includes options to control commonly used software settings. Each executable also provides detailed, built-in help describing its function, available options, and settings.

- [Chapter 3, Tcl Scripting](#)

This chapter discusses developing and running Tcl scripts in the Quartus II software to allow you to perform a wide range of functions, such as compiling a design or automating common tasks. This chapter includes sample Tcl scripts for automating the Quartus II software. You can modify these example scripts for use with your own designs.

- [Chapter 4, Managing Quartus II Projects](#)

This chapter discusses the best ways to manage Quartus II projects. This chapter also discusses how to migrate your projects from one computing platform to another; create and compare revisions; and copy, archive and restore projects.

This chapter discusses the various tools and methods for constraining and re-constraining Quartus II designs in different design flows, both with the Quartus II GUI and with Tcl to facilitate a scripted flow.

Constraints, sometimes known as assignments or logic options, control the way the Quartus II software implements a design for an FPGA. Constraints are also central in the way that the TimeQuest Timing Analyzer and the PowerPlay Power Analyzer inform synthesis, placement, and routing. There are several types of constraints:

- Global design constraints and software settings, such as device family selection, package type, and pin count.
- Entity-level constraints, such as logic options and placement assignments.
- Instance-level constraints.
- Pin assignments and I/O constraints.

User-created constraints are contained in one of two files: the Quartus II Settings File (**.qsf**) or, in the case of timing constraints, the Synopsys Design Constraints file (**.sdc**). Constraints and assignments made with the **Device** dialog box, **Settings** dialog box, Assignment Editor, Chip Planner, and Pin Planner are contained in the Quartus II Settings File. The **.qsf** file contains project-wide and instance-level assignments for the current revision of the project in Tcl syntax. You can create separate revisions of your project with different settings, and there is a separate **.qsf** file for each revision.

The TimeQuest Timing Analyzer uses industry-standard Synopsys Design Constraints, also using Tcl syntax, that are contained in Synopsys Design Constraints (**.sdc**) files. The TimeQuest Timing Analyzer GUI is a tool for making timing constraints and viewing the results of subsequent analysis.

There are several ways to constrain a design, each potentially more appropriate than the others, depending on your tool chain and design flow. You can constrain designs for compilation and analysis in the Quartus II software using the GUI, as well as using Tcl syntax and scripting. By combining the Tcl syntax of the **.qsf** files and the **.sdc** files with procedural Tcl, you can automate iteration over several different settings, changing constraints and recompiling.

Constraining Designs with the Quartus II GUI

In the Quartus II GUI, the New Project Wizard, **Device** dialog box, and **Settings** dialog box allow you to make global constraints and software settings. The Assignment Editor and Pin Planner are spreadsheet-style interfaces for constraining your design at the instance or entity level. The Assignment Editor and Pin Planner make constraint types and values available based on global design characteristics such as the targeted device. These tools help you verify that your constraints are valid before compilation by allowing you to pick only from valid values for each constraint.

The TimeQuest Timing Analyzer GUI allows you to make timing constraints in SDC format and view the effects of those constraints on the timing in your design. Before running the TimeQuest timing analyzer, you must specify initial timing constraints that describe the clock characteristics, timing exceptions, and external signal arrival and required times. The Quartus II Fitter optimizes the placement of logic in the device to meet your specified constraints.

- ❓ For more information about timing constraints and the TimeQuest Timing Analyzer, refer to *About TimeQuest Timing Analysis* in Quartus II Help.

Global Constraints

Global constraints affect the entire Quartus II project and all of the applicable logic in the design. Many of these constraints are simply project settings, such as the targeted device selected for the design. Synthesis optimizations and global timing and power analysis settings can also be applied with globally. Global constraints are often made when running the New Project Wizard, or in the **Device** dialog box or the **Settings** dialog box, early project development.

The following are the most common types of global constraints:

- Target device specification
- Top-level entity of your design, and the names of the design files included in the project
- Operating temperature limits and conditions
- Physical synthesis optimizations
- Analysis and synthesis options and optimization techniques
- Verilog HDL and VHDL language versions used in your project
- Fitter effort and timing driven compilation settings
- .sdc files for the TimeQuest timing analyzer to use during analysis as part of a full compilation flow

Settings that direct compilation and analysis flows in the Quartus II software are also stored in the Quartus II Settings File for your project, including the following global software settings:

- Early Timing Estimate mode
- Settings for EDA tool integration such as third-party synthesis tools, simulation tools, timing analysis tools, and formal verification tools.
- Settings and settings file specifications for the Quartus II Assembler, SignalTap II Logic Analyzer, PowerPlay power analyzer, and SSN Analyzer.

Global constraints and software settings stored in the Quartus II settings file are specific to each revision of your design, allowing you to control the operation of the software differently for different revisions. For example, different revisions can specify different operating temperatures and different devices, so that you can compare results.

Only the valid assignments made in the Assignment Editor are saved in the Quartus II Settings File, which is located in the project directory. When you make a design constraint, the new assignment is placed on a new line at the end of the file.

When you create or update a constraint in the GUI, the Quartus II software displays the equivalent Tcl command in the **System** tab of the Messages window. You can use the displayed messages as references when making assignments using Tcl commands.

 For more information about specifying initial global constraints and software settings, refer to *Setting up and Running a Compilation* in Quartus II Help.

 For more information about how the Quartus II software uses Quartus II Settings Files, refer to the *Managing Quartus II Projects* chapter in volume 2 of the *Quartus II Handbook*.

Node, Entity, and Instance-Level Constraints

Node, entity, and instance-level constraints constrain a particular segment of the design hierarchy, as opposed to the entire design. In the Quartus II software GUI, most instance-level constraints are made with the Assignment Editor, Pin Planner, and Chip Planner. Both the Assignment Editor and Pin Planner aid you in correctly constraining your design, both passively, through device-and-assignment-determined pick lists, and actively, through live I/O checking.

You can assign logic functions to physical resources on the device, using location assignments with the Assignment Editor or the Chip Planner. Node, entity, and instance-level constraints take precedence over any global constraints that affect the same sections of the design hierarchy. You can edit and view all node and entity-level constraints you created in the Assignment Editor, or you can filter the assignments by choosing to view assignments only for specific locations, such as DSP blocks.

The Pin Planner provides a graphical representation of the target device, which allows you to easily plan, view, create, and edit pin assignments in terms of where the pins actually exist on the targeted device package. With the Pin Planner, you can visually identify I/O banks, VREF groups, edges, and differential pin pairings to assist you in the pin planning process. You can verify the legality of new and existing pin assignments with the live I/O check feature and view the results in the Live I/O Check Status window.

The Chip Planner allows you to view the device from a variety of different perspectives, and you can make precise assignments to specific floorplan locations. With the Chip Planner, you can adjust existing assignments to device resources, such as pins, logic cells, and LABs using drag and drop features and a graphical interface. You can also view equations and routing information, and demote assignments by dragging and dropping assignments to various regions in the Regions window.

 For more information about the Assignment Editor, refer to *About the Assignment Editor* in Quartus II Help. For more information about the Chip Planner, refer to *About the Chip Planner* in Quartus II Help. For more information about the Pin Planner, refer to *About the Pin Planner* in Quartus II Help.

Probing Between Components of the Quartus II GUI

The Assignment Editor, Chip Planner, and Pin Planner let you locate nodes and instances in the source files for your design in other Quartus II viewers. You can select a cell in the Assignment Editor spreadsheet and locate the corresponding item in another applicable Quartus II software window, such as the Chip Planner. To locate an item from the Assignment Editor in another window, right-click the item of interest in the spreadsheet, point to **Locate**, and click the appropriate command.

You can also locate nodes in the Assignment Editor and other constraint tools from other windows within the Quartus II software. First, select the node or nodes in the appropriate window. For example, select an entity in the **Entity** list in the **Hierarchy** tab in the Project Navigator, or select nodes in the Chip Planner. Next, right-click the selected object, point to **Locate**, and click **Locate in Assignment Editor**. The Assignment Editor opens, or it is brought to the foreground if it is already open.

- ② For more information about the Assignment Editor, refer to *About the Assignment Editor* in Quartus II Help. For more information about the Chip Planner, refer to *About the Chip Planner* in Quartus II Help. For more information about the Pin Planner, refer to *About the Pin Planner* in Quartus II Help.

SDC and the TimeQuest Timing Analyzer

You can make individual timing constraints for individual entities, nodes, and pins with the Constraints menu of the TimeQuest Timing Analyzer. The TimeQuest Timing Analyzer GUI provides easy access to timing constraints, and reporting, without requiring knowledge of SDC syntax. As you specify commands and options in the GUI, the corresponding SDC or Tcl command appears in the Console. This lets you know exactly what constraint you have added to your Synopsys Design Constraints file, and also enables you to learn SDC syntax for use in scripted flows. The GUI also provides enhanced graphical reporting features.

Individual timing assignments override project-wide requirements. You can also assign timing exceptions to nodes and paths to avoid reporting of incorrect or irrelevant timing violations. The TimeQuest timing analyzer supports point-to-point timing constraints, wildcards to identify specific nodes when making constraints, and assignment groups to make individual constraints to groups of nodes.

- ② For more information about timing constraints and the TimeQuest Timing Analyzer, refer to *About TimeQuest Timing Analysis* in Quartus II Help.

Constraining Designs with Tcl

Because **.sdc** files and **.qsf** files are both in Tcl syntax, you can modify these files to be part of a scripted constraint and compilation flow. With Quartus II Tcl packages, Tcl scripts can open projects, make the assignments procedurally that would otherwise be specified in a **.qsf** file, compile a design, and compare compilation results against known goals and benchmarks for the design. Such a script can further automate the iterative process by modifying design constraints and recompiling the design.

- ② For more information about controlling the Quartus II software with Tcl, refer to *About Quartus II Tcl Scripting* in Quartus II Help.

Quartus II Settings Files and Tcl

QSF files use Tcl syntax, but, unmodified, are not executable scripts. However, you can embed QSF constraints in a scripted iterative compilation flow, where the script that automates compilation and custom results reporting also contains the design constraints. [Example 1-1](#) shows an example QSF file with boilerplate comments removed.

Example 1-1. Quartus II Settings File

```
set_global_assignment -name FAMILY "Cyclone II"
set_global_assignment -name DEVICE EP2C35F672C6
set_global_assignment -name TOP_LEVEL_ENTITY chiptrip
set_global_assignment -name ORIGINAL_QUARTUS_VERSION 10.0
set_global_assignment -name PROJECT_CREATION_TIME_DATE "11:45:02 JUNE 08, 2010"
set_global_assignment -name LAST_QUARTUS_VERSION 10.0
set_global_assignment -name MIN_CORE_JUNCTION_TEMP 0
set_global_assignment -name MAX_CORE_JUNCTION_TEMP 85
set_instance_assignment -name PARTITION_HIERARCHY root_partition -to | -section_id Top
set_global_assignment -name PARTITION_NETLIST_TYPE SOURCE -section_id Top
set_global_assignment -name PARTITION_FITTER_PRESERVATION_LEVEL PLACEMENT_AND_ROUTING \
-section_id Top
set_global_assignment -name PARTITION_COLOR 16764057 -section_id Top
set_global_assignment -name LL_ROOT_REGION ON -section_id "Root Region"
set_global_assignment -name LL_MEMBER_STATE LOCKED -section_id "Root Region"
set_global_assignment -name STRATIX_DEVICE_IO_STANDARD "3.3-V LVTTTL"
set_location_assignment PIN_P2 -to clk2
set_location_assignment PIN_AE4 -to ticket[0]
set_location_assignment PIN_J23 -to ticket[2]
set_location_assignment PIN_Y12 -to timeo[1]
set_location_assignment PIN_N2 -to reset
set_location_assignment PIN_R2 -to timeo[7]
set_location_assignment PIN_P1 -to clk1
set_location_assignment PIN_M3 -to ticket[1]
set_location_assignment PIN_AE24 -to ~LVDS150p/nCEO~
set_location_assignment PIN_C2 -to accel
set_location_assignment PIN_K4 -to ticket[3]
set_location_assignment PIN_B3 -to stf
set_location_assignment PIN_T9 -to timeo[0]
set_location_assignment PIN_M5 -to timeo[6]
set_location_assignment PIN_J8 -to dir[1]
set_location_assignment PIN_C5 -to timeo[5]
set_location_assignment PIN_F6 -to gt1
set_location_assignment PIN_P24 -to timeo[2]
set_location_assignment PIN_B2 -to at_altera
set_location_assignment PIN_P3 -to timeo[4]
set_location_assignment PIN_M4 -to enable
set_location_assignment PIN_E3 -to ~ASDO~
set_location_assignment PIN_E5 -to dir[0]
set_location_assignment PIN_R25 -to timeo[3]
set_location_assignment PIN_D3 -to ~nCSO~
set_location_assignment PIN_G4 -to gt2
set_global_assignment -name MISC_FILE "D:/altera/chiptrip/chiptrip.dpf"
set_global_assignment -name USE_TIMEQUEST_TIMING_ANALYZER ON
set_global_assignment -name POWER_PRESET_COOLING_SOLUTION \
"23 MM HEAT SINK WITH 200 LFPM AIRFLOW"
set_global_assignment -name POWER_BOARD_THERMAL_MODEL "NONE (CONSERVATIVE)"
set_global_assignment -name SDC_FILE chiptrip.sdc
```

[Example 1-1](#) shows the way that the `set_global_assignment` Quartus II Tcl command makes all global constraints and software settings, with `set_location_assignment` constraining each I/O node in the design to a physical pin on the device.

However, after you initially create the Quartus II Settings File for your design, you can export the contents to a procedural, executable Tcl (.tcl) file. You can then use that generated script to restore certain settings after experimenting with other constraints. You can also use the generated Tcl script to archive your assignments instead of archiving the Quartus II Settings file itself.

To export your constraints as an executable Tcl script, on the Project menu, click **Generate Tcl File for Project**. [Example 1-2](#) shows the constraints in [Example 1-1](#) converted to an executable Tcl script.

Example 1-2. Generated Tcl Script for a Quartus II Project (Part 1 of 2)

```
# Quartus II: Generate Tcl File for Project
# File: chiptrip.tcl
# Generated on: Tue Jun 08 13:08:48 2010

# Load Quartus II Tcl Project package
package require ::quartus::project

set need_to_close_project 0
set make_assignments 1

# Check that the right project is open
if {[is_project_open]} {
    if {[string compare $quartus(project) "chiptrip"]} {
        puts "Project chiptrip is not open"
        set make_assignments 0
    }
} else {
    # Only open if not already open
    if {[project_exists chiptrip]} {
        project_open -revision chiptrip chiptrip
    } else {
        project_new -revision chiptrip chiptrip
    }
    set need_to_close_project 1
}

# Make assignments
if {$make_assignments} {
    set_global_assignment -name FAMILY "Cyclone II"
    set_global_assignment -name DEVICE EP2C35F672C6
    set_global_assignment -name TOP_LEVEL_ENTITY chiptrip
    set_global_assignment -name ORIGINAL_QUARTUS_VERSION 10.0
    set_global_assignment -name PROJECT_CREATION_TIME_DATE "11:45:02 JUNE 08, 2010"
    set_global_assignment -name LAST_QUARTUS_VERSION 10.0
    set_global_assignment -name MIN_CORE_JUNCTION_TEMP 0
    set_global_assignment -name MAX_CORE_JUNCTION_TEMP 85
    set_instance_assignment -name PARTITION_HIERARCHY root_partition -to | -section_id Top
    set_global_assignment -name PARTITION_NETLIST_TYPE SOURCE -section_id Top
    set_global_assignment -name PARTITION_FITTER_PRESERVATION_LEVEL PLACEMENT_AND_ROUTING \
-section_id Top
```

Example 1-2. Generated Tcl Script for a Quartus II Project (Part 2 of 2)

```
set_global_assignment -name PARTITION_COLOR 16764057 -section_id Top
set_global_assignment -name LL_ROOT_REGION ON -section_id "Root Region"
set_global_assignment -name LL_MEMBER_STATE LOCKED -section_id "Root Region"
set_global_assignment -name STRATIX_DEVICE_IO_STANDARD "3.3-V LVTTL"
set_location_assignment PIN_P2 -to clk2
set_location_assignment PIN_AE4 -to ticket[0]
set_location_assignment PIN_J23 -to ticket[2]
set_location_assignment PIN_Y12 -to timeo[1]
set_location_assignment PIN_N2 -to reset
set_location_assignment PIN_R2 -to timeo[7]
set_location_assignment PIN_P1 -to clk1
set_location_assignment PIN_M3 -to ticket[1]
set_location_assignment PIN_AE24 -to ~LVDS150p/nCEO~
set_location_assignment PIN_C2 -to accel
set_location_assignment PIN_K4 -to ticket[3]
set_location_assignment PIN_B3 -to stf
set_location_assignment PIN_T9 -to timeo[0]
set_location_assignment PIN_M5 -to timeo[6]
set_location_assignment PIN_J8 -to dir[1]
set_location_assignment PIN_C5 -to timeo[5]
set_location_assignment PIN_F6 -to gt1
set_location_assignment PIN_P24 -to timeo[2]
set_location_assignment PIN_B2 -to at_altera
set_location_assignment PIN_P3 -to timeo[4]
set_location_assignment PIN_M4 -to enable
set_location_assignment PIN_E3 -to ~ASDO~
set_location_assignment PIN_E5 -to dir[0]
set_location_assignment PIN_R25 -to timeo[3]
set_location_assignment PIN_D3 -to ~nCSO~
set_location_assignment PIN_G4 -to gt2
set_global_assignment -name MISC_FILE "D:/altera/chiptrip/chiptrip.dpf"
set_global_assignment -name USE_TIMEQUEST_TIMING_ANALYZER ON
set_global_assignment -name POWER_PRESET_COOLING_SOLUTION \
"23 MM HEAT SINK WITH 200 LFPM AIRFLOW"
set_global_assignment -name POWER_BOARD_THERMAL_MODEL "NONE (CONSERVATIVE)"
set_global_assignment -name SDC_FILE chiptrip.sdc

# Commit assignments
export_assignments

# Close project
if {$need_to_close_project} {
    project_close
}
}
```

After setting initial values for variables to control constraint creation and whether or not the project needs to be closed at the end of the script, the generated script checks to see if a project is open. If a project is open but it is not the correct project, in this case, **chiptrip**, the script prints Project chiptrip is not open to the console and does nothing else.

If no project is open, the script determines if **chiptrip** exists in the current directory. If the project exists, the script opens the project. If the project does not exist, the script creates a new project and opens the project.

The script then creates the constraints. After creating the constraints, the script writes the constraints to the Quartus II Settings File and then closes the project.

Timing Analysis with Synopsys Design Constraints and Tcl

Timing constraints used in analysis by the Quartus II TimeQuest Timing Analyzer are stored in `.sdc` files. Because they use Tcl syntax, the constraints in `.sdc` files can be incorporated into other scripts for iterative timing analysis. [Example 1-3](#) shows a basic `.sdc` file for the `chiptrip` project.

Example 1-3. Initial `.sdc` file for the `chiptrip` Project

```
# -----
set_time_unit ns
set_decimal_places 3

# -----
#
create_clock -period 10.0 -waveform { 0 5.0 } clk2 -name clk2
create_clock -period 4.0 -waveform { 0 2.0 } clk1 -name clk1

# clk1 -> dir* : INPUT_MAX_DELAY = 1 ns
set_input_delay -max 1ns -clock clk1 [get_ports dir*]
# clk2 -> time* : OUTPUT_MAX_DELAY = -2 ns
set_output_delay -max -2ns -clock clk2 [get_ports time*]
```

Similar to the constraints in the Quartus II Settings File, you can make the SDC constraints in [Example 1-3](#) part of an executable timing analysis script, as shown in [example Example 1-4](#).

Example 1-4. Tcl Script Making Basic Timing Constraints and Performing Multi-Corner Timing Analysis

```
project_open chiptrip
create_timing_netlist

#
# Create Constraints
#
create_clock -period 10.0 -waveform { 0 5.0 } clk2 -name clk2
create_clock -period 4.0 -waveform { 0 2.0 } clk1 -name clk1

# clk1 -> dir* : INPUT_MAX_DELAY = 1 ns
set_input_delay -max 1ns -clock clk1 [get_ports dir*]
# clk2 -> time* : OUTPUT_MAX_DELAY = -2 ns
set_output_delay -max -2ns -clock clk2 [get_ports time*]

#
# Perform timing analysis for several different sets of operating conditions
#
foreach_in_collection oc [get_available_operating_conditions] {
    set_operating_conditions $oc
    update_timing_netlist

    report_timing -setup -npaths 1
    report_timing -hold -npaths 1
    report_timing -recovery -npaths 1
    report_timing -removal -npaths 1
    report_min_pulse_width -nworst 1
}

delete_timing_netlist
project_close
```

The script in [Example 1-4](#) opens the project, creates a timing netlist, then constrains the two clocks in the design and applies input and output delay constraints. The clock settings and delay constraints are identical to those in the `.sdc` file shown in [Example 1-3](#). The next section of the script updates the timing netlist for the constraints and performs multi-corner timing analysis on the design.

A Fully Iterative Scripted Flow

You can use the `::quartus::flow` Tcl package and other packages in the Quartus II Tcl API to add flow control to modify constraints and recompile your design in an automated flow. You can combine your timing constraints with the other constraints for your design, and embed them in an executable Tcl script that also iteratively compiles your design as different constraints are applied.

Each time such a modified generated script is run, it can modify the `.qsf` file and `.sdc` file for your project based on the results of iterative compilations, effectively replacing these files for the purposes of archiving and version control using industry-standard source control methods and practices.

This type of scripted flow can include automated compilation of a design, modification of design constraints, and recompilation of the design, based on how you foresee results and pre-determine next-step constraint changes in response to those results.

-  For more information about the Quartus II Tcl API, refer to [API Functions for Tcl](#) in Quartus II Help. For more information about controlling the Quartus II software with Tcl scripts, refer to [About Quartus II Tcl Scripting](#) in Quartus II Help.

Document Revision History

[Table 1-1](#) shows the revision history for this chapter.

Table 1-1. Document Revision History

Date	Version	Changes
December 2010	10.0.1	Template update.
July 2010	10.0.0	Rewrote chapter to more broadly cover all design constraint methods. Removed procedural steps and user interface details, and replaced with links to Quartus II Help.
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Added two notes. ■ Minor text edits.
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Revised and reorganized the entire chapter. ■ Added section “Probing to Source Design Files and Other Quartus II Windows” on page 1-2. ■ Added description of node type icons (Table 1-3). ■ Added explanation of wildcard characters.
November 2008	8.1.0	Changed to 8½” × 11” page size. No change to content.
May 2008	8.0.0	Updated Quartus II software 8.0 revision and date.

-  For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

FPGA design software that easily integrates into your design flow saves time and improves productivity. The Altera® Quartus® II software provides you with a command-line executable for each step of the FPGA design flow to make the design process customizable and flexible.

The benefits provided by command-line executables include:

- Command-line control over each step of the design flow
- Easy integration with scripted design flows including makefiles
- Reduced memory requirements
- Improved performance

The command-line executables are also completely compatible with the Quartus II GUI, allowing you to use the exact combination of tools that you prefer.

This chapter describes how to take advantage of Quartus II command-line executables, and provides several examples of scripts that automate different segments of the FPGA design flow. This chapter includes the following topics:

- “Benefits of Command-Line Executables”
- “Introductory Example” on page 2-2
- “Compilation with `quartus_sh --flow`” on page 2-7
- “The MegaWizard Plug-In Manager” on page 2-11
- “Command-Line Scripting Examples” on page 2-17

Benefits of Command-Line Executables

The Quartus II command-line executables provide control over each step of the design flow. Each executable includes options to control commonly used software settings. Each executable also provides detailed, built-in help describing its function, available options, and settings.

Command-line executables allow for easy integration with scripted design flows. You can easily create scripts in any language with a series of commands. These scripts can be batch-processed, allowing for integration with distributed computing in server farms. You can also integrate the Quartus II command-line executables in makefile-based design flows. These features enhance the ease of integration between the Quartus II software and other EDA synthesis, simulation, and verification software.

Command-line executables add flexibility without sacrificing the ease-of-use of the Quartus II GUI. You can use the Quartus II GUI and command-line executables at different stages in the design flow. For example, you might use the Quartus II GUI to edit the floorplan for the design, use the command-line executables to perform place-and-route, and return to the Quartus II GUI to perform debugging with the Chip Editor.

Command-line executables reduce the amount of memory required during each step in the design flow. Because each executable targets only one step in the design flow, the executables themselves are relatively compact, both in file size and the amount of memory used when running. This memory usage reduction improves performance, and is particularly beneficial in design environments where computer networks or workstations are heavily used with reduced memory.

- ② For a complete list of the Quartus II command-line executables, refer to *Using the Quartus II Executables in Shell Scripts* in Quartus II Help.

Introductory Example

The following introduction to command-line executables demonstrates how to create a project, fit the design, and generate programming files.

The tutorial design included with the Quartus II software is used to demonstrate this functionality. If installed, the tutorial design is found in the `<Quartus II directory>/qdesigns/fir_filter` directory.

Before making changes, copy the tutorial directory and type the four commands shown in [Example 2-1](#) at a command prompt in the new project directory.



The `<Quartus II directory>/quartus/bin` directory must be in your PATH environment variable.

Example 2-1. Introductory Example

```
quartus_map filtref --source=filtref.bdf --family=CYCLONE ←
quartus_fit filtref --part=EP1C12F256C6 --fmax=80MHz ←
quartus_asm filtref ←
quartus_sta filtref ←
```

The `quartus_map filtref --source=filtref.bdf --family=CYCLONE` command creates a new Quartus II project called **filtref** with **filtref.bdf** as the top-level file. It targets the Cyclone® device family and performs logic synthesis and technology mapping on the design files.

The `quartus_fit filtref --part=EP1C12Q240C6 --fmax=80MHz` command performs fitting on the **filtref** project. This command specifies an EP1C12Q240C6 device and the Fitter attempts to meet a global f_{MAX} requirement of 80 MHz and a global t_{SU} requirement of 8 ns.

The `quartus_asm filtref` command creates programming files for the **filtref** project.

The `quartus_sta filtref` command performs basic timing analysis on the **filtref** project using the Quartus II TimeQuest Timing Analyzer, reporting worst-case setup slack, worst-case hold slack, and other measurements.

 The TimeQuest Timing Analyzer employs Synopsys Design Constraints to fully analyze the timing of your design. For more information about using all of the features of the `quartus_sta` executable, refer to the *TimeQuest Timing Analyzer Quick Start Tutorial*.

You can put the four commands from [Example 2-1](#) into a batch file or script file, and run them. For example, you can create a simple UNIX shell script called `compile.sh`, which includes the code shown in [Example 2-2](#).

Example 2-2. UNIX Shell Script: `compile.sh`

```
#!/bin/sh
PROJECT=filtref
TOP_LEVEL_FILE=filtref.bdf
FAMILY=Cyclone
PART=EP1C12F256C6
FMAX=80MHz
quartus_map $PROJECT --source=$TOP_LEVEL_FILE --family=$FAMILY
quartus_fit $PROJECT --part=$PART --fmax=$FMAX
quartus_asm $PROJECT
quartus_sta $PROJECT
```

Edit the script as necessary and compile your project.

Command-Line Scripting Help

Help for command-line executables is available through different methods. You can access help built in to the executables with command-line options. You can use the Quartus II Command-Line and Tcl API Help browser for an easy graphical view of the help information.

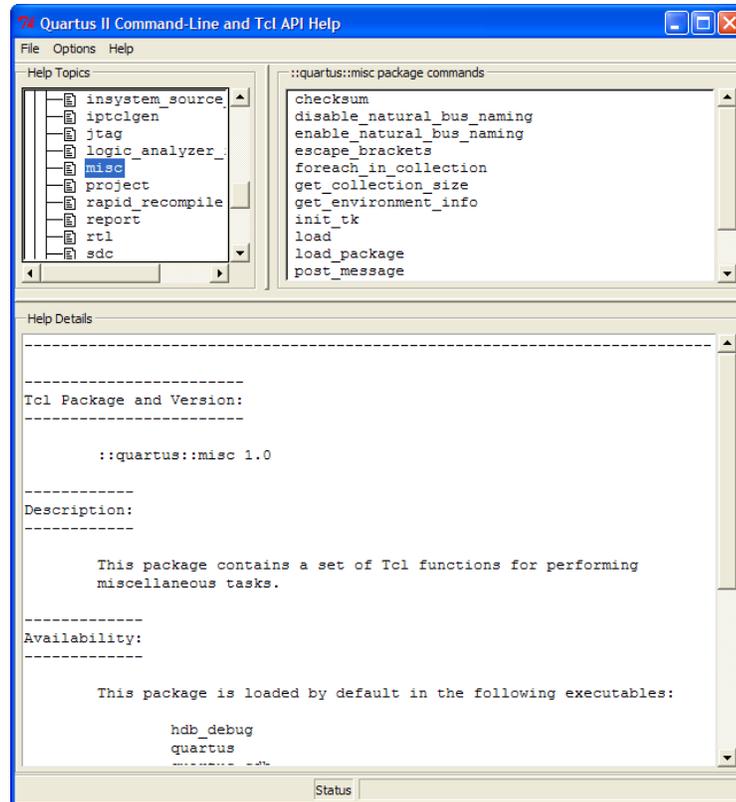
To use the Quartus II Command-Line and Tcl API Help browser, type the following command:

```
quartus_sh --qhelp ←
```

This command starts the Quartus II Command-Line and Tcl API Help browser, a viewer for information about the Quartus II Command-Line executables and Tcl API ([Figure 2-1](#)).

Use the `-h` option with any of the Quartus II Command-Line executables to get a description and list of supported options. Use the `--help=<option name>` option for detailed information about each option.

Figure 2-1. Quartus II Command-Line and Tcl API Help Browser



Command-Line Option Details

Command-line options are provided for many common global project settings and for performing common tasks. You can use either of two methods to make assignments to an individual entity. If the project exists, open the project in the Quartus II GUI, change the assignment, and close the project. The changed assignment is updated in the Quartus II Settings File. Any command-line executables that are run after this update use the updated assignment. For more information refer to “[Option Precedence](#)” on page 2-5. You can also make assignments using the Quartus II Tcl scripting API. If you want to completely script the creation of a Quartus II project, choose this method.

- For more information about the Quartus II Tcl scripting API, refer to the [Tcl Scripting](#) chapter in volume 2 of the *Quartus II Handbook*. For more information about Quartus II project settings and assignments, refer to the [QSF Reference Manual](#).

Option Precedence

If you use command-line executables, you must be aware of the precedence of various project assignments and how to control the precedence. Assignments for a particular project exist in the Quartus II Settings File for the project. Assignments for a project can also be made with command-line options. Project assignments are reflected in compiler database files that hold intermediate compilation results and reflect assignments made in the previous project compilation.

All command-line options override any conflicting assignments found in the Quartus II Settings File or the compiler database files. There are two command-line options to specify whether the Quartus II Settings File or compiler database files take precedence for any assignments not specified as command-line options.



Any assignment not specified as a command-line option or found in the Quartus II Settings File or compiler database file is set to its default value.

The file precedence command-line options are `--read_settings_files` and `--write_settings_files`.

By default, the `--read_settings_files` and `--write_settings_files` options are turned on. Turning on the `--read_settings_files` option causes a command-line executable to read assignments from the Quartus II Settings File instead of from the compiler database files. Turning on the `--write_settings_files` option causes a command-line executable to update the Quartus II Settings File to reflect any specified options, as happens when you close a project in the Quartus II GUI.

If you use command-line executables, be aware of the precedence of various project assignments and how to control the precedence. Assignments for a particular project can exist in three places:

- The Quartus II Settings File for the project
- The result of the last compilation, in the `/db` directory, which reflects the assignments that existed when the project was compiled
- Command-line options

Table 2-1 lists the precedence for reading assignments depending on the value of the `--read_settings_files` option.

Table 2-1. Precedence for Reading Assignments

Option Specified	Precedence for Reading Assignments
<code>--read_settings_files = on</code> (Default)	<ol style="list-style-type: none"> 1. Command-line options 2. Quartus II Settings File 3. Project database (db directory, if it exists) 4. Quartus II software defaults
<code>--read_settings_files = off</code>	<ol style="list-style-type: none"> 1. Command-line options 2. Project database (db directory, if it exists) 3. Quartus II software defaults

Table 2–2 lists the locations to which assignments are written, depending on the value of the `--write_settings_files` command-line option.

Table 2–2. Location for Writing Assignments

Option Specified	Location for Writing Assignments
<code>--write_settings_files = on</code> (Default)	Quartus II Settings File and compiler database
<code>--write_settings_files = off</code>	Compiler database

Example 2–3 assumes that a project named `fir_filter` exists, and that the analysis and synthesis step has been performed (using the `quartus_map` executable).

Example 2–3. Write Settings Files

```
quartus_fit fir_filter --fmax=80MHz ←
quartus_sta fir_filter ←
mv fir_filter_sta.rpt fir_filter_1_sta.rpt ←
quartus_fit fir_filter --fmax=100MHz --write_settings_files=off ←
quartus_sta fir_filter ←
mv fir_filter_sta.rpt fir_filter_2_sta.rpt ←
```

The first command, `quartus_fit fir_filter --fmax=80MHz`, runs the `quartus_fit` executable and specifies a global f_{MAX} requirement of 80 MHz.

The second command, `quartus_sta fir_filter`, runs timing analysis for the results of the previous fit.

The third command uses the UNIX `mv` command to copy the report file output from `quartus_sta` to a file with a new name, so that the results are not overwritten by subsequent timing analysis.

The fourth command reruns the Fitter with a global f_{MAX} requirement of 100 MHz. By specifying the `--write_settings_files=off` option, the command-line executable does not update the Quartus II Settings File to reflect the changed f_{MAX} requirement. The compiler database files reflect the changed f_{MAX} requirement. If the `--write_settings_files=off` option is not specified, the command-line executable updates the Quartus II Settings File to reflect the 100-MHz global f_{MAX} requirement.

The fifth command reruns timing analysis, and the sixth command renames the report file, so that it is not overwritten by subsequent timing analysis.

Use the options `--read_settings_files=off` and `--write_settings_files=off` (where appropriate) to optimize the way that the Quartus II software reads and updates settings files. Example 2–4 shows how to avoid unnecessary reading and writing.

Example 2–4. Avoiding Unnecessary Reading and Writing

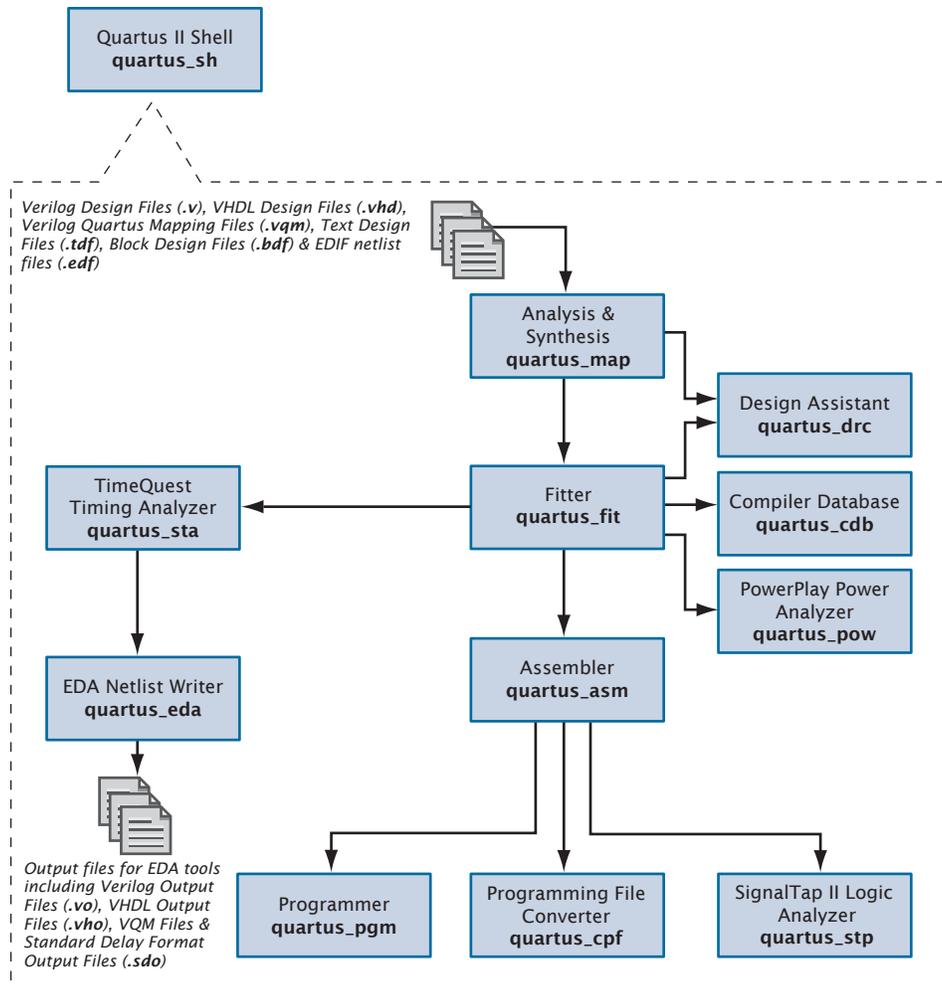
```
quartus_map filtref --source=filtref --part=EP1C12F256C6 ←
quartus_fit filtref --fmax=100MHz --read_settings_files=off ←
quartus_asm filtref --read_settings_files=off --write_settings_files=off ←
```

In Example 2–4, the `quartus_asm` executable does not read or write settings files because they do not change any settings in the project.

Compilation with quartus_sh --flow

Figure 2-2 shows a typical Quartus II FPGA design flow.

Figure 2-2. Typical Design Flow



Use the `quartus_sh` executable with the `--flow` option to perform a complete compilation flow with a single command. (For information about specialized flows, type `quartus_sh --help=flow` at a command prompt.) The `--flow` option supports the smart recompile feature and efficiently sets command-line arguments for each executable in the flow.

The following example runs compilation, timing analysis, and programming file generation with a single command:

```
quartus_sh --flow compile filtref
```

Text-Based Report Files

Each command-line executable creates a text report file when it is run. These files report success or failure, and contain information about the processing performed by the executable.

Report file names contain the revision name and the short-form name of the executable that generated the report file: `<revision>.<executable>.rpt`. For example, using the `quartus_fit` executable to place and route a project with the revision name **design_top** generates a report file named **design_top.fit.rpt**. Similarly, using the `quartus_sta` executable to perform timing analysis on a project with the revision name **fir_filter** generates a report file named **fir_filter.sta.rpt**.

- ❓ As an alternative to parsing text-based report files, you can use the `::quartus::report` Tcl package. For more information about this package, refer to `::quartus::report` in Quartus II Help.

You can use command-line executables in scripts that control a design flow that uses other software in addition to the Quartus II software. For example, if your design flow uses other synthesis or simulation software, and you can run the other software at a command prompt, you can include it in a single script. The Quartus II command-line executables include options for common global project settings and operations, but you must use a Tcl script or the Quartus II GUI to set up a new project and apply individual constraints, such as pin location assignments and timing requirements. Command-line executables are very useful for working with existing projects, for making common global settings, and for performing common operations. For more flexibility in a flow, use a Tcl script, which makes it easier to pass data between different stages of the design flow and have more control during the flow.

-  For more information about Tcl scripts, refer to the *Tcl Scripting* chapter in volume 2 of the *Quartus II Handbook*, or *About Quartus II Tcl Scripting* in Quartus II Help.

For example, your script could run other synthesis software, then place-and-route the design in the Quartus II software, then generate output netlists for other simulation software. [Example 2-5](#) shows how to do this with a UNIX shell script for a design that targets a Cyclone II device.

Example 2-5. Script for End-to-End Flow (Part 1 of 2)

```
#!/bin/sh
# Run synthesis first.
# This example assumes you use Synplify software
synplify -batch synthesize.tcl

# If your Quartus II project exists already, you can just
# recompile the design.
# You can also use the script described in a later example to
# create a new project from scratch
quartus_sh --flow compile myproject

# Use the quartus_sta executable to do fast and slow-model
# timing analysis
quartus_sta myproject --model=slow
quartus_sta myproject --model=fast

# Use the quartus_eda executable to write out a gate-level
# Verilog simulation netlist for ModelSim
quartus_eda my_project --simulation --tool=modelsim --format=verilog
```

Example 2-5. Script for End-to-End Flow (Part 2 of 2)

```
# Perform the simulation with the ModelSim software
vlib cycloneii_ver
vlog -work cycloneii_ver /opt/quartusii/eda/sim_lib/cycloneii_atoms.v
vlib work
vlog -work work my_project.vo
vsim -L cycloneii_ver -t lps work.my_project
```

Makefile Implementation

You can use the Quartus II command-line executables in conjunction with the `make` utility to automatically update files when other files they depend on change. The file dependencies and commands used to update files are specified in a text file called a makefile.

To facilitate easier development of efficient makefiles, the following “smart action” scripting command is provided with the Quartus II software:

```
quartus_sh --determine_smart_action ←
```

Because assignments for a Quartus II project are stored in the Quartus II Settings File (`.qsf`), including it in every rule results in unnecessary processing steps. For example, updating a setting related to programming file generation (which requires re-running only `quartus_asm`) modifies the Quartus II Settings File, requiring a complete recompilation if the Quartus II Settings File is included in every rule.

The smart action command determines the earliest command-line executable in the compilation flow that must be run based on the current Quartus II Settings File, and generates a change file corresponding to that executable. For a given command-line executable named `quartus_<executable>`, the change file is named with the format `<executable>.chg`. For example, if `quartus_map` must be re-run, the smart action command creates or updates a file named `map.chg`. Thus, rather than including the Quartus II Settings File in each makefile rule, include only the appropriate change file.

Example 2-6 uses change files and the smart action command. You can copy and modify it for your own use. A copy of this example is included in the help for the makefile option, which is available by typing:

```
quartus_sh --help=makefiles ←
```

Example 2-6. Sample Makefile (Part 1 of 2)

```
#####
# Project Configuration:
#
# Specify the name of the design (project), the Quartus II Settings
# File (.qsf), and the list of source files used.
#####

PROJECT = chiptrip
SOURCE_FILES = auto_max.v chiptrip.v speed_ch.v tick_cnt.v time_cnt.v
ASSIGNMENT_FILES = chiptrip.qpf chiptrip.qsf

#####
# Main Targets
#
# all: build everything
# clean: remove output files and database
#####
all: smart.log $(PROJECT).asm.rpt $(PROJECT).sta.rpt

clean:
    rm -rf *.rpt *.chg smart.log *.htm *.eqn *.pin *.sof *.pof db

map: smart.log $(PROJECT).map.rpt
fit: smart.log $(PROJECT).fit.rpt
asm: smart.log $(PROJECT).asm.rpt
sta: smart.log $(PROJECT).sta.rpt
smart: smart.log
#####
# Executable Configuration
#####

MAP_ARGS = --family=Stratix
FIT_ARGS = --part=EP1S20F484C6
ASM_ARGS =
STA_ARGS =

#####
# Target implementations
#####

STAMP = echo done >

$(PROJECT).map.rpt: map.chg $(SOURCE_FILES)
    quartus_map $(MAP_ARGS) $(PROJECT)
    $(STAMP) fit.chg

$(PROJECT).fit.rpt: fit.chg $(PROJECT).map.rpt
    quartus_fit $(FIT_ARGS) $(PROJECT)
    $(STAMP) asm.chg
    $(STAMP) sta.chg

$(PROJECT).asm.rpt: asm.chg $(PROJECT).fit.rpt
    quartus_asm $(ASM_ARGS) $(PROJECT)

$(PROJECT).sta.rpt: sta.chg $(PROJECT).fit.rpt
    quartus_sta $(STA_ARGS) $(PROJECT)
smart.log: $(ASSIGNMENT_FILES)
    quartus_sh --determine_smart_action $(PROJECT) > smart.log
```

Example 2-6. Sample Makefile (Part 2 of 2)

```
#####  
# Project initialization  
#####  
  
$(ASSIGNMENT_FILES):  
    quartus_sh --prepare $(PROJECT)  
  
map.chg:  
    $(STAMP) map.chg  
fit.chg:  
    $(STAMP) fit.chg  
sta.chg:  
    $(STAMP) sta.chg  
asm.chg:  
    $(STAMP) asm.chg
```

A Tcl script is provided with the Quartus II software to create or modify files that are specified as dependencies in the make rules, assisting you in makefile development. Complete information about this Tcl script and how to integrate it with makefiles is available by running the following command:

```
quartus_sh --help=determine_smart_action ←
```

The MegaWizard Plug-In Manager

The MegaWizard™ Plug-In Manager provides a GUI-based flow to configure megafunction and IP variation files. However, you can use command-line options for the **qmegawiz** executable to modify, update, or create variation files without using the GUI. This capability is useful in a fully scripted design flow, or in cases where you want to generate variation files without using the wizard GUI flow.

The MegaWizard Plug-In Manager has three functions:

- Providing an interface for you to select the output file or files
- Running a specific MegaWizard Plug-In
- Creating output files (such as variation files, symbol files, and simulation netlist files)

Each MegaWizard Plug-In provides a user interface for configuring the variation, and performs validation and error checking of your selected ports and parameters. When you create or update a variation with the GUI, the parameters and values are entered through the GUI provided by the Plug-In. When you create a Plug-In variation with the command line, you provide the parameters and values as command-line options.

Example 2-7 shows how to create a new variation file at a system command prompt.

Example 2-7. MegaWizard Plug-In Manager Command-Line Executable

```
qmegawiz [options] [module=<module name>|wizard=<wizard name>] [<param>=<value> ...  
    <port>=<used|unused> ...] [OPTIONAL_FILES=<optional files>] <variation file name>
```

When you use **qmegawiz** to update an existing variation file, the module or wizard name is not required.

If a megafunction changes between software versions, the variation files must be regenerated. To do this, use `qmegawiz -silent <variation file name>`. For example, if your design contains a variation file called `myrom.v`, type the following command:

```
qmegawiz -silent myrom.v ←
```

For more information on updating megafunction variation files as part of a scripted flow, refer to “Regenerating Megafunctions After Updating the Quartus II Software” on page 2-23.

Table 2-3 describes the supported options.

Table 2-3. qmegawiz Options

Option	Description
<code>-silent</code>	Run the MegaWizard Plug-In Manager in command-line mode, without displaying the GUI.
<code>-f: <param file></code>	A file that contains all options for the <code>qmegawiz</code> command. Refer to “Parameter File” on page 2-16.
<code>-p: <working directory></code>	Sets the default working directory. Refer to “Working Directory” on page 2-17.

For information about specifying the module name or wizard name, refer to “Module and Wizard Names” on page 2-13.

For information about specifying ports and parameters, refer to “Ports and Parameters” on page 2-14.

For information about generating optional files, refer to “Optional Files” on page 2-15.

For information about specifying the variation file name, refer to “Variation File Name” on page 2-17.

Command-Line Support

Only the MegaWizard Plug-Ins listed in Table 2-4 support creation and update in command-line mode. For Plug-Ins not listed in the table, you must use the MegaWizard Plug-In Manager GUI for creation and updates.

Table 2-4. MegaWizard Plug-Ins with Command Line Support (Part 1 of 2)

MegaWizard Plug-In	Wizard Name	Module Name
<code>alt2gxb</code>	ALT2GXB	<code>alt2gxb</code>
<code>alt4gxb</code>	ALTGX	<code>alt4gxb</code>
<code>altasmi_parallel</code>	ALTASMI_PARALLEL	<code>altasmi_parallel</code>
<code>altclkctrl</code>	ALTCLKCTRL	<code>altclkctrl</code>
<code>altddio_bidir</code>	ALTDIO_BIDIR	<code>altddio_bidir</code>
<code>altddio_in</code>	ALTDIO_IN	<code>altddio_in</code>
<code>altddio_out</code>	ALTDIO_OUT	<code>altddio_out</code>
<code>altecc_decoder</code>	ALTECC	<code>altecc_decoder</code>
<code>altecc_encoder</code>		<code>altecc_encoder</code>
<code>altfp_abs</code>	ALTFP_ABS	<code>altfp_abs</code>

Table 2-4. MegaWizard Plug-Ins with Command Line Support (Part 2 of 2)

MegaWizard Plug-In	Wizard Name	Module Name
altfp_add_sub	ALTFP_ADD_SUB	altfp_add_sub
altfp_compare	ALTFP_COMPARE	altfp_compare
altfp_convert	ALTFP_CONVERT	altfp_convert
altfp_div	ALTFP_DIV	altfp_div
altfp_exp	ALTFP_EXP	altfp_exp
altfp_inv_sqrt	ALTFP_INV_SQRT	altfp_inv_sqrt
altfp_inv	ALTFP_INV	altfp_inv
altfp_log	ALTFP_LOG	altfp_log
altfp_matrix_inv	ALTFP_MATRIX_INV	altfp_matrix_inv
altfp_matrix_mult	ALTFP_MATRIX_MULT	altfp_matrix_mult
altfp_mult	ALTFP_MULT	altfp_mult
altfp_sincos	ALTFP_SINCOS	altfp_sincos
altfp_sqrt	ALTFP_SQRT	altfp_sqrt
altiobuf_bidir	ALTIobuf	altiobuf_bidir
altiobuf_in		altiobuf_in
altiobuf_out		altiobuf_out
altlvds_rx	ALTLVDS	altlvds_rx
altlvds_tx		altlvds_tx
altmult_accum	ALTMULT_ACCUM (MAC)	altmult_accum
altmult_complex	ALTMULT_COMPLEX	altmult_complex
altotp	ALTOtp	altotp
altpll_reconfig	ALTPLL_RECONFIG	altpll_reconfig
altpll	ALTPLL	altpll
altremote_update	ALTRemOTE_UPDATE	altremote_update
altshift_taps	ALTSHIFT_TAPS	altshift_taps
altsyncram	RAM: 2-PORT	altsyncram
	RAM: 1-PORT	
	ROM: 1-PORT	
alttemp_sense	ALTTEMP_SENSE	alttemp_sense
alt_c3gxb	ALT_C3GXB	alt_c3gxb
dcfifo	FIFO	dcfifo
scfifo		scfifo

Module and Wizard Names

You must specify the wizard or module name, shown in [Table 2-4](#), as a command-line option when you create a variation file. Use the option `module=<module name>` to specify the module, or use the option `wizard=<wizard name>` to specify the wizard. If there are spaces in the wizard or module name, enclose the name in double quotes, for example:

```
wizard="RAM: 2-PORT"
```

When there is a one-to-one mapping between the MegaWizard Plug-In and the wizard name and the module name, you can use either the wizard option or the module option.

When there are multiple wizard names that correspond to one module name, you should use the wizard option to specify one wizard.

When there are multiple module names that correspond to one wizard name, you should use the module option to specify one module. For example, use the module option if you create a FIFO because one wizard is common to both modules. However, you should use the wizard option if you create a RAM, because one module is common to three wizards.

If you edit or update an existing variation file, the wizard or module option is not necessary, because information about the wizard or module is already in the variation file.

Ports and Parameters

Ports and parameters for each MegaWizard Plug-In are described in Quartus II Help, and in the [Megafunction User Guides](#) on the Altera website. You should use these references to determine appropriate values for each port and parameter required for a particular variation configuration. Refer to “[Strategies to Determine Port and Parameter Values](#)” for more information. You do not have to specify every port and parameter supported by a Plug-In. The MegaWizard Plug-In Manager uses default values for any port or parameter you do not specify.

Specify ports as used or unused, for example:

```
<port>=used  
<port>=unused
```

You can specify port names in any order. Grouping does not matter. Separate port configuration options from each other with spaces.

Specify a value for a parameter with the equal sign, for example:

```
<parameter>=<value>
```

You can specify parameters in any order. Grouping does not matter. Separate parameter configuration options from each other with spaces. You can specify port names and parameter names in upper or lower case; case does not matter.

All MegaWizard Plug-Ins allow you to specify the target device family with the `INTENDED_DEVICE_FAMILY` parameter, as shown in the following example:

```
qmegawiz wizard=<wizard> INTENDED_DEVICE_FAMILY="Cyclone III" <file>
```

You must specify enough ports and parameters to create a legal configuration of the Plug-In. When you use the GUI flow, each MegaWizard Plug-In performs validation and error checking for the particular ports and parameters you choose. When you use command-line options to specify ports and parameters, you must ensure that the ports and parameters you use are complete and valid for your particular configuration.

For example, when you use a RAM Plug-In to configure a RAM to be 32 words deep, the Plug-In automatically configures an address port that is five bits wide. If you use the command-line flow to configure a RAM that is 32 words deep, you must use one option to specify the depth of the RAM, then calculate the width of the address port and specify that width with another option.

Invalid Configurations

If the combination of default and specified ports and parameters is not complete to create a legal configuration of the Plug-In, `qmegawiz` generates an error message that indicates what is missing and what values are supported. If the combination of default and specified ports and parameters results in an illegal configuration of the Plug-In, `qmegawiz` generates an error message that indicates what is illegal, and displays the legal values.

Strategies to Determine Port and Parameter Values

For simple Plug-In variations, it is often easy to determine appropriate port and parameter values with the information in Quartus II Help and other megafunction documentation. For example, determining that a 32-word-deep RAM requires an address port that is five bits wide is straightforward. For complex Plug-In variations, an option in the GUI might affect multiple port and parameter settings, so it can be difficult to determine a complete set of ports and parameters. In this case, you should use the GUI to generate a variation file that includes the ports and parameters for your desired configuration. Open the variation file in a text editor and use the port and parameter values in the variation file as command-line options.

Optional Files

In addition to the variation file, the MegaWizard Plug-In Manager can generate other files, such as instantiation templates, simulation netlists, and symbols for graphic design entry. Use the `OPTIONAL_FILES` parameter to control whether the MegaWizard Plug-In Manager generates optional files. [Table 2-5](#) lists valid arguments for the `OPTIONAL_FILES` parameter.

Table 2-5. Arguments for the `OPTIONAL_FILES` Parameter

Argument	Description
INST	Controls the generation of the <code><variation>_inst.v</code> file.
INC	Controls the generation of the <code><variation>.inc</code> file.
CMP	Controls the generation of the <code><variation>.cmp</code> file.
BSF	Controls the generation of the <code><variation>.bsf</code> file.
BB	Controls the generation of the <code><variation>_bb.v</code> file.
SIM_NETLIST	Controls the generation of the simulation netlist file, wherever there is wizard support.
SYNTH_NETLIST	Controls the generation of the synthesis netlist file, wherever there is wizard support.
ALL	Generates all applicable optional files.
NONE	Disables the generation of all optional files.

Specify a single optional file, for example:

```
OPTIONAL_FILES=<argument>
```

Specify multiple optional files separated by a vertical bar character, for example:

```
OPTIONAL_FILES=<argument 1>|...|<argument n>
```

If you prefix an argument with a dash (for example, `-BB`), it is excluded from the generated optional files. If any of the optional files exist when you run `qmegawiz` and they are excluded in the `OPTIONAL_FILES` parameter (with the `NONE` argument, or prefixed with a dash), they are deleted.

You can combine the `ALL` argument with other excluded arguments to generate “all files except *<excluded files>*.” You can combine the `NONE` argument with other included arguments to generate “no files except *<files>*.”

When you combine multiple arguments, they are processed from left to right, and arguments evaluated later have precedence over arguments evaluated earlier. Therefore, the `ALL` or `NONE` argument should be the first in a combination of multiple arguments. When `ALL` is the first argument, all optional files are generated before exclusions are processed (deleted). When `NONE` is the first argument, none of the optional files are generated (in other words, any that exist are deleted), then any files you subsequently specify are generated.

Table 2-6 shows examples for the `OPTIONAL_FILES` parameter and describes the result of each example.

Table 2-6. Examples of Different Optional File Arguments

Example Values for OPTIONAL_FILES	Description
BB	The optional file <i><variation>_bb.v</i> is generated, and no optional files are deleted
BB INST	The optional file <i><variation>_bb.v</i> is generated, then the optional file <i><variation>_inst.v</i> is generated, and no optional files are deleted.
NONE	No optional files are generated, and any existing optional files are deleted.
NONE INC BSF	Any existing optional files are deleted, then the optional file <i><variation>.inc</i> is generated, then the optional file <i><variation>.bsf</i> is generated.
ALL -INST	All optional files are generated, then <i><variation>_inst.v</i> is deleted if it exists.
-BB	The optional file <i><variation>_bb.v</i> is deleted if it exists.
-BB INST	The optional file <i><variation>_bb.v</i> is deleted if it exists, then the optional file <i><variation>_inst.v</i> is generated.

The `qmegawiz` command accepts the `ALL` argument combined with other included file arguments, for example, `ALL|BB`, but that combination is equivalent to `ALL` because first all optional files are generated, then the file *<variation>_bb.v* is generated (again). Additionally, the software accepts the `NONE` argument combined with other excluded file arguments, for example, `NONE|-BB`, but that combination is equivalent to `NONE` because no optional files are generated (any that exist are deleted), then the file *<variation>_bb.v* is deleted if it exists.

Parameter File

You can put all parameter values and port values in a file, and pass the file name as an argument to `qmegawiz` with the `-f:<parameter file>` option. For example, the following command specifies a parameter file named `rom_params.txt`:

```
qmegawiz -silent module=altsyncram -f:rom_params.txt myrom.v ←
```

The `rom_params.txt` parameter file can include options similar to the following:

```
RAM_BLOCK_TYPE=M4K DEVICE_FAMILY=Stratix WIDTH_A=5 WIDTHAD_A=5
NUMWORDS_A=32 INIT_FILE=rom.hex OPERATION_MODE=ROM
```



If you use the `-f` option and the `-p` option together, the MegaWizard Plug-In Manager sources the parameter file in a directory specified with the `-p` option, or in a directory relative to that directory. For example, if you specify `C:\project\work` with the `-p` option and `work\params.txt` with the `-f` option, the MegaWizard Plug-In Manager attempts to source the file `params.txt` in `C:\project\work\work`.

Working Directory

You can change the working directory that `qmegawiz` uses when it generates files. By default, the working directory is the current directory when you execute the `qmegawiz` command. Use the `-p` option to specify a different working directory, for example:

```
-p:<working directory>
```

You can specify the working directory with an absolute or relative path. Specify an alternative working directory any time you do not want files generated in the current directory. The alternative working directory can be useful if you generate multiple variations in a batch script, and keep generated files for the different Plug-In variations in separate directories.

Variation File Name

The language used for a variation file depends on the file extension of the variation file name. The MegaWizard Plug-In Manager creates HDL output files in a language based on the file name extension. Therefore, you must always specify a complete file name, including file extension, as the last argument to the `qmegawiz` command.

Table 2-7 shows the file extension that corresponds to supported HDL types.

Table 2-7. Variation File Extensions

Variation File HDL Type	Required File Extension
Verilog HDL	.v
VHDL	.vhd
AHDL	.tdf

Command-Line Scripting Examples

This section presents various examples of command-line executable use.

Create a Project and Apply Constraints

The command-line executables include options for common global project settings and commands. To apply constraints such as pin locations and timing assignments, run a Tcl script with the constraints in it. You can write a Tcl constraint file yourself, or generate one for an existing project. From the Project menu, click **Generate Tcl File for Project**.

Example 2-8 creates a project with a Tcl script and applies project constraints using the tutorial design files in the `<Quartus II installation directory>/qdesigns/fir_filter/` directory.

Example 2-8. Tcl Script to Create Project and Apply Constraints

```
project_new filtref -overwrite
# Assign family, device, and top-level file
set_global_assignment -name FAMILY Cyclone
set_global_assignment -name DEVICE EP1C12F256C6
set_global_assignment -name BDF_FILE filtref.bdf
# Assign pins
set_location_assignment -to clk Pin_28
set_location_assignment -to clkx2 Pin_29
set_location_assignment -to d[0] Pin_139
set_location_assignment -to d[1] Pin_140
# Other assignments could follow
project_close
```

Save the script in a file called `setup_proj.tcl` and type the commands illustrated in **Example 2-9** at a command prompt to create the design, apply constraints, compile the design, and perform fast-corner and slow-corner timing analysis. Timing analysis results are saved in two files, `filtref_sta_1.rpt` and `filtref_sta_2.rpt`.

Example 2-9. Script to Create and Compile a Project

```
quartus_sh -t setup_proj.tcl ←
quartus_map filtref ←
quartus_fit filtref ←
quartus_asm filtref ←
quartus_sta filtref --model=fast --export_settings=off ←
mv filtref_sta.rpt filtref_sta_1.rpt ←
quartus_sta filtref --export_settings=off ←
mv filtref_sta.rpt filtref_sta_2.rpt ←
```

Type the following commands to create the design, apply constraints, and compile the design, without performing timing analysis:

```
quartus_sh -t setup_proj.tcl ←
quartus_sh --flow compile filtref ←
```

The `quartus_sh --flow compile` command performs a full compilation, and is equivalent to clicking the **Start Compilation** button in the toolbar.

Check Design File Syntax

The UNIX shell script example shown in **Example 2-10** assumes that the Quartus II software `fir_filter` tutorial project exists in the current directory. You can find the `fir_filter` project in the `<Quartus II directory>/qdesigns/fir_filter` directory unless the Quartus II software tutorial files are not installed.

The `--analyze_file` option performs a syntax check on each file. The script checks the exit code of the `quartus_map` executable to determine whether there is an error during the syntax check. Files with syntax errors are added to the `FILES_WITH_ERRORS` variable, and when all files are checked, the script prints a message indicating syntax errors. When options are not specified, the executable uses the project database values. If not specified in the project database, the executable uses the Quartus II software default values. For example, the `fir_filter` project is set to target the Cyclone device family, so it is not necessary to specify the `--family` option.

Example 2-10. Shell Script to Check Design File Syntax

```
#!/bin/sh
FILES_WITH_ERRORS=""
# Iterate over each file with a .bdf or .v extension
for filename in `ls *.bdf *.v`
do
# Perform a syntax check on the specified file
  quartus_map fir_filter --analyze_file=$filename
  # If the exit code is non-zero, the file has a syntax error
  if [ $? -ne 0 ]
  then
    FILES_WITH_ERRORS="$FILES_WITH_ERRORS $filename"
  fi
done
if [ -z "$FILES_WITH_ERRORS" ]
then
  echo "All files passed the syntax check"
  exit 0
else
  echo "There were syntax errors in the following file(s)"
  echo $FILES_WITH_ERRORS
  exit 1
fi
```

Create a Project and Synthesize a Netlist Using Netlist Optimizations

This example creates a new Quartus II project with a file `top.edf` as the top-level entity. The `--enable_register_retiming=on` and `--enable_wysiwyg_resynthesis=on` options allow the technology mapper to optimize the design using gate-level register retiming and technology remapping.

- ② For more information about register retiming, WYSIWYG primitive resynthesis, and other netlist optimization options, refer to Quartus II Help.

The `--part` option tells the technology mapper to target an EP20K600EBC652-1X device. To create the project and synthesize it using the netlist optimizations described above, type the command shown in [Example 2-11](#) at a command prompt.

Example 2-11. Creating a Project and Synthesizing a Netlist Using Netlist Optimizations

```
quartus_map top --source=top.edf --enable_register_retiming=on
--enable_wysiwyg_resynthesis=on --part=EP20K600EBC652-1X ↵
```

Archive and Restore Projects

You can archive or restore a Quartus II project with a single command. This makes it easy to take snapshots of projects when you use batch files or shell scripts for compilation and project management. Use the `--archive` or `--restore` options for `quartus_sh` as appropriate. Type the command shown in [Example 2-12](#) at a command prompt to archive your project.

Example 2-12. Archiving a Project

```
quartus_sh --archive <project name> ↵
```

The archive file is automatically named `<project name>.qar`. If you want to use a different name, rename the archive after it has been created. This command overwrites any existing archive with the same name.

To restore a project archive, type the command shown in [Example 2-13](#) at a command prompt.

Example 2-13. Restoring a Project Archive

```
quartus_sh --restore <archive name> ↵
```

The command restores the project archive to the current directory and overwrites existing files.



For more information about archiving and restoring projects, refer to the [Managing Quartus II Projects](#) chapter in volume 2 of the *Quartus II Handbook*.

Perform I/O Assignment Analysis

You can perform I/O assignment analysis with a single command. I/O assignment analysis checks pin assignments to ensure they do not violate board layout guidelines. I/O assignment analysis does not require a complete place and route, so it is a quick way to ensure your pin assignments are correct. The command shown in [Example 2-14](#) performs I/O assignment analysis for the specified project and revision.

Example 2-14. Performing I/O Assignment Analysis

```
quartus_fit --check_ios <project name> --rev=<revision name> ↵
```

Update Memory Contents Without Recompiling

You can use two commands to update the contents of memory blocks in your design without recompiling. Use the `quartus_cdb` executable with the `--update_mif` option to update memory contents from `.mif` or `.hexout` files. Then, rerun the assembler with the `quartus_asm` executable to regenerate the `.sof`, `.pof`, and any other programming files.

Example 2-15 shows these two commands.

Example 2-15. Commands to Update Memory Contents Without Recompiling

```
quartus_cdb --update_mif <project name> [--rev=<revision name>] ←  
quartus_asm <project name> [--rev=<revision name>] ←
```

Example 2-16 shows the commands for a DOS batch file for this example. You can paste the following lines into a DOS batch file called `update_memory.bat`.

Example 2-16. Batch file to Update Memory Contents Without Recompiling

```
quartus_cdb --update_mif %1 --rev=%2  
quartus_asm %1 --rev=%2
```

Type the following command at a command prompt:

```
update_memory.bat <project name> <revision name> ←
```

Create a Compressed Configuration File

You can create a compressed configuration file in two ways. The first way is to run `quartus_cpf` with an option file that turns on compression. The second way is to run `quartus_cpf` with a Conversion Setup File (`.cof`).

To create an option file that turns on compression, type the following command at a command prompt:

```
quartus_cpf -w <filename>.opt ←
```

This interactive command guides you through some questions, then creates an option file based on your answers. Use the `--option` option to `quartus_cpf` to specify the option file you just created. For example, the following command creates a compressed `.pof` that targets an EPCS64 device:

```
quartus_cpf --convert --option=<filename>.opt --device=EPCS64 <file>.sof <file>.pof ←
```

Alternatively, you can use the Convert Programming Files utility in the Quartus II software to create a conversion setup file. Configure any options you want, including compression, then save the conversion setup. Use the following command to run the conversion setup you specified.

```
quartus_cpf <file>.cof ←
```

Fit a Design as Quickly as Possible

This example assumes that a project called `top` exists in the current directory, and that the name of the top-level entity is `top`. The `--effort=fast` option causes the Fitter to use the fast fit algorithm to increase compilation speed, possibly at the expense of reduced f_{MAX} performance. The `--one_fit_attempt=on` option restricts the Fitter to only one fitting attempt for the design.

To attempt to fit the project called `top` as quickly as possible, type the command shown in Example 2-17 at a command prompt.

Example 2-17. Fitting a Project Quickly

```
quartus_fit top --effort=fast --one_fit_attempt=on ←
```

Fit a Design Using Multiple Seeds

This shell script example assumes that the Quartus II software tutorial project called **fir_filter** exists in the current directory (defined in the file **fir_filter.qpf**). If the tutorial files are installed on your system, this project exists in the *<Quartus II directory>/qdesigns<quartus_version_number> /fir_filter* directory. Because the top-level entity in the project does not have the same name as the project, you must specify the revision name for the top-level entity with the `--rev` option. The `--seed` option specifies the seeds to use for fitting.

A seed is a parameter that affects the random initial placement of the Quartus II Fitter. Varying the seed can result in better performance for some designs.

After each fitting attempt, the script creates new directories for the results of each fitting attempt and copies the complete project to the new directory so that the results are available for viewing and debugging after the script has completed.

[Example 2-18](#) is designed for use on UNIX systems using **sh** (the shell).

Example 2-18. Shell Script to Fit a Design Using Multiple Seeds

```
#!/bin/sh
ERROR_SEEDS=""
quartus_map fir_filter --rev=filtref
# Iterate over a number of seeds
for seed in 1 2 3 4 5
do
echo "Starting fit with seed=$seed"
# Perform a fitting attempt with the specified seed
quartus_fit fir_filter --seed=$seed --rev=filtref
# If the exit-code is non-zero, the fitting attempt was
# successful, so copy the project to a new directory
if [ $? -eq 0 ]
then
        mkdir ../fir_filter-seed_$seed
        mkdir ../fir_filter-seed_$seed/db
        cp * ../fir_filter-seed_$seed
        cp db/* ../fir_filter-seed_$seed/db
else
        ERROR_SEEDS="$ERROR_SEEDS $seed"
fi
done
if [ -z "$ERROR_SEEDS" ]
then
echo "Seed sweeping was successful"
exit 0
else
echo "There were errors with the following seed(s)"
echo $ERROR_SEEDS
exit 1
fi
```

 Use the Design Space Explorer (DSE) included with the Quartus II software script (by typing `quartus_sh --dse` at a command prompt) to improve design performance by performing automated seed sweeping.

 For more information about the DSE, type `quartus_sh --help=dse` at a command prompt, or refer to *Design Space Explorer* in Quartus II Help.

Regenerating Megafunctions After Updating the Quartus II Software

Some megafunction variations may require regeneration when you update your installation of the Quartus II software. Read the release notes for the Quartus II software and any new documentation for the IP functions used in your design to determine if regeneration is necessary.

If regeneration is necessary, you can use a Tcl script to run the **qmegawiz** executable to update each function, allowing you to avoid regenerating each function in the Megawizard Plug-In Manager GUI.

Wizard-generated files are identified in the Source Files Used report panel (contained in `<project name>.map.rpt`) in the File Type column as “Auto-Found Wizard-Generated File”. In a Tcl script, use the commands in the **::quartus::report** package from the Quartus II Tcl API to recover the list of files. Use the `qexec` command in a loop to run **qmegawiz** for each variation file:

```
qexec "qmegawiz -silent <variation file name>"
```

For example, if your script determines that your design contains a variation file called `myrom.v`, in one iteration of the loop in your script, a combination of strings and variables passed to the `qexec` command would be equivalent to the following command:

```
qexec "qmegawiz -silent myrom.v"
```

If your design flow incorporates parameter files, those can be included in the `qmegawiz` call in the same way you would include them from a command prompt:

```
qexec "qmegawiz -silent -f:<parameter file>.txt <variation file name>"
```

 For more information about the **::quartus::report** Tcl package, refer to [::quartus::report](#) in Quartus II Help.

 For more information about the Quartus II Tcl scripting API, refer to the [Tcl Scripting](#) chapter in volume 2 of the *Quartus II Handbook*.

The QFlow Script

A Tcl/Tk-based graphical interface called QFlow is included with the command-line executables. You can use the QFlow interface to open projects, launch some of the command-line executables, view report files, and make some global project assignments. The QFlow interface can run the following command-line executables:

- `quartus_map` (Analysis and Synthesis)
- `quartus_fit` (Fitter)
- `quartus_sta` (TimeQuest timing analyzer)
- `quartus_asm` (Assembler)
- `quartus_eda` (EDA Netlist Writer)

To view floorplans or perform other GUI-intensive tasks, launch the Quartus II software.

Start QFlow by typing the following command at a command prompt:

```
quartus_sh -g ↵
```

 The QFlow script is located in the `<Quartus II directory>/common/tcl/apps/qflow/` directory.

Document Revision History

Table 2-8 shows the revision history for this chapter.

Table 2-8. Document Revision History

Date	Version	Changes
December 2012	10.1.0	<ul style="list-style-type: none"> Template update. Added section on using a script to regenerate megafunction variations. Removed references to the Classic Timing Analyzer (<code>quartus_tan</code>). Removed Qflow illustration.
July 2010	10.0.0	Updated script examples to use <code>quartus_sta</code> instead of <code>quartus_tan</code> , and other minor updates throughout document.
November 2009	9.1.0	Updated Table 2-1 to add <code>quartus_jli</code> and <code>quartus_jbcc</code> executables and descriptions, and other minor updates throughout document.
March 2009	9.0.0	No change to content.
November 2008	8.1.0	<ul style="list-style-type: none"> Added the following sections: <ul style="list-style-type: none"> ■ “The MegaWizard Plug-In Manager” on page 2-11 ■ “Command-Line Support” on page 2-12 ■ “Module and Wizard Names” on page 2-13 ■ “Ports and Parameters” on page 2-14 ■ “Invalid Configurations” on page 2-15 ■ “Strategies to Determine Port and Parameter Values” on page 2-15 ■ “Optional Files” on page 2-15 ■ “Parameter File” on page 2-16 ■ “Working Directory” on page 2-17 ■ “Variation File Name” on page 2-17 ■ “Create a Compressed Configuration File” on page 2-21 ■ Updated “Option Precedence” on page 2-5 to clarify how to control precedence ■ Corrected Example 2-5 on page 2-8 ■ Changed Example 2-1, Example 2-2, Example 2-4, and Example 2-7 to use the EP1C12F256C6 device ■ Minor editorial updates ■ Updated entire chapter using 8½” × 11” chapter template
May 2008	8.0.0	<ul style="list-style-type: none"> ■ Updated “Referenced Documents” on page 2-20. ■ Updated references in document.

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

Introduction

Developing and running Tcl scripts to control the Altera® Quartus® II software allows you to perform a wide range of functions, such as compiling a design or writing procedures to automate common tasks.

You can use Tcl scripts to manage a Quartus II project, make assignments, define design constraints, make device assignments, run compilations, perform timing analysis, import LogicLock™ region assignments, use the Quartus II Chip Editor, and access reports. You can automate your Quartus II assignments using Tcl scripts so that you do not have to create them individually. Tcl scripts also facilitate project or assignment migration. For example, when using the same prototype or development board for different projects, you can automate reassignment of pin locations in each new project. The Quartus II software can also generate a Tcl script based on all the current assignments in the project, which aids in switching assignments to another project.

The Quartus II software Tcl commands follow the EDA industry Tcl application programming interface (API) standards for using command-line options to specify arguments. This simplifies learning and using Tcl commands. If you encounter an error using a command argument, the Tcl interpreter gives help information showing correct usage.

This chapter includes sample Tcl scripts for automating the Quartus II software. You can modify these example scripts for use with your own designs. You can find more Tcl scripts in the Design Examples section of the Support area of Altera's website.

This chapter includes the following topics:

- "Quartus II Tcl Packages" on page 3-2
- "Quartus II Tcl API Help" on page 3-3
- "Command-Line Options: -t, -s, and --tcl_eval" on page 3-6
- "End-to-End Design Flows" on page 3-7
- "Creating Projects and Making Assignments" on page 3-8
- "Compiling Designs" on page 3-9
- "Reporting" on page 3-9
- "Timing Analysis" on page 3-12
- "Automating Script Execution" on page 3-12
- "Other Scripting Features" on page 3-14
- "Using the Quartus II Tcl Shell in Interactive Mode" on page 3-19
- "Using the tclsh Shell" on page 3-20
- "Using the tclsh Shell" on page 3-20

- “Tcl Scripting Basics” on page 3–20

What is Tcl?

Tcl (pronounced “tickle”) is a popular scripting language that is similar to many shell scripting and high-level programming languages. It provides support for control structures, variables, network socket access, and APIs. Tcl is the EDA industry-standard scripting language used by Synopsys, Mentor Graphics®, and Altera software. It allows you to create custom commands and works seamlessly across most development platforms. For a list of recommended literature on Tcl, refer to “External References” on page 3–26.

You can create your own procedures by writing scripts containing basic Tcl commands, user-defined procedures, and Quartus II API functions. You can then automate your design flow, run the Quartus II software in batch mode, or execute the individual Tcl commands interactively in the Quartus II Tcl interactive shell.

If you are unfamiliar with Tcl scripting, or are a Tcl beginner, refer to “Tcl Scripting Basics” on page 3–20 for an introduction to Tcl scripting.

The Quartus II software, beginning with version 4.1, supports Tcl/Tk version 8.4, supplied by the Tcl DeveloperXchange at tcl.activestate.com.

Quartus II Tcl Packages

The Quartus II Tcl commands are grouped in packages by function. Table 3–1 describes each Tcl package.

Table 3–1. Tcl Packages (Part 1 of 2)

Package Name	Package Description
backannotate	Back annotate assignments
chip_planner	Identify and modify resource usage and routing with the Chip Editor
database_manager	Manage version-compatible database files
device	Get device and family information from the device database
flow	Compile a project, run command-line executables and other common flows
incremental compilation	Manipulate design partitions and LogicLock regions, and settings related to incremental compilation
insystem_memory_edit	Read and edit memory contents in Altera devices
insystem_source_probe	interact with the In-System Sources and Probes tool in an Altera device
jtag	Control the JTAG chain
logic_analyzer_interface	Query and modify the logic analyzer interface output pin state
misc	Perform miscellaneous tasks
project	Create and manage projects and revisions, make any project assignments including timing assignments
rapid_recompile	Manipulate Quartus II Rapid Recompile features
report	Get information from report tables, create custom reports
rtl	Traversing and querying the RTL netlist of your design
sdc	Specifies constraints and exceptions to the TimeQuest Timing Analyzer

Table 3-1. Tcl Packages (Part 2 of 2)

Package Name	Package Description
<code>sdc_ext</code>	Altera-specific SDC commands
<code>simulator</code>	Configure and perform simulations
<code>sta</code>	Contains the set of Tcl functions for obtaining advanced information from the Quartus II TimeQuest Timing Analyzer
<code>stp</code>	Run the SignalTap® II Logic Analyzer

By default, only the minimum number of packages is loaded automatically with each Quartus II executable. This keeps the memory requirement for each executable as low as possible. Because the minimum number of packages is automatically loaded, you must load other packages before you can run commands in those packages.

Because different packages are available in different executables, you must run your scripts with executables that include the packages you use in the scripts. For example, if you use commands in the `sdc_ext` package, you must use the `quartus_sta` executable to run the script because the `quartus_sta` executable is the only one with support for the `timing` package.

The following command prints lists of the packages loaded or available to load for an executable, to the console:

```
<executable name> --tcl_eval help ←
```

For example, type the following command to list the packages loaded or available to load by the `quartus_fit` executable:

```
quartus_fit --tcl_eval help ←
```

Loading Packages

To load a Quartus II Tcl package, use the `load_package` command as follows:

```
load_package [-version <version number>] <package name>
```

This command is similar to the `package require` Tcl command (described in [Table 3-2 on page 3-4](#)), but you can easily alternate between different versions of a Quartus II Tcl package with the `load_package` command.



For additional information about these and other Quartus II command-line executables, refer to the *Command-Line Scripting* chapter in volume 2 of the *Quartus II Handbook*.

Quartus II Tcl API Help

Access the Quartus II Tcl API Help reference by typing the following command at a system command prompt:

```
quartus_sh --qhelp ←
```

This command runs the Quartus II Command-Line and Tcl API help browser, which documents all commands and options in the Quartus II Tcl API. It includes detailed descriptions and examples for each command.

Quartus II Tcl help allows easy access to information about the Quartus II Tcl commands. To access the help information, type `help` at a Tcl prompt, as shown in [Example 3-1](#).

Example 3-1. Help Output

```
tcl> help
-----
---
-----
Available Quartus II Tcl Packages:
-----

Loaded                                Not Loaded
-----                                -----
::quartus::misc                       ::quartus::device
::quartus::old_api                    ::quartus::backannotate
::quartus::project                    ::quartus::flow
::quartus::timing_assignment          ::quartus::logiclock
::quartus::timing_report              ::quartus::report

* Type "help -tcl"
to get an overview on Quartus II Tcl usages.
```

Using the `-tcl` option with `help` displays an introduction to the Quartus II Tcl API that focuses on how to get help for Tcl commands (short help and long help) and Tcl packages.



The Tcl API help is also available in Quartus II online help. Search for the command or package name to find details about that command or package.

[Table 3-2](#) summarizes the help options available in the Tcl environment.

Table 3-2. Help Options Available in the Quartus II Tcl Environment (Part 1 of 2)

Help Command	Description
<code>help</code>	To view a list of available Quartus II Tcl packages, loaded and not loaded.
<code>help -tcl</code>	To view a list of commands used to load Tcl packages and access command-line help.
<code>help -pkg <package_name></code> <code>[-version <version number>]</code>	To view help for a specified Quartus II package that includes the list of available Tcl commands. For convenience, you can omit the <code>::quartus::</code> package prefix, and type <code>help -pkg <package name></code> ↵. If you do not specify the <code>-version</code> option, help for the currently loaded package is displayed by default. If the package for which you want help is not loaded, help for the latest version of the package is displayed by default. Examples: <code>help -pkg ::quartus::project</code> ↵ <code>help -pkg project</code> ↵ <code>help -pkg project -version 1.0</code> ↵

Table 3-2. Help Options Available in the Quartus II Tcl Environment (Part 2 of 2)

Help Command	Description
<pre><command_name> -h or <command_name> -help</pre>	<p>To view short help for a Quartus II Tcl command for which the package is loaded.</p> <p>Examples:</p> <pre>project_open -h ↵ project_open -help ↵</pre>
<pre>package require ::quartus::<package [<version>]<="" name>="" pre=""> </package></pre>	<p>To load a Quartus II Tcl package with the specified version. If <i><version></i> is not specified, the latest version of the package is loaded by default.</p> <p>Example:</p> <pre>package require ::quartus::project 1.0 ↵</pre> <p>This command is similar to the load_package command.</p> <p>The advantage of using load_package is that you can alternate freely between different versions of the same package.</p> <p>Type <i><package name></i> [-version <i><version number></i>]↵ to load a Quartus II Tcl package with the specified version. If the -version option is not specified, the latest version of the package is loaded by default.</p> <p>Example:</p> <pre>load_package ::quartus::project -version 1.0 ↵</pre>
<pre>help -cmd <command_name> [-version <version>] or <command_name> -long_help</pre>	<p>To view long help for a Quartus II Tcl command. Only <i><command name></i> -long_help requires that the associated Tcl package is loaded.</p> <p>If you do not specify the -version option, help for the currently loaded package is displayed by default.</p> <p>If the package for which you want help is not loaded, help for the latest version of the package is displayed by default.</p> <p>Examples:</p> <pre>project_open -long_help ↵ help -cmd project_open ↵ help -cmd project_open -version 1.0 ↵</pre>
<pre>help -examples</pre>	<p>To view examples of Quartus II Tcl usage.</p>
<pre>help -quartus</pre>	<p>To view help on the predefined global Tcl array that can be accessed to view information about the Quartus II executable that is currently running.</p>
<pre>quartus_sh --qhelp</pre>	<p>To launch the Tk viewer for Quartus II command-line help and display help for the command-line executables and Tcl API packages.</p> <p>For more information about this utility, refer to the <i>Command-Line Scripting</i> chapter in volume 2 of the <i>Quartus II Handbook</i>.</p>

Command-Line Options: -t, -s, and --tcl_eval

Table 3-3 lists three command-line options you can use with executables that support Tcl.

Table 3-3. Command-Line Options Supporting Tcl Scripting

Command-Line Option	Description
-t <script file> [<script args>]	Run the specified Tcl script with optional arguments.
-s	Open the executable in the interactive Tcl shell mode.
--tcl_eval <tcl command>	Evaluate the remaining command-line arguments as Tcl commands. For example, the following command displays help for the project package: quartus_sh --tcl_eval help -pkg project

Run a Tcl Script

Running an executable with the -t option runs the specified Tcl script. You can also specify arguments to the script. Access the arguments through the argv variable, or use a package such as cmdline, which supports arguments of the following form:

```
-<argument name> <argument value>
```

The `cmdline` package is included in the <Quartus II directory>/common/tcl/packages directory.

For example, to run a script called `myscript.tcl` with one argument, `Stratix`, type the following command at a system command prompt:

```
quartus_sh -t myscript.tcl Stratix ←
```



Beginning with version 4.1, the Quartus II software supports the argv variable. In previous software versions, script arguments are accessed in the `quartus(args)` global variable.

Refer to “[Accessing Command-Line Arguments](#)” on page 3-17 for more information.

Interactive Shell Mode

Running an executable with the -s option starts an interactive Tcl shell. For example, to open the Quartus II TimeQuest Timing Analyzer executable in interactive shell mode, type the following command:

```
quartus_sta -s ←
```

Commands you type in the Tcl shell are interpreted when you click **Enter**. You can run a Tcl script in the interactive shell with the following command:

```
source <script name> ←
```

If a command is not recognized by the shell, it is assumed to be an external command and executed with the `exec` command.

Evaluate as Tcl

Running an executable with the --tcl_eval option causes the executable to immediately evaluate the remaining command-line arguments as Tcl commands. This can be useful if you want to run simple Tcl commands from other scripting languages.

For example, the following command runs the Tcl command that prints out the commands available in the project package.

```
quartus_sh --tcl_eval help -pkg project ↵
```

Using the Quartus II Tcl Console Window

You can run Tcl commands directly in the Quartus II Tcl Console window. On the View menu, click **Utility Windows**. By default, the Tcl Console window is docked in the bottom-right corner of the Quartus II GUI. Everything typed in the Tcl Console is interpreted by the Quartus II Tcl shell.



The Quartus II Tcl Console window supports the Tcl API used in the Quartus II software version 3.0 and earlier for backward compatibility with older designs and EDA tools.

Tcl messages appear in the **System** tab (Messages window). Errors and messages written to `stdout` and `stderr` also are shown in the Quartus II Tcl Console window.

End-to-End Design Flows

You can use Tcl scripts to control all aspects of the design flow, including controlling other software if it includes a scripting interface.

Typically, EDA tools include their own script interpreters that extend core language functionality with tool-specific commands. For example, the Quartus II Tcl interpreter supports all core Tcl commands, and adds numerous commands specific to the Quartus II software. You can include commands in one Tcl script to run another script, which allows you to combine or chain together scripts to control different tools. Because scripts for different tools must be executed with different Tcl interpreters, it is difficult to pass information between the scripts unless one script writes information into a file and another script reads it.

Within the Quartus II software, you can perform many different operations in a design flow (such as synthesis, fitting, and timing analysis) from a single script, making it easy to maintain global state information and pass data between the operations. However, there are some limitations on the operations you can perform in a single script due to the various packages supported by each executable.

There are no limitations on running flows from any executable. Flows include operations found in the Start section of the Processing menu in the Quartus II GUI, and are also documented with the `execute_flow` Tcl command. If you can make settings in the Quartus II software and run a flow to get your desired result, you can make the same settings and run the same flow in any command-line executable.

Creating Projects and Making Assignments

A benefit of the Tcl scripting API is that you can easily create a script that makes all the assignments for an existing project. You can use the script at any time to restore your project settings to a known state. From the Project menu, click **Generate Tcl File for Project** to automatically generate a `.tcl` file with all of your assignments. You can source this file to recreate your project, and you can edit the file to add other commands, such as compiling the design. The file is a good starting point to learn about project management commands and assignment commands.

 Refer to “[Interactive Shell Mode](#)” on page 3-6 for information about sourcing a script. Scripting information for all Quartus II project settings and assignments is located in the *QSF Reference Manual*.

[Example 3-2](#) shows how to create a project, make assignments, and compile the project. It uses the `fir_filter` tutorial design files in the `qdesigns` installation directory. Run this script in the `fir_filter` directory, with the `quartus_sh` executable.

Example 3-2. Create and Compile a Project

```
load_package flow

# Create the project and overwrite any settings
# files that exist
project_new fir_filter -revision filtref -overwrite
# Set the device, the name of the top-level BDF,
# and the name of the top level entity
set_global_assignment -name FAMILY Cyclone
set_global_assignment -name DEVICE EP1C6F256C6
set_global_assignment -name BDF_FILE filtref.bdf
set_global_assignment -name TOP_LEVEL_ENTITY filtref
# Add other pin assignments here
set_location_assignment -to clk Pin_G1
# compile the project
execute_flow -compile
project_close
```

 The assignments created or modified while a project is open are not committed to the Quartus II Settings Files (`.qsf`) unless you explicitly call `export_assignments` or `project_close` (unless `-dont_export_assignments` is specified). In some cases, such as when running `execute_flow`, the Quartus II software automatically commits the changes.

HardCopy Device Design

 For information about using a scripted design flow for HardCopy II designs, refer to the *Script-Based Design for HardCopy II Devices* chapter of the *HardCopy Handbook*. A separate chapter in the *HardCopy Handbook* called *Timing Constraints for HardCopy II Devices* also contains information about script-based design for HardCopy II devices, with an emphasis on timing constraints.

Compiling Designs

You can run the Quartus II command-line executables from Tcl scripts. Use the included **flow** package to run various Quartus II compilation flows, or run each executable directly.

The flow Package

The **flow** package includes two commands for running Quartus II command-line executables, either individually or together in standard compilation sequence. The `execute_module` command allows you to run an individual Quartus II command-line executable. The `execute_flow` command allows you to run some or all of the modules in commonly-used combinations. Use the **flow** package instead of using system calls to run Quartus II executables from scripts.

Another way to run a Quartus II executable from Tcl environment is with the `qexec` Tcl command, a Quartus II implementation of the Tcl `exec` command. For example, to run Quartus II Analysis and Synthesis on a given project, use:

```
qexec "quartus_map <project_name>" ←
```

When you use the `qexec` command to compile a design, assignments made in the Tcl script (or from the Tcl shell) are not exported to the project database and settings file before compilation. Use the `export_assignments` command or compile the project with commands in the **flow** package to ensure assignments are exported to the project database and settings file.

You can also run executables directly in a Tcl shell, using the same syntax as at the system command prompt. For example, to run the Quartus II technology mapper on a given project, type the following at the Tcl shell prompt:

```
quartus_map <project_name> ←
```

Compile All Revisions

You can use a simple Tcl script to compile all revisions in your project. Save the script shown in [Example 3-3](#) in a file called **compile_revisions.tcl** and type the following to run it:

```
quartus_sh -t compile_revisions.tcl <project name> ←
```

Example 3-3. Compile All Revisions

```
load_package flow
project_open [lindex $quartus(args) 0]
set original_revision [get_current_revision]
foreach revision [get_project_revisions] {
    set_current_revision $revision
    execute flow -compile
}
set_current_revision $original_revision
project_close
```

Reporting

It is sometimes necessary to extract information from the Compilation Report to evaluate results. The Quartus II Tcl API provides easy access to report data so you do not have to write scripts to parse the text report files.

If you know the exact cell or cells you want to access, use the `get_report_panel_data` command and specify the row and column names (or x and y coordinates) and the name of the appropriate report panel. You can often search for data in a report panel. To do this, use a loop that reads the report one row at a time with the `get_report_panel_row` command.

Column headings in report panels are in row 0. If you use a loop that reads the report one row at a time, you can start with row 1 to skip the row with column headings. The `get_number_of_rows` command returns the number of rows in the report panel, including the column heading row. Because the number of rows includes the column heading row, your loop should continue as long as the loop index is less than the number of rows, as illustrated in [Example 3-5](#).

Report panels are hierarchically arranged and each level of hierarchy is denoted by the string “|” in the panel name. For example, the name of the Fitter Settings report panel is **Fitter|Fitter Settings** because it is in the Fitter folder. Panels at the highest hierarchy level do not use the “|” string. For example, the Flow Settings report panel is named **Flow Settings**.

The code in [Example 3-4](#) prints a list of all report panel names in your project. You can run this code with any executable that includes support for the report package.

Example 3-4. Print All Report Panel Names

```
load_package report
project_open myproject
load_report
set panel_names [get_report_panel_names]
foreach panel_name $panel_names {
post_message "$panel_name"
}
```

Example 3-5 prints the number of failing paths in each clock domain in your design. It uses a loop to access each row of the Timing Analyzer Summary report panel. Clock domains are listed in the column named **Type** with the format `Clock Setup: '<clock name>'`. Other summary information is listed in the **Type** column as well. If the **Type** column matches the pattern `"Clock Setup*"`, the script prints the number of failing paths listed in the column named **Failed Paths**. You can run this script example with any executable that supports the report package.

Example 3-5. Print Number of Failing Paths per Clock

```
load_package report
project_open my-project
load_report
set report_panel_name "Timing Analyzer||Timing Analyzer Summary"
set num_rows [get_number_of_rows -name $report_panel_name]

# Get the column indices for the Type and Failed Paths columns
set type_column [get_report_panel_column_index -name \
    $report_panel_name "Type"]
set failed_paths_column [get_report_panel_column_index -name \
    $report_panel_name "Failed Paths"]

# Go through each line in the report panel
for {set i 1} {$i < $num_rows} {incr i}{

    # Get the row of data, then the type of summary
    # information in the row, and the number of failed paths
    set report_row [get_report_panel_row -name \
        $report_panel_name -row $i]
    set row_type [lindex $report_row $type_column]
    set failed_paths [lindex $report_row $failed_paths_column]
    if { [string match "Clock Setup*" $row_type] } {
        puts "$row_type has $failed_paths failing paths"
    }
}
unload_report
```

Creating .csv Files for Excel

The Microsoft Excel software is sometimes used to view or manipulate timing analysis results. You can create a `.csv` file to import into Excel with data from any Quartus II report. **Example 3-6** shows a simple way to create a `.csv` file with data from a timing analysis panel in the report. You could modify the script to use command-line arguments to pass in the name of the project, report panel, and output file to use. You can run this script example with any executable that supports the report package.

Example 3-6. Create .csv Files from Reports

```

load_package report
project_open my-project

load_report

# This is the name of the report panel to save as a CSV file
set panel_name "Timing Analyzer|Clock Setup: 'clk'"
set csv_file "output.csv"

set fh [open $csv_file w]
set num_rows [get_number_of_rows -name $panel_name]

# Go through all the rows in the report file, including the
# row with headings, and write out the comma-separated data
for { set i 0 } { $i < $num_rows } { incr i } {
    set row_data [get_report_panel_row -name $panel_name \
        -row $i]
    puts $fh [join $row_data ","]
}

close $fh
unload_report

```

Timing Analysis

The Quartus II software includes a comprehensive Tcl APIs for the TimeQuest Timing Analyzer in the **sta**, **sdc**, and **sdc_ext** packages. This section includes introductory scripting information about the TimeQuest Tcl API.

The Quartus II TimeQuest Timing Analyzer includes support for SDC commands in the **sdc** package.



Refer to the *Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook* for detailed information about how to perform timing analysis with the Quartus II TimeQuest Timing Analyzer.

Automating Script Execution

You can configure scripts to run automatically at various points during compilation. Use this capability to automatically run scripts that perform custom reporting, make specific assignments, and perform many other tasks.

The following three global assignments control when a script is run automatically:

- `PRE_FLOW_SCRIPT_FILE` —before a flow starts
- `POST_MODULE_SCRIPT_FILE` —after a module finishes
- `POST_FLOW_SCRIPT_FILE` —after a flow finishes

A module is another term for a Quartus II executable that performs one step in a flow. For example, two modules are Analysis and Synthesis (**quartus_map**), and timing analysis (**quartus_sta**).

A flow is a series of modules that the Quartus II software runs with predefined options. For example, compiling a design is a flow that typically consists of the following steps (performed by the indicated module):

1. Analysis and synthesis (**quartus_map**)
2. Fitter (**quartus_fit**)
3. Assembler (**quartus_asm**)
4. Timing Analyzer (**quartus_sta**)

Other flows are described in the help for the `execute_flow` Tcl command. In addition, many commands in the Processing menu of the Quartus II GUI correspond to this design flow.

To make an assignment automatically run a script, add an assignment with the following form to your project's `.qsf` file:

```
set_global_assignment -name <assignment name> <executable>:<script name>
```

The Quartus II software runs the scripts as shown in [Example 3-7](#).

Example 3-7.

```
<executable> -t <script name> <flow or module name> <project name> <revision name>
```

The first argument passed in the `argv` variable (or `quartus(args)` variable) is the name of the flow or module being executed, depending on the assignment you use. The second argument is the name of the project and the third argument is the name of the revision.

When you use the `POST_MODULE_SCRIPT_FILE` assignment, the specified script is automatically run after every executable in a flow. You can use a string comparison with the module name (the first argument passed in to the script) to isolate script processing to certain modules.

Execution Example

[Example 3-8](#) illustrates how automatic script execution works in a complete flow, assuming you have a project called **top** with a current revision called **rev_1**, and you have the following assignments in the `.qsf` file for your project.

Example 3-8.

```
set_global_assignment -name PRE_FLOW_SCRIPT_FILE quartus_sh:first.tcl
set_global_assignment -name POST_MODULE_SCRIPT_FILE quartus_sh:next.tcl
set_global_assignment -name POST_FLOW_SCRIPT_FILE quartus_sh:last.tcl
```

When you compile your project, the `PRE_FLOW_SCRIPT_FILE` assignment causes the following command to be run before compilation begins:

```
quartus_sh -t first.tcl compile top rev_1
```

Next, the Quartus II software starts compilation with analysis and synthesis, performed by the `quartus_map` executable. After the analysis and synthesis finishes, the `POST_MODULE_SCRIPT_FILE` assignment causes the following command to run:

```
quartus_sh -t next.tcl quartus_map top rev_1
```

Then, the Quartus II software continues compilation with the Fitter, performed by the `quartus_fit` executable. After the Fitter finishes, the `POST_MODULE_SCRIPT_FILE` assignment runs the following command:

```
quartus_sh -t next.tcl quartus_fit top rev_1
```

Corresponding commands are run after the other stages of the compilation. When the compilation is over, the `POST_FLOW_SCRIPT_FILE` assignment runs the following command:

```
quartus_sh -t last.tcl compile top rev_1
```

Controlling Processing

The `POST_MODULE_SCRIPT_FILE` assignment causes a script to run after every module. Because the same script is run after every module, you might have to include some conditional statements that restrict processing in your script to certain modules.

For example, if you want a script to run only after timing analysis, you should include a conditional test like the one shown in [Example 3-9](#). It checks the flow or module name passed as the first argument to the script and executes code when the module is `quartus_sta`.

Example 3-9. Restrict Processing to a Single Module

```
set module [lindex $quartus(args) 0]

if [string match "quartus_sta" $module] {

    # Include commands here that are run
    # after timing analysis
    # Use the post-message command to display
    # messages
    post_message "Running after timing analysis"
}
```

Displaying Messages

Because of the way the Quartus II software runs the scripts automatically, you must use the `post_message` command to display messages, instead of the `puts` command. This requirement applies only to scripts that are run by the three assignments listed in [“Automating Script Execution”](#) on page 3-12.



Refer to [“Using the `post_message` Command”](#) on page 3-16 for more information about this command.

Other Scripting Features

The Quartus II Tcl API includes other general-purpose commands and features described in this section.

Natural Bus Naming

The Quartus II software supports natural bus naming. Natural bus naming means that square brackets used to specify bus indexes in HDL do not have to be escaped to prevent Tcl from interpreting them as commands. For example, one signal in a bus named address can be identified as address[0] instead of address\[0\]. You can take advantage of natural bus naming when making assignments, as in [Example 3-10](#).

Example 3-10. Natural Bus Naming

```
set_location_assignment -to address[10] Pin_M20
```

The Quartus II software defaults to natural bus naming. You can turn off natural bus naming with the `disable_natural_bus_naming` command. For more information about natural bus naming, type the following at a Quartus II Tcl prompt:

```
enable_natural_bus_naming -h ←
```

Short Option Names

You can use short versions of command options, as long as they are unambiguous. For example, the `project_open` command supports two options: `-current_revision` and `-revision`. You can use any of the following abbreviations of the `-revision` option: `-r`, `-re`, `-rev`, `-revi`, `-revis`, and `-revisio`. You can use an option as short as `-r` because in the case of the `project_open` command no other option starts with the letter `r`. However, the `report_timing` command includes the options `-recovery` and `-removal`. You cannot use `-r` or `-re` to shorten either of those options, because the abbreviation would not be unique to only one option.

Using Collection Commands

Some Quartus II Tcl functions return very large sets of data that would be inefficient as Tcl lists. These data structures are referred to as collections. The Quartus II Tcl API uses a collection ID to access the collection. There are two Quartus II Tcl commands for working with collections, `foreach_in_collection` and `get_collection_size`. Use the `set` command to assign a collection ID to a variable.

- ② For information about which Quartus II Tcl commands return collection IDs, refer to [foreach_in_collection](#) in Quartus II Help.

The `foreach_in_collection` Command

The `foreach_in_collection` command is similar to the `foreach` Tcl command. Use it to iterate through all elements in a collection. [Example 3-11](#) prints all instance assignments in an open project.

Example 3-11. Using Collection Commands

```
set all_instance_assignments [get_all_instance_assignments -name *]
foreach_in_collection asgn $all_instance_assignments {
    # Information about each assignment is
    # returned in a list. For information
    # about the list elements, refer to Help
    # for the get-all-instance-assignments command.
    set to [lindex $asgn 2]
    set name [lindex $asgn 3]
    set value [lindex $asgn 4]
    puts "Assignment to $to: $name = $value"
}
```

The `get_collection_size` Command

Use the `get_collection_size` command to get the number of elements in a collection. [Example 3-12](#) prints the number of global assignments in an open project.

Example 3-12. `get_collection_size` Command

```
set all_global_assignments [get_all_global_assignments -name *]
set num_global_assignments [get_collection_size $all_global_assignments]
puts "There are $num_global_assignments global assignments in your project"
```

Using the `post_message` Command

To print messages that are formatted like Quartus II software messages, use the `post_message` command. Messages printed by the `post_message` command appear in the **System** tab of the Messages window in the Quartus II GUI, and are written to standard at when scripts are run. Arguments for the `post_message` command include an optional message type and a required message string.

The message type can be one of the following:

- `info` (default)
- `extra_info`
- `warning`
- `critical_warning`
- `error`

If you do not specify a type, Quartus II software defaults to `info`.

When you are using the Quartus II software in Windows, you can color code messages displayed at the system command prompt with the `post_message` command. Add the following line to your `quartus2.ini` file:

```
DISPLAY_COMMAND_LINE_MESSAGES_IN_COLOR = on
```

[Example 3-13](#) shows how to use the `post_message` command.

Example 3-13. `post_message` command

```
post_message -type warning "Design has gated clocks"
```

Accessing Command-Line Arguments

Many Tcl scripts are designed to accept command-line arguments, such as the name of a project or revision. The global variable `quartus(args)` is a list of the arguments typed on the command-line following the name of the Tcl script. [Example 3-14](#) shows code that prints all of the arguments in the `quartus(args)` variable.

Example 3-14. Simple Command-Line Argument Access

```
set i 0
foreach arg $quartus(args) {
    puts "The value at index $i is $arg"
    incr i
}
```

If you copy the script in the previous example to a file named `print_args.tcl`, it displays the following output when you type the command shown in [Example 3-15](#) at a command prompt.

Example 3-15. Passing Command-Line Arguments to Scripts

```
quartus_sh -t print_args.tcl my_project 100MHz ←
The value at index 0 is my_project
The value at index 1 is 100MHz
```

Using the `cmdline` Package

You can use the `cmdline` package included with the Quartus II software for more robust and self-documenting command-line argument passing. The `cmdline` package supports command-line arguments with the form `-<option> <value>`.

[Example 3-16](#) uses the `cmdline` package.

Example 3-16. `cmdline` Package

```
package require cmdline
variable ::argv0 $::quartus(args)
set options {
    { "project.arg" "" "Project name" }
    { "frequency.arg" "" "Frequency" }
}
set usage "You need to specify options and values"

array set optshash [::cmdline::getoptions ::argv $options $usage]
puts "The project name is $optshash(project)"
puts "The frequency is $optshash(frequency)"
```

If you save those commands in a Tcl script called `print_cmd_args.tcl` you see the following output when you type the command shown in [Example 3-23](#) at a command prompt.

Example 3-17. Passing Command-Line Arguments for Scripts

```
quartus_sh -t print_cmd_args.tcl -project my_project -frequency 100MHz ←
The project name is my_project
The frequency is 100MHz
```

Virtually all Quartus II Tcl scripts must open a project. [Example 3-18](#) opens a project, and you can optionally specify a revision name. The example checks whether the specified project exists. If it does, the example opens the current revision, or the revision you specify.

Example 3-18. Full-Featured Method to Open Projects

```
package require cmdline
variable ::argv0 $::quartus(args)
set options { \
{ "project.arg" "" "Project Name" } \
{ "revision.arg" "" "Revision Name" } \
}
array set optshash [::cmdline::getoptions ::argv0 $options]

# Ensure the project exists before trying to open it
if {[project_exists $optshash(project)]} {

    if {[string equal "" $optshash(revision)]} {

        # There is no revision name specified, so default
        # to the current revision
        project_open $optshash(project) -current_revision
    } else {

        # There is a revision name specified, so open the
        # project with that revision
        project_open $optshash(project) -revision \
            $optshash(revision)
    }
} else {
    puts "Project $optshash(project) does not exist"
    exit 1
}
# The rest of your script goes here
```

If you do not require this flexibility or error checking, you can use just the `project_open` command, as shown in [Example 3-19](#).

Example 3-19. Simple Method to Open Projects

```
set proj_name [lindex $argv 0]
project_open $proj_name
```

Using the Quartus II Tcl Shell in Interactive Mode

This section presents an example of using the `quartus_sh` interactive shell to make some project assignments and compile the finite impulse response (FIR) filter tutorial project. This example assumes that you already have the FIR filter tutorial design files in a project directory.

To begin, type the following at the system command prompt to run the interactive Tcl shell:

```
quartus_sh -s ↵
```

Create a new project called **fir_filter**, with a revision called **filtref** by typing the following command at a Tcl prompt:

```
project_new -revision filtref fir_filter ↵
```

 If the project file and project name are the same, the Quartus II software gives the revision the same name as the project.

Because the revision named **filtref** matches the top-level file, all design files are automatically picked up from the hierarchy tree.

Next, set a global assignment for the device with the following command:

```
set_global_assignment -name family Cyclone ↵
```

 To learn more about assignment names that you can use with the `-name` option, refer to Quartus II Help.

 For assignment values that contain spaces, the value should be enclosed in quotation marks.

To quickly compile a design, use the `::quartus::flow` package, which properly exports the new project assignments and compiles the design using the proper sequence of the command-line executables. First, load the package:

```
load_package flow ↵
```

It returns the following:

```
1.0
```

To perform a full compilation of the FIR filter design, use the `execute_flow` command with the `-compile` option:

```
execute_flow -compile ↵
```

This command compiles the FIR filter tutorial project, exporting the project assignments and running `quartus_map`, `quartus_fit`, `quartus_asm`, and `quartus_sta`. This sequence of events is the same as selecting **Start Compilation** from the Processing menu in the Quartus II GUI.

When you are finished with a project, close it using the `project_close` command as shown in [Example 3-20](#).

Example 3-20.

```
project_close ↵
```

To exit the interactive Tcl shell, type `exit`  at a Tcl prompt.

Using the tclsh Shell

On the UNIX and Linux operating systems, the tclsh shell included with the Quartus II software is initialized with a minimal `PATH` environment variable. As a result, system commands may not be available within the tclsh shell because certain directories are not in the `PATH` environment variable. To include other directories in the path searched by the tclsh shell, set the `QUARTUS_INIT_PATH` environment variable before running the tclsh shell. Directories in the `QUARTUS_INIT_PATH` environment variable are searched by the tclsh shell when you execute a system command.

Tcl Scripting Basics

The core Tcl commands support variables, control structures, and procedures. Additionally, there are commands for accessing the file system and network sockets, and running other programs. You can create platform-independent graphical interfaces with the Tk widget set.

Tcl commands are executed immediately as they are typed in an interactive Tcl shell. You can also create scripts (including the examples in this chapter) in files and run them with the Quartus II executables or with the tclsh shell.

Hello World Example

The following shows the basic “Hello world” example in Tcl:

```
puts "Hello world" 
```

Use double quotation marks to group the words `hello` and `world` as one argument. Double quotation marks allow substitutions to occur in the group. Substitutions can be simple variable substitutions, or the result of running a nested command, described in “[Substitutions](#)” on page 3-21. Use curly braces `{ }` for grouping when you want to prevent substitutions.

Variables

Use the `set` command to assign a value to a variable. You do not have to declare a variable before using it. Tcl variable names are case-sensitive. [Example 3-21](#) assigns the value `1` to the variable named `a`.

Example 3-21. Assigning Variables

```
set a 1
```

To access the contents of a variable, use a dollar sign before the variable name. [Example 3-22](#) prints “Hello world” in a different way.

Example 3-22. Accessing Variables

```
set a Hello  
set b world  
puts "$a $b"
```

Substitutions

Tcl performs three types of substitution:

- Variable value substitution
- Nested command substitution
- Backslash substitution

Variable Value Substitution

Variable value substitution, as shown in [Example 3-22](#), refers to accessing the value stored in a variable by using a dollar sign (“\$”) before the variable name.

Nested Command Substitution

Nested command substitution refers to how the Tcl interpreter evaluates Tcl code in square brackets. The Tcl interpreter evaluates nested commands, starting with the innermost nested command, and commands nested at the same level from left to right. Each nested command result is substituted in the outer command.

[Example 3-23](#) sets `a` to the length of the string `foo`.

Example 3-23. Command Substitution

```
set a [string length foo]
```

Backslash Substitution

Backslash substitution allows you to quote reserved characters in Tcl, such as dollar signs (“\$”) and braces (“[]”). You can also specify other special ASCII characters like tabs and new lines with backslash substitutions. The backslash character is the Tcl line continuation character, used when a Tcl command wraps to more than one line.

[Example 3-24](#) shows how to use the backslash character for line continuation.

Example 3-24. Backslash Substitution

```
set this_is_a_long_variable_name [string length "Hello \  
world."]
```

Arithmetic

Use the `expr` command to perform arithmetic calculations. Using curly braces (“{ }”) to group the arguments of this command makes arithmetic calculations more efficient and preserves numeric precision. [Example 3-25](#) sets `b` to the sum of the value in the variable `a` and the square root of 2.

Example 3-25. Arithmetic with the expr Command

```
set a 5  
set b [expr { $a + sqrt(2) }]
```

Tcl also supports boolean operators such as `&&` (AND), `||` (OR), `!` (NOT), and comparison operators such as `<` (less than), `>` (greater than), and `==` (equal to).

Lists

A Tcl list is a series of values. Supported list operations include creating lists, appending lists, extracting list elements, computing the length of a list, sorting a list, and more. [Example 3-26](#) sets `a` to a list with three numbers in it.

Example 3-26. Creating Simple Lists

```
set a { 1 2 3 }
```

You can use the `lindex` command to extract information at a specific index in a list. Indexes are zero-based. You can use the `index end` to specify the last element in the list, or the `index end-<n>` to count from the end of the list. [Example 3-27](#) prints the second element (at index 1) in the list stored in `a`.

Example 3-27. Accessing List Elements

```
puts [lindex $a 1]
```

The `llength` command returns the length of a list. [Example 3-28](#) prints the length of the list stored in `a`.

Example 3-28. List Length

```
puts [llength $a]
```

The `lappend` command appends elements to a list. If a list does not already exist, the list you specify is created. The list variable name is not specified with a dollar sign. [Example 3-29](#) appends some elements to the list stored in `a`.

Example 3-29. Appending to a List

```
lappend a 4 5 6
```

Arrays

Arrays are similar to lists except that they use a string-based index. Tcl arrays are implemented as hash tables. You can create arrays by setting each element individually or by using the `array set` command. To set an element with an index of `Mon` to a value of `Monday` in an array called `days`, use the following command:

```
set days(Mon) Monday
```

The `array set` command requires a list of index/value pairs. This example sets the array called `days`:

```
array set days { Sun Sunday Mon Monday Tue Tuesday \
  Wed Wednesday Thu Thursday Fri Friday Sat Saturday }
```

[Example 3-30](#) shows how to access the value for a particular index.

Example 3-30. Accessing Array Elements

```
set day_abbreviation Mon
puts $days($day_abbreviation)
```

Use the `array names` command to get a list of all the indexes in a particular array. The index values are not returned in any specified order. [Example 3–31](#) shows one way to iterate over all the values in an array.

Example 3–31. Iterating Over Arrays

```
foreach day [array names days] {  
    puts "The abbreviation $day corresponds to the day \  
name $days($day)"  
}
```

Arrays are a very flexible way of storing information in a Tcl script and are a good way to build complex data structures.

Control Structures

Tcl supports common control structures, including if-then-else conditions and `for`, `foreach`, and `while` loops. The position of the curly braces as shown in the following examples ensures the control structure commands are executed efficiently and correctly. [Example 3–32](#) prints whether the value of variable `a` is positive, negative, or zero.

Example 3–32. If-Then-Else Structure

```
if { $a > 0 } {  
    puts "The value is positive"  
} elseif { $a < 0 } {  
    puts "The value is negative"  
} else {  
    puts "The value is zero"  
}
```

[Example 3–33](#) uses a `for` loop to print each element in a list.

Example 3–33. For Loop

```
set a { 1 2 3 }  
for { set i 0 } { $i < [llength $a] } { incr i } {  
    puts "The list element at index $i is [lindex $a $i]"  
}
```

[Example 3–34](#) uses a `foreach` loop to print each element in a list.

Example 3–34. foreach Loop

```
set a { 1 2 3 }  
foreach element $a {  
    puts "The list element is $element"  
}
```

[Example 3-35](#) uses a while loop to print each element in a list.

Example 3-35. while Loop

```
set a { 1 2 3 }
set i 0
while { $i < [llength $a] } {
    puts "The list element at index $i is [lindex $a $i]"
    incr i
}
```

You do not have to use the `expr` command in boolean expressions in control structure commands because they invoke the `expr` command automatically.

Procedures

Use the `proc` command to define a Tcl procedure (known as a subroutine or function in other scripting and programming languages). The scope of variables in a procedure is local to the procedure. If the procedure returns a value, use the `return` command to return the value from the procedure. [Example 3-36](#) defines a procedure that multiplies two numbers and returns the result.

Example 3-36. Simple Procedure

```
proc multiply { x y } {
    set product [expr { $x * $y }]
    return $product
}
```

[Example 3-37](#) shows how to use the `multiply` procedure in your code. You must define a procedure before your script calls it.

Example 3-37. Using a Procedure

```
proc multiply { x y } {
    set product [expr { $x * $y }]
    return $product
}
set a 1
set b 2
puts [multiply $a $b]
```

You should define procedures near the beginning of a script. If you want to access global variables in a procedure, use the `global` command in each procedure that uses a global variable. [Example 3-38](#) defines a procedure that prints an element in a global list of numbers, then calls the procedure.

Example 3-38. Accessing Global Variables

```
proc print_global_list_element { i } {
    global my_data
    puts "The list element at index $i is [lindex $my_data $i]"
}
set my_data { 1 2 3 }
print_global_list_element 0
```

File I/O

Tcl includes commands to read from and write to files. You must open a file before you can read from or write to it, and close it when the read and write operations are done. To open a file, use the `open` command; to close a file, use the `close` command. When you open a file, specify its name and the mode in which to open it. If you do not specify a mode, Tcl defaults to read mode. To write to a file, specify `w` for write mode as shown in [Example 3-39](#).

Example 3-39. Open a File for Writing

```
set output [open myfile.txt w]
```

Tcl supports other modes, including appending to existing files and reading from and writing to the same file.

The `open` command returns a file handle to use for read or write access. You can use the `puts` command to write to a file by specifying a filehandle, as shown in [Example 3-40](#).

Example 3-40. Write to a File

```
set output [open myfile.txt w]
puts $output "This text is written to the file."
close $output
```

You can read a file one line at a time with the `gets` command. [Example 3-41](#) uses the `gets` command to read each line of the file and then prints it out with its line number.

Example 3-41. Read from a File

```
set input [open myfile.txt]
set line_num 1
while { [gets $input line] >= 0 } {
    # Process the line of text here
    puts "$line_num: $line"
    incr line_num
}
close $input
```

Syntax and Comments

Arguments to Tcl commands are separated by white space, and Tcl commands are terminated by a newline character or a semicolon. As shown in [“Substitutions” on page 3-21](#), you must use backslashes when a Tcl command extends more than one line.

Tcl uses the hash or pound character (`#`) to begin comments. The `#` character must begin a comment. If you prefer to include comments on the same line as a command, be sure to terminate the command with a semicolon before the `#` character.

[Example 3-42](#) is a valid line of code that includes a `set` command and a comment.

Example 3-42. Comments

```
set a 1;# Initializes a
```

Without the semicolon, it would be an invalid command because the set command would not terminate until the new line after the comment.

The Tcl interpreter counts curly braces inside comments, which can lead to errors that are difficult to track down. [Example 3-43](#) causes an error because of unbalanced curly braces.

Example 3-43. Unbalanced Braces in Comments

```
# if { $x > 0 } {
if { $y > 0 } {
    # code here
}
```

External References

 For more information about using Tcl, refer to the following sources:

- *Practical Programming in Tcl and Tk*, Brent B. Welch
- *Tcl and the TK Toolkit*, John Ousterhout
- *Effective Tcl/TK Programming*, Michael McLennan and Mark Harrison
- Quartus II Tcl example scripts at www.altera.com/support/examples/tcl/tcl.html
- Tcl Developer Xchange at tcl.activestate.com

Document Revision History

[Table 3-4](#) shows the revision history for this chapter.

Table 3-4. Document Revision History

Date	Version	Changes
December 2010	10.1.0	Template update Updated to remove tcl packages used by the Classic Timing Analyzer
July 2010	10.0.0	Minor updates throughout document.
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Removed LogicLock example. ■ Added the incremental_compilation, insystem_source_probe, and rtl packages to Table 3-1 and Table 3-2. ■ Added quartus_map to table 3-2.
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Removed the “EDA Tool Assignments” section ■ Added the section “Compile All Revisions” on page 3-9 ■ Added the section “Using the tclsh Shell” on page 3-20
November 2008	8.1.0	Changed to 8½” × 11” page size. No change to content.
May 2008	8.0.0	Updated references.

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter discusses how to create and manage projects, and how to migrate them from one computing platform to another.

Today's larger and more sophisticated FPGA designs are often developed by several engineers and are constantly changing throughout the project lifecycle. Designers must track their project changes to ensure efficient design coordination.

This chapter discusses the following topics:

- “Managing Your Quartus II Projects” on page 4–1
- “Exporting and Importing Version-Compatible Database Files” on page 4–6
- “Managing Projects in a Team-Based Design Environment” on page 4–15
- “Scripting Support” on page 4–16

Managing Your Quartus II Projects

To help you manage your FPGA designs, the Quartus® II software provides tools that assist you with your design tasks, including creating a project, creating assignments, managing revisions, and archiving projects.

A Quartus II project contains all your design files, setting files, and other files necessary for the successful compilation of your design.

- ❓ For more information about creating and opening a project, adding files to and removing files from a project, modifying project settings, saving project changes, and specifying the top-level entity, refer to *Managing Files in a Project* in Quartus II Help.

For more information about libraries, refer to “*Specifying Libraries*” on page 4–10.

-  On the **General** page of the **Options** dialog box, you can also specify a default directory that automatically stores all project files.

After you create a new project, the Quartus II software automatically generates various project files necessary for successful compilation, including the Quartus II Project File (.qpf) and Quartus II Settings File (.qsf).

- ❓ For more information about the .qpf and .qsf, refer to *Quartus II Project File (.qpf)* and *Quartus II Setting File (.qsf)* in Quartus II Help.
-  For a list of supported Quartus II project files and design file types, refer to *Introduction to the Quartus II Software* manual.

File Association

Quartus II project files are files associated with a Quartus II project, but are not design files in the project hierarchy. Most project files do not contain design logic. The Quartus II software supports project files such as **.qpf**, Quartus II IP File (**.qip**), and **.qsf**, among others.

Design files are files that contain logic for a Quartus II project. The Compiler compiles the Quartus II design files. The Quartus II software also supports designs created from EDIF Input Files (**.edf**) or Verilog Quartus Mapping Files (**.vqm**) generated by EDA design entry and synthesis tools. You can also create Verilog HDL or VHDL designs in the Quartus II software and EDA design entry tools and either generate EDIF Input Files (**.edf**) and Verilog Quartus Mapping File (**.vqm**), or use the Verilog HDL or VHDL design files directly in Quartus II projects. The Quartus II software also supports use of Quartus II Exported Partition Files (**.qxp**) as source files containing entities you can add to your design.

The Quartus II software sets file type association when you run the Quartus II software version 9.1 or earlier; however, in the Quartus II software versions 10.0 and later, the Quartus II software sets file type association after installation, which can be overwritten if you run prior versions of the Quartus II software after installing Quartus II software versions 10.0 and later. If your files are associated with a different version of the Quartus II software, and you want to associate the files with the Quartus II software version 10.0, you can manually associate the files to the Quartus II software version 10.0.

Example 4–1 shows how to associate files with the current version of the Quartus II software manually:

Example 4–1. Command to Associate Files

```
<path to installation directory>\quartus\bin\qreg.exe --file_assoc ←
```

Editing Text-Based Designs with the Quartus II Text Editor

You can use any text editor with the Quartus II software; however, the Quartus II Text Editor allows you to take advantage of features available only in the Quartus II software, error location, and predefined templates to help you with coding.

The Quartus II software provides templates that allow you to insert predefined code directly into your design file; you can choose from several design languages, and you can directly add TimeQuest analyzer design constraints and megafunction information. You can also create and save your own templates.

- ❓ For more information about editing Quartus II Text Editor files, refer to *Editing Quartus II Text Editor Files* in Quartus II Help. For more information about text editors, refer to *About the Quartus II Text Editor* in Quartus II Help. For more information about the Quartus II Text Editor options and setting a preferred text editor, refer to *Setting Quartus II Text Editor Options* in Quartus II Help.

- 🔧 For more information about the Quartus II language template feature, refer to the “Quartus II Language Templates” section in the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*.

Creating Assignments

Assignments control a variety of different functions of the Quartus II software and are an important part of an efficient and effective design. When used in conjunction with a good design practices, assignments can help you successfully compile your design. You can create assignments with different editors and dialog boxes in the Quartus II software or with Tcl scripts. Assignments are logic functions you assign to a physical resource on the device, or compilation resources you assign to logic functions.

Quartus II Settings File

As you create assignments in the Quartus II software, you can choose either to store the assignments in memory temporarily or write the assignment out to the **.qsf** with the **Update assignments to disk during design processing only** option located on the **Processing** page of the **Options** dialog box. You can open the **Options** dialog box by clicking **Options** on the Tools menu. If you turn on the **Update assignments to disk during design processing only** option, the Quartus II software stores all assignments in memory and writes to the **.qsf** when a compilation starts or when you save or close the project. The performance of the software improves when you save assignments in memory. You can view this performance improvement when the Quartus II software stores the project files on a remote data disk.



For more information about the **.qsf**, refer to the *Quartus II Settings File Manual*.

Preserving QSF Format

When you create new assignments, the Quartus II software appends the assignments to the end of the **.qsf**. If you modify an assignment, the corresponding line in the **.qsf** changes to maintain the order of assignments in the **.qsf**, unless you add and remove project source files, or when you add, remove, and exclude members from an assignment group. In these cases, the Quartus II software appends all assignments to the end of the **.qsf**. For example, if you add a new design file to the project, the Quartus II software appends the list of all your design files to the end of the **.qsf**.

The Quartus II software preserves all spaces and tabs for all unmodified assignments and comments. When you create a new assignment or modify an existing assignment in the GUI, the Quartus II software writes the assignment with the default formatting.

Quartus II Default Settings File

You can ensure consistent results when defaults change between versions of the Quartus II software with the **assignment_defaults.qdf**, located in the **bin** or **bin64** directory of the Quartus II installation path.

The Quartus II software reads assignments from various files and stores the assignments in memory. The Quartus II software reads settings files in the following order and assignments in subsequent files take precedence over earlier ones:

1. **assignment_defaults.qdf** from *<Quartus II Installation directory>/bin* or *bin64*
2. **assignment_defaults.qdf** from the project directory
3. *<revision name>_assignment_defaults.qdf* from the project directory
4. *<revision name>.qsf* from the project directory

As the Quartus II software reads each new file, if an existing assignment from an existing project file matches, following rules of case sensitivity, multivalued fields, and other rules, the Quartus II software replaces the old value with a new value. For example, if the first file has an assignment A=1, and the second file has A=2, the software replaces assignment A=1 with assignment A=2.

Creating Timing Assignments

If you create timing assignments with the TimeQuest Timing Analyzer, the Quartus II software creates a Synopsys Design Constraints File (.sdc) that contains your SDC commands.



For more information about TimeQuest analyzer and SDC constraints, refer to the *Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

Creating Revisions

In the Quartus II software, a revision is a set of assignments and settings. A project may have multiple revisions, and each revision has its own set of assignments and settings. You can create multiple revisions in a project, and you can create a unique revision based on an existing revision. Creating a unique revision allows you to optimize a design for different results; creating a revision based on an existing revision allows you to try new settings and assignments and then compare the revisions.

Creating a revision of your design allows you to create a new set of assignments and settings for a set of design files without losing your previous assignments and settings. You can perform the following tasks with revisions:

- Create a revision not based on a previous revision. Creating a unique revision allows you to optimize a design for different fundamental reasons, such as to optimize by area in one revision and then optimize for f_{MAX} in another revision. When you create a unique revision, the Quartus II software uses all default settings.
- Create a revision based on an existing revision, but try new settings and assignments in the new revision. A new revision includes all the assignments and settings in the existing revision. You can revert from the new revision to the original revision. You can compare revisions manually, or with features in the Quartus II software.

Managing Project Revisions

The **Revisions** dialog box manages your revisions by allowing you to create and delete a revision, specify the current revision, and compare revisions.

Each time you create a new revision of a project, the Quartus II software creates a new .qsf and adds the name of the new revision to the list of revisions in the .qsf. The name of a new .qsf matches the revision name.

You can compare the compilation results of multiple revisions side by side with the **Compare Revisions** dialog box. The **Compare Revisions** dialog box compares the compilation results of each revision in three categories:

- Analysis & Synthesis
- Fitter
- TimeQuest Timing Analyzer

In addition to viewing the compilation results of each revision, you can also compare the assignments for each revision. Comparing the compilation results and assignments for each revision allows you to gain a better understanding of how different optimization options affect your design.

- ❓ For more information about creating, deleting, specifying, and comparing revisions, refer to *Managing Project Revisions* in Quartus II Help.

Creating New Copies of Your Design

If your design requires that you have two separate copies of your project, rather than just a separate revision, you can create a second copy of your project with the **Copy Project** command. For example, if you have a design that is compatible with a 32-bit data bus and you require a new copy of your design to interface with a 64-bit data bus, you may want a separate copy of the project.

The Quartus II software provides utilities to copy and save different copies of your project. Creating a copy of your project with the **Copy Project** command directs the Quartus II software to copy all your design files, your **.qsf**, and all your associated revisions.

If you are creating a new copy of a project that contains an **.edf** or a **.vqm** from a third-party EDA synthesis tool, first create a copy of your project and then replace any **.edf** or **.vqm** files with the newly generated **.edf** or **.vqm**.

- ❓ For more information about the **Copy Project** command, refer to *Copy Project Dialog Box* in Quartus II Help.

Archiving and Restoring Projects

To share large projects between engineers or to transfer your project to a new version of the Quartus II software, you can archive your project. Archiving your project creates a single compressed Quartus II Archive File (**.qar**) that contains all your design, project, and settings files. The **.qar** contains all the **.qdf** files required to compile your design and restore the original compilation results. When you restore the archive in a different version of the Quartus II software, you must include the **.qdf** in the archive to preserve previous compilation results. For more information about the **.qdf**, refer to “*Quartus II Default Settings File*” on page 4-3.

-  You can copy files listed in the **Source Control** file set for the Archiver with the **Copy Project** command. If you cannot find your source file in the **Source Control** file set, add the source file to your project before copying.

- ② For more information about archiving a project and restoring an archived project, refer to *About Archiving Projects* and *Archiving Projects* in Quartus II Help.

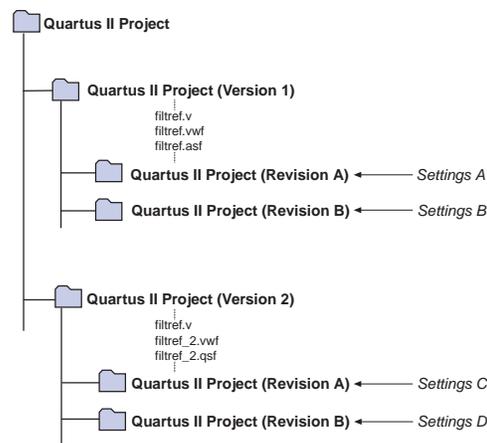
Exporting and Importing Version-Compatible Database Files

The Quartus II software generates version-compatible database files that are a representation of the internal database files. The Quartus II software allows you to export and import version-compatible database files for a project to use compilation databases in different versions of the Quartus II software. If you export version-compatible database files for a project, you can import these files in a future version of the Quartus II software. By importing version-compatible database files and rerunning timing analysis, you can check a project's fitting and timing results in newer versions of the Quartus II software.

Version-compatible databases allows you to use the same project database when you upgrade to a newer version of the Quartus II software, eliminating the need to recompile your project, which saves design time.

Figure 4-1 shows the Quartus II software version-compatible database structure.

Figure 4-1. Quartus II Version-Compatible Database Structure



- ② For more information about exporting and importing version-compatible database files, including device support, refer to *Exporting and Importing Version-Compatible Database Files* in Quartus II Help.

If you require the database files to reproduce the compilation results in the same Quartus II software version, you can use the command-line option to archive a full compilation database. For more information, refer to “*Archiving and Restoring Projects*” on page 4-5.

Migrating to a New Version of the Quartus II Software

To migrate your design to a newer version of the Quartus II software, follow these steps:

1. On the File menu, click **Open Project** and browse to select the Quartus II project file to open the older version of the Quartus II software project.
2. On the Project menu, click **Copy Project** to create a new copy of the project. The older version closes and the copied project opens.
3. Before exporting the database, you must run Analysis and Synthesis for the new version. On the Project menu, click **Export Database**. By default, the Quartus II software exports the database to the **export_db** directory of the copied project. You can also create a new directory.
4. Open the copied project from the new version of the Quartus II software. The Quartus II software deletes the existing database but not the exported database.
5. On the Project menu, click **Import Database**. By default, the Quartus II software selects the directory that contains the exported database you just created. Select the exported database and the Quartus II software imports the version-compatible database files.

Saving the Database in a Version-Compatible Format

To save the database in a version-compatible format during a full compilation, follow these steps:

1. On the Assignments menu, click **Settings**. The **Settings** dialog box appears.
2. In the **Category** list, select **Compilation Process Settings**. The **Compilation Process Settings** page appears.
3. Turn on the **Export version compatible database** option.
4. Browse to the directory in which you want to save the database.
5. Click **OK**.

You can also export a project database as version-compatible database files during a full compilation.

- ④ For more information about importing and exporting version-compatible databases, refer to *Exporting and Importing Version-Compatible Database Files* in Quartus II Help.

Quartus II Project Platform Migration

When moving your project from one computing platform to another, you must consider the following cross-platform issues:

- “File Names and Hierarchies”
- “Specifying Libraries”
- “Quartus II Search Path Precedence Rules”
- “Quartus II-Generated Files for Third-Party EDA Tools”
- “Migrating Database Files Between Platforms”

File Names and Hierarchies

To ensure a successful migration across platforms, consider the following differences between operating systems when naming source files, especially when interacting with the operating systems from the command prompt or a Tcl script:

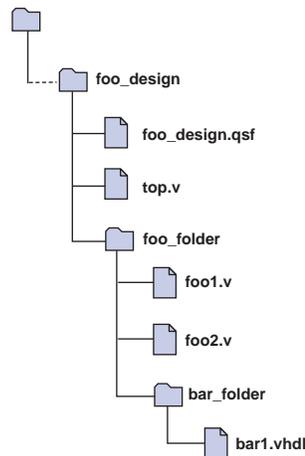
- Some operating system file systems are case sensitive. When writing scripts, ensure that you specify paths exactly, even if the current operating system is not case sensitive. Use lowercase letters when naming files.
- Use a character set common to all the used platforms.
- Do not change the forward-slash (/) and back-slash (\) path separators in the .qsf because the Quartus II software changes all back-slash (\) path separators to forward-slashes (/).
- Observe the shortest file name length limit of the different operating systems you are using.



Altera recommends that you avoid using spaces in the name of the project directory. You can rename the directory with a symbol such as the underscore (_) as a place holder instead of spaces (for example, “my_design” instead of “my design”).

You can specify files and directories inside a Quartus II project as paths relative to the project directory. For example, for a project titled **foo_design** with a directory structure shown in [Figure 4-2](#), specify the source files as: **top.v**, **foo_folder/foo1.v**, **foo_folder/foo2.v**, and **foo_folder/bar_folder/bar1.vhdl**.

Figure 4-2. All Inclusive Project Directory Structure

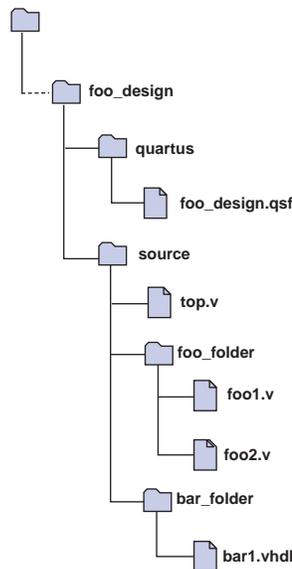


If the **.qsf** is in a directory that is separate from the source files, you can specify paths using the relative and absolute paths and libraries options.

Relative Paths

If the source files are very near to the Quartus II project directory, you can express relative paths using the **..** notation. For example, in the directory structure shown in [Figure 4-3](#), you can specify **top.v** as **../source/top.v** and **foo1.v** as **../source/foo_folder/foo1.v**.

Figure 4-3. Quartus II Project Directory Separate from Design Files



 When you copy a directory structure to a different platform, ensure that all the subdirectories are in the same hierarchical structure and relative path as in the original platform.

Specifying Libraries

You can specify the directory containing source files as a library that the Quartus II software searches when you compile your project. A Quartus II library is a directory containing your Quartus II project design files. You can specify the following libraries in the Quartus II software:

- Project libraries—Apply to a specific project
- Global libraries—Apply to all projects

 The project directory takes precedence over the project libraries.

All files in your libraries are relative to the libraries. For example, if you specify the **user_lib1** directory as a project library and you want to add the **/user_lib1/foo1.v** file to the library, you can specify the **foo1.v** file in the **.qsf** as **foo1.v**. The Quartus II software searches the file in directories that the Quartus II software specifies as libraries.

- ② For more information about libraries, refer to *Libraries Page (Settings Dialog Box)* in Quartus Help.

Specifying Project Libraries

To specify project libraries from the GUI, on the Assignments menu, click **Settings** and select **Libraries**. Type the name of the directory in the **Project Library name** box, or browse to the name of the directory. The **.qsf** of the current revision stores project libraries.

You can also specify project libraries in the **Libraries** page in the **General** category in the **Options** dialog box.

Specifying Global Libraries

To specify global libraries from the GUI, on the Tools menu, click **Options** and select **Libraries**. Type the name of the directory in the **Global Library name** box, or browse to the name of the directory. The **quartus2.ini** file stores global libraries.

To specify libraries from the GUI, on the Assignments menu, click **Settings** and select **Libraries**.

For Windows, the Quartus II software searches for the **quartus2.ini** file in the following directories and order:

1. USERPROFILE, for example, **C:\Documents and Settings\<user name>**
2. Directory specified by the **TMP** environmental variable
3. Directory specified by the **TEMP** environmental variable
4. Root directory, for example, **C:**

For Linux, the Quartus II software creates the file in the **altera.quartus** directory under the **<home>** directory.

-  If the **altera.quartus** directory does not exist, the Quartus II software creates the file in the *<home>* directory.
-  Whenever you specify a directory name in the GUI or in Tcl, the Quartus II software maintains the directory name you use in the **.qsf** rather than resolved to an absolute path.

If the directory is outside of the project directory, the path returned in the dialog box is an absolute path. You can use the **Browse** button in either the **Settings** dialog box or the **Options** dialog box to select a directory. You can change the absolute path to a relative path by editing the absolute path displayed in the **library name** field to create a relative path before you click **Add** to put the directory in the **Libraries** list. Alternatively, you can also select from the **Libraries** list and double-click to edit the path.

When copying projects that specify project libraries, you must either copy your project library files along with the project directory or ensure that your project library files exist in the target platform.

Quartus II Search Path Precedence Rules

If two files have the same file name, the Quartus II software's search path precedence rules determine the found file. The Quartus II software resolves relative paths by searching for the file in the following directories and order:

1. The project directory.
2. The project's database (**db**) directory.
3. Project libraries are searched in the order specified by the **SEARCH_PATH** setting of the **.qsf** for the current revision.

 Altera recommends that you use the **SEARCH_PATH** assignment to define the project libraries. You can have multiple **SEARCH_PATH** assignments. However, you can specify only one source directory for each **SEARCH_PATH** assignment. For more information about **SEARCH_PATH** assignments, refer to [Example 4-18 on page 4-19](#).

4. Global user libraries are searched in the order specified by the **SEARCH_PATH** setting on the **Global User Libraries** page in the **Options** dialog box.
5. The Quartus II software **libraries** directory, for example, *<Quartus II Software Installation directory>\libraries*. For more information about libraries, refer to ["Specifying Libraries Using Scripts" on page 4-19](#).

Quartus II-Generated Files for Third-Party EDA Tools

The project archive and copy features in the Quartus II software do not include Quartus II generated files for third-party EDA tools such as:

- Verilog Output Files (.vo)
- VHDL Output Files (.vho)
- Standard Delay Format Output Files (.sdo) output netlist files
- Stamp model files
- PartMiner XML-Format Files (.xml)
- IBIS Output Files (.ibs)

When you archive your design project, you can save the database in a version-compatible format during a full compilation and include the version-compatible database files in your project archive.

 Version-compatible databases may not be available for all device families because the archive does not include the compilation database. If you require the database files to reproduce the compilation results in the same Quartus II software version, you can use the command-line option to archive a full database. For more information, refer to [“Archiving and Restoring Projects” on page 4-5](#).

For more information about saving the database in a version-compatible format and archiving projects, refer to [“Saving the Database in a Version-Compatible Format” on page 4-7](#) and [“Archiving Projects” on page 4-17](#).

To copy your project to another platform, you can regenerate the output netlist or output files by following these steps:

1. Import the version-compatible database. For more information, refer to [“Migrating to a New Version of the Quartus II Software” on page 4-7](#).
2. From the Tools menu, run the TimeQuest analyzer.
3. Run the EDA Netlist Writer.

To restore your project, you can regenerate the output netlist or output files by performing the following steps:

1. Restore your design project. For more information about restoring an archived project, refer to [“Archiving and Restoring Projects” on page 4-5](#).
2. Import the version-compatible database. For more information about migrating to a new version, refer to [“Migrating to a New Version of the Quartus II Software” on page 4-7](#).
3. From the Processing menu, run the TimeQuest analyzer.
4. Run the EDA Netlist Writer.

 When you create version-compatible databases, you do not need to recompile your design as you move across platforms.

Migrating Database Files Between Platforms

There is nothing inherent in the file format and syntax of the exported version-compatible database files that might cause problems when migrating the files to other platforms. However, the contents of the database can cause problems for platform migration. For example, using the absolute paths in version-compatible database files generated by the Quartus II software can cause problems for migration. Altera recommends that you change the absolute paths to relative paths before migrating files whenever possible.

Working with Messages

The Quartus II software generates various types of messages, including Information, Warning, Extra Info, Critical Warning, and Error messages. Some messages include information about software status during a compilation and alert you to possible problems with your design. The Messages box in the Quartus II GUI displays messages, and these messages are written to `stdout` when you use command-line executables. In both cases, Quartus II report files write messages.

You can right-click a message in the Message window to get help for the message, locate the source of the message of your design, and manage messages.

Messages provide useful information if you take time to review them after each compilation. The following sections describe the Quartus II software features to help you manage messages.

Messages Window

The Messages window displays nine message tabs, enabling you to review all messages of a certain type. The **Info**, **Extra Info**, **Warning**, **Critical Warning**, and **Error** tabs display messages by type.

- ② For more information about the Messages window and message tabs, refer to [About the Messages Window](#) in Quartus II Help. For more information about managing messages in the Messages window, refer to [Managing Messages in the Messages Window](#) in Quartus II Help.

Message Suppression

You can use message suppression to reduce the number of messages after a compilation by preventing individual messages and entire categories of messages from displaying. For example, if you review a particular message and determine that your design is not the cause of the message, you can suppress the message for subsequent compilations. Message suppression saves time because you see only new messages during subsequent compilations.

Adding a suppressed message creates a suppression rule. Suppressing exact selected messages adds patterns that are exact strings to the suppression rules. Suppressing all similar messages adds patterns with wildcards to the suppression rules.

Furthermore, you can suppress all messages of a particular type in a particular stage of the compilation flow. On the Tools menu, click **Options**. In the **Category** list, select **Suppression** in the **Messages** section.

Suppressing individual messages is controlled in two locations in the Quartus II GUI. You can right-click on a message in the Messages window and choose commands in the Suppress sub-menu entry. To open the Message Suppression Manager, right-click in the Messages window. From the Suppress sub-menu, click **Message Suppression Manager**. For more information about the Message Suppression Manager, refer to “[Message Suppression Manager](#)” on page 4-15.

Message Suppression Methods

You can use the following methods to create suppression rules:

- **Suppress Exact Selected Messages**
- **Suppress All Similar Messages**
- **Suppress All Flagged Messages**

If you suppress a message with the **Suppress Exact Selected Messages** option, the Quartus II software suppresses only messages that match the exact text during subsequent compilations. The **Suppress All Similar Messages** option behaves like a wildcard pattern on variable fields in messages and the **Suppress All Flagged Messages** option only suppresses flagged messages.

[Example 4-2](#) shows an example of suppressing common Info type of messages:

Example 4-2. Example of Suppressing Common Info Type Message

```
Info: Found 1 design units, including 1 entities, in source file mult.v.
```

This Info type of message is common during synthesis. The Quartus II software displays the message for each processed source file with varying information about the number of design units, entities, and source file name.

[Example 4-3](#) shows an example of this message in Help:

Example 4-3. Example of Suppressing Common Info Type Message

```
Found <number> design units, including <number> entities, in source file <name>.
```

Choosing to suppress all similar messages effectively replaces the variable parts of that message (<number>, <number>, and <name>) with wildcards.

[Example 4-4](#) shows the suppression rule to suppress common Info type of messages:

Example 4-4. Suppression Rule to Suppress Common Info Type of Messages

```
Info: Found * design units, including * entities, in source file *.
```

The Quartus II software suppresses all messages that match the pattern.

Message Suppression Details and Limitations

The following rules describe which messages that you can suppress and how to suppress them:

- You cannot suppress error messages or messages with information about Altera legal agreements.
- Suppressing a message also suppresses all its submessages, if any.
- Suppressing a submessage causes the Quartus II software to suppress matching submessages only if the parent messages are the same.
- You cannot create your own custom wildcards to suppress messages.
- You must use the Quartus II GUI to manage message suppression, including choosing messages to suppress. The Quartus II software suppresses these messages during compilation in the GUI and when using command-line executables.
- The Quartus II software suppresses the messages on a per-revision basis, not for an entire project. The Quartus II software stores information about which messages to suppress in a file called `<revision>.srf`. If you create a revision based on a suppressed messages revision, the Quartus II software copies the suppression rules file to the new revision. You cannot make all revisions in one project using the same suppression rules file.
- You cannot remove messages or modify message suppression rules while a compilation is running.

Message Suppression Manager

You can use the Message Suppression Manager to view and suppress messages, view and delete suppression rules, and view suppressed messages. The Message Suppression Manager has three tabs labeled **Suppressible Messages**, **Suppression Rules**, and **Suppressed Messages**.

- ② For more information about the Message Suppression Manager, refer to [About Message Suppression](#) in Quartus II Help.

Managing Projects in a Team-Based Design Environment

The Quartus II software provides several methods to help you manage efficient design coordination across multiple designers and design iterations. You can use the following features to preserve and track project changes:

- Creating Revisions
- Managing Different Design Versions
- Creating Version-Compatible Databases
- Archiving Projects

- ② For more information about creating revisions, managing different design versions, creating version-compatible databases, and archiving projects in a team-based design environment, refer to [About Project Management](#) in Quartus II Help.

The Quartus II software supports top-down incremental compilation flows. With top-down compilation, one designer or project lead compiles the entire design in the software. Different designers or IP providers can design and verify different parts of the design, and the project lead can add design entities to the project as they are completed. However, the project lead compiles and optimizes the top-level project as a whole. Completed parts of the design can have fitting results and performance fixed as other parts of the design change.

 For more information about incremental compilation for team-based design, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*. For more information about best practices for incremental compilation partitions and floorplan assignments, refer to *Best Practices for Incremental Compilation Partitions and Floorplan Assignments* chapter in volume 1 of the *Quartus II Handbook*.

Scripting Support

You can run procedures and create settings described in this chapter in a Tcl script. You can also run some procedures at a command prompt. For more information about scripting command options, refer to the Quartus II Command-Line and Tcl API Help browser.

[Example 4-5](#) shows the command to run the Help browser:

Example 4-5. Command to Run the Help Browser

```
quartus_sh --qhelp ↵
```

Managing Revisions

You can use the following commands to create and manage revisions. For more information about managing revisions, including creating and deleting revisions, setting the current revision, and getting a list of revisions, refer to “[Creating Revisions](#)” on page 4-4.

Creating Revisions

The `-based_on` and `-set_current` options are optional. You can also use `-copy_results` option to copy results from the “`based_on`” revision.

[Example 4-6](#) shows a Tcl command to create a new revision called `speed_ch`, based on a revision called `chiptrip`, and sets the new revision as the current revision:

Example 4-6. Creating Revisions Command

```
create_revision speed_ch -based_on chiptrip -set_current
```

Setting the Current Revision

The `-force` option enables you to open the revision that you specify under revision name and overwrite the compilation database if the database version is incompatible.

[Example 4-7](#) shows the Tcl command to specify the current revision:

Example 4-7. Specifying the Current Revision Command

```
set_current_revision -force <revision name>
```

Getting a List of Revisions

[Example 4-8](#) shows the Tcl command to get a list of revisions in the opened project:

Example 4-8. Getting a List of Revisions in an Opened Project Command

```
get_project_revisions <project_name>
```

Deleting Revisions

[Example 4-9](#) shows the Tcl command to delete a revision:

Example 4-9. Deleting a Revision Command

```
delete_revision <revision name>
```

Archiving Projects

You can archive projects with a Tcl command or a command run at the system command prompt.

[Example 4-10](#) shows the Tcl command to create a project archive with the default settings, overwriting the existing specified archived file:

Example 4-10. Creating a Project Archive with the Default Settings Command

```
project_archive archive.qar -overwrite ↵
```

You can change default settings with the `project_archive` command with options such as:

- `-all_revisions`
- `-include_libraries`
- `-include_outputs`
- `-use_file_set <file_set>`
- `-version_compatible_database`



For new device families, a version-compatible database might not be available because the archive does not include the compilation database. If you require the database files to reproduce the compilation results in the same Quartus II software version, you can use the `-use_file_set full_db` command-line option to archive a full database. For more information, refer to [“Archiving and Restoring Projects” on page 4-5](#).

[Example 4-11](#) shows the command to create a project archive called **top**:

Example 4-11. Creating a Project Archive Command

```
quartus_sh --archive top ↵
```

You can overwrite the existing archive file with the `-overwrite` option.

Restoring Archived Projects

You can restore archived projects with a Tcl command or with a command run at a command prompt. For more information about restoring archived projects, refer to [“Archiving and Restoring Projects”](#) on page 4-5.

[Example 4-12](#) shows the Tcl command to restore the project archive named **archive.qar** in the **restored** subdirectory and overwrite existing files:

Example 4-12. Restoring a Project Archive Command

```
project_restore archive.qar -destination restored -overwrite ↵
```

[Example 4-13](#) shows the command to restore a project archive:

Example 4-13. Restoring a Project Archive Command

```
quartus_sh --restore archive.qar ↵
```

Importing and Exporting Version-Compatible Databases

You can import and export version-compatible databases with either a Tcl command or a command run at a command prompt. For more information about importing and exporting version-compatible databases, refer to [“Exporting and Importing Version-Compatible Database Files”](#) on page 4-6.



The `flow` and `database_manager` packages contain commands to manage version-compatible databases.

[Example 4-14](#) shows the Tcl command to import or export version-compatible databases from the `database_manager` package.

Example 4-14. Importing and Exporting Version-Compatible Databases Command

```
export_database <directory> ↵
import_database <directory> ↵
```

[Example 4-15](#) shows the Tcl commands from the `flow` package to import or export version-compatible databases. If you use the `flow` package, you must specify the database directory variable name.

Example 4-15. Importing and Exporting Version-Compatible Databases from the flow Package Command

```
set_global_assignment -name VER_COMPATIBLE_DB_DIR <directory>
execute_flow -flow export_database
execute_flow -flow import_database
```

[Example 4-16](#) shows the Tcl commands to generate version-compatible databases after every compilation.

Example 4-16. Generating Version-Compatible Databases After Every Compilation Command

```
set_global_assignment -name AUTO_EXPORT_VER_COMPATIBLE_DB ON
set_global_assignment -name VER_COMPATIBLE_DB_DIR <directory>
```

[Example 4-17](#) shows the `quartus_cdb` and the `quartus_sh` executables to manage version-compatible databases:

Example 4-17. quartus_cdb and quartus_sh Executable

```
quartus_cdb <project> -c <revision>--export_database=<directory> ←
quartus_cdb <project> -c <revision> --import_database=<directory>←
quartus_sh -flow export_database <project> -c \ <revision> ←
quartus_sh -flow import_database <project> -c \ <revision> ←
```

Specifying Libraries Using Scripts

In Tcl, use commands in the `::quartus::project` package to specify project libraries. To specify project libraries, use the `set_global_assignment` command.

[Example 4-18](#) shows the typical usage of the `set_global_assignment` command:

Example 4-18. Commands to Specify Project Libraries Using the SEARCH_PATH Assignment

```
set_global_assignment -name SEARCH_PATH "../other_dir/library1"
set_global_assignment -name SEARCH_PATH "../other_dir/library2"
set_global_assignment -name SEARCH_PATH "../other_dir/library3"
```

To report any project libraries specified for a project and any global libraries specified for the current installation of the Quartus II software, use the `get_global_assignment` and `get_user_option` Tcl commands.

[Example 4-19](#) shows that the Tcl script outputs the user paths and global libraries for an open Quartus II project:

Example 4-19. Commands to Report Specified Project Libraries

```
get_global_assignment -name SEARCH_PATH
get_user_option -name SEARCH_PATH
```

 For more information about Tcl scripting, refer to the [Tcl Scripting](#) chapter in volume 2 of the *Quartus II Handbook*. For more information about all settings and constraints in the Quartus II software, refer to the [Quartus II Settings File Manual](#). For more information about command-line scripting, refer to the [Command-Line Scripting](#) chapter in volume 2 of the *Quartus II Handbook*.

 For more information about Tcl scripting, refer to the [API Functions for Tcl](#) in Quartus II Help.

Conclusion

Designers often try different settings and versions of their designs throughout the development process. The Quartus II project revisions facilitate the creation and management of different assignments and settings. Project archives are useful to save your results, or pass designs between different members of a team. In addition, understanding how to migrate your projects from one computing platform to another, controlling messages, and reducing compilation time are important as well. The Quartus II software facilitates efficient management of your design to accommodate today's sophisticated FPGA designs.

Document Revision History

Table 4-1 shows the revision history for this chapter.

Table 4-1. Document Revision History (Part 1 of 2)

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Changed to new document template. ■ Removed Figure 4-1, Figure 4-6, Table 4-2. ■ Moved "Hiding Messages" to Help. Moved information in "Message Suppression Manager" on page 4-15 to Help. ■ Removed references about the <code>set_user_option</code> command in "Specifying Libraries Using Scripts" on page 4-19. ■ Removed Classic Timing Analyzer references.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Major reorganization done to this chapter. ■ Updated "Working with Messages" on page 4-17. Added a link to Help. Removed Figure 4-2 on page 4-7, Figure 4-11 on page 23, and Figure 4-12 on page. ■ Updated "Specifying Libraries" on page 4-14 section. Changed "User Libraries" to "Libraries". Removed "Reducing Compilation Time" on page 4-26. ■ Added "Managing Projects in a Team-Based Design Environment" on page 4-22 and "File Association" on page 4-2. ■ Updated Figure 4-1 on page 4-6, Figure 4-2 on page 4-8, Figure 4-6 on page 4-18, Figure 4-6 on page 4-19, and Figure 4-7 on page 4-21.
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Updated "Creating a New Project" on page 4-4, "Archiving a Project" on page 4-9, "Restoring an Archived Project" on page 4-11. ■ Added "Quartus II Text Editor" on page 4-2, "Reducing Compilation Time" on page 4-32. ■ Updated Table 4-1 on page 4-10, Table 4-2 on page 4-20. ■ Updated Figure 4-4 on page 4-9, Figure 4-7 on page 4-19.

Table 4-1. Document Revision History (Part 2 of 2)

Date	Version	Changes
April 2009	9.0.0	Updated to fix “Document Revision History” for version 9.0.0.
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Updated “Managing Quartus II Projects” on page 4-1, “Creating a New Project” on page 4-2, “Using Revisions with Your Design” on page 4-3, “Creating and Deleting Revisions” on page 4-4, “Creating New Copies of Your Design” on page 4-6, “Version-Compatible Databases” on page 4-11, “Quartus II Project Platform Migration” on page 4-12, “Filenames and Hierarchies” on page 4-12, “Quartus II Search Path Precedence Rules” on page 4-15, “Quartus II-Generated Files for Third-Party EDA Tools” on page 4-15, “Migrating Database Files between Platforms” on page 4-16, “Message Suppression” on page 4-20, “Quartus II Settings File” on page 4-24, “Quartus II Default Settings File” on page 4-25, “Managing Revisions” on page 4-26, “Archiving Projects” on page 4-26 and “Archiving Projects with the Quartus II Archive Project Feature” on page 4-7, “Importing and Exporting Version-Compatible Databases” on page 4-27, “Specifying Libraries Using Scripts” on page 4-28, “Conclusion” on page 4-30. ■ Updated Figure 4-1, Figure 4-7, Figure 4-8, and Figure 4-11. ■ Updated Table 4-1 and Table 4-2. ■ Updated Example 4-3, Example 4-4, Example 4-5, and Example 4-6.

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

This section provides an overview of the I/O planning process, Altera FPGA pin terminology, as well as the various methods for importing, exporting, creating, and validating pin-related assignments using the Quartus[®] II software. This section also describes ways to use the Quartus II software to analyze signal integrity, including simultaneous switching noise, as well as interfaces with third-party PCB design tools.

This section includes the following chapters:

■ **Chapter 5, I/O Management**

This chapter provides an overview of the I/O planning process, Altera FPGA pin terminology, and the various methods for importing, exporting, creating, and validating pin-related assignments.

■ **Chapter 6, Simultaneous Switching Noise (SSN) Analysis and Optimizations**

This chapter describes the tools in the Quartus II software that allow you to estimate the SSN performance of your design both early in the design cycle and when your PCB is complete.

■ **Chapter 7, Signal Integrity Analysis with Third-Party Tools**

This chapter is intended for logic designers and board designers, and describes simulation and how to adjust designs to improve board-level timing and signal integrity. Also included is information about how to create accurate models of your design with the Quartus II software for use in simulation software.

■ **Chapter 8, Mentor Graphics PCB Design Tools Support**

This chapter discusses how the Quartus II software interacts with the Mentor Graphics I/O Designer software and the DxDesigner software to provide a completely cyclical FPGA-to-board integration design workflow.

■ **Chapter 9, Cadence PCB Design Tools Support**

This chapter addresses how the Quartus II software interacts with the Cadence Allegro Design Entry HDL software and the Allegro Design Entry CIS (Component Information System) software (also known as OrCAD Capture CIS) to provide a complete FPGA-to-board integration design workflow.

The process of managing I/O assignments involves more than fitting design pins into a package. The increasing complexity of I/O standards and pin placement guidelines are just some of the factors that influence pin-related assignments. Both I/O capabilities of the target device and board layout guidelines influence pin location and other types of assignments. Therefore, it is necessary to begin I/O planning and PCB development even before starting to design for the target device.

Altera provides many resources for I/O planning. This chapter provides information on how to make create assignments, how to enter I/O interface information in the Pin Planner, how to create I/O-based top-level HDL files, how to validate your pin assignments, and how to generate a valid pin-out file for use with third-party PCB tools. You can consult the device-specific pin connection guidelines available on the Altera® website for your board layout. You can also benefit from the Pin Advisors available in the Quartus® II software.

-  For more information about the Altera resources available for I/O planning, refer to the [I/O Management](#), [Board Development Support](#), and [Signal Integrity Analysis Resource Center](#) of the Altera website.
-  For more information about PCB designs for Altera high-speed FPGAs, refer to [AN 315: Guidelines for Designing High-Speed FPGA PCBs](#), and the [Board Design Resource Center](#) of the Altera website.

This chapter includes the following topics:

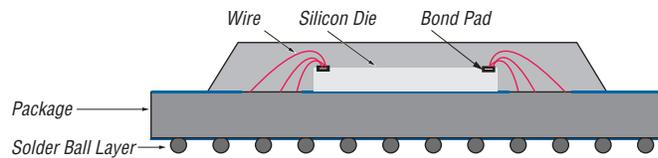
- [“Understanding Altera Pin Terminology”](#) on page 5–1
- [“I/O Planning Overview”](#) on page 5–5
- [“Performing Early I/O Planning with the Pin Planner”](#) on page 5–8
- [“Importing and Exporting Pin Assignments”](#) on page 5–12
- [“Creating Pin-Related Assignments”](#) on page 5–13
- [“Validating Pin Assignments”](#) on page 5–23
- [“Performing I/O Timing Analysis”](#) on page 5–32
- [“Incorporating PCB Design Tools”](#) on page 5–37

Understanding Altera Pin Terminology

Altera devices are available in a variety of package types. To describe pin terminology, this chapter uses a wire bond ball-grid array (BGA) package in its examples. On the top surface of the silicon die, there is a ring of bond pads that connect to the silicon to the I/O pins. In a wire bond BGA package, the device is placed in the package and copper wires connect the bond pads to the solder balls of the package. [Figure 5–1](#) shows a cross section of a wire bond BGA package.

For more information about the BGA packages available for each Altera device, refer to the [Altera Device Package Information Data Sheet](#).

Figure 5-1. Wire Bond BGA

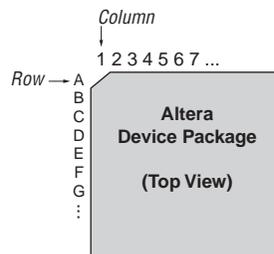


Package Pins

The pins of a BGA package are small solder balls arranged in a grid pattern on the bottom of the package. In the Quartus II software, the package pins are identified with pin numbers. The pin numbers are determined by pin locations using a coordinate system, with letters and numbers identifying the row and column of the pins, respectively.

Figure 5-2 shows the coordinate system used to identify pin locations. The top row of pins is labeled *A* and continues alphabetically as you move down. The left-most column of pins is labeled *1* and increments by one as you move right. For example, pin number *B4* represents the pin located in row *B* and column 4. The letters *I*, *O*, *Q*, *S*, *X*, and *Z* are never used in pin numbers. If the device contains more rows than letters of the alphabet, the alphabet is repeated, prefixed with the letter *A*.

Figure 5-2. Row and Column Labeling

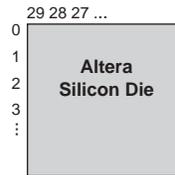


For more information about the pin numbers for your Altera device, refer to the [Pin-Out Files for Altera Devices](#) page of the Altera website.

Pads

Package pins are connected to pads located on the perimeter of the top metal layer of the silicon die (refer to [Figure 5-1](#)). [Figure 5-3](#) shows the numbering scheme for pads on the device. Each pad is identified by a pad ID, which is numbered starting at 0, and increments by one in a counterclockwise direction around the device.

Figure 5-3. Pad Number Ordering



To prevent signal integrity issues, the Quartus II software uses pin placement rules to validate your pin placements and pin-related assignments. You must understand the pad locations to which your pins were assigned, because some pin placement rules describe pad placement restrictions. For example, to ensure signal integrity, in certain devices there is a restriction on the number of I/O pins supported by a VREF pad. There are also restrictions on the number of pads between single-ended input or output pins and a differential pin. The Quartus II software performs pin placement analysis, and if pins are not placed according to the pin placement rules, the design compilation fails and the Quartus II software reports an error.

 For more information about pin placement guidelines, refer to the appropriate device handbook available on the [Literature and Technical Documentation](#) page of the Altera website.

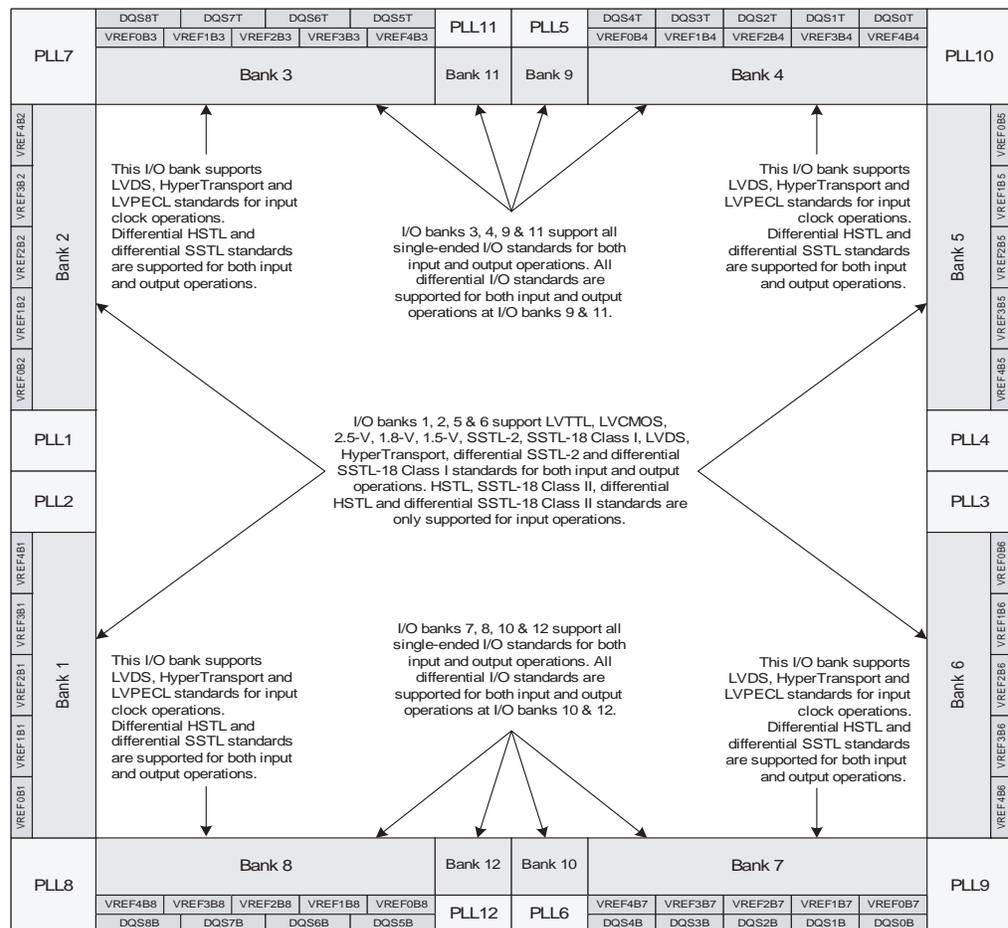
I/O Banks

I/O pins are organized into I/O banks designed to facilitate various supported I/O standards. Each I/O bank is numbered and has its own voltage source pins, called V_{CCIO} pins, to offer the highest I/O performance. Depending on the device and I/O standards for the pins in the I/O bank, the specified voltage of the V_{CCIO} pin is between 1.5 V and 3.3 V. Each I/O bank can support multiple pins with different I/O standards, however the pins must use the same V_{CCIO} signal.

[Figure 5-4](#) shows the I/O banks of a Stratix[®] II device. The pins in the I/O banks on the left and right side support high-speed I/O standards such as LVDS, whereas the pins on the top and bottom I/O banks support all single-ended I/O standards, including data strobe signaling (DQS). Pins belonging to the same I/O bank must use the same V_{CCIO} signal.

For more information about the capabilities of each I/O bank, refer to the appropriate device handbook available on the [Literature and Technical Documentation](#) page of the Altera website.

Figure 5-4. Stratix II I/O Banks (Note 1), (2), (3), (4)



Notes to Figure 5-4:

- (1) This figure shows a top view of the silicon die that corresponds to a reverse view for flip chip packages. It is a graphical representation only.
- (2) Depending on the size of the device, different device members have a different number of VREF groups. For more information, refer to the pin list and the Quartus II software for exact locations.
- (3) Banks 9 through 12 are enhanced phase-locked loop (PLL) external clock output banks.
- (4) Horizontal I/O banks feature serializer/deserializer (SERDES) and dynamic phase alignment (DPA) circuitry for high-speed differential I/O standards. For more information about differential I/O standards, refer to the *High-Speed Differential I/O Interfaces with DPA in Stratix II and Stratix II GX Devices* chapter in volume 2 of the *Stratix II Device Handbook*.

VREF Groups

A VREF group is a group of pins that includes one dedicated VREF pin as required by voltage-referenced I/O standards. A VREF group is made up of a small number of pins, as compared to the I/O bank, to maintain the signal integrity of the VREF pin. One or more VREF groups exist in an I/O bank. The pins in a VREF group share the same V_{CCIO} and V_{REF} voltages.

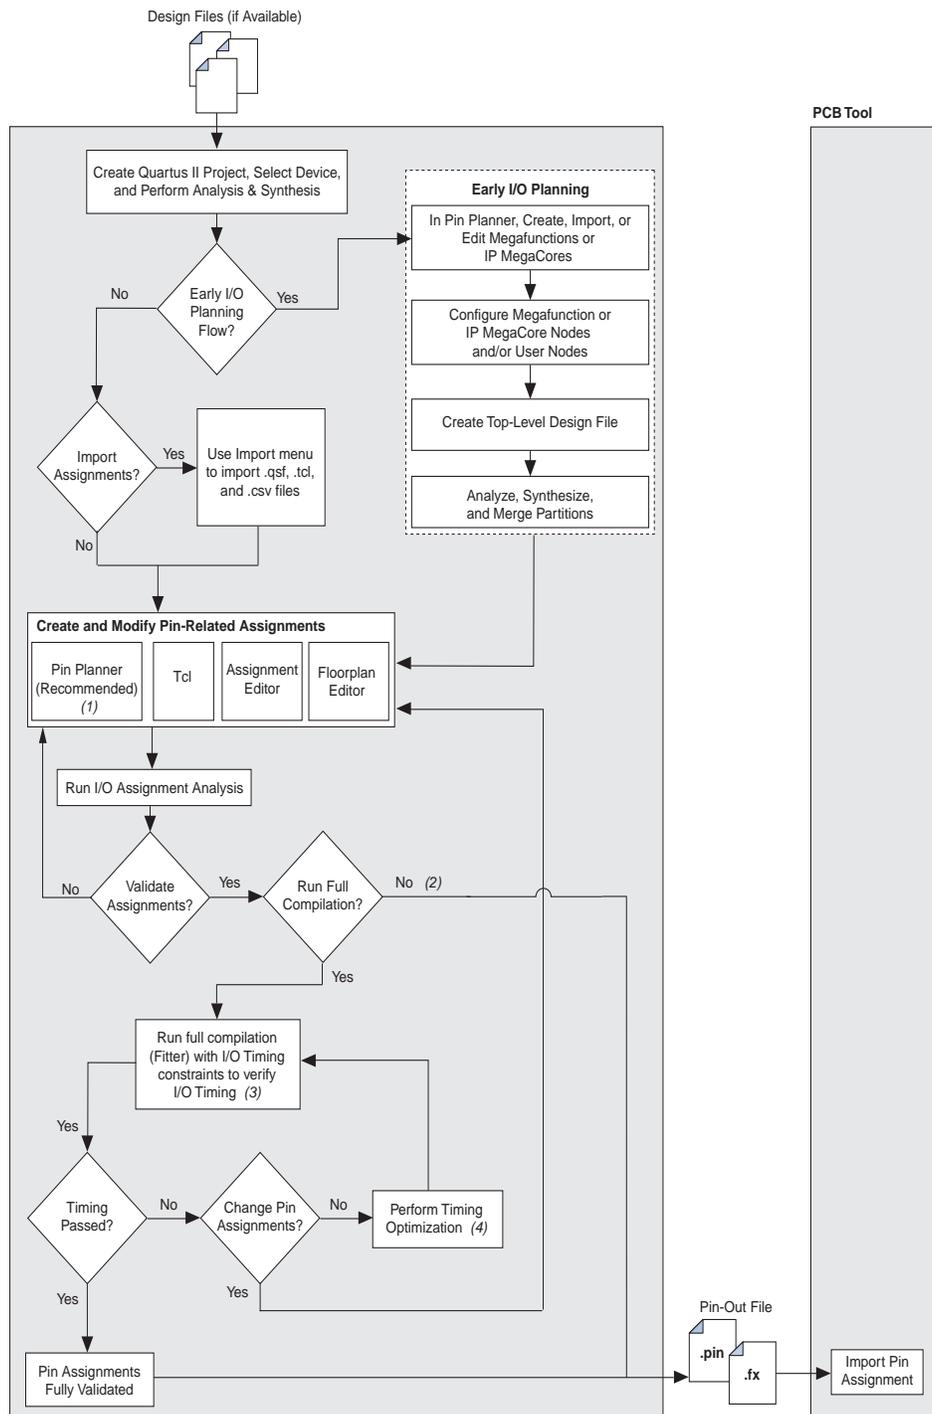
- For more information about I/O banks, VREF groups, and supported I/O standards, refer to the appropriate device handbook available on the [Literature and Technical Documentation](#) page of the Altera website.

I/O Planning Overview

The most important step in I/O planning is to create, modify, complete, and validate pin-related assignments. The Quartus II software includes the Pin Planner and the I/O assignment analysis feature to assist you in I/O planning. The method you use to create I/O assignments depends on your design requirements. [Figure 5-5](#) shows the recommended design flow for creating I/O assignments for a design. I/O planning for your design in the Quartus II software can include the following tasks:

- Selecting a device that meets your logic and I/O requirements, based on the supported I/O standards for the device, I/O bank structure, supply voltage requirements such as V_{REF} and V_{CCIO} requirements in I/O banks, available pins for user I/O, power supply requirements, and other design requirements.
- Preparing your design files. Design files contain the top-level ports or top-level interface information; if you do not have the design files, you can use the Early I/O Planning flow (refer to [Figure 5-5](#)) to generate a top-level HDL wrapper file.
- Importing any existing assignments from a Tcl Script File (.tcl), Comma-Separated Value File (.csv), or Quartus II Settings File (.qsf).
- Creating, modifying, and completing all pin-related assignments, including pin location assignments, I/O standard assignments, output loading assignments for output and bidirectional pins, slew rate assignments, current strength assignments.
- Validating your pin-related assignments. You can perform preliminary I/O assignment validation while creating the assignments with the live I/O check feature; you can perform a more thorough validation with the I/O assignment analysis feature; and finally you can perform complete I/O assignment validation by running the Fitter with timing constraints.
- Generating a validated Pin-Out File (.pin) for third-party PCB tools.

Figure 5-5. Quartus II Software I/O Planning Flow

**Notes to Figure 5-5:**

- (1) Use the live I/O check feature in the Pin Planner to validate pin assignments as you create them.
- (2) To create the FPGA Xchange file (.fx), on the Processing menu, point to **Start** and click **EDA Netlist Writer**. The .pin is created at the <project_dir> level. The .fx is created at the <project_dir/board/.../> level.
- (3) You must complete your design files and constraints before you begin full compilation. To learn how to create I/O timing constraints, refer to *The Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.
- (4) For more information, refer to the *Area and Timing Optimization* chapter in volume 2 of the *Quartus II Handbook*.

Selecting a Device

Before you begin I/O planning or I/O assignment analysis you must first select a device family that has an appropriate I/O structure, supports the correct I/O standards, has enough available pins for user I/O, supports the correct clocking schemes and options, and has an appropriate I/O bank structure for your design. Then, you must choose an appropriate device from a supported device family for your design.

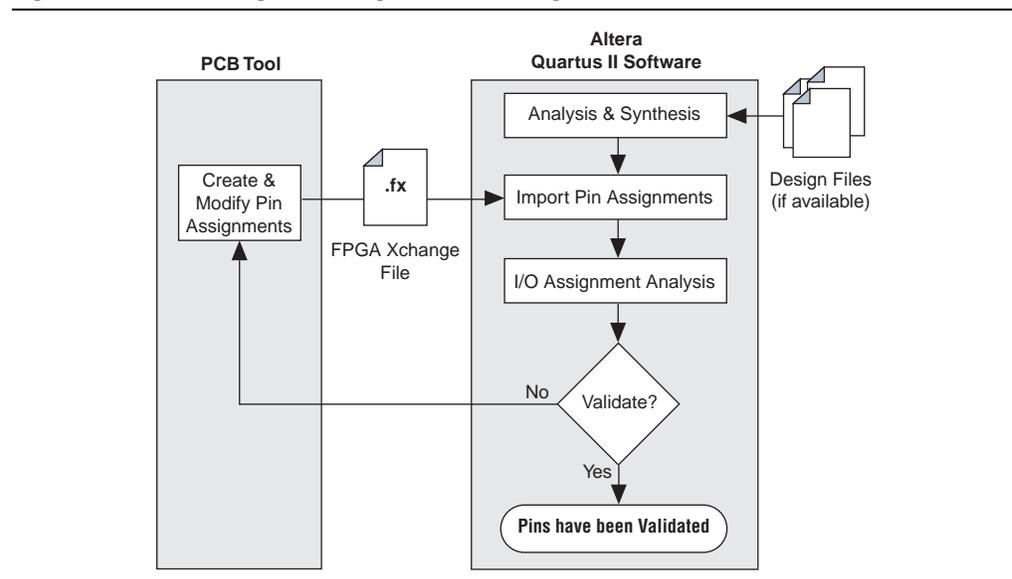
-  For more information, refer to the various device handbooks available on the [Literature and Technical Documentation](#) page of the Altera website.
-  For more information about selecting a device in the Quartus II software, refer to [Setting Up and Running a Compilation](#) in Quartus II Help.

Working with Third-Party PCB Tools

If you have not designed your PCB, the recommended I/O planning flow (refer to [Figure 5-5](#)) indicates that you should first create and validate your I/O assignments in the Quartus II software, and then export them to the PCB tool. If your PCB is partially designed, [Figure 5-6](#) shows the recommend I/O planning flow to ensure that your pin assignments are valid. First, you create your device assignments in your PCB tool and then import them into the Quartus II software for validation.

-  Currently, only the Mentor Graphics® I/O Designer PCB tool and the Cadence Allegro PCB tool are supported in this reverse I/O planning flow.

Figure 5-6. I/O Planning Flow Using an FPGA Xchange File from a PCB Tool



Performing Early I/O Planning with the Pin Planner

In design planning, typically you create design elements in an HDL such as Verilog HDL or VHDL, or in a schematic editor such as the Quartus II Block Editor. Central to the design is a top-level file that instantiates the next level of hierarchy and includes port names and their direction. [Example 5-1](#) shows a top-level file, written in Verilog HDL, that lists the input and output ports in a design.

Example 5-1. Top-Level Design File

```
module top (      clk_in,
                  rst,
                  a,
                  z,
                  b,
                  c_in,
                  d,
                  e);

input clk_in;
input rst_;
input a;
input z;
input [7:0] b;
input [7:0] c_in;
input [7:0] d;
output reg [7:0] e;

/* Instantiations of sub blocks */

endmodule
```

Top-level design files often contain interfaces for memory, high-speed I/O, device configuration, and debugging tools. Listing the ports in an HDL or drawing them in the schematic can be extremely time-consuming. To reduce the time spent in creating top-level design files, you can use the Pin Planner to create a top-level design file if the design files for your entire project are not available or are incomplete. The interfaces between your target device and other devices are determined and documented in design specifications. By adding the interfaces required to connect your target device and other devices in your design directly in the Pin Planner, you can plan your I/O assignments efficiently without design files, and generate a top-level module in Verilog HDL or VHDL. By importing or creating MegaCore® functions or Altera megafunctions in the Pin Planner, as well as creating or adding additional top-level I/O information, the generated top-level design file accurately anticipates the rest of the HDL to come.

The following sections describe the typical steps of the early I/O planning flow:

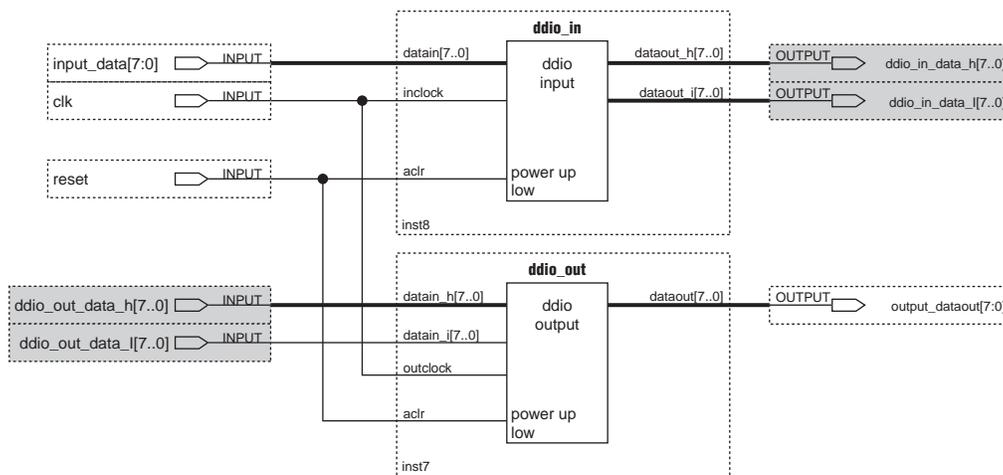
- [“Instantiating or Importing IP Cores in the Pin Planner”](#)
- [“Adding and Connecting Nodes” on page 5-9](#)
- [“Setting Up and Creating the Top-Level Design File” on page 5-10](#)

Instantiating or Importing IP Cores in the Pin Planner

You can use the MegaWizard™ Plug-In Manager from the Pin Planner to instantiate or import custom IP cores to assist you in I/O planning. Adding interface information directly in the Pin Planner allows you to assign required pins without manually creating each pin individually.

You can create complex interfaces from the Pin Planner. Figure 5-7 shows the megafunctions `ddio_in` and `ddio_out` in the schematic created by invoking the MegaWizard Plug-In Manager from the Pin Planner. The ports shown in gray are ports specified as internal ports, which are connected later in the design.

Figure 5-7. Connections between Input and Output Megafunctions and User Nodes



- For more information about creating or customizing IP cores with the Pin Planner, refer to *Managing Pins and Groups in the Pin Planner* and *Create/Import Megafunction Dialog Box* in Quartus II Help.

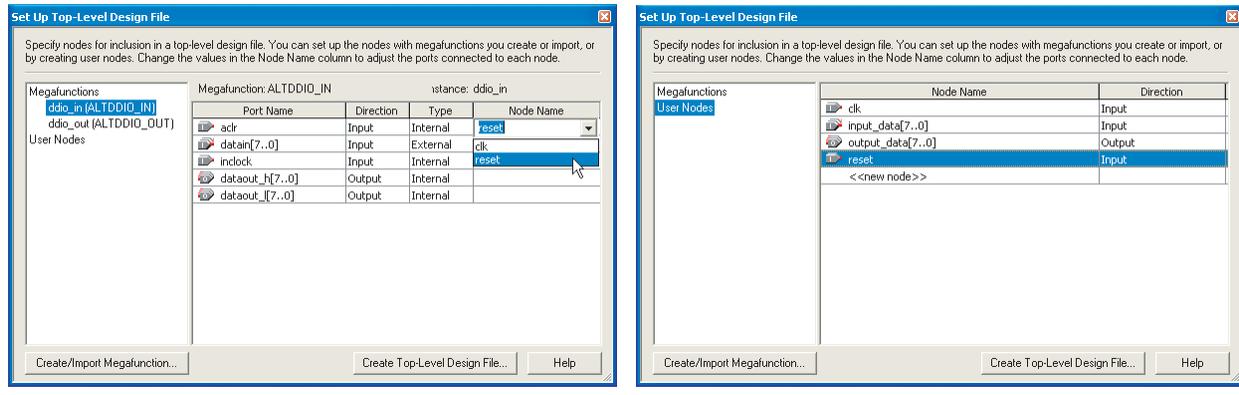
Adding and Connecting Nodes

After you instantiate an IP core in your design, you can set up the user nodes in the design and customize information such as port direction and type in the **Set Up Top-Level File** window. Before you create a top-level design file, you must connect the user nodes and ports to each other and to the rest of the design. You also can add new nodes to the design in the **Set Up Top-Level File** window.

The IP core nodes you specify as internal nodes in the **Set Up Top-Level Design File** dialog box are declared as virtual pins when you generate the top-level design file. The Pin Planner makes virtual pin assignments to internal nodes so that internal nodes are not assigned to device pins during compilation. These virtual pins are not shown in the **All Pins** list or **Groups** list in the Pin Planner because they are not actual external ports of the design. Any new nodes you add to the design in the **Set Up Top-Level Design File** dialog box are declared as ports in the top-level design file and are shown in the **All Pins** list and **Groups** list.

Figure 5-8 shows how to make the user node connection between the **reset** signal and the megafunction reset input port shown in Figure 5-7.

Figure 5-8. Connecting a User Node to a Megafunction Port



- ❓ For more information about connecting nodes from IP cores, refer to *Generating a Top-Level Design File Based on Pin Planner Megafunctions and User Nodes* in Quartus II Help.

Setting Up and Creating the Top-Level Design File

You can create a top-level design file after you create pin connections and add or modify user nodes and IP core nodes with the Pin Planner. If the internal logic is incomplete, generating the top-level design file allows you to validate your I/O assignments and provides a basis on which to build the rest of your design.

- 👉 You must update the top-level design file whenever you change the top-level ports of the design, including any node changes made in the **Set Up Top-Level Design File** window.

Example 5-2 shows a sample of a top-level HDL wrapper file representing the design in Figure 5-7.

Example 5-2. HDL Wrapper File Generated with the Early I/O Planning Flow

```
module top
(
    reset,
    input_data,
    clk,
    output_data,
    ddio_in_dataout_h, // Internal
    ddio_in_dataout_l, // Internal
    ddio_out_datain_h, // Internal
    ddio_out_datain_l  // Internal
);

input reset;
input[7:0]input_data;
input clk;
output[7:0]output_data;

output[7:0]ddio_in_dataout_h /* synthesis altera_attribute="-name VIRTUAL_PIN ON" */;
output[7:0]ddio_in_dataout_l /* synthesis altera_attribute="-name VIRTUAL_PIN ON" */;
input[7:0]ddio_out_datain_h /* synthesis altera_attribute="-name VIRTUAL_PIN ON" */;
input[7:0]ddio_out_datain_l /* synthesis altera_attribute="-name VIRTUAL_PIN ON" */;

ddio_in ddio_in_inst
(
    .aclr(reset),
    .datain(input_data),
    .inclock(clk),
    .dataout_h(ddio_in_dataout_h),
    .dataout_l(ddio_in_dataout_l)
);

ddio_out ddio_out_inst
(
    .aclr(reset),
    .datain_h(ddio_out_datain_h),
    .datain_l(ddio_out_datain_l),
    .outclock(clk),
    .dataout(output_data)
);

endmodule
```

- ❓ For more information about generating a top-level design file, refer to *Generating a Top-Level Design File Based on Pin Planner Megafunctions and User Nodes* in Quartus II Help.

After you generate the top-level design file and compile the design, use I/O assignment analysis as described in “Validating Pin Assignments with I/O Assignment Analysis” on page 5-25 and continue with your design flow by modifying or creating pin assignments with the Pin Planner.

Importing and Exporting Pin Assignments

If you have existing pin assignments in the file formats `.tcl`, `.csv`, `.qsf`, `.fx`, or `.pin` from a different Quartus II project or from third-party PCB tools, you can transfer these assignments between the Quartus II software and other tools.

Importing and Exporting Assignments with the Quartus II Software

You can import and export pin-related assignments contained in `.tcl`, `.csv`, and `.qsf` files with the Quartus II software. You can import assignments from `.tcl`, `.csv`, and `.qsf` files and then view and change the assignments in the Pin Planner.

When you create pin assignments with the Pin Planner all your pin-related assignments are written to the `.qsf` as Tcl commands. You can import and export `.qsf` files between projects to transfer all assignments, including pin assignments.

 For more information, refer to *Importing and Exporting Assignments* in Quartus II Help.

 For more information about `.qsf` files, refer to the *Managing Quartus II Projects* chapter in volume 2 of the *Quartus II Handbook*.

When you export pin assignments from the Pin Planner as a `.csv`, the row and column headings in the exported file is in the same order and format as the columns displayed in the **All Pins** list in the Pin Planner. Do not modify the row or column headings if you plan to import the `.csv` file later.

When you export pin assignments as Tcl commands in a `.tcl`, you create a script you can later run to add the assignments as part of a scripted compilation flow.

 For more information about importing and exporting pin assignments as `.tcl` and `.csv` files, refer to *Assigning Pins* in Quartus II Help.

 For more information about Quartus II scripting support, including examples, refer to the *Tcl Scripting* and *Command-Line Scripting* chapters in volume 2 of the *Quartus II Handbook*.

Importing and Exporting Assignments with Third-Party PCB Tools

After creating and validating your pin assignments, you can export device and pin-related information from the Quartus II software to third-party PCB tools for board development. You must generate a `.pin` with the Quartus II software to export the correct pin locations and other important pin information.

The device `.pin` files display the pin name and pin number as well as detailed properties about each pin. [Table 5-1](#) describes the columns in a `.pin`.

Table 5-1. .pin Header Description (Part 1 of 2)

Column Name	Description
Pin Name/Usage	The name of the design pin, or whether the pin is GND or V_{CC} pin
Location	The pin number of the location on the device package
Dir	The direction of the pin

Table 5-1. .pin Header Description (Part 2 of 2)

Column Name	Description
I/O Standard	The name of the I/O standard to which the pin is configured
Voltage	The voltage level that is required to be connected to the pin
I/O Bank	The I/O bank to which the pin belongs
User Assignment	Y or N indicating if the location assignment for the design pin was user assigned (Y) or assigned by the Fitter (N)

 For more information, refer to the [Pin-Out Files for Altera Devices](#) page of the Altera website.

To transfer device and pin-related information between the Quartus II software and the Mentor Graphics I/O Designer software you must generate an **.fx**. Importing assignments into the I/O Designer software requires both the **.fx** and a **.pin** generated by the Quartus II software. However, the Quartus II software requires only the **.fx** to import pin assignments back from the I/O Designer software.

 For more information, refer to [Importing and Exporting Assignments](#) in Quartus II Help.

 For more information about the I/O Designer software and the DxDesigner interface, refer to the [Mentor Graphics PCB Tools Support](#) chapter in volume 2 of the *Quartus II Handbook*.

Creating Pin-Related Assignments

A pin-related assignment is any assignment applied to a top-level pin. For example, a pin location assignment assigns a top-level port or node to a pin number (location) on the targeted device. Other examples of pin-related assignments include assigning an I/O standard, assigning drive strength, or assigning a slew rate to a pin.

When making pin assignments, if you do not have complete information for all the top-level pins, you can reserve certain device pins to temporarily represent your top-level design I/O pins until the I/O pins are defined in your design files. Reserved pins are pins you intend to use in the future, but do not currently perform a function in your design. Reserved pins require a unique pin name and pin location. Using reserved pins as place holders for future design pins increases the accuracy of I/O assignment analysis.

 For more information about reserving pins with the Pin Planner, refer to [Assigning Pins](#) in Quartus II Help.

Table 5-2 describes the tools and features in the Quartus II software for creating reserved pins and other pin-related assignments. Each tool and feature is described in more detail in the following sections. Altera recommends that you use the Pin Planner to create and edit pin-related assignments; however, depending on your design flow, you may find some of the other tools useful for working with pin-related assignments.

Table 5-2. Overview of Quartus II Tools and Features to Create Pin-Related Assignments

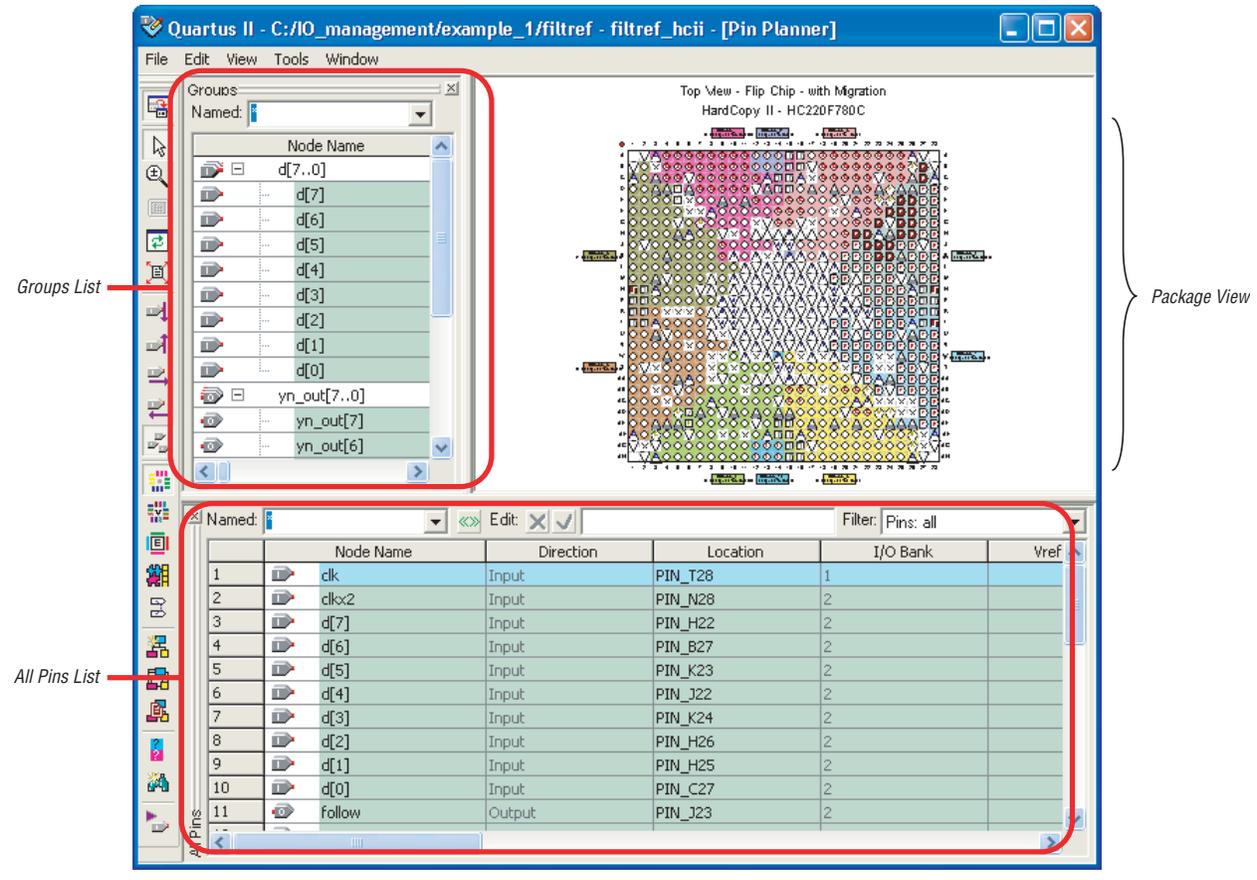
Feature	Overview
Pin Planner	<ul style="list-style-type: none"> ■ Create pin location assignments to one or more node names by dragging and dropping unassigned pins into the package view ■ Edit pin location assignments for one or more node names by dragging and dropping groups of pins in the package view ■ Visually analyze pin resources in the package view ■ Display I/O banks, VREF groups, and differential pin pairs ■ View the function of package pins with the Pin Legend window ■ Make correct pin location decisions by referring to the Pad View window ■ Create, import, and edit IP cores for early I/O planning ■ Generate a top-level wrapper file without design files based on early I/O assignments ■ Configure board trace model assignments, instead of capacitive loading assignments, to generate Advanced I/O Timing results
Tcl Scripts	<ul style="list-style-type: none"> ■ Create any pin-related assignments for multiple pins ■ Store and reapply all pin-related assignments with Tcl scripts ■ Create assignments from the command line
Chip Planner	<ul style="list-style-type: none"> ■ Create and change pin locations by dragging and dropping pins into the floorplan ■ Make correct pin location decisions by referring to the pad ID number and spacing ■ Display I/O banks, VREF groups, and differential pin pairing information
Synthesis Attributes	<ul style="list-style-type: none"> ■ Embed pin-related assignments with attributes in the design files to pass assignments to the Quartus II software

Creating Pin Assignments With the Pin Planner

During I/O planning, it can be cumbersome to try to correlate pin numbers with their relative location on the package and their pin properties. The Pin Planner (refer to Figure 5-9) is an interface for creating and editing pin-related assignments. The Pin Planner provides an intuitive graphical representation of the target device to make it easy to plan your I/Os, create reserved pins, and create pin location assignments. With the Pin Planner, you can identify I/O banks, VREF groups, and differential pin pairings to help you through the I/O planning process. When deciding on a pin location, use the Pin Planner to gather information about available resources, as well as the functionality of each individual pin, I/O bank, and VREF group.

- For more information about the Pin Planner, refer to *About the Pin Planner* in Quartus II Help. For more information about creating pin assignments with the Pin Planner, refer to *Assigning Pins* in Quartus II Help.

Figure 5-9. Pin Planner



Finding Compatible Pin Locations with the Pin Finder

As device pin-counts and I/O capabilities continue to increase, it becomes more difficult to understand the capabilities of each I/O pin and to correctly assign your design I/Os. As you edit pin assignments to help your design fit or to achieve timing goals, you can use the tools in the Pin Planner to help you find appropriate pin locations for your design I/O nodes.

- For more information, refer to *Managing Pins and Groups in the Pin Planner* in Quartus II Help.

Verifying Pin Migration Compatibility

You can use the Pin Migration View window in Pin Planner to assist you in verifying whether your pin assignments migrate to a different device successfully. You can vertically migrate to a device with a different density while using the same device package, or migrate between packages with different densities and ball counts.

- ❓ For more information, refer to *Viewing Pin Migration Compatibility* in Quartus II Help.

When you select migration devices early in the design process, the Pin Planner displays only the pins that are available in the current device and in all migration devices. If you select migration devices later in your design cycle, there may be assignments for I/O nodes in your original design that do not have corresponding pins in a migration device. If no corresponding pin exists, the Compiler cannot honor the assignment and an error occurs when you try to recompile the design.

The Pin Migration View window helps you identify the difference in pins that can exist between migration devices. For example, [Figure 5-10](#) shows the highlighted pin AC24 existed in the target EP2S30 device, but does not exist in one of the migration devices, resulting in a No Connect (NC).

Figure 5-10. Pin Migration View

Pin Number	Pin	Migration Result				Migration Devices											
		Pin Function	I/O Bank	VREF Group	Clock Pin	EP2S30F672C4				EP2S15F672C4				EP2S60F672C4			
						Pin Function	I/O Bank	VREF Group	Clock Pin	Pin Function	I/O Bank	VREF Group	Clock Pin	Pin Function	I/O Bank	VREF Group	Clock Pin
87	PIN_AC11	VREFB7N0	7	B7_N0		VREFB7N0	7	B7_N0		VREFB7N0	7	B7_N0		VREFB7N0	7	B7_N0	
88	PIN_AC12	Column I/O	10	B7_N0	Yes	Column I/O	10	B7_N0	Yes	Column I/O	10	B7_N0	Yes	Column I/O	10	B7_N0	Yes
89	PIN_AC13	Column I/O	7	B7_N0	Yes	Column I/O	7	B7_N0	Yes	Column I/O	7	B7_N0	Yes	Column I/O	7	B7_N0	Yes
90	PIN_AC14	Column I/O	8	B8_N1	Yes	Column I/O	8	B8_N1	Yes	Column I/O	8	B8_N1	Yes	Column I/O	8	B8_N2	Yes
91	PIN_AC15	NC				Column I/O	8	B8_N1		NC				Column I/O	12	B8_N2	Yes
92	PIN_AC16	VREFB8N1	8	B8_N1		VREFB8N1	8	B8_N1		VREFB8N1	8	B8_N1		VREFB8N2	8	B8_N2	
93	PIN_AC17	Column I/O	8	B8_N1		Column I/O	8	B8_N1		Column I/O	8	B8_N1		Column I/O	8	B8_N1	
94	PIN_AC18	Column I/O	8	B8_N0		Column I/O	8	B8_N0		Column I/O	8	B8_N1		Column I/O	8	B8_N0	
95	PIN_AC19	Column I/O	8	B8_N0		Column I/O	8	B8_N0		Column I/O	8	B8_N0		Column I/O	8	B8_N0	
96	PIN_AC20	Column I/O	8	B8_N0		Column I/O	8	B8_N0		Column I/O	8	B8_N0		Column I/O	8	B8_N0	
97	PIN_AC21	Column I/O	8	B8_N0		Column I/O	8	B8_N0		Column I/O	8	B8_N0		Column I/O	8	B8_N0	
98	PIN_AC22	VREFB8N0	8	B8_N0		VREFB8N0	8	B8_N0		VREFB8N0	8	B8_N0		VREFB8N0	8	B8_N0	
99	PIN_AC23	VREFB1N2	1	B1_N2		Column I/O	8	B8_N0		NC				VREFB1N2	1	B1_N2	
100	PIN_AC24	NC				Row I/O	1	B1_N1		NC				Row I/O	1	B1_N1	
101	PIN_AC25	NC				Row I/O	1	B1_N1		NC				Row I/O	1	B1_N1	
102	PIN_AC26	VCCIO1	1			VCCIO1	1			VCCIO1	1			VCCIO1	1		
103	PIN_AD1	NC				Row I/O	6	B6_N0		NC				Row I/O	6	B6_N1	
104	PIN_AD2	NC				Row I/O	6	B6_N0		NC				Row I/O	6	B6_N1	
105	PIN_AD3	Column I/O	7	B7_N1		Column I/O	7	B7_N1		Column I/O	7	B7_N1		Column I/O	7	B7_N2	
106	PIN_AD4	Column I/O	7	B7_N1		Column I/O	7	B7_N1		Column I/O	7	B7_N1		Column I/O	7	B7_N2	
107	PIN_AD5	Column I/O	7	B7_N1		Column I/O	7	B7_N1		Column I/O	7	B7_N1		Column I/O	7	B7_N2	
108	PIN_AD6	Column I/O	7	B7_N1		Column I/O	7	B7_N1		Column I/O	7	B7_N1		Column I/O	7	B7_N2	
109	PIN_AD7	Column I/O	7	B7_N1		Column I/O	7	B7_N1		Column I/O	7	B7_N1		Column I/O	7	B7_N1	
110	PIN_AD8	Column I/O	7	B7_N0		Column I/O	7	B7_N0		Column I/O	7	B7_N0		Column I/O	7	B7_N1	
111	PIN_AD9	Column I/O	7	B7_N0		Column I/O	7	B7_N0		Column I/O	7	B7_N0		Column I/O	7	B7_N1	
112	PIN_AD10	Column I/O	7	B7_N0		Column I/O	7	B7_N0		Column I/O	7	B7_N0		Column I/O	7	B7_N0	
113	PIN_AD11	Column I/O	7	B7_N0		Column I/O	7	B7_N0		Column I/O	7	B7_N0		Column I/O	7	B7_N0	
114	PIN_AD12	Column I/O	10	B7_N0	Yes	Column I/O	10	B7_N0	Yes	Column I/O	10	B7_N0	Yes	Column I/O	10	B7_N0	Yes
115	PIN_AD13	Column I/O	10	B7_N0	Yes	Column I/O	10	B7_N0	Yes	Column I/O	10	B7_N0	Yes	Column I/O	10	B7_N0	Yes
116	PIN_AD14	Column I/O	7	B7_N0	Yes	Column I/O	7	B7_N0	Yes	Column I/O	7	B7_N0	Yes	Column I/O	7	B7_N0	Yes

The migration result for the pin function of highlighted PIN_AC23 is not an NC but a voltage reference VREFB1N2 even though the pin is an NC in one of the migration devices. VREF standards have a higher priority than an NC, thus the migration result display the voltage reference. Even if you do not use that pin for a port connection in your design, you must use the VREF standard for I/O standards that require it on the actual board for the migration device.

If one of the migration devices has pins intended for connection to V_{CC} or GND and these same pins are I/O pins on a different device in the migration path, the Quartus II software ensures these pins are not used for I/O. Ensure that these pins are connected to the correct PCB plane.

When migrating between two devices in the same package, pins that are not connected to the smaller die may be intended to connect to V_{CC} or GND on the larger die. To facilitate migration, you can connect these pins to V_{CC} or GND in your original design because the pins are not physically connected to the smaller die.

-  For more information about migration, refer to *AN90: SameFrame Pin-Out Design for FineLine BGA Packages*. For more information about designing for HardCopy series devices, refer to the *Quartus II Support for HardCopy Series Devices* chapter in volume 1 of the *Quartus II Handbook*.

Viewing Simultaneous Switching Noise (SSN) Results

You can perform and SSN analysis of your design to estimate the voltage noise for each pin in the design. You can view the SSN results in the Pin Planner and adjust your I/O assignments to avoid potential signal integrity issues.

-  For more information about running the SSN Analyzer and viewing the results in the Pin Planner, refer to *Running the SSN Analyzer* in Quartus II Help.
-  For more information about the SSN Analyzer, refer to the *Simultaneous Switching Noise (SSN) Analysis and Optimization* chapter in volume 2 of the *Quartus II Handbook*.

Creating Location Assignments

You can create the following types of location assignments for your design and its reserved pins:

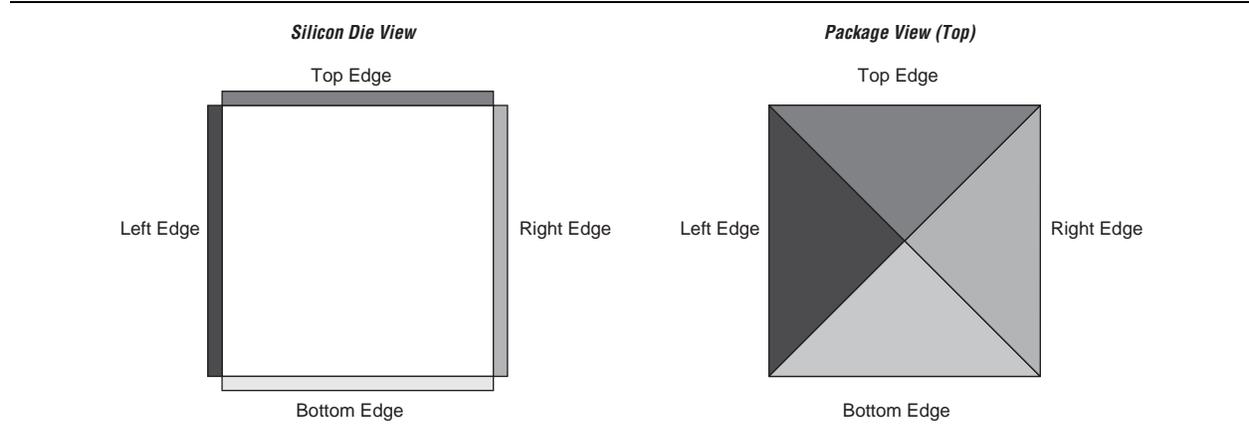
- Pin number
 - I/O bank
 - VREF group
 - Edge
-  For more information about device support for I/O bank, VREF group, and edge location assignments, refer to *I/O bank*, *VREF group*, and *edge* in Quartus II Help.

You can assign your pins to a location with the Pin Planner. It is common to place a group of pins (or bus) with compatible I/O standards in the same I/O bank or VREF group. For example, two buses with two compatible I/O standards, such as **2.5-V** and **SSTL-II Class I**, can be placed in the same I/O bank.

If your design contains a large bus that exceeds the pins available in a particular I/O bank, you can use edge location assignments to place the bus. Edge location assignments improve the circuit board routing ability of large buses, because they are close together near an edge. [Figure 5-11](#) shows Altera device package edges.

- For more information about creating pin assignments with the Pin Planner, refer to *Assigning Pins* in Quartus II Help.

Figure 5-11. Die View and Package View of the Four Edges on an Altera Device



Creating Exclusive I/O Group Assignments

You can create exclusive groups comprised of pins by creating and modifying the **Exclusive I/O Group** logic option with the Pin Planner. When you create exclusive I/O groups in your design and use the Quartus II software to map the signals onto device pins, the Fitter does not place the I/O pins belonging to one exclusive group in an I/O bank if the pins belong to another exclusive I/O group. To understand this, consider an example in which you have a set of signals assigned exclusively to a group called `group_a`, and another set of signals assigned to `group_b`. In both exclusive groups you have pins with different I/O standards. When you create these groups, the Quartus II software maps the pins of both groups in such a way that they are placed in different I/O banks.

- For more information about creating logic option assignments with the Pin Planner, refer to *Assigning Pins* in Quartus II Help.

Changing the Slew Rate and Drive Strength

As part of I/O planning you can set both the slew rate and drive strength of pins. Both slew rate and drive strength affect the outgoing signal integrity of the pin. Depending on your target device family, you can adjust the slew rate of a pin with either the **Slew Rate** or **Slow Slew Rate** logic option; you can adjust the drive strength of a pin with the **Current Strength** logic option. The settings you create for slew rate and drive strength are honored during live I/O check, I/O assignment analysis, and full compilation.

- For more information about slew rate support, refer to the appropriate device handbook available on the [Literature and Technical Documentation](#) page of the Altera website.

- For more information about creating logic option assignments with the Pin Planner, refer to *Assigning Pins* in Quartus II Help.

If you want to make changes to slew rate or drive strength after you compile your design, you can use the Resource Property Editor to perform engineering change orders (ECOs). ECOs allow you to compile the changes to your design, without changing the synthesis or fitting results.

? For more information about making post-compilation changes, refer to *About Post-Compilation Changes* in Quartus II Help.

For more information about ECOs, refer to the *Engineering Change Management with the Chip Planner* chapter in volume 2 of the *Quartus II Handbook*. For more information about the effect of I/O settings on signal integrity on the board, refer to *AN 476: Impact of I/O Settings on Signal Integrity in Stratix III Devices*.

Assigning Locations for Differential Pins

When you use the Pin Planner to assign a differential I/O standard to a single-ended top-level pin in your design, it automatically recognizes the negative pin as part of the differential pin pair assignment and creates the negative pin for you. The Quartus II software writes the location assignment for the negative pin to the .qsf file; however, the I/O standard assignment is not added to the .qsf file for the negative pin of the differential pair.

For example, [Figure 5-12](#) shows a design in which you have a top-level pin defined as `lvds_in` to which you assign a differential I/O standard. The Pin Planner automatically creates the differential pin, `lvds_in(n)`, to complete the differential pin pair.

☞ If you have a single-ended clock that feeds a PLL, assign the pin only to the positive clock pin of a differential pair in the target device. Single-ended pins that feed a PLL and are assigned to the negative clock pin device cause the design to not fit.

? For more information about identifying and assigning differential pins with the Pin Planner, refer to *Assigning Pins* in Quartus II Help.

For more information about assigning locations for differential pins in HDL code with low-level I/O primitives, refer to “[Creating Pin Assignments with Low-Level I/O Primitives](#)” on page 5-22.

Figure 5-12. Creating a Differential Pin Pair in the Pin Planner

Node Name	Differential Pair	I/O Standard	Direction
input_data[4]		3.3-V LVTTTL (default)	Input
input_data[3]		3.3-V LVTTTL (default)	Input
input_data[2]		3.3-V LVTTTL (default)	Input
input_data[1]		3.3-V LVTTTL (default)	Input
input_data[0]		3.3-V LVTTTL (default)	Input
lvds_in	lvds_in(n)	LVDS	Input
output_data[7]		3.3-V LVTTTL (default)	Output
output_data[6]		3.3-V LVTTTL (default)	Output
output_data[5]		3.3-V LVTTTL (default)	Output
output_data[4]		3.3-V LVTTTL (default)	Output
output_data[3]		3.3-V LVTTTL (default)	Output
output_data[2]		3.3-V LVTTTL (default)	Output
output_data[1]		3.3-V LVTTTL (default)	Output
output_data[0]		3.3-V LVTTTL (default)	Output
reset		3.3-V LVTTTL (default)	Input

Creating Pin Assignments with the Chip Planner

The floorplan of the device shows the pins in the same order as the pads of the device. Understanding the relative distance between a pad and related logic can help you meet your timing requirements. You can view the floorplan of the device in the Chip Planner and determine the distances between user I/O pads and V_{CC} , GND, and V_{REF} pads to avoid signal integrity issues.

 For more information about creating pin location assignments in the Chip Planner, refer to *Working with Assignments in the Chip Planner* in Quartus II Help.

 For more information about pin placement guidelines, refer to the appropriate device handbook available on the [Literature and Technical Documentation](#) page of the Altera website.

Creating Pin Assignments with Tcl Scripts

You can use Tcl scripts to create pin-related assignments as part of a script-based compilation flow. You can enter individual Tcl commands in the Tcl Console window of the Quartus II software. To run a Tcl script with the Quartus II software, type the following command at a system prompt:

```
quartus_sh -t my_tcl_script.tcl ↵
```

Example 5-3 shows the `set_location_assignment` and `set_instance_assignment` Tcl commands used to create pin-related assignments to the input pin `address[10]`.

Example 5-3. Tcl Commands to Create Pin-Related Assignments

```
set_location_assignment PIN M20 -to address[10] -comment "Address pin to Second FPGA"
set_instance_assignment -name IO_STANDARD "2.5 V" -to address[10]
set_instance_assignment -name CURRENT_STRENGTH_NEW "MAXIMUM CURRENT" -to address[10]
```

 For more information about creating and running Tcl scripts with the Quartus II software, refer to *Creating and Running Tcl Scripts* in Quartus II Help.

 For more information about using Tcl scripts to create pin-related assignments, refer to the *Tcl Scripting* chapter in volume 2 of the *Quartus II Handbook* and *API Functions for Tcl* in Quartus II Help.

Creating Pin Assignments in HDL Code

You can use synthesis attributes or low-level I/O primitives to embed pin-related assignments directly in your HDL code. When you analyze and synthesize your HDL code, the information in the HDL code is converted into the appropriate assignments. There are two ways to specify pin-related assignments with HDL code:

- Using synthesis attributes for signal names that are top-level pins
- Using low-level I/O primitives, such as `ALT_BUF_IN`, to specify input, output, and differential buffers, and for setting parameters or attributes

Synthesis Attributes

Synthesis attributes allow you to embed pin-related assignments in your HDL code. During Analysis and Synthesis, the Quartus II software reads these synthesis attributes and translates them into assignments. The assignments are then populated in the Pin Planner. If you modify or delete these pin assignments in the Pin Planner and then recompile your design, any changes made in the Pin Planner take precedence over the assignments you made with synthesis attributes in your HDL code. Quartus II integrated synthesis supports the `chip_pin`, `useioff`, and `altera_attribute` synthesis attributes.

- For more information about integrated synthesis, synthesis attributes, and syntax, refer to the *Quartus II Integrated Synthesis* chapter in volume 1 of the *Quartus II Handbook*. For more information about synthesis attributes supported by third-party synthesis tools, contact your vendor.

`chip_pin` and `useioff`

Use the `chip_pin` and `useioff` synthesis attributes to create pin location assignments and to create **Fast Input Register**, **Fast Output Register**, and **Fast Output Enable Register** logic option assignments, respectively, in your HDL code. For all other assignments, including pin-related assignments, use the `altera_attribute` synthesis attribute as discussed in the “`altera_attribute`” section.

- For more information, refer to *useioff VHDL Synthesis Attribute*, *useioff Verilog HDL Synthesis Attribute*, *chip_pin VHDL Synthesis Attribute*, and *chip_pin Verilog HDL Synthesis Attribute* in Quartus II Help.

Example 5-4 and **Example 5-5** use the `chip_pin` and `useioff` attributes to embed location and **Fast Input Register** logic option assignments in both a Verilog HDL and VHDL design file using the synthesis attributes.

Example 5-4. Verilog HDL Example

```
input my_pin1 /* synthesis chip_pin = "C1" useioff = 1 */;
```

Example 5-5. VHDL Example

```
entity my_entity is
    port(
        my_pin1: in std_logic
    );
end my_entity;

architecture rtl of my_entity is
    attribute useioff : boolean;
    attribute useioff of my_pin1 : signal is true;
    attribute chip_pin : string;
    attribute chip_pin of my_pin1 : signal is "C1";
begin -- The architecture body
end rtl;
```

altera_attribute

Use the `altera_attribute` synthesis attribute to create other pin-related assignments in your HDL code. The `altera_attribute` attribute is understood only by Quartus II integrated synthesis and supports all types of instance assignments.

- For more information, refer to *altera_attribute VHDL Synthesis Attribute*, and *altera_attribute Verilog HDL Synthesis Attribute* in Quartus II Help.

Example 5-6 and **Example 5-7** use the `altera_attribute` attribute to embed **Fast Input Register** logic option assignments and I/O standard assignments in both a Verilog HDL and a VHDL design file.

Example 5-6. Verilog HDL Example

```
input my_pin1 /* synthesis altera_attribute = "-name FAST_INPUT_REGISTER ON; -name
IO_STANDARD \"2.5 V\" " */ ;
```

Example 5-7. VHDL Example

```
entity my_entity is
    port(
        my_pin1: in std_logic
    );
end my_entity;
architecture rtl of my_entity is
begin

attribute altera_attribute : string;
attribute altera_attribute of my_pin1: signal is "-name FAST_INPUT_REGISTER ON;
-- The architecture body
end rtl;
```

Creating Pin Assignments with Low-Level I/O Primitives

You can create pin-related assignments for your top-level nodes using low-level I/O primitives, which allow you to create pin location assignments and set I/O standards, drive strengths, slew rates, and on-chip termination (OCT) value assignments. You can also use low-level differential I/O primitives to define both positive and negative pins of a differential pair in the HDL code for your design.

The pin-related assignments made with primitives do not automatically appear in the Pin Planner. When you use low-level I/O primitives to define various pin-related I/O assignments, the assignments are honored only after you perform a full compilation. After performing a full compilation, you can populate these assignments in the Pin Planner by back-annotating pin assignments.

- For more information about back-annotating assignments, refer to *Back-Annotating Assignments for A Project* in Quartus II Help. For more information about differential I/O primitives, refer to *Primitives* in Quartus II Help.

-  For more information about using low-level I/O primitives in your design, refer to the *Designing with Low-Level Primitives User Guide*.

Validating Pin Assignments

The Quartus II software includes predefined I/O rules to guide you in pin placement and checks your pin-related assignments against these rules during pin planning. You must validate all pin-related assignments in your design. During a full compilation, the Quartus II software does not report illegal pin assignments until the Fitter stage. To preliminarily validate pin-related assignments against the predefined I/O rules, you can use the live I/O check feature or run I/O assignment analysis after performing analysis and synthesis. Typically, the I/O assignment analysis completes quickly. By preliminarily validating your pin-related assignments before fully compiling your design, you can avoid recompiling your design to fix pin-related assignment errors, thus reducing your overall design time. To fully validate pin-related assignments against all I/O timing checks, you must perform a full compilation.

Table 5-3 and Table 5-4 list a subset of the I/O rule checks performed when you run I/O assignment analysis.

 For more information about each I/O rule, including which devices support which rules, refer to the device handbooks available on the [Literature and Technical Documentation](#) page of the Altera website.

Table 5-3. Examples of I/O Rule Checks (Part 1 of 2)

Rule	Description	HDL Required?
I/O bank capacity	Checks the number of pins assigned to an I/O bank against the number of pins allowed in the I/O bank.	No
I/O bank V_{CCIO} voltage compatibility	Checks that no more than one V_{CCIO} is required for the pins assigned to the I/O bank.	No
I/O bank VREF voltage compatibility	Checks that no more than one VREF is required for the pins assigned to the I/O bank.	No
I/O standard and location conflicts	Checks whether the pin location supports the assigned I/O standard.	No
I/O standard and signal direction conflicts	Checks whether the pin location supports the assigned I/O standard and direction. For example, certain I/O standards on a particular pin location can only support output pins.	No
Differential I/O standards cannot have open drain turned on	Checks that open drain is turned off for all pins with a differential I/O standard.	No
I/O standard and drive strength conflicts	Checks whether the drive strength assignments are within the specifications of the I/O standard.	No
Drive strength and location conflicts	Checks whether the pin location supports the assigned drive strength.	No
BUSHOLD and location conflicts	Checks whether the pin location supports BUSHOLD. For example, dedicated clock pins do not support BUSHOLD.	No
WEAK_PULLUP and location conflicts	Checks whether the pin location supports WEAK_PULLUP (for example, dedicated clock pins do not support WEAK_PULLUP)	No
Electromigration check	Checks whether combined drive strength of consecutive pads exceeds a certain limit. For example, the total current drive for 10 consecutive pads on a Stratix II device cannot exceed 200 mA.	No

Table 5-3. Examples of I/O Rule Checks (Part 2 of 2)

Rule	Description	HDL Required?
PCI_IO clamp diode, location, and I/O standard conflicts	Checks whether the pin location along with the I/O standard assigned supports PCI_IO clamp diode.	No
SERDES and I/O pin location compatibility check	Checks that all pins connected to a SERDES in your design are assigned to dedicated SERDES pin locations.	Yes
PLL and I/O pin location compatibility check	Checks whether pins connected to a PLL are assigned to the dedicated PLL pin locations.	Yes

Table 5-4. SSN-Related Rules

Rule	Description	HDL Required?
I/O bank can not have single-ended I/O when DPA exists	Checks that no single-ended I/O pin exists in the same I/O bank as a DPA.	No
A PLL I/O bank does not support both a single-ended I/O and a differential signal simultaneously	Checks that there are no single-ended I/O pins present in the PLL I/O Bank when a differential signal exists.	No
Single-ended output is required to be a certain distance away from a differential I/O pin	Checks whether single-ended output pins are a certain distance away from a differential I/O pin.	No
Single-ended output has to be a certain distance away from a VREF pad	Checks whether single-ended output pins are a certain distance away from a VREF pad.	No
Single-ended input is required to be a certain distance away from a differential I/O pin	Checks whether single-ended input pins are a certain distance away from a differential I/O pin.	No
Too many outputs or bidirectional pins in a VREFGROUP when a VREF is used	Checks that there are no more than a certain number of outputs or bidirectional pins in a VREFGROUP when a VREF is used.	No
Too many outputs in a VREFGROUP	Checks whether too many outputs are in a VREFGROUP.	No

Validating Pin Assignments with the Live I/O Check Feature

The live I/O check feature provides live I/O rule checking capability to prevent you from creating pin placements that violate I/O fitting rules. When the live I/O check feature is turned on, pin-related assignment error and warning messages appear immediately in the Quartus II Messages window as you create pin-related assignments in the Pin Planner. This feature enhances your productivity by showing you warnings and errors as you create pin-related assignments so you can immediately correct basic errors before you proceed to the next step in your design flow.

The most basic I/O rules are the I/O buffer rules. The I/O buffer rules checked by the live I/O check feature include:

- V_{CCIO} and V_{REF} voltage compatibility rules
- Electromigration (current density) rules
- Simultaneous Switching Output (SSO) rules

- I/O property compatibility rules, such as drive strength compatibility, I/O standard compatibility, PCI_IO clamp diode compatibility, and I/O direction compatibility

When the live I/O check feature is turned on, the Quartus II software prevents you from assigning pins to unavailable locations. The following are examples of unassignable locations:

- An I/O bank or VREF group with no available pins
- The negative pin of a differential pair if the positive pin of the differential pair is assigned with a node name with a differential I/O standard
- Pin locations that do not support the I/O standard assigned to the selected node name
- For HSTL- and SSTL-type I/O standards, VREF groups of a different V_{REF} voltage than the selected node name

You can turn on or turn off the live I/O check feature at any time. By default, the live I/O check feature is turned off. When the live I/O check feature is turned on, the Quartus II software immediately checks whether your new pin-related assignments pass the basic I/O buffer rules. The Live I/O Check Status window displays the total numbers of errors and warnings while you create and edit pin-related assignments. The Messages window shows detailed messages about any errors or warnings.

Although the live I/O check feature checks all the basic I/O buffer rules, you must run I/O assignment analysis to validate your pin-related assignments against the complete set of I/O system rules.

- ❓ For more information about using the live I/O check feature to validate pin assignments, refer to *Assigning Pins* in Quartus II Help.

Validating Pin Assignments with I/O Assignment Analysis

Performing I/O assignment analysis allows you to validate your I/O assignments and surrounding logic for illegal assignments and violations of board layout rules early in the design process. The I/O assignment analysis feature also checks blocks that directly feed or are fed by resources such as a PLLs, LVDS, or gigabit transceiver blocks. You can check the legality of pin assignments before you compile your design, or allow the process to run automatically during compilation. If design files are available, you can perform more thorough legality checks on the I/O pins and surrounding logic in your design. Legality checks include proper VREF pin use, valid pin location assignments, and acceptable mixed I/O standards. Altera recommends that you run I/O assignment analysis each time you add or modify a pin-related assignment.

- 👉 If you have partial or complete design files, you must perform Analysis and Synthesis to generate a synthesized (mapped) netlist before you can perform I/O assignment analysis.

Performing I/O assignment analysis directs the Fitter to read assignments from your mapped netlist and the .qsf to determine the legality of your pin-related assignments. These pin-related assignments include pin settings such as I/O standards, drive strength, and location assignments.

- ② For more information about performing I/O assignment analysis, refer to *Assigning Pins* in Quartus II Help.

Incomplete I/O assignments trigger warnings during I/O assignment analysis. You can view the I/O Assignment Warnings report to find and resolve warnings generated during I/O assignment analysis. For example, you may receive a warning that some of the pins in the design are missing a drive strength or slew rate. Single-ended output and bidirectional pins default to the non-calibrated OCT setting if you do not assign drive strength and slew rate options to the pins, or if other OCT options are assigned to the pins. To resolve this issue, you can either assign drive strength or slew rate settings to the pins with the **Current Strength** or **Slew Rate** or **Slow Slew Rate** logic options, or assign the **Termination** logic option to the pins with a series setting. You cannot use drive strength and slew rate settings when a pin is assigned an OCT setting.

- ② For more information about creating logic option assignments with the Pin Planner, refer to *Assigning Pins* in Quartus II Help.

During I/O assignment analysis the Fitter automatically assigns suggested pin locations to unassigned pins in your design, based on your design constraints, so it can perform pin legality checks. For example, if you assign an edge location to a group of LVDS pins, the Fitter assigns pin locations for each LVDS pin in the specified edge location and then performs legality checks. To display the Fitter-placed pins use the Show Fitter Placements feature in the Pin Planner. To accept these suggested pin locations, you must back-annotate your pin assignments.

- ② For more information about the Show Fitter Placements feature, refer to the *Show Commands* in Quartus II Help. For more information about back-annotating assignments, refer to *Back-Annotating Assignments for A Project* in Quartus II Help.

The following design flows show two different circumstances in which you can use I/O assignment analysis:

- Use the flow shown in [Figure 5-13](#) if you must complete board layout before starting the design. This flow does not require design files and checks the legality of your pin assignments.
- Use the flow shown in [Figure 5-14](#) if your design is complete. This flow thoroughly checks the legality of your pin assignments against any design files provided.

Each flow involves creating pin assignments, running analysis, and reviewing the report file.

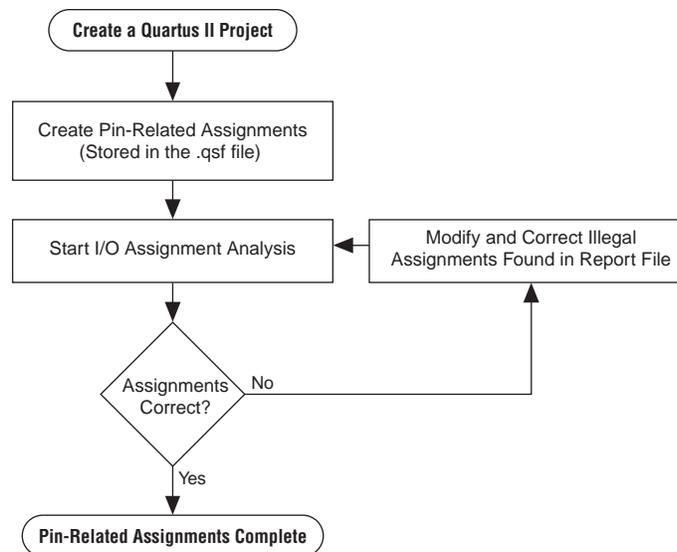
Running I/O Assignment Analysis without Design Files

During the early stages of developing a device, board layout engineers may request preliminary or final pin-outs. It is time consuming to manually check whether the pin-outs violate any design rules. To quickly perform basic checks on the legality of your pin assignments, you can use the Quartus II software to perform I/O assignment analysis.

 If you create pin-related assignments in Mentor Graphics I/O Designer software, you can import an .fx into the Quartus II software.

Running I/O assignment analysis performs limited checks on pin assignments made in a design in which you specified a device, but does not yet include any HDL design files. For example, you can create a Quartus II project with only a target device specified and create pin-related assignments based on circuit board layout considerations that are already determined. Even though the Quartus II project does not yet contain any design files, you can reserve input and output pins and create pin-related assignments for each pin with the Pin Planner. After you assign an I/O standard to each reserved pin, run I/O assignment analysis to ensure that there are no I/O standard conflicts in each I/O bank. [Figure 5-13](#) shows the work flow for assigning and analyzing pin-outs without design files.

Figure 5-13. Assigning and Analyzing Pin-Outs without Design Files



When you make and analyze pin-related assignments without design files, make sure you reserve the pins you intend to use as I/O pins, so the Fitter can determine each pin type. After performing I/O assignment analysis, correct any errors reported by the Fitter and rerun I/O assignment analysis until all errors are corrected.

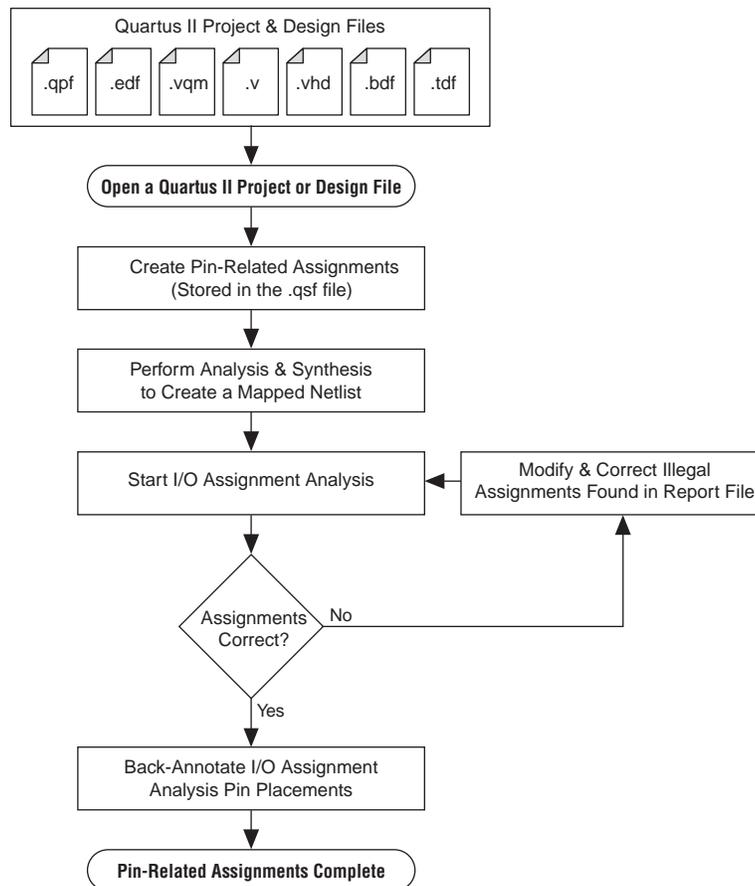
 For more information about reserving pins with the Pin Planner, refer to [Assigning Pins](#) in Quartus II Help.

 Without a complete design, running I/O assignment analysis performs limited checks and cannot guarantee that your assignments do not violate design rules.

Running I/O Assignment Analysis with Design Files

If you have preliminary or complete design files, you can run I/O assignment analysis to help you determine if your design will fit. The rules checked during I/O assignment analysis depend on the completeness of the design. If you have a complete design, the legality of all pin-related assignments are thoroughly checked during I/O assignment analysis. If you have a partial design, which can be just the top-level wrapper file, the legality of those pin-related assignments for which there is enough information are checked during I/O assignment analysis. Figure 5-14 shows the work flow for assigning and analyzing pin-outs without design files.

Figure 5-14. Assigning and Analyzing Pin-Outs with Design Files

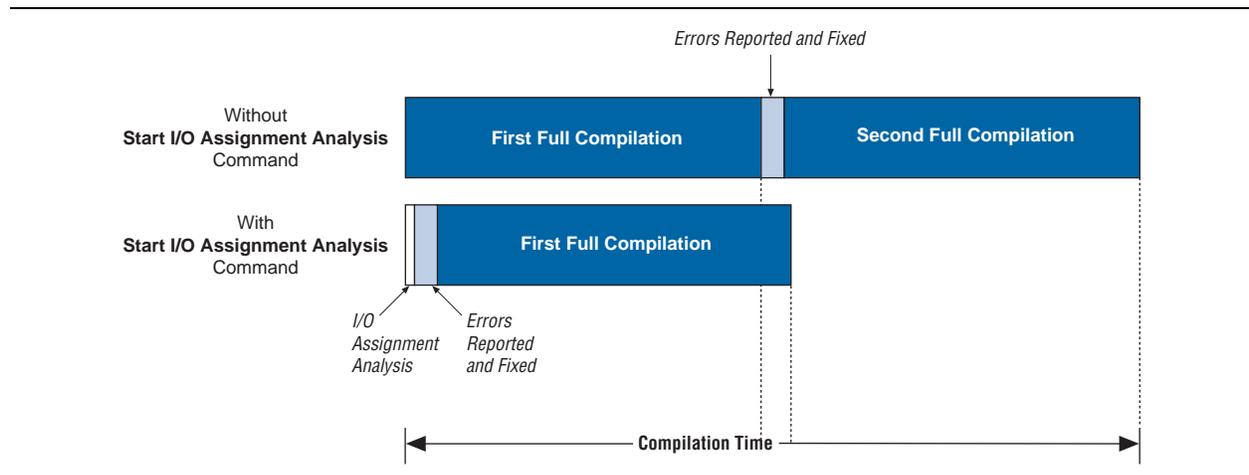


If you run I/O assignment analysis on incomplete design files, you may still encounter errors during full compilation. For example, you may assign a clock to a user I/O pin instead of assigning it to a dedicated clock pin, or design the clock to drive a PLL that you have not yet instantiated in the design. The checks run during I/O assignment analysis do not account for the logic that the pin drives, and do not check that only a dedicated clock input pin can drive the clock port of a PLL. To obtain better coverage, analyze as much of the design as possible, especially logic that

connects to pins. For example, if your design includes PLLs or LVDS blocks, include these MegaWizard Plug-In Manager-generated files. To assign and analyze pin-related assignments successfully, after performing I/O assignment analysis, correct any errors reported by the Fitter and rerun I/O assignment analysis until all errors are corrected.

Figure 5-15 shows the compilation time benefit of performing I/O assignment analysis before running a full compilation.

Figure 5-15. Saving Compilation Time with the I/O Assignment Analysis



Optimizing I/O Assignment Analysis with Output Enable Group Logic Option Assignments

Each device has a certain number of VREF pins, and each VREF pin supports a certain number of I/O pins. A VREF pin and its supported I/O pins are called a VREF bank. The VREF pins are used only for inputs with VREF I/O standards, such as HSTL- and SSTL-type I/O standards. VREF outputs do not require the VREF pin. When a voltage-referenced input is present in a VREF bank, only a certain number of outputs can be present in that VREF bank. For example, for devices in the Stratix II flip chip package, only 20 outputs can be present in a VREF bank when a VREF I/O standard input is present in that bank.

 For more information about device VREF pins and their associated I/O pins, refer to the [Pin-Out Files for Altera Devices](#) page of the Altera website.

For interfaces that use bidirectional VREF I/O pins, your design must meet the output restriction for each I/O bank when the pins are driving in either direction. If a set of bidirectional signals are controlled by different output enables, they are treated as independent output enables during I/O assignment analysis, thus creating a situation in which the Fitter may determine that your design violates VREF restrictions. To treat the set of bidirectional signals as a single output enable group so that the Fitter does not determine that the design violates the requirements for the maximum number of pins driving out of a VREF group, assign the **Output Enable Group** logic option assignment to the bidirectional signals. Assign an integer value for the **Output Enable Group** logic option assignment and assign the same integer value to all sets of signals that are driving in the same direction. Assigning this logic option to groups of signals is important in the case of external memory interfaces.

For example, in the case of a DDR2 interface in a Stratix II device, the device can have 30 pins in a VREF group. Each byte lane for a $\times 8$ DDR2 interface has one DQS pin and eight DQ pins, for a total of nine pins per byte lane. The DDR2 interface uses **SSTL-18 Class I** as its I/O standard, which is a VREF I/O standard. In typical interfaces, each byte lane has its own output enable. In this example, the DDR2 interface has four byte lanes. Using 30 I/O pins in a VREF group, there are three byte lanes and an extra byte lane that supports the three remaining pins. If you do not use the **Output Enable Group** logic option assignment, the Fitter analyzes each byte lane as an independent group driven by a unique output enable during I/O assignment analysis. With this arrangement, the worst-case scenario is when the three pins are inputs, and the other 27 pins are outputs. In this case, the 27 output pins violate the 20 output pin limit.

In a DDR2 interface, all DQS and DQ pins are always driven in the same direction. Therefore, the Fitter reports an error that is not applicable to your design. Assigning the **Output Enable Group** logic option assignment to the DQS and DQ pins forces the Fitter to check these pins as a group driven by a common output enable during I/O assignment analysis. When you use the **Output Enable Group** logic option assignment, the DQS and DQ pins are checked as all input pins or all output pins and are not in violation of the I/O rules.

You can also use the **Output Enable Group** logic option assignment with pins that are driven only at certain times. For example, the data mask signal in DDR2 interfaces is an output signal, but it is driven only when the DDR2 is writing (bidirectional signals are outputs). To avoid errors during I/O assignment analysis, use the **Output Enable Group** logic option assignment to assign the data mask to the same value as the DQ and DQS signals.

You can also assign the **Output Enable Group** logic option to VREF input pins. If the VREF input pins are not active during the time the outputs are driving, add the VREF input pins to the output enable group, thus removing the VREF input pins from the VREF analysis. For example, the QVLD signal for an RLDRAM II interface is active only during a read. During a write, the QVLD pin is not active and does not count as an active VREF input pin in the VREF group. You can place the QVLD pins in the same output enable group as the RLDRAM II data pins.

Understanding the I/O Assignment Analysis Report

When I/O assignment analysis is complete, you can view detailed analysis reports and a **.pin**. The detailed messages in the reports help you quickly understand and resolve pin assignment errors. Each message includes a related node name and a description of the problem.

The Fitter section of the Compilation report contains information generated during I/O assignment analysis, including the following reports:

- I/O Assignment Warnings
- Resource Section
- I/O Rules Section

The I/O Assignment Warnings report provides a list of pins and the Fitter warnings generated for the pins during I/O assignment analysis

The Resource Section contains reports that categorize the pins as input pins, output pins, and bidirectional pins. You can view the utilization of each I/O bank in your device in the I/O Bank Usage report.

The I/O Rules Section includes detailed information about the I/O rules tested during I/O assignment analysis and contains the following reports:

- I/O Rules Summary report
- I/O Rules Details report
- I/O Rules Matrix report

The I/O Rules Summary report provides a quick summary of the number of I/O rules tested and how many applicable rules passed, failed, or were not checked due to other failing rules.

The I/O Rules Details report provides detailed information on all I/O rules. The **Status** column indicates whether applicable rules passed, failed, or could not be checked. All rules are given a level of severity rating to indicate their level of importance for an effective analysis.

The I/O Rules Matrix report (refer to [Figure 5-16](#)) provides a list of the I/O rules checked by the Fitter for each pin in the design. Rules that apply to the target device family either pass or fail for each pin. Rules marked **Inapplicable** are rules that do not apply to the target device family. You can ignore any rule marked **Inapplicable**.

Figure 5-16. I/O Rules Matrix

Pin/Rules	I/O_000001	I/O_000002	I/O_000003	I/O_000004	I/O_000005	I/O_000006	I/O_000007	I/O_000008	I/O_000009	I/O_000010	I/O_000011
1 Total Pass	21	0	21	0	0	21	21	0	21	21	20
2 Total Unchecked	1	0	1	0	0	1	1	0	1	1	1
3 Total Inapplicable	0	22	0	22	22	0	0	22	0	0	0
4 Total Fail	0	0	0	0	0	0	0	0	0	0	1
5 yvalid	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
6 follow	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Fail
7 yn_out[7]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
8 yn_out[6]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
9 yn_out[5]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
10 yn_out[4]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
11 yn_out[3]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
12 yn_out[2]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
13 yn_out[1]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
14 yn_out[0]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
15 clk	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
16 reset	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
17 clkx2	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
18 newt	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
19 d[7]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
20 d[6]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
21 d[5]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
22 d[4]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
23 d[3]	Unchecked	Inapplicable	Unchecked	Inapplicable	Inapplicable	Unchecked	Unchecked	Inapplicable	Unchecked	Unchecked	Unchecked
24 d[2]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
25 d[1]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass
26 d[0]	Pass	Inapplicable	Pass	Inapplicable	Inapplicable	Pass	Pass	Inapplicable	Pass	Pass	Pass

Validating Pin Assignments with Full Compilation

After performing preliminary pin assignment validation with the live I/O check feature and running I/O assignment analysis, you must perform a final I/O timing check of your design by performing a full compilation. To avoid costly board respins, you must include complete design files and constraints. With timing information, the Fitter makes intelligent placement and routing to achieve optimal timing performance in your design. Use the TimeQuest Timing Analyzer to create timing constraints for input, output, and bidirectional pins.

 For more information about the TimeQuest analyzer, refer to the *The Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

Performing I/O Timing Analysis

Timing analysis is usually run during a full compilation of your design or when you perform an early timing estimate. You can also run timing analysis independently after you fully compile the design. For example, if you change the slew rates or drive strengths of some I/O pins with ECOs, you do not have to recompile the entire design, but only run timing analysis to verify the timing of your design.

As part of I/O planning, especially with high-speed designs, take board-level signal integrity and timing into account. When adding a device with high-speed interfaces to a board design, the quality of the signal at the far end of the board route, as well as the propagation delay, is vital for proper system operation

As part of I/O planning, you must understand the I/O timing results reported by the Quartus II software after performing timing analysis on your design. If all your design files are complete and you have fully compiled your design, all the timing checks related to I/O timing are performed during timing analysis. Static timing analysis is performed when you compile your design in the Quartus II software. You must understand I/O timing and what factors affect I/O timing paths in your design. One important factor in I/O timing results is how accurately you specify the output loads of the output and bidirectional pins in your design. Incomplete I/O constraints can affect your I/O timing results.

The Quartus II software supports three different methods of I/O timing analysis:

- Advanced I/O timing using a user-defined board trace model to produce enhanced timing reports from accurate, “board-aware”, simulation models

Advanced I/O timing allows you to configure a complete board trace model for each I/O standard or pin used in your design. With advanced I/O timing, the TimeQuest analyzer uses the results of simulations of the I/O buffer, package, and board trace model to generate more accurate I/O delays and extra reports to give insight into signal behavior at the system level. You can use these advanced timing reports as a guide to make changes to your I/O assignments and board design to improve timing and signal integrity.

 For more information about advanced I/O timing, including device support, refer to *About Advanced I/O Timing* in Quartus II Help.

- I/O timing using a default or user-specified capacitive load without signal integrity analysis

The TimeQuest analyzer creates timing reports that measure t_{CO} to an I/O pin using a default or user-specified value for a capacitive load.

- Full board routing simulation in third-party tools using Altera-provided or Quartus II software-generated IBIS or HSPICE I/O models

The IBIS and HSPICE Writers the simulation model files for use by third-party board simulation tools. The IBIS and HSPICE Writers in the Quartus II software can export accurate simulation models for use in applications such as Mentor Graphics HyperLynx and Synopsys HSPICE.



For devices that support advanced I/O timing, it is the default method of I/O timing analysis. For all other devices, you must use a default or user-specified capacitive load assignment to determine t_{CO} and power measurements.

-  For more information about advanced I/O timing support, refer to the appropriate device handbook available on the [Literature and Technical Documentation](#) page of the Altera website. For more information about board-level signal integrity and tips on how to improve signal integrity in your high-speed designs, refer to the [Altera Signal Integrity Center](#) page of the Altera website.
-  For information about creating IBIS and HSPICE models with the Quartus II software and integrating those models into HyperLynx and HSPICE simulations, refer to the [Signal Integrity Analysis with Third-Party Tools](#) chapter in volume 2 of the *Quartus II Handbook*.

Enabling and Configuring Advanced I/O Timing

With the advanced I/O timing feature, you can expand upon the basic timing and power measurements made with the capacitive loading settings. The advanced I/O timing feature gives you the ability to fully define not only the capacitive load, but also any termination components and trace impedances in the board routing for any output pin or bidirectional pin in output mode. You can configure an overall board trace model for each I/O standard as well as customize the model for specific pins.

When you use the advanced I/O timing feature, the board trace model replaces any capacitive load setting you made because the load is included in the model. For timing measurements, the entire board trace model is taken into account when calculating I/O delays. For power measurements, an effective capacitive load is used based on the sum of the capacitive elements in the model, including the **Near capacitance**, **Far capacitance**, and **Transmission line distributed capacitance** elements of the model.

Defining Overall Board Trace Models

You can define an overall board trace model for each I/O standard in your design that is the default model for all pins that use a particular I/O standard. After configuring the overall board trace model, you can customize the model for specific pins using the Board Trace Model window in the Pin Planner.

-  Custom component value changes you make to selected pins in the Pin Planner take priority and are not affected by subsequent changes to a specific components for the entire design. Similarly, any changes you make to specific pins do not affect the component settings for the entire design.

When you define an overall board trace model you can specify the board trace, termination, and capacitive load parameters for each I/O standard. The default settings for components in the model for each I/O standard are device-specific and match the default test model used for calculating delay without advanced I/O timing. For differential I/O standards, the component values you set are used for both the positive and negative signals of a differential pin pair.

All the assignments for board trace models you specify are saved to the `.qsf`. You can also use Tcl commands to create board trace model assignments. [Example 5-8](#) shows Tcl commands for specifying board trace model assignments.

Example 5-8. Specifying Board Trace Models

```
## setting the near end series resistance model of sel_p output pin to 25 ohms
set_instance_assignment -name BOARD_MODEL_NEAR_SERIES_R 25 -to sel_p
## Setting the far end capacitance model for sel_p output signal to 6 picofarads
set_instance_assignment -name BOARD_MODEL_FAR_C 6P -to sel_p
```

 For more information about defining a board trace model for your entire design, refer to [Using Advanced I/O Timing](#) in Quartus II Help. For more information about configuring specific pins in your board trace model, refer to [Creating Board Trace Model Assignments in the Pin Planner](#) in Quartus II Help. For more information about configuring component values for a board trace model, including a complete list of the supported unit prefixes and setting the values with Tcl scripts, refer to [Board Trace Model](#) Quartus II Help.

 For more information about the default models used for measuring I/O delay, refer to the appropriate device handbook available on the [Literature and Technical Documentation](#) page of the Altera website.

Customizing the Board Trace Model in the Pin Planner

In the Pin Planner, you can view a graphical representation of the board trace model you configured with the Board Trace Model window. Initially, the settings you create for the overall board trace model match the settings in the Pin Planner. For differential signals, the Board Trace Model window displays the routing and components for both the positive and negative signals of the differential pair. Any changes you make for a differential signal pair must be performed on the positive signal of the pair. The settings must match between the positive and negative signals of a differential pair, so the changes are automatically reflected in the settings for the negative signal.

When editing board trace model assignments, for numerical values, use standard unit prefixes such as *p*, *n*, and *k* to represent pico, nano, and kilo, respectively. To short a series component or have an open circuit for a parallel component, select **short** or **open**, respectively, for the component value.

Configuring Board Trace Models

The Quartus II software provides board trace model templates for various I/O standards in which you can fill in various parameters. [Figure 5-17](#) shows the template for a 2.5-V I/O standard. This model consists of near-end and far-end board component parameters.

Modeling of the near-end of the board trace includes the elements which are close to the device and modeling of the far-end includes the elements which are at the receiver end of the link, closer to the receiving device. The topology represented in the Quartus II board trace model is conceptual and does not necessarily match the board trace component for component. For example near-end model parameters can

represent device-end discrete termination and breakout traces. Far-end modeling can represent the bulk of the board trace to discrete external memory components, and the far end termination network. The same circuit can be analyzed with near-end modeling of the entire board, including memory component termination, and far-end modeling of the actual memory component.

Figure 5-17. Board Trace Model

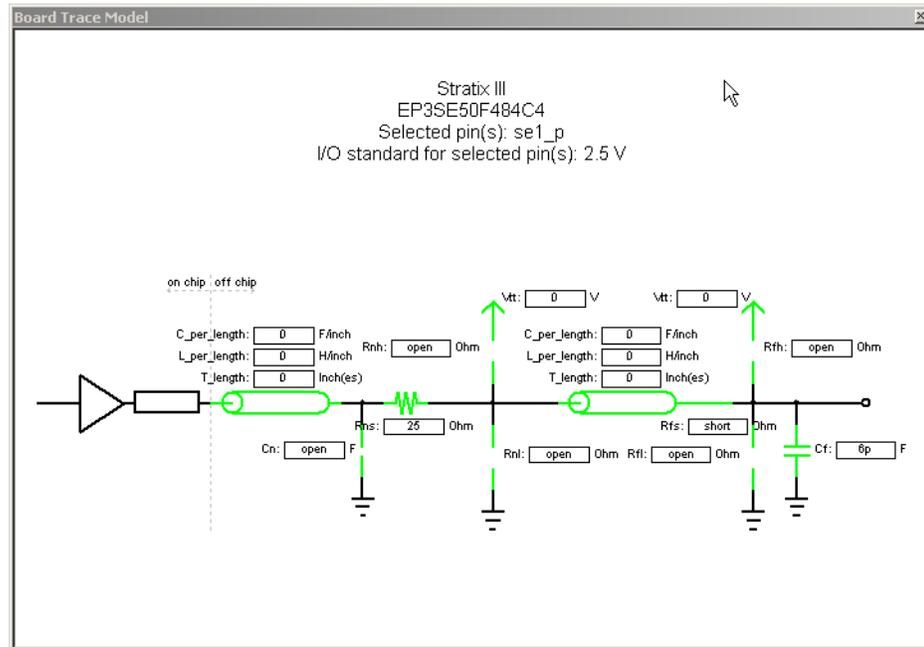
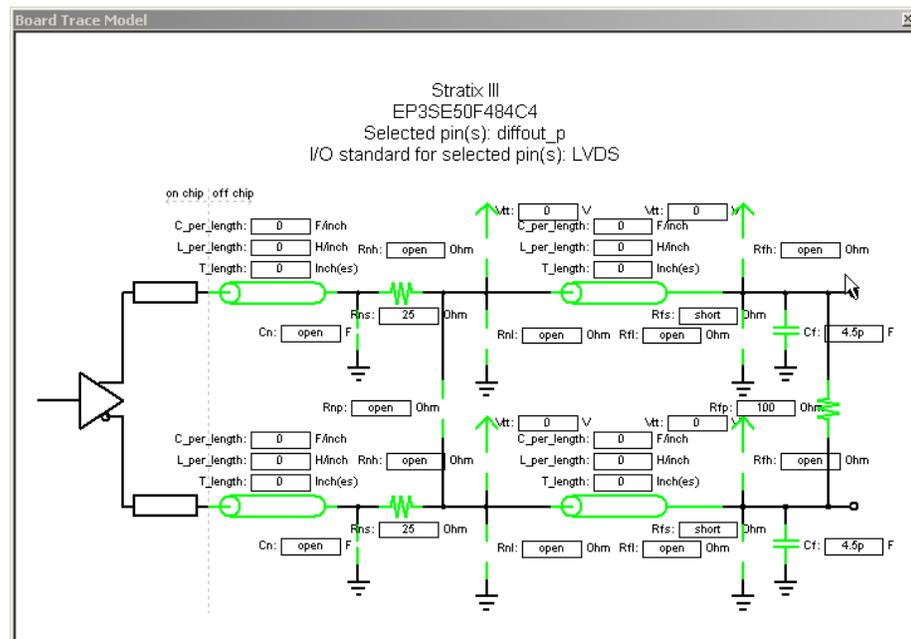


Figure 5-18 shows the template for the LVDS I/O standard. The far-end capacitance (C_f) represents the external-device or multiple-device capacitive load. If you have multiple devices on the far-end, you must find the equivalent capacitance at the far-end, taking into account all receiver capacitances. The far-end capacitance can be the sum of all the receiver capacitances.

For more information about the specifications for external device capacitance values, refer to the appropriate device handbook available on the [Literature and Technical Documentation](#) page of the Altera website.

The Quartus II software models lossless transmission lines, and does not require a transmission-line resistance value. Only distributed inductance (L) and capacitance (C) values are needed. The distributed L and C values of transmission lines must be entered on a per-inch basis, and can be obtained from the PCB vendor or manufacturer, the CAD Design tool, or a signal integrity tool such as the Mentor Graphics Hyperlynx software.

Figure 5–18. Differential Board Trace Model

Specifying Near-End vs Far-End Timing Analysis

With advanced I/O timing analysis, you have the option of selecting a near-end or far-end point for your I/O timing. With near-end timing, the timing is analyzed to the device pin. [Figure 5–17](#) shows near-end timing ending at the vertical dashed line separating the device I/O pin and off-chip components.

By default, advanced I/O timing analysis analyzes output I/O timing to the device pin. When you choose a near-end endpoint, you can use the `set_output_delay` SDC timing constraint to account for the delay across the board. However, when you choose a far-end I/O timing endpoint, then advanced I/O timing analysis analyzes timing to the external device input, at the far end of the board trace. Whether you choose a near-end or far-end timing endpoint, the board trace models are taken into account during timing analysis.

- ❓ For more information about calculating I/O timing to the near-end or far-end of the board trace, refer to [Using Advanced I/O Timing](#) in Quartus II Help.

Understanding Advanced I/O Timing Analysis Reports

When you perform advanced I/O analysis, the TimeQuest analyzer creates reports that signal integrity reports that provide board delay estimates and signal integrity data.

The TimeQuest analyzer section of the Compilation report contains information generated during advanced I/O timing analysis, including the following reports:

- Board Trace Model Assignments
- Signal Integrity Metrics

The Board Trace Model Assignments report summarizes the board trace model component settings for each output and bidirectional signal.

The Signal Integrity Metrics report contains all the signal integrity metrics calculated during advanced I/O timing analysis based on the board trace model settings for each output or bidirectional pin. The reports contain many metrics, including measurements at both the FPGA pin and at the far-end load of board delay, steady state voltages, and rise and fall times.

-  By default, the TimeQuest analyzer generates the Slow-Corner Signal Integrity Metrics report. To generate a Fast-Corner Signal Integrity Metrics report you must change the delay model.
-  For more information about the reports generated during advanced I/O timing analysis, refer to *Advanced I/O Timing Reports* in Quartus II Help. For more information about the metrics calculated during advanced I/O timing analysis, including diagrams illustrating the metrics on output waveforms, refer to *Signal Integrity Metrics* in Quartus II Help. For more information about changing the delay model, refer to the *Create Timing Netlist Dialog Box* in Quartus II Help.
-  For information about the configuration and use of the TimeQuest analyzer, refer to *The Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

Adjusting I/O Timing and Power with Capacitive Loading

When calculating t_{CO} and power for output and bidirectional pins, the TimeQuest analyzer and the PowerPlay Power Analyzer use a bulk capacitive load. You can adjust the value of the capacitive load per I/O standard to obtain t_{CO} and power measurements that more accurately reflect the behavior of the output or bidirectional net on your PCB. Input pins ignore any capacitive load settings. You can adjust the capacitive load settings per I/O standard, in picofarads (pF), for your entire design. During compilation, the Compiler measures power and t_{CO} measurements based on your settings. You can also adjust the capacitive load on an individual pin with the **Output Pin Load** logic option.

-  For more information about adjusting the value of the capacitive load for your entire design, refer to *Setting Up and Running a Compilation* in Quartus II Help. For more information about adjusting the value of the capacitive load on an individual pin, refer to *Assigning Pins* in Quartus II Help.

Incorporating PCB Design Tools

Signal and pin assignments are initially made by the chip designer and it is up to the board designer to correctly transfer these assignments to the symbols in their system circuit schematics and board layout. As the board design progresses, pin reassignments may be requested or required to optimize the layout. These reassignments must in turn be communicated to the chip designer, so that the new assignments can be validated during I/O assignment analysis and processed through an updated placement-and-routing of the device.

The Quartus II software interacts with board layout tools by importing and exporting pin information files, including the `.qsf`, `.pin`, and `.fx` files.

 For more information about incorporating PCB design tools, refer to the *Cadence PCB Design Tools Support* and *Mentor Graphics PCB Design Tools Support* chapters in volume 2 of the *Quartus II Handbook*.

Scripting Support

A Tcl script allows you to run procedures and determine settings described in this chapter. You can also run some of these procedures at a command prompt.

For detailed information about specific scripting command options and Tcl API packages, type the following command at a system command prompt to run the Quartus II Command-Line and Tcl API Help browser:

```
quartus_sh --qhelp ←
```

 For more information about Quartus II scripting support, including examples, refer to the *Tcl Scripting* and *Command-Line Scripting* chapters in volume 2 of the *Quartus II Handbook*.

Running I/O Assignment Analysis

You can run I/O assignment analysis with a Tcl command or with a command-line command.

Enter the following in the Tcl console or a Tcl script:

```
execute_flow -check_ios ←
```

Type the following at a system command prompt:

```
quartus_fit <project name> --check_ios ←
```

For more information about running I/O assignment analysis, refer to “[Understanding the I/O Assignment Analysis Report](#)” on page 5-30.

Generating a Mapped Netlist

You can generate a mapped netlist with a Tcl command or with a command-line command.

Enter the following in the Tcl console or in a Tcl script:

```
execute_module -tool map ←
```

The `execute_module` command is in the flow package.

Type the following at a system command prompt:

```
quartus_map <project name> ←
```

Reserving Pins

You can reserve pins with a Tcl command.

Use the following Tcl command to reserve a pin:

```
set_instance_assignment -name RESERVE_PIN <value> -to <signal name>
```

Use one of the following valid reserved pin values:

- "AS BIDIRECTIONAL"
- "AS INPUT TRI-STATE"
- "AS OUTPUT DRIVING AN UNSPECIFIED SIGNAL"
- "AS OUTPUT DRIVING GROUND"
- "AS SIGNALPROBE OUTPUT"



Ensure you include the quotation marks when specifying the reserved pin value.

Creating Location Assignments

You can create location assignments with a Tcl command.

Use the following Tcl command to assign a signal to a pin or device location:

```
set_location_assignment <location> -to <signal name> ←
```

Valid locations are pin locations, I/O bank locations, or edge locations. Pin locations include pin names, such as PIN_A3. I/O bank locations include IOBANK_1 up to IOBANK_*n*, in which *n* is the number of I/O banks in the device.

Use one of the following valid edge location values:

- EDGE_BOTTOM
- EDGE_LEFT
- EDGE_TOP
- EDGE_RIGHT

For more information about location assignments, refer to [“Creating Location Assignments”](#) on page 5-17.



For more information about I/O banks in your device, refer to the appropriate device handbook available on the [Literature and Technical Documentation](#) page of the Altera website.

Creating Exclusive I/O Group Assignments

You can create exclusive I/O group assignments with a Tcl command.

Use the following Tcl command to create an exclusive I/O group assignments:

```
set_instance_assignment -name "EXCLUSIVE_IO_GROUP" -to pin
```

For more information about exclusive I/O group assignments, refer to [“Creating Exclusive I/O Group Assignments”](#) on page 5-18.

Changing the Slew Rate and Drive Strength

You can create slew rate and drive strength assignments with a Tcl command.

Use the following Tcl commands to create an slew rate and drive strength assignments:

```
set_instance_assignment -name CURRENT_STRENGTH_NEW 8MA -to e[0]
set_instance_assignment -name SLEW_RATE 2 -to e[0]
```

For more information about slew rate and drive strength assignments, refer to [“Creating Pin Assignments with the Chip Planner”](#) on page 5-20.

Conclusion

The Quartus II software provides many tools and features to help you with I/O planning, including the ability to validate pin assignments in all design stages, even before the development of your design. The ability to import and export assignments between the Quartus II software and other PCB tools allows you to make iterative changes efficiently. Finally, the ability to enter a board trace model and create advanced timing reports based on how I/O signals are routed on a board truly makes the Quartus II software board-aware.

Document Revision History

Table 5-5 shows the revision history for this chapter.

Table 5-5. Document Revision History

Date	Version	Changes
December 2010	10.0.1	Template update
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Reorganized and edited the chapter ■ Added links to Quartus II Help for procedural information previously included in the chapter ■ Added information on rules marked Inapplicable in the I/O Rules Matrix Report ■ Added information on assigning slew rate and drive strength settings to pins to fix I/O assignment warnings
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Reorganized entire chapter to include links to Quartus II help for procedural information previously included in the chapter ■ Added documentation on near-end and far-end advanced I/O timing
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Updated “Pad View Window” on page 5-20 ■ Added new figures: <ul style="list-style-type: none"> ■ Figure 5-15 ■ Figure 5-16 ■ Added new section “Viewing Simultaneous Switching Noise (SSN) Results” on page 5-17 ■ Added new section “Creating Exclusive I/O Group Assignments” on page 5-18

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

FPGA design has evolved from small programmable circuits to designs that compete with multimillion-gate ASICs. At the same time, the I/O counts on FPGAs and logic density requirements of designs have increased exponentially. The higher-speed interfaces in FPGAs, including high-speed serial interfaces and memory interfaces, require careful interface design on the PCB. Designers must address the timing and signal integrity requirements of these interfaces early in the design cycle. Simultaneous switching noise (SSN) often leads to the degradation of signal integrity by causing signal distortion, thereby reducing the noise margin of a system.

Today's complex FPGA system design is incomplete without addressing the integrity of signals coming in to and out of the FPGA. Altera recommends that you perform SSN analysis early in your FPGA design and prior to the layout of your PCB with complete SSN analysis of your FPGA in the Quartus® II software. This chapter describes the Quartus II SSN Analyzer tool and covers the following topics:

- “Definitions”
- “Understanding SSN” on page 6–2
- “SSN Estimation Tools” on page 6–5
- “SSN Analysis Overview” on page 6–5
- “Optimizing Your Design for SSN Analysis” on page 6–8
- “Performing SSN Analysis and Viewing Results” on page 6–15
- “Decreasing Processing Time for SSN Analysis” on page 6–17

Definitions

The terminology used in this chapter includes the following terms:

Aggressor: An output or bidirectional signal that contributes to the noise for a victim I/O pin

PDN: Power distribution network

QH: Quiet high signal level on a pin

QHN: Quiet high noise on a pin, measured in volts

QL: Quiet low signal level on a pin

QLN: Quiet low noise on a pin, measured in volts

SI: Signal integrity (a superset of SSN, covering all noise sources)

SSN: Simultaneous switching noise

SSO: Simultaneous switching output (which are either the output or bidirectional pins)

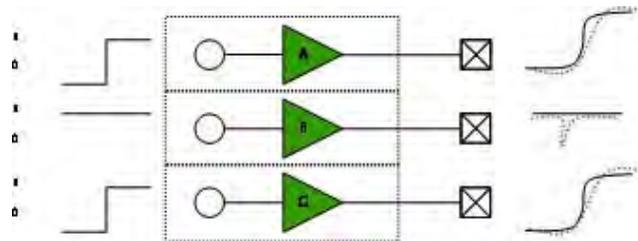
Victim: An input, output, or bidirectional pin that is analyzed during SSN analysis. During SSN analysis, each pin is analyzed as a victim. If a pin is an output or bidirectional pin, the same pin acts as an aggressor signal for other pins.

Understanding SSN

SSN is defined as a noise voltage induced onto a single victim I/O pin on a device due to the switching behavior of other aggressor I/O pins on the device. SSN can be divided into two types of noise: voltage noise and timing noise.

Figure 6-1 shows a system with three pins. Two of the pins (A and C) are switching, while one pin (B) is quiet. If the pins are driven in isolation, the voltage waveforms at the output of the buffers appear without noise interference, as shown by the solid curves at the left of the figure. However, when the pins are switched simultaneously, the noise generated by pins A and C switching is injected onto the other pins, manifesting itself as a voltage noise on pin B and timing noise on pins A and C, as shown by the dotted curves in the figure.

Figure 6-1. System with Three Pins



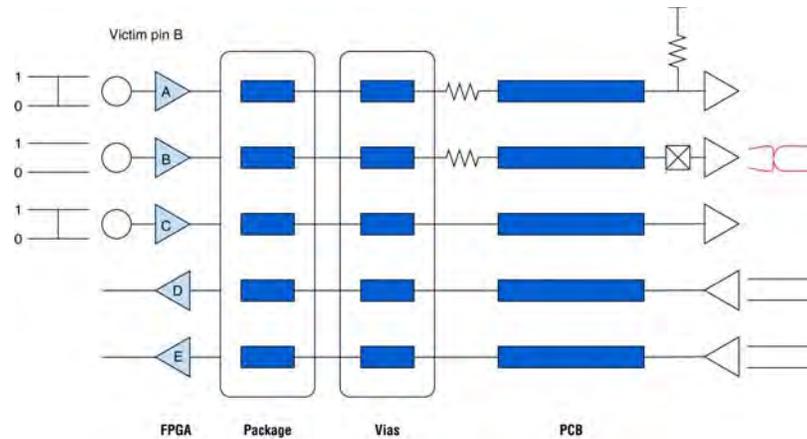
Voltage noise is measured as the worst-case change in voltage of a signal due to SSN. When a signal is QH, it is measured as the change in voltage toward 0 V. When a signal is QL, it is measured as the change in voltage toward V_{CC} .

In the Quartus II software, only voltage noise is analyzed. Voltage noise can be caused by SSOs under two worst-case conditions:

- The victim pin is high and the aggressor pins (SSOs) are switching from low to high
- The victim pin is low and the aggressor pins (SSOs) are switching from high to low

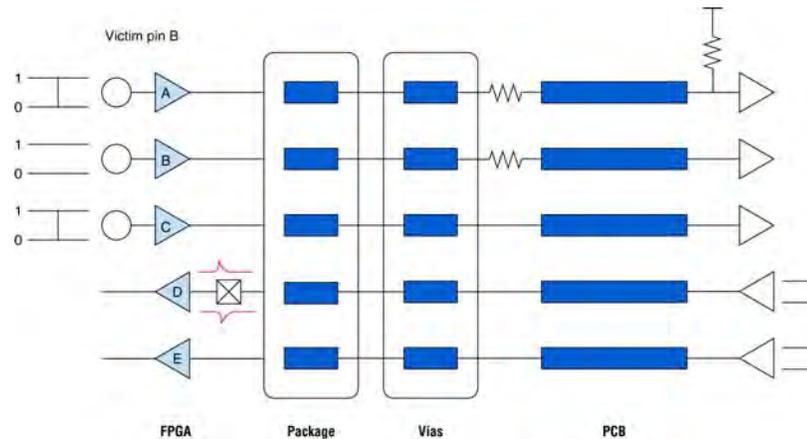
For outputs, the noise is computed at the far-end receiver for pin B (refer to Figure 6-2).

Figure 6-2. Quiet High Output Noise Estimation



For inputs, the noise is computed at the FPGA bumps as shown in for pin D (refer to Figure 6-3).

Figure 6-3. Quiet Low Input Noise Estimation



SSN can occur in any system, but the induced noise does not always result in failures. Voltage functional errors are caused by SSN on quiet victim pins only when the voltage values on the quiet pins change by a large enough voltage that the logic listening to that signal reads a change in the logic value. For QH signals, a voltage functional error occurs when noise events cause the voltage to fall below V_{IH} . Similarly, for QL signals, a voltage functional error occurs when noise events cause the voltage to rise above V_{IL} (refer to Figure 6-4). Because V_{IH} and V_{IL} are different for

different I/O standards, and because signals have different quiet voltage values, the absolute amount of SSN, measured in volts, cannot be used to determine if a voltage failure occurs. Instead, to quantify whether an SSN event will cause a voltage error, the Quartus II software uses the amount of noise as a percent of signal margin when reporting noise margins in SSN analysis (refer to Figure 6-4).

Figure 6-4. Reporting Noise Margins

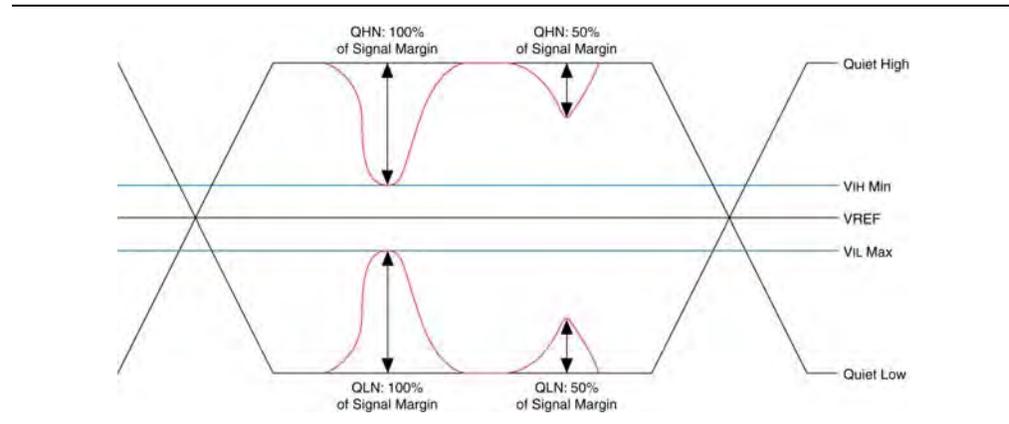
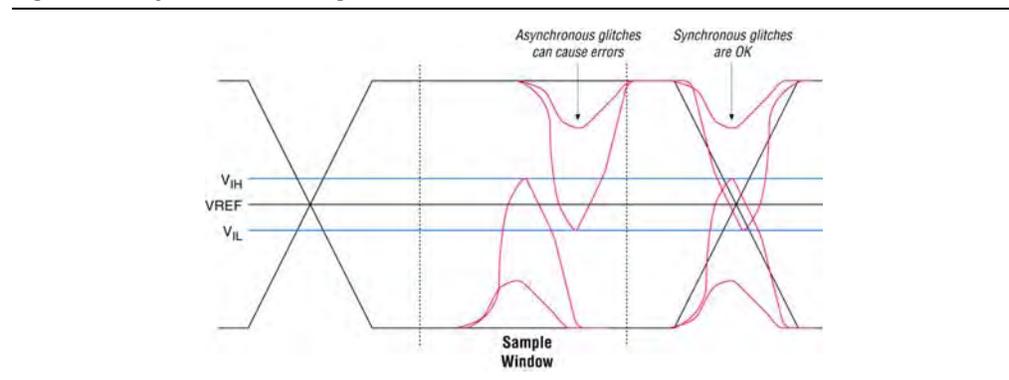


Figure 6-4 shows four noise events, two on QH signals and two on QL signals. The two noise events on the right-side of the figure consume 50 percent of the signal margin and do not cause voltage functional errors. However, the two noise events on the left side of the figure consume 100 percent of the signal margin and can cause a voltage functional error.

Figure 6-5 illustrates a synchronous voltage noise event that does not result in a voltage functional error. Noise or glitches caused by aggressor signals are synchronously related to the victim pin outside of the sampling window of a receiver. The noise or glitches affect the switching time of a victim pin, but are not considered an input threshold violation failure.

Figure 6-5. Synchronous Voltage Noise



For more information about the design factors that affect the noise margins during SSN analysis in the Quartus II software, refer to “[SSN Analysis Overview](#)”.

SSN Estimation Tools

Addressing SSN early in your FPGA design and PCB layout can help you avoid costly board respins and lost time, both of which can impact your time-to-market. Altera provides many tools for SSN analysis and estimation, including the following tools:

- SSN characterization reports
- An early SSN estimation (ESE) tool
- The SSN Analyzer in the Quartus II software

 For more information about the SSN characterization reports and the ESE tool, including device support information, refer to the [Signal Integrity Center](#) page of the Altera website.

 For more information about the devices for which you can run the SSN Analyzer, refer to [About the SSN Analyzer](#) in Quartus II Help.

The ESE tool is useful for preliminary SSN analysis of your FPGA design; for more accurate results, however, you must use the SSN Analyzer in the Quartus II software. [Table 6–1](#) compares some of the differences between the ESE tool and the SSN Analyzer.

Table 6–1. Comparison of ESE Tool and SSN Analyzer Tool

ESE Tool	SSN Analyzer
Is not integrated with the Quartus II software.	Integrated with the Quartus II software, allowing you to perform preliminary SSN analysis while making I/O assignment changes in the Quartus II software.
QL and QH levels are computed assuming a worst-case pattern of I/O placements.	QL and QH levels are computed based on the I/O placements in your design.
No support for entering board information.	Supports board trace models and board layer information, resulting in a more accurate SSN analysis.
No graphical representation.	Integrated with the Quartus II Pin Planner, in which an SSN map shows the QL and QH levels on victim pins.
Good for doing an early SSN estimate. Does not require you to use the Quartus II software.	Requires you to create a Quartus II software project and provide the top-level port information.

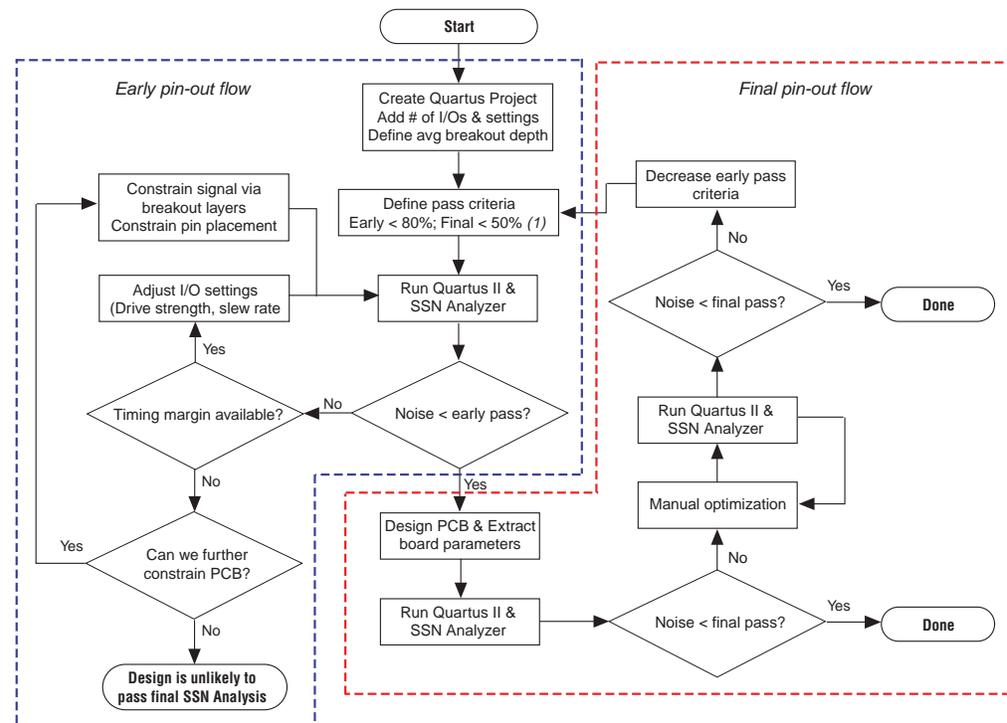
SSN Analysis Overview

You can run the SSN Analyzer at different stages in your design cycle to obtain SSN results. The accuracy of the results depends on the completeness of your design information. Altera recommends that you start SSN analysis early in the design cycle to obtain preliminary results and make adjustments to your I/O assignments, and iterate through the design cycle to finally perform a fully constrained SSN analysis with complete information about your board.

[Figure 6–6](#) shows the flows for both early pin-out and final pin-out SSN analysis. The early pin-out flow assumes conservative design rules initially, and then lets you analyze the design and iteratively apply tighter design rules until SSN analysis indicates your design meets SSN constraints. You must define pass criteria for SSN analysis as a percentage of signal margin in both the early pin-out flow and the final

pin-out flow. The pass criteria you define is specific to your design requirements. For example, a pass criterion you might define is a condition that verifies you have sufficient SSN margins in your design. You may require that the acceptable voltage noise on a pin must be below 70% of the voltage level for that pin. The pass criteria for the early-pin out flow may be higher than the final pin-out flow criteria, so that you do not spend too much time optimizing the on-FPGA portions of your design when the SSN metrics for the design may improve after the design is fully specified.

Figure 6-6. Pin-Out Analysis (Note 1)



Note to Figure 6-6:

(1) Pass criteria determined by customer requirements.

Performing Early Pin-Out SSN Analysis

In the early stages of your design cycle, before you create pin location for your design, use the early pin-out flow (refer to Figure 6-6) to obtain preliminary SSN analysis results. In order to obtain useful SSN results, you must define the top-level ports of your design, but your design files do not have to be complete.

Performing Early Pin-Out SSN Analysis with the ESE Tool

If you know the I/O standards and signaling standards for your design, you can use the ESE tool to perform an initial SSN evaluation.



For more information about the ESE tool, refer to the [Signal Integrity Center](#) page of the Altera website.

Performing Early Pin-Out SSN Analysis with the SSN Analyzer

If you have complete information for the top-level ports of your design, you can use the SSN Analyzer to perform an initial SSN evaluation. Use the following steps to perform early pin-out SSN analysis:

1. Create a project in the Quartus II software.
2. Specify your top-level design information either in schematic form or in HDL code.
3. Perform Analysis and Synthesis.
4. Create I/O assignments, such as I/O standard assignments, for the top-level ports in your design.



Do not create pin location assignments. The Fitter automatically creates optimized pin location assignments.

5. If you do not have completed design files and timing constraints, run I/O assignment analysis.



During I/O assignment analysis, the Fitter places all the unplaced pins on the device, and checks all the I/O placement rules.

6. Run the SSN Analyzer.



For more information about creating and managing projects, refer to the *Managing Quartus II Projects* chapter in volume 2 of the *Quartus II Handbook*. For more about generating a top-level design file in the Quartus II software and I/O assignment analysis, refer to the *I/O Management* chapter in volume 2 of the *Quartus II Handbook*.

In the early stages of your design cycle, you may not have complete board information, such as board trace parameters, layer information, and the signal breakout layers. If you run the SSN Analyzer without this specific information, it uses default board trace models and board layer information for SSN analysis, and as a result the SSN Analyzer confidence level is low. If the noise amounts are larger than the pass criteria for early pin-out SSN analysis, verify whether the SSN noise violations are true failures or false failures. For example, sometimes the SSN Analyzer can determine whether pins are switching synchronously and use that information to filter false positives; however, it may not be able to determine all the synchronous groups. You can improve the SSN analysis results by adjusting your I/O assignments and other design settings. After you optimize your design such that it meets the pass criteria for the early pin-out flow, you can then begin to design your PCB.

For more information, refer to “*Optimizing Your Design for SSN Analysis*”.

Performing Final Pin-Out SSN Analysis

You perform final pin-out SSN analysis after you place all the pins in your design, or the Fitter places them for you, and you have complete information about the board trace models and PCB layers. Even if your design achieves sufficient SSN results during early pin-out SSN analysis, you should run SSN analysis with the complete PCB information to ensure that SSN does not cause failures in your final design. You must specify the board parameters in the Quartus II software, including the PCB layer thicknesses, the signal breakout layers, and the board trace models, before you can run SSN analysis on your final assignments.

For more information, refer to [“Optimizing Your Design for SSN Analysis”](#).

If the SSN analysis results meet the pass criteria for final pin-out SSN analysis, SSN analysis is complete. If the SSN analysis results do not meet the pass criteria, you must further optimize your design by changing the board and design parameters and then rerun the SSN Analyzer. If the design still does not meet the pass criteria, reduce the pass criteria for early pin-out SSN analysis, and restart the process. By reducing the pass criteria for early pin-out SSN analysis, you place a greater emphasis on reducing SSN through I/O settings and I/O placement. Changing the drive strength and slew rate of output and bidirectional pins, as well as adjusting the placement of different SSOs, can affect SSN results. Adjusting I/O settings and placement allows the design to meet the pass criteria for final pin-out SSN analysis after you specify the actual PCB board parameters.

Design Factors Affecting SSN Results

There are many factors that affect the SSN levels in your design. The two main factors are the drive strength and slew rate settings of the output and bidirectional pins in your design.



For more information about the factors that contribute to SSN voltage noise in your FPGA design and managing SSN in your system, refer to [AN 472: Stratix II GX SSN Design Guidelines](#), [AN 508: Cyclone III Simultaneous Switching Noise \(SSN\) Design Guidelines](#), and the [Signal Integrity Center](#) page of the Altera website.

Optimizing Your Design for SSN Analysis

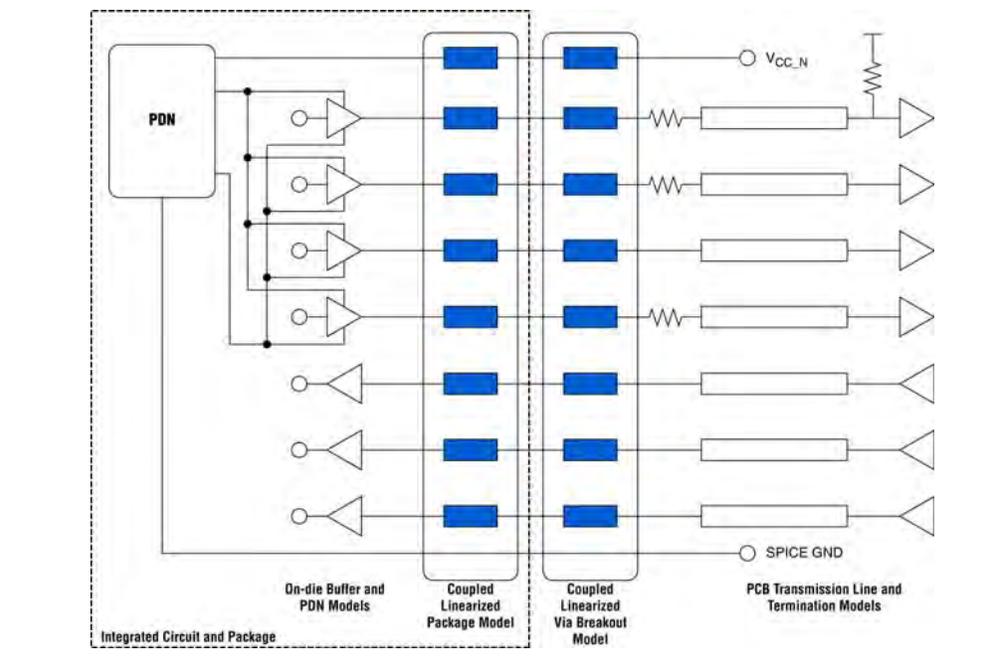
The SSN Analyzer gives you flexibility to precisely define your system to obtain accurate SSN results. The SSN Analyzer produces a voltage noise estimate for each input, output, and bidirectional pin in the design. It allows you to estimate the SSN levels, comprised of QLN and QHN levels, for your FPGA pins. Performing SSN analysis helps you optimize your design for SSN during compilation.

Because the SSN Analyzer is integrated into the Quartus II software, it can automatically set up a system topology that matches your design. The SSN Analyzer accounts for different I/O standards and slew rate settings for each buffer in the design and models different board traces for each signal. Also, it correctly models the state of the unused pins in the design. The SSN Analyzer leverages any custom board trace assignments you set up for use by the advanced I/O timing feature.

The SSN Analyzer also models the package and vias in the design. Models for the different packages that Altera devices support are integrated into the Quartus II software. In the Quartus II software, you can specify different layers on which signals break out, each with its own thickness, and then specify which signal breaks out on which layer.

Figure 6-7 shows the circuit topology the SSN Analyzer automatically constructs. After constructing the circuit topology, the SSN Analyzer uses a simulation-based methodology to determine the SSN for each victim pin in the design.

Figure 6-7. Circuit Topology for SSN Analysis



Optimizing Pin Placements for Signal Integrity

You can take advantage of a built-in SSN optimization feature in the Quartus II software with the **SSN Optimization** logic option.

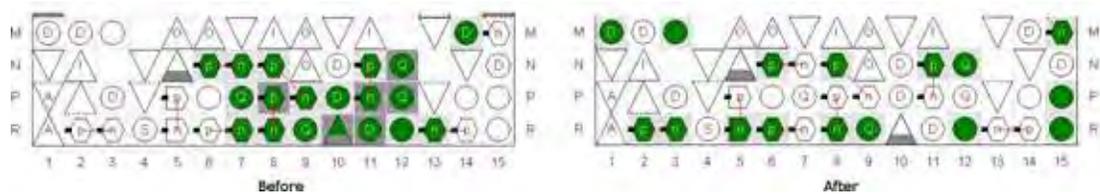
The I/O placements in your design may be affected when you use this option. Setting this option to **Normal compilation** does not affect the f_{MAX} of your design during compilation, however setting this option to **Extra effort** level may impact your design f_{MAX} .



In order to use the **SSN Optimization** logic option, Altera recommends that you do not create location assignments for your pins; instead, let the Fitter place the pins during compilation so that it places the pins to meet the timing performance of your design. To display the Fitter-placed pins use the Show Fitter Placements feature in the Pin Planner. To accept these suggested pin locations, you must back-annotate your pin assignments.

Figure 6-8 shows the results of turning on the **SSN Optimization** logic option for a design. The image on the left shows the placement of the pins without the **SSN Optimization** logic option, and the image on the right shows the adjustments the Fitter made to pin placements to reduce the amount of SSN in the design when the **SSN Optimization** logic option is turned on.

Figure 6-8. SSN Analysis Results Before and After Using the SSN Optimization Logic Option



- ? For more information about creating project-wide logic option assignments, refer to *Setting Up and Running the Fitter* in Quartus II Help. For more information about the Show Fitter Placements feature, refer to *Show Commands* in Quartus II Help. For more information about back-annotating assignments, refer to *Back-Annotating Assignments for A Project* in Quartus II Help.

? For more information about design optimization features, refer to the *Area, Timing, and Compilation Time Optimization* section in volume 2 of the *Quartus II Handbook*.

Specifying Board Trace Model Settings

The SSN Analyzer uses circuit models to determine voltage noise during SSN analysis. The circuit topology (refer to Figure 6-7) is incomplete without board trace information and PCB layer information. You must describe the board trace and PCB layer parameters in your design to accurately compute the SSN in your FPGA device. However, if you do not specify some or all of the board trace parameters and PCB layer information, the SSN Analyzer uses default parameters during SSN analysis. When you use the default parameters, the SSN confidence level is low.

For more information about the default parameters used by the SSN Analyzer and SSN confidence levels, refer to “*Confidence Metric Details Report*” on page 6-16.

The board trace models required for the SSN Analyzer include the board trace termination resistors, pin loads (capacitance), and transmission line parameters. You can define the board circuit models, which are also known as board trace models, in the Quartus II software. The board trace model settings are shared with the models used during advanced I/O timing.

? For more information about defining board trace models and advanced I/O timing, refer to the *I/O Management* chapter in volume 2 of the *Quartus II Handbook*.

You can define an overall board trace model for each I/O standard in your design; this overall board trace model is the default model for all pins that use a particular I/O standard. After configuring the overall board trace model, you can customize the model for specific pins. The parameters you specify for the board trace model are also used in during advanced I/O timing analysis with the TimeQuest Timing Analyzer. If you already specified the board trace models as part of your advanced I/O timing assignments, the same parameters are used during SSN analysis.

- ② For more information about defining a board trace model for your entire design, refer to *Using Advanced I/O Timing* in Quartus II Help. For more information about configuring specific pins in your board trace model, refer to *Creating Board Trace Model Assignments in the Pin Planner* in Quartus II Help. For more information about configuring component values for a board trace model, including a complete list of the supported unit prefixes and setting the values with Tcl scripts, refer to *Board Trace Model* in Quartus II Help.

All the assignments for board trace models you specify are saved to the `.qsf`. You can also use Tcl commands to create board trace model assignments. [Example 6-1](#) shows Tcl commands for specifying transmission line parameters.

Example 6-1. Specifying Board Trace Models

```
set_instance_assignment -name BOARD_MODEL_TLINE_L_PER_LENGTH "3.041E-7" -to e[0]
set_instance_assignment -name BOARD_MODEL_TLINE_LENGTH 0.1391 -to e[0]
set_instance_assignment -name BOARD_MODEL_TLINE_C_PER_LENGTH "1.463E-10" -to e[0]
```

The best way to calculate transmission line parameters is to use a two-dimensional solver to estimate the inductance per inch and capacitance per inch for the transmission line. The termination resistor topology information can be obtained from the PCB schematics. The near-end and far-end pin load (capacitance) values can be obtained from the PCB schematic and other device data sheets. For example, if you know that an FPGA pin is driving a DIMM, you can obtain the far-end loading information in the data sheet for your target device.

-  For more information, refer to the *Device Family Data Sheet* in the appropriate device handbook available on the [Literature and Technical Documentation](#) page of the Altera website.

Defining PCB Layers and PCB Layer Thickness

Every PCB is fabricated using a number of layers. To remove some of the pessimism from your SSN results, Altera recommends that you create assignments describing your PCB layers in the Quartus II software. You can specify the number of layers on you PCB, and their thickness. The PCB layer information is used only during SSN analysis and is not used in other processes run by the Quartus II software. If a custom PCB breakout region is not described you can select the default thickness, which directs the SSN Analyzer to use a single-layer PCB breakout region during SSN analysis.

- ② For more information about specifying PCB layer information, refer to *Running the SSN Analyzer* in Quartus II Help.

All the assignments you create for the PCB layers are saved to the `.qsf`. You can also use Tcl commands to create PCB layer assignments. You can create any number of PCB layers, however, the layers must be consecutive. [Example 6-2](#) shows Tcl commands for specifying PCB layer assignments.

Example 6-2. Specifying PCB Layer Assignments

```
set_global_assignment -name PCB_LAYER_THICKNESS 0.00099822M -section_id 1
set_global_assignment -name PCB_LAYER_THICKNESS 0.00034036M -section_id 2
set_global_assignment -name PCB_LAYER_THICKNESS 0.00034036M -section_id 3
```

[Figure 6-9](#) shows the layout cross-section of a PCB in the Cadence Allegro PCB tool. The cross-section shows the stackup information of a PCB, which tells you the number of layers used in your PCB. The PCB shown in this example consists of various signal and circuit layers on which FPGA pins are routed, as well as the power and ground layers.

Figure 6-9. Snapshot of Stackup of a PCB Shown in the Allegro Board Design Environment

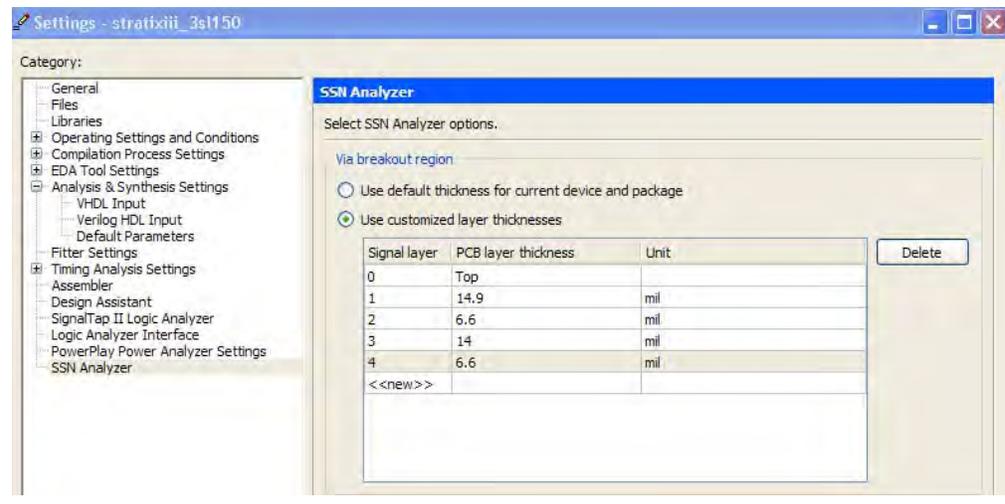
	Subclass Name	Type	Material	Thickness (MIL)
1		SURFACE	AIR	
2	TOP	CONDUCTOR	PLATED_COPPER_FOIL	1.9
3		DIELECTRIC	FR-4	3.6
4	L2-PWR	PLANE	COPPER	1.2
5		DIELECTRIC	FR-4	3
6	L3-GND	PLANE	COPPER	1.2
7		DIELECTRIC	FR-4	4
8	L4-SIGNAL	CONDUCTOR	COPPER	0.6
9		DIELECTRIC	FR-4	6
10	L5-SIGNAL	CONDUCTOR	COPPER	0.6
11		DIELECTRIC	FR-4	4
12	L6-GND	PLANE	COPPER	1.2
13		DIELECTRIC	FR-4	3
14	L7-PWR	PLANE	COPPER	1.2
15		DIELECTRIC	FR-4	4
16	L8-SIGNAL	CONDUCTOR	COPPER	0.6
17		DIELECTRIC	FR-4	6
18	L9-SIGNAL	CONDUCTOR	COPPER	0.6
19		DIELECTRIC	FR-4	4
20	L10-GND	PLANE	COPPER	1.2
21		DIELECTRIC	FR-4	3

In this example, each of the four signal layers are a different thickness, with the depths shown in the **Thickness (MIL)** column. The layer thickness for each signal layer is computed as follows:

- Signal Layer 1 is the L4-SIGNAL, at thickness $(1.9+3.6+1.2+3+1.2+4=)$ 14.9 mils
- Signal Layer 2 is the L5-SIGNAL, at thickness $(0.6+6=)$ 6.6 mils
- Signal Layer 3 is the L8-SIGNAL, at thickness $(0.6+4+1.2+3+1.2+4=)$ 14 mils
- Signal Layer 4 is the L9-SIGNAL, at thickness $(0.6+6=)$ 6.6 mils

Figure 6-10 shows the results in the Quartus II software after you enter these PCB signal layers and thickness assignments.

Figure 6-10. PCB Layers Specified in the Quartus II Software



Specifying Signal Breakout Layers

Each user I/O pin in your FPGA device can break out at different layers on your PCB. In the Pin Planner, you can specify on which layers the I/O pins in your design break out. The breakout layer information is used only during SSN analysis and is not used in other processes run by the Quartus II software. To assign a pin to PCB layer, follow these steps:

1. On the Assignments menu, click **Pin Planner**.
2. If necessary, perform Analysis & Elaboration, Analysis & Synthesis, or fully compile the design to populate the Pin Planner with the node names in the design.
3. Right-click anywhere in the **All Pins** or **Groups** list, and then click **Customize Columns**.
4. Select the **PCB layer** column and move it from the **Available columns** list to the **Show these columns in this order** list.
5. Click **OK**.
6. In the **PCB layer** column, specify the PCB layer to which you want to connect the signal.
7. On the File menu, click **Save Project** to save the changes.

 When you create PCB breakout layer assignments in the Pin Planner, you can assign the pin to any layer, even if you did not yet define the PCB layer.

Creating I/O Assignments

I/O assignments are required in FPGA design and are also used during SSN analysis to estimate voltage noise. Each input, output, or bidirectional signal in your design is assigned a physical pin location on the device using pin location assignments. Each signal has a physical I/O buffer that has a specific I/O standard, pin location, drive strength, and slew rate. The SSN Analyzer supports most I/O standards in a device family, such as the LVTTTL and LVCMOS I/O standards.



The SSN Analyzer does not support differential I/O standards, such as the LVDS I/O standard and its variations, because differential I/O standards contribute a small amount of SSN.



For more information about supported I/O standards, refer to the appropriate device handbook available on the [Literature and Technical Documentation](#) page of the Altera website.



For more information about creating and managing I/O assignments, refer to the *I/O Management* chapter in volume 2 of the *Quartus II Handbook*.

Decreasing Pessimism in SSN Analysis

In the absence of specific timing information, the SSN Analyzer analyzes your design under worst-case conditions. Worst-case conditions include all pins acting as aggressor signals on all possible victim pins and all aggressor pins switching with the worst possible timing relationship. The results of SSN analysis under worst-case conditions are very pessimistic. You can improve the results of SSN Analysis by creating group assignments for specific types of pins. Use the following group assignments to decrease the pessimism in SSN analysis results:

- Assign pins to an output enable group—All pins in an output enable group must be either all input pins or all output pins. If all the pins in a group are always either all inputs or all outputs, it is impossible for an output pin in the group to cause SSN noise on an input pin in the group. You can assign pins to an output enable group with the **Output Enable Group** logic option.
- Assign pins to a synchronous group—I/O pins that are part of a synchronous group (signals that switch at the same time) may cause SSN, but do not result in any failures because the noise glitch occurs during the switching period of the signal. The noise, therefore, does not occur in the sampling window of that signal. You can assign pins to an output enable group with the **Synchronous Group** logic option. For example, in your design you have a bus with 32 pins that all belong to the same group. In a real operation, the bus switches at the same time, so any voltage noise induced by a pin on its groupmates does not matter, because it does not fall in the sampling window. If you do not assign the bus to a synchronous group, the other 31 pins can act as aggressors for the first pin in that group, leading to higher QL and QH noise levels during SSN analysis.

In some cases, the SSN Analyzer can detect the grouping for bidirectional pins by looking at the output enable signal of the bidirectional pins. However, Altera recommends that you explicitly specify the bidirectional groups and output groups in your design.

- ② For more information about creating logic option assignments, refer to [Assigning Pins](#) in Quartus II Help.

Excluding Pins as Aggressor Signals

The SSN Analyzer uses the following conditions to exclude pins as aggressor signals for a specific victim pin:

- A pin that is a complement of the victim pin. For example, any pin that is assigned a differential I/O standard cannot be an aggressor pin.
 - A programming pin or JTAG pin because these pins are not active in user mode.
 - Pins that have the same output enable signal as a bidirectional victim pin that the SSN Analyzer analyzes as an input pin. Pins with the same output enable signal also act as input pins and therefore cannot be aggressor pins at the same time. For information about grouping bidirectional pins, refer to [“Performing SSN Analysis and Viewing Results”](#).
 - Pins in the same synchronous group as a victim output pin. For information about grouping output pins, refer to [“Performing SSN Analysis and Viewing Results”](#).
 - A pin assigned the **I/O Maximum Toggle Rate** logic option with a frequency setting of zero. The SSN Analyzer does not consider pins with this setting as aggressor pins.
- ② For more information about creating pin assignments with the Pin Planner, refer to [Assigning Pins](#) in Quartus II Help.

Performing SSN Analysis and Viewing Results

You can perform SSN analysis either on your entire design, or you can limit the analysis to specific I/O banks.

If you know the problem area for SSN is within one I/O bank and you are changing pin assignments only in that bank, you can run SSN analysis for just that one I/O bank to reduce analysis time.

- ② For more information, refer to [Running the SSN Analyzer](#) in Quartus II Help.
-  For more information about I/O bank numbering, refer to the appropriate device handbook available on the [Literature and Technical Documentation](#) page of the Altera website.

Understanding the SSN Reports

When SSN analysis is complete, you can view detailed analysis reports. The detailed messages in the reports help you understand and resolve SSN problems.

The SSN Analyzer section of the Compilation report contains information generated during SSN analysis, including the following reports:

- Summary
- Output Pins

- Input Pins
- Unanalyzed Pins
- Confidence Metric Details

Summary Report

The Summary report summarizes the SSN Analyzer status and rates the SSN Analyzer confidence level as low, medium, or high. The confidence level depends on the completeness of your board trace model assignments. The more assignments you complete, the higher the confidence level. However, the confidence level does not always contribute to the accuracy of the QL and QH noise levels on a victim pin. The accuracy of QH and QL noise levels depends the accuracy of your board trace model assignments.

Output Pins and Input Pins Reports

The Output Pins report lists all of the output pins and bidirectional pins that are treated as output pins during SSN analysis. The Input Pins report lists all of the input pins and bidirectional pins that are treated as inputs during SSN analysis. Both reports list the location assignments for the pins treated as SSN outputs or inputs during SSN analysis, the QL and QH noise in volts, and what percentage the QL and QH margins are for the I/O standard used for that signal. The QH and QL noise margins that fall in the critical range (> 90%) are shown in red. The QH and QL noise margins that fall in the range of 70% to 90% are shown in gray.

Unanalyzed Pins Report

Not all pins are analyzed for SSN analysis. The following pins are not analyzed and are reported in the Unanalyzed Pins report:

- Pins assigned the LVDS I/O standard or any LVDS variations, such as the **mini-LVDS** I/O standard
- Pins created in the migration flow that cover power and supply pins in other packages
- The negative terminals of pseudo-differential I/O standards; the noise on differential standards is reported as the differential noise and is reported on the positive terminal

Confidence Metric Details Report

The Confidence Metric Details Report lists the values used during SSN Analysis for unspecified I/O, board, and PCB assignments.

Viewing SSN Analysis Results in the Pin Planner

After SSN analysis completes, you can analyze the results in the Pin Planner. In the Pin Planner you can identify the SSN hotspots in your device, as well as the QL and QH noise levels. The QL and QH results for each pin are displayed with a different color that represents whether the pin is below the warning threshold, below the critical threshold, or above the critical threshold. This color representation is also referred to as the SSN map of your FPGA device.

When you view the SSN map, you can customize which details to display, including input pins, output pins, QH signals, QL signals, and noise levels. You can also adjust the threshold levels for QH and QL noise voltages. Adjusting the threshold levels in the Pin Planner does not change the threshold levels reported during SSN analysis and does not change the data in any of the SSN reports.

You can also you change I/O assignments and board trace information and rerun the SSN Analyzer to view the SSN analysis results based on those modified settings.

- ② For more information, refer to *Show SSN Analyzer Results* and *Running the SSN Analyzer* in Quartus II Help.

Decreasing Processing Time for SSN Analysis

FPGA designs are getting larger in density, logic, and I/O count. The time it takes to complete SSN analysis and other Quartus II software processes affects your development time. Faster processing times can reduce your design cycle time. Use the following guidelines to reduce processing time:

- Direct the Quartus II software to use more than one processor for parallel executables, including the SSN Analyzer
- Perform SSN analysis after I/O assignment analysis if your design files and constraints are complete, and you are interested in generating the SSN results early in the design process and want to adjust I/O placements to see if you can obtain better results
- Perform SSN analysis after fitting if you want to view preliminary SSN results that do not take into account complete I/O assignment and I/O timing results
- Perform engineering change orders (ECOs) on your design, rather than recompiling the entire design, if you want to rerun SSN analysis after changing I/O assignments

- ② For more information about using parallel processors, refer to *Setting Up and Running Analysis and Synthesis* and *Compilation Process Settings Page* in Quartus II Help. For more information about performing I/O assignment analysis, refer to *Assigning Pins* in Quartus II Help. For more information about running the Fitter, refer to *Setting Up and Running the Fitter* in Quartus II Help.

-  For more information about performing ECOs on your design, refer to the *Engineering Change Management with the Chip Planner* chapter in volume 2 of the *Quartus II Handbook*.

Scripting Support

A Tcl script allows you to run procedures and determine settings described in this chapter. You can also run some of these procedures at a command prompt. The Quartus II software provides several packages to compile your design and create I/O assignments for analysis and fitting. You can create a custom Tcl script that maps the design and runs SSN analysis on your design.

For detailed information about specific scripting command options and Tcl API packages, type the following command at a system command prompt to run the Quartus II Command-Line and Tcl API Help browser:

```
quartus_sh --qhelp ←
```



For more information about Quartus II scripting support, including examples, refer to the *Tcl Scripting* and *Command-Line Scripting* chapters in volume 2 of the *Quartus II Handbook* and *API Functions for Tcl* in Quartus II Help.

Optimizing Pin Placements for Signal Integrity

You can create an assignment that directs the Fitter to optimize pin placements for signal integrity with a Tcl command.

The following Tcl command directs the Fitter to optimize pin placement for signal integrity without affecting design f_{MAX} :

```
set_global_assignment -name OPTIMIZE_SIGNAL_INTEGRITY "Normal  
Compilation"
```

For more information, refer to “[Optimizing Pin Placements for Signal Integrity](#)” on page 6-9.

Defining PCB Layers and PCB Layer Thickness

You can create PCB layer and thickness assignments with a Tcl command. [shows](#) Tcl commands for specifying PCB layer assignments.

Example 6-3. Specifying PCB Layer Assignments

```
set_global_assignment -name PCB_LAYER_THICKNESS 0.00099822M -section_id 1  
set_global_assignment -name PCB_LAYER_THICKNESS 0.00034036M -section_id 2  
set_global_assignment -name PCB_LAYER_THICKNESS 0.00034036M -section_id 3  
set_global_assignment -name PCB_LAYER_THICKNESS 0.00055372M -section_id 4  
set_global_assignment -name PCB_LAYER_THICKNESS 0.00034036M -section_id 5  
set_global_assignment -name PCB_LAYER_THICKNESS 0.00034036M -section_id 6  
set_global_assignment -name PCB_LAYER_THICKNESS 0.00082042M -section_id 7
```

These Tcl commands specify that there are seven PCB layers in the design, each with a different thickness. In each assignment, the letter M indicates the unit of measurement is millimeters. When you specify PCB layer assignments with Tcl commands, you must list the layers in consecutive order. For example, you would receive an error during SSN Analysis if your Tcl commands created the following assignments:

```
set_global_assignment -name PCB_LAYER_THICKNESS 0.00099822M -section_id 1  
set_global_assignment -name PCB_LAYER_THICKNESS 0.00082042M -section_id 7
```

To create assignments with the unit of measurement in mils, refer to the syntax in the following Tcl commands. These Tcl commands specify the same settings as shown in [Figure 6-10 on page 6-13](#).

```
set_global_assignment -name PCB_LAYER_THICKNESS 14.9MIL -section_id 1  
set_global_assignment -name PCB_LAYER_THICKNESS 6.6MIL -section_id 2  
set_global_assignment -name PCB_LAYER_THICKNESS 14MIL -section_id 3  
set_global_assignment -name PCB_LAYER_THICKNESS 6.6MIL -section_id 4
```

For more information, refer to [“Defining PCB Layers and PCB Layer Thickness”](#) on page 6-11.

Specifying Signal Breakout Layers

You can create signal breakout layer assignments with a Tcl command. [Example 6-4](#) shows Tcl commands for specifying signal breakout layer assignments:

Example 6-4. Specifying Signal Breakout Layer Assignments

```
set_instance_assignment -name PCB_LAYER 10 -to e[2]  
set_instance_assignment -name PCB_LAYER 3 -to e[3]
```

When you create PCB breakout layer assignments with Tcl commands, if you do not specify a PCB layer, or if you specify a PCB layer that does not exist, the SSN Analyzer breaks out the signal at the bottommost PCB layer.



If you create a PCB layer breakout assignment to a layer that does not exist, the SSN Analyzer will generate a warning message.

For more information, refer to [“Specifying Signal Breakout Layers”](#) on page 6-13.

Decreasing Pessimism in SSN Analysis

You can create output enable group and synchronous group assignments to help decrease pessimism during SSN Analysis with a Tcl command.

The following Tcl command assigns the bidirectional bus DATAINOUT to an output enable group:

```
set_instance_assignment -name OUTPUT_ENABLE_GROUP 1 -to DATAINOUT
```

The following Tcl command assigns the bus PCI_ADD_io to a synchronous group:

```
set_instance_assignment -name SYNCHRONOUS_GROUP 1 -to PCI_AD_io
```

For more information, refer to [“Decreasing Pessimism in SSN Analysis”](#) on page 6-14.

Performing SSN Analysis

You can perform SSN analysis with a command-line command. Use the `quartus_si` package that is provided with the Quartus II software.

Type the following command at a system command prompt to start the SSN Analyzer:

```
quartus_si <project name> ␣
```

To analyze just one I/O bank, type the following command at a system command prompt:

```
quartus_si <project revision> <--bank = bank id> ␣
```

For example, to run analyze the I/O bank 2A type the following command:

```
quartus_si counter --bank=2A ␣
```

For more information, refer to “Performing SSN Analysis and Viewing Results” on page 6-15.



For more information about the `quartus_si` package, type `quartus_si -h` at a system command prompt.

Conclusion

To assist you with SSN Analysis, you can use the fast and accurate SSN Analyzer to help you estimate the SSN performance of your FPGA both early in the design cycle and when your PCB is complete. The SSN methodology discussed in this chapter gives you the tools you need to ensure your FPGA design meets your SSN requirements.

Document Revision History

Table 6-2 shows the revision history for this chapter.

Table 6-2. Document Revision History

Date	Version	Changes
December 2010	10.0.1	Template update
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Reorganized and edited the chapter ■ Added links to Quartus II Help for procedural information previously included in the chapter
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Added “Figure 6-9 shows the layout cross-section of a PCB in the Cadence Allegro PCB tool. The cross-section shows the stackup information of a PCB, which tells you the number of layers used in your PCB. The PCB shown in this example consists of various signal and circuit layers on which FPGA pins are routed, as well as the power and ground layers.” on page 6-12 ■ Updated for the Quartus II software 9.1 release
March 2009	9.0.0	Initial release



For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).



Take an [online survey](#) to provide feedback about this handbook chapter.

Introduction

With the ever-increasing operating speed of interfaces in traditional FPGA design, the timing and signal integrity margins between the FPGA and other devices on the board must be within specification and tolerance before a single PCB is built. If the board trace is designed poorly or the route is too heavily loaded, noise in the signal can cause data corruption, while overshoot and undershoot can potentially damage input buffers over time.

As FPGA devices are used in high-speed applications, signal integrity and timing margin between the FPGA and other devices on the printed circuit board (PCB) are important aspects to consider to ensure proper system operation. To avoid time-consuming redesigns and expensive board respins, the topology and routing of critical signals must be simulated. The high-speed interfaces available on current FPGA devices must be modeled accurately and integrated into timing models and board-level signal integrity simulations. The tools used in the design of an FPGA and its integration into a PCB must be “board-aware”—able to take into account properties of the board routing and the connected devices on the board.

This chapter contains the following topics:

- “I/O Model Selection: IBIS or HSPICE” on page 7–3
- “FPGA to Board Signal Integrity Analysis Flow” on page 7–4
- “Simulation with IBIS Models” on page 7–7
- “Simulation with HSPICE Models” on page 7–16

The Quartus® II software provides methodologies, resources, and tools to ensure good signal integrity and timing margin between Altera® FPGA devices and other components on the board. Three types of analysis are possible with the Quartus II software:

- I/O timing with a default or user-specified capacitive load and no signal integrity analysis (default)
- The Quartus II **Enable Advanced I/O Timing** option utilizing a user-defined board trace model to produce enhanced timing reports from accurate “board-aware” simulation models
- Full board routing simulation in third-party tools using Altera-provided or generated Input/Output Buffer Information Specification (IBIS) or HSPICE I/O models

I/O timing using a specified capacitive test load requires no special configuration other than setting the size of the load. I/O timing reports from the Quartus II TimeQuest or the Quartus II Classic Timing Analyzer are generated based only on point-to-point delays within the I/O buffer and assume the presence of the capacitive test load with no other details about the board specified. The default size of the load is based on the I/O standard selected for the pin. Timing is measured to the FPGA pin with no signal integrity analysis details.

The **Enable Advanced I/O Timing** option expands the details in I/O timing reports by taking board topology and termination components into account. A complete point-to-point board trace model is defined and accounted for in the timing analysis. This ability to define a board trace model is an example of how the Quartus II software is “board-aware.”

In this case, timing and signal integrity metrics between the I/O buffer and the defined far end load are analyzed and reported in enhanced reports generated by the Quartus II TimeQuest Timing Analyzer.

 For more information about defining capacitive test loads or how to use the **Enable Advanced I/O Timing** option to configure a board trace model, refer to the *I/O Management* chapter in volume 2 of the *Quartus II Handbook*.

This chapter focuses on the third type of analysis. The Quartus II software can export accurate HSPICE models with the built-in HSPICE Writer. You can run signal integrity simulations with these complete HSPICE models in Synopsys HSPICE. IBIS models of the FPGA I/O buffers are also created easily with the Quartus II IBIS Writer. You can integrate IBIS models into any third-party simulation tool that supports them, such as the Mentor Graphics® Hyperlynx software. With the ability to create industry-standard model definition files quickly, you can build accurate simulations that can provide data to help improve board-level signal integrity.

The I/O's IBIS and HSPICE model creation available in the Quartus II software can help prevent problems before a costly board respin is required. In general, creating and running accurate simulations is difficult and time consuming. The tools in the Quartus II software automate the I/O model setup and creation process by configuring the models specifically for your design. With these tools, you can set up and run accurate simulations quickly and acquire data that helps guide your FPGA and board design.

The information about signal integrity in this chapter refers to board-level signal integrity based on I/O buffer configuration and board parameters, not simultaneous switching noise (SSN), also known as ground bounce or V_{CC} sag. SSN is a product of multiple output drivers switching at the same time, causing an overall drop in the voltage of the chip's power supply. This can cause temporary glitches in the specified level of ground or V_{CC} for the device.

 For a more information about SSN and ways to prevent it, refer to *AN 315: Guidelines for Designing High-Speed FPGA PCBs*.

This chapter is intended for FPGA and board designers and includes details about the concepts and steps involved in getting designs simulated and how to adjust designs to improve board-level timing and signal integrity. Also included is information about how to create accurate models from the Quartus II software and how to use those models in simulation software.

The information in this chapter is meant for those who are familiar with the Quartus II software and basic concepts of signal integrity and the design techniques and components in good PCB design. Finally, you should know how to set up simulations and use your selected third-party simulation tool.

 For information about basic signal integrity concepts and signal integrity details pertaining to Altera FPGA devices, refer to the [Altera Signal Integrity Center](#).

I/O Model Selection: IBIS or HSPICE

The Quartus II software can export two different types of I/O models that are useful for different simulation situations. IBIS models define the behavior of input or output buffers through the use of voltage-current (V-I) and voltage-time (V-t) data tables. HSPICE models, often referred to as HSPICE decks, include complete physical descriptions of the transistors and parasitic capacitances that make up an I/O buffer along with all the parameter settings required to run a simulation. The HSPICE decks generated by the Quartus II software are preconfigured with the I/O standard, voltage, and pin loading settings for each pin in your design.

The choice of I/O model type is based on many factors. [Table 7-1](#) shows a detailed comparison of the two I/O model types and information and examples of situations in which they might be used.

Table 7-1. IBIS and HSPICE Model Comparison

Feature	IBIS Model	HSPICE Model
I/O Buffer Description	Behavioral —I/O buffers are described by voltage-current and voltage-time tables in typical, minimum, and maximum supply voltage cases.	Physical —I/O buffers and all components in a circuit are described by their physical properties, such as transistor characteristics and parasitic capacitances, as well as their connections to one another.
Model Customization	Simple and limited —The model completely describes the I/O buffer and does not usually have to be customized.	Fully customizable —Unless connected to an arbitrary board description, the description of the board trace model must be customized in the model file. All parameters of the simulation are also adjustable.
Simulation Set Up and Run Time	Fast —Simulations run quickly after set up correctly.	Slow —Simulations take time to set up and take longer to run and complete.
Simulation Accuracy	Good —For most simulations, accuracy is sufficient to make useful adjustments to the FPGA and/or board design to improve signal integrity.	Excellent —Simulations are highly accurate, making HSPICE simulation almost a requirement for any high-speed design where signal integrity and timing margins are tight.
Third-Party Tool Support	Excellent —Almost all third-party board simulation tools support IBIS.	Good —Most third-party tools that support SPICE support HSPICE. However, Synopsys HSPICE is required for simulations of Altera's encrypted HSPICE models.

 For more information about IBIS files created by the Quartus II IBIS Writer and IBIS files in general, as well as links to websites with detailed information, refer to [AN 283: Simulating Altera Devices with IBIS Models](#).

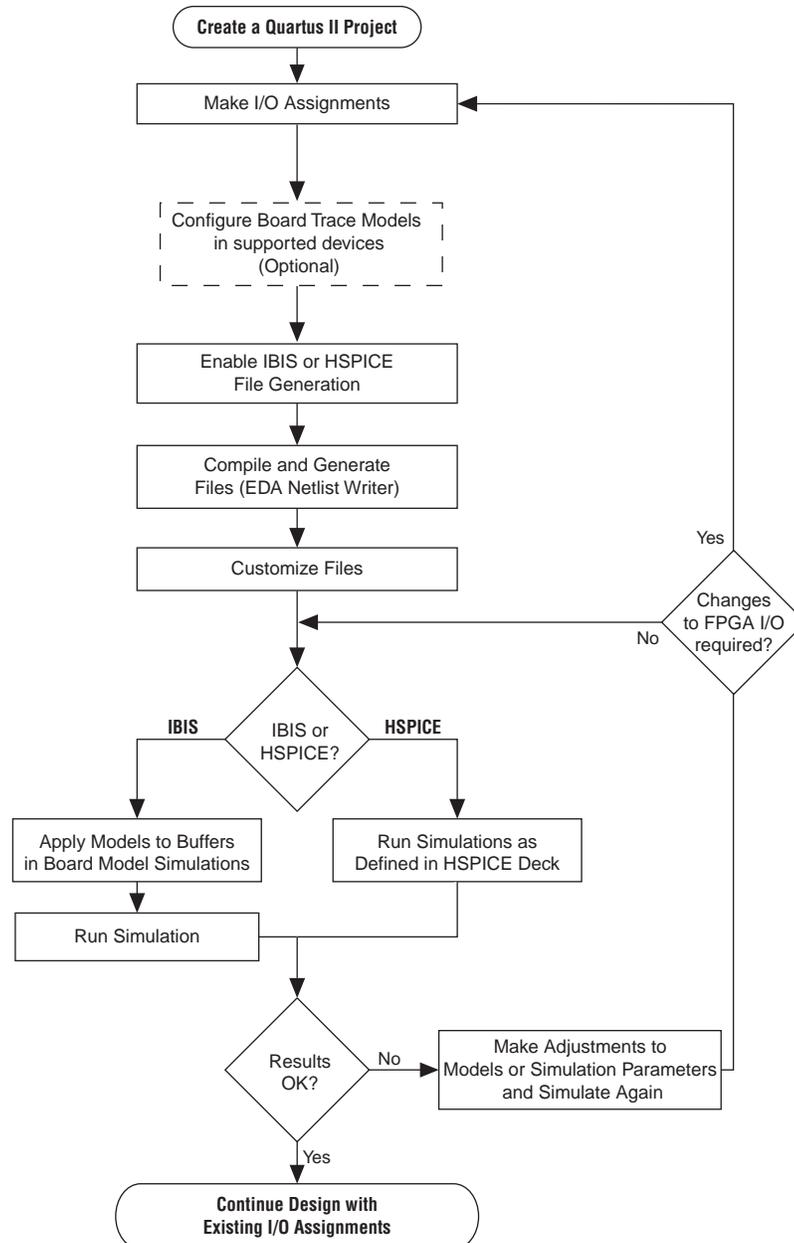
FPGA to Board Signal Integrity Analysis Flow

Board signal integrity analysis can take place at any point in the FPGA design process and is often performed before and after board layout. If it is performed early in the process as part of a pre-PCB layout analysis, the models used for simulations can be more generic and can be changed as much as required to see how adjustments improve timing or signal integrity and help with the design and routing of the PCB. Simulations and the resulting changes made at this stage allow you to analyze “what if” scenarios to plan and implement your design better. To assist with early board signal integrity analysis, you can download generic IBIS model files for each device family and obtain HSPICE buffer simulation kits from the “Board Level Tools” section of the [EDA Tool Support Resource Center](#).

Typically, if board signal integrity analysis is performed late in the design, it is used for a post-layout verification. The inputs and outputs of the FPGA are defined, and required board routing topologies and constraints are known. Simulations can help you find problems that might still exist in the FPGA or board design before fabrication and assembly. In either case, a simple process flow illustrates how to create accurate IBIS and HSPICE models from a design in the Quartus II software and transfer them to third-party simulation tools. [Figure 7-1](#) shows this flow.

 This chapter is organized around the type of model, IBIS or HSPICE, that you use for your simulations. When you understand the steps in the analysis flow, refer to the section of this chapter that corresponds to the model type you are using.

Figure 7-1. Third-Party Board Signal Integrity Analysis Flow



Create I/O and Board Trace Model Assignments

If your design uses a Stratix® III, Stratix II, or Cyclone® III device, you can configure a board trace model for output signals or for bidirectional signals in output mode and automatically transfer its description to HSPICE decks generated by the HSPICE Writer. This helps improve simulation accuracy.

To configure a board trace model, in the **Settings** dialog box, in the **TimeQuest Timing Analyzer** page, turn on the **Enable Advanced I/O Timing** option and configure the board trace model assignment settings for each I/O standard used in your design. You can add series or parallel termination, specify the transmission line length, and set the value of the far-end capacitive load. You can configure these parameters either in the Board Trace Model view of the Pin Planner, or click **Device and Pin Options** in the **Device** page of the **Settings** dialog box.



For information about how to use the **Enable Advanced I/O Timing** option and configure board trace models for the I/O standards used in your design, refer to the *I/O Management* chapter in volume 2 of the *Quartus II Handbook*.

The Quartus II software can generate IBIS models and HSPICE decks without having to configure a board trace model with the **Enable Advanced I/O Timing** option. In fact, IBIS models ignore any board trace model settings other than the far-end capacitive load. If any load value is set other than the default, the delay given by IBIS models generated by the IBIS Writer cannot be used to account correctly for the double counting problem. The load value mismatch between the IBIS delay and the t_{CO} measurement of the Quartus II software prevents the delays from being safely added together. Warning messages displayed when the EDA Netlist Writer runs indicate when this mismatch occurs.

Output File Generation

IBIS and HSPICE model files are not generated by the Quartus II software by default. To generate or update the files automatically during each project compilation, select the type of file to generate and a location where to save the file in the project settings. These settings can also be specified with commands in a Tcl script.

The IBIS and HSPICE Writers in the Quartus II software are run as part of the EDA Netlist Writer during normal project compilation. If either writer is turned on in the project settings, IBIS or HSPICE files are created and stored in the specified location. For IBIS, a single file is generated containing information about all assigned pins. HSPICE file generation creates separate files for each assigned pin. You can run the EDA Netlist Writer separately from a full compilation in the Quartus II software or at the command line. However, you must fully compile the project or perform I/O Assignment Analysis at least once for the IBIS and HSPICE Writers to have information about the I/O assignments and settings in the design.

Customize the Output Files

The files generated by either the IBIS or HSPICE Writer are text files that you can edit and customize easily for design or experimentation purposes. IBIS files downloaded from the Altera website must be customized with the correct RLC values for the specific device package you have selected for your design. IBIS files generated by the IBIS Writer do not require this customization because they are configured automatically with the RLC values for your selected device. HSPICE decks require modification to include a detailed description of your board. With **Enable Advanced I/O Timing** turned on and a board trace model defined in the Quartus II software, generated HSPICE decks automatically include that model's parameters. However, Altera recommends that you replace that model with a more detailed model that

describes your board design more accurately. A default simulation included in the generated HSPICE decks measures delay between the FPGA and the far-end device. You can make additions or adjustments to the default simulation in the generated files to change the parameters of the default simulation or to perform additional measurements.

Set Up and Run Simulations in Third-Party Tools

When you have generated the files, you can use them to perform simulations in your selected simulation tool. With IBIS models, you can apply them to input, output, or bidirectional buffer entities and quickly set up and run simulations. For HSPICE decks, the simulation parameters are included in the files. Open the files in Synopsys HSPICE and run simulations for each pin as required.

With HSPICE decks generated from the HSPICE Writer, the double counting problem is accounted for, which ensures that your simulations are accurate. Simulations that involve IBIS models created with anything other than the default loading settings in the Quartus II software must take the change in the size of the load between the IBIS delay and the Quartus II t_{CO} measurement into account. Warning messages during compilation alert you to this change.

Interpret Simulation Results

If you encounter timing or signal integrity issues with your high-speed signals after running simulations, you can make adjustments to I/O assignment settings in the Quartus II software. These could include such things as drive strength or I/O standard, or making changes to your board routing or topology. After regenerating models in the Quartus II software based on the changes you have made, rerun the simulations to check whether your changes corrected the problem.

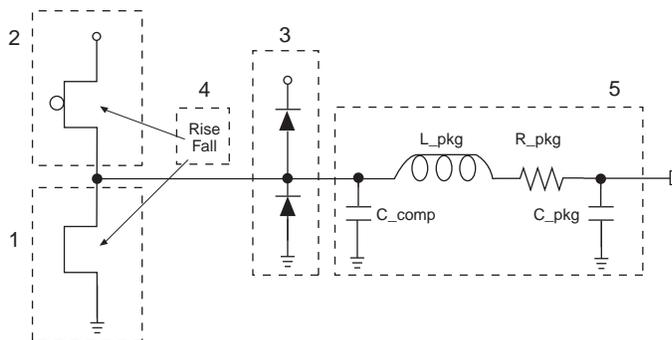
Simulation with IBIS Models

IBIS models provide a way to run accurate signal integrity simulations quickly. IBIS models describe the behavior of I/O buffers with voltage-current and voltage-time data curves. Because of their behavioral nature, IBIS models do not have to include any information about the internal circuit design of the I/O buffer. Most component manufacturers, including Altera, provide IBIS models for free download and use in signal integrity analysis simulation tools. You can download generic device family IBIS models from the Altera website for early design simulation or use the IBIS Writer to create custom IBIS models for your existing design.

Elements of an IBIS Model

An IBIS model file (**.ibs**) is a text file that describes the behavior of an I/O buffer across minimum, typical, and maximum temperature and voltage ranges with a specified test load. The tables and values specified in the IBIS file describe five basic elements of the I/O buffer. [Figure 7-2](#) highlights each of these elements in the I/O buffer model.

Figure 7-2. Five Basic Elements in IBIS Models



The following elements correspond to each numbered block in [Figure 7-2](#).

1. **Pulldown**—A voltage-current table describes the current when the buffer is driven low based on a pull-down voltage range of $-V_{CC}$ to $2 V_{CC}$.
2. **Pullup**—A voltage-current table describes the current when the buffer is driven high based on a pull-up voltage range of $-V_{CC}$ to V_{CC} .
3. **Ground and Power Clamps**—Voltage-current tables describe the current when clamping diodes for electrostatic discharge (ESD) are present. The ground clamp voltage range is $-V_{CC}$ to V_{CC} , and the power clamp voltage range is $-V_{CC}$ to ground.
4. **Ramp and Rising/Falling Waveform**—A voltage-time (dv/dt) ratio describes the rise and fall time of the buffer during a logic transition. Optional rising and falling waveform tables can be added to more accurately describe the characteristics of the rising and falling transitions.
5. **Total Output Capacitance and Package RLC**—The total output capacitance includes the parasitic capacitances of the output pad, clamp diodes (if present), and input transistors. The package RLC is device package-specific and defines the resistance, inductance, and capacitance of the bond wire and pin of the I/O.



For more information about IBIS models and Altera-specific features, including links to the official IBIS specification, refer to [AN 283: Simulating Altera Devices with IBIS Models](#).

Creating Accurate IBIS Models

There are two methods to obtain Altera device IBIS files for your board-level signal integrity simulations. You can download generic IBIS models from the Altera website or you can use the IBIS writer in the Quartus II software to create design-specific models.

Download IBIS Models

Altera provides IBIS models for almost all FPGA and FPGA configuration devices. Check the [Altera IBIS Models](#) page for information about whether models for your selected device are available. You can use the IBIS models from the website to perform early simulations of the I/O buffers you expect to use in your design as part of a pre-layout analysis.

Downloaded IBIS models have the RLC package values set to one particular device in each device family. To simulate your design with the model accurately, you must adjust the RLC values in the IBIS model file to match the values for your particular device package by performing the following steps:

1. Download and expand the ZIP file (.zip) of the IBIS model for the device family you are using for your design. The .zip file contains the .ibs file along with an IBIS model user guide and a model data correlation report.
2. Download the Package RLC Values spreadsheet for the same device family.
3. Open the spreadsheet and locate the row that describes the device package used in your design.
4. From the package's **I/O** row, copy the minimum, maximum, and typical values of resistance, inductance, and capacitance for your device package.
5. Open the .ibs file in a text editor and locate the [Package] section of the file.
6. Overwrite the listed values copied with the values from the spreadsheet and save the file.

The .ibs file is now customized for your device package and can be used for any simulation. IBIS models downloaded and used for simulations in this manner are generic. They describe only a certain set of models listed for each device on the [Altera IBIS Models](#) page of the Altera website. To create customized models for your design, use the IBIS Writer as described in the next section.

Generate Custom IBIS Models with the IBIS Writer

If you have started your FPGA design and have created custom I/O assignments, such as drive strength settings or the enabling of clamping diodes for ESD protection, you can use the Quartus II IBIS Writer to create custom IBIS models to accurately reflect your assignments. IBIS models created with the IBIS Writer take I/O assignment settings into account.

If the **Enable Advanced I/O Timing** option is turned off, the generated .ibs files are based on the load value setting for each I/O standard on the **Capacitive Loading** page of the **Device and Pin Options** dialog box in the **Device** dialog box. With the **Enable Advanced I/O Timing** option turned on, IBIS models use an effective capacitive load based on settings found in the board trace model on the **Board Trace Model** page in the **Device and Pin Options** dialog box or the **Board Trace Model** view in the Pin Planner. The effective capacitive load is based on the sum of the **Near capacitance**, **Transmission line distributed capacitance**, and the **Far capacitance** settings in the board trace model. Resistances and transmission line inductance values are ignored.

-  If you made any changes from the default load settings, the delay in the generated IBIS model cannot safely be added to the Quartus II t_{CO} measurement to account for the double counting problem. This is because the load values between the two delay measurements do not match. When this happens, the Quartus II software displays warning messages when the EDA Netlist Writer runs to remind you about the load value mismatch.
-  For step-by-step instructions on how to generate IBIS models with the Quartus II software, refer to *Generating IBIS Output Files with the Quartus II Software* in Quartus II Help.
-  For more information about IBIS model generation, refer to the *AN 283: Simulating Altera Devices with IBIS Models* or to the Quartus II Help.

Design Simulation Using the Mentor Graphics HyperLynx Software

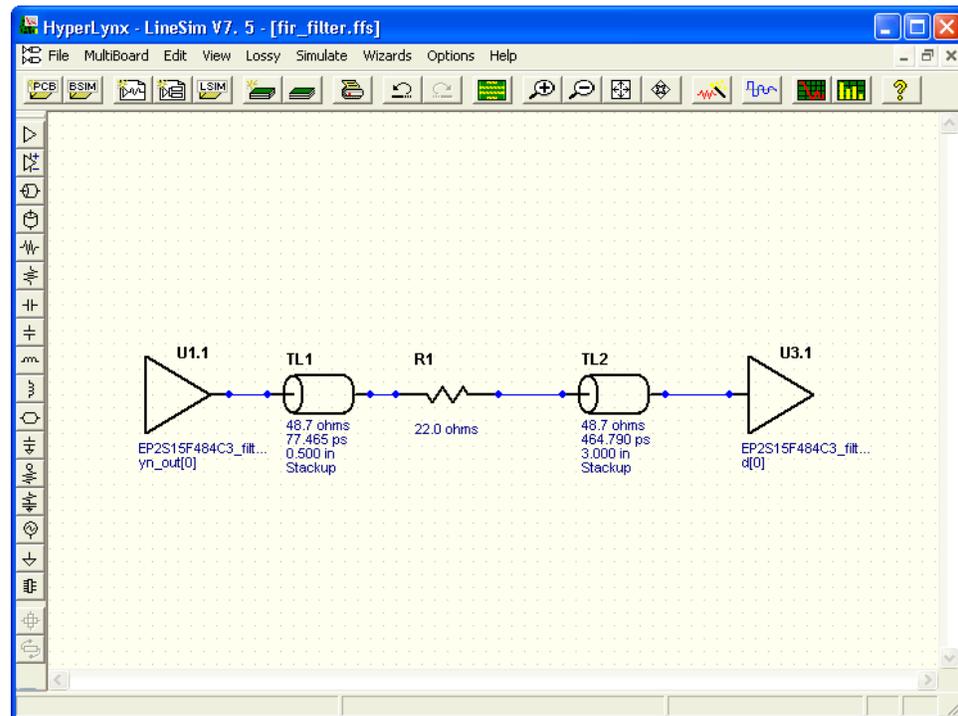
You must integrate IBIS models downloaded from the Altera website (www.altera.com) or created with the Quartus II IBIS Writer into board design simulations to accurately model timing and signal integrity. The HyperLynx software from Mentor Graphics is one of the most popular tools for design simulation. The HyperLynx software makes it easy to integrate IBIS models into simulations.

The HyperLynx software is a PCB analysis and simulation tool for high-speed designs, consisting of two products, LineSim and BoardSim. LineSim is an early simulation tool. Before any board routing takes place, LineSim is used to simulate “what if” scenarios to assist in creating routing rules and defining board parameters. BoardSim is a post-layout tool used to analyze existing board routing. Specific nets are selected from a board layout file and simulated in a manner similar to LineSim. With board and routing parameters, and surrounding signal routing known, highly accurate simulations of the final fabricated PCB are possible. This section focuses on LineSim. Because the process of creating and running simulations is very similar for both LineSim and BoardSim, the details of IBIS model use in LineSim applies to simulations in BoardSim.

Simulations in LineSim are configured using a schematic GUI to create connections and topologies between I/O buffers, route trace segments, and termination components. LineSim provides two methods for creating routing schematics: cell-based and free-form. Cell-based schematics are based on fixed cells consisting of typical placements of buffers, trace impedances, and components. Parts of the grid-based cells are filled with the desired objects to create the topology. A topology in a cell-based schematic is limited by the available connections within and between the cells.

A more robust and expandable way to create a circuit schematic for simulation is to use the free-form schematic format in LineSim as shown in Figure 7-3. The free-form schematic format makes it easy to place parts into any configuration and edit them as required. This section describes the use of IBIS models with free-form schematics, but the process is nearly identical for cell-based schematics.

Figure 7-3. HyperLynx LineSim Free-Form Schematic Editor



When you use HyperLynx software to perform simulations, you typically perform the following steps:

1. Create a new LineSim free-form schematic document and set up the board stackup for your PCB using the Stackup Editor. In this editor, specify board layer properties including layer thickness, dielectric constant, and trace width.
2. Create a circuit schematic for the net you want to simulate. The schematic represents all the parts of the routed net including source and destination I/O buffers, termination components, transmission line segments, and representations of impedance discontinuities such as vias or connectors.
3. Assign IBIS models to the source and destination I/O buffers to represent their behavior during operation.
4. Attach probes from the digital oscilloscope that is built in to LineSim to points in the circuit that you want to monitor during simulation. Typically, at least one probe is attached to the pin of a destination I/O buffer. For differential signals, you can attach a differential probe to both the positive and negative pins at the destination.
5. Configure and run the simulation. You can simulate a rising or falling edge and test the circuit under different drive strength conditions.

6. Interpret the results and make adjustments. Based on the waveforms captured in the digital oscilloscope, you can adjust anything in the circuit schematic to correct any signal integrity issues, such as overshoot or ringing. If necessary, you can make I/O assignment changes in the Quartus II software, regenerate the IBIS file with the IBIS Writer, and apply the updated IBIS model to the buffers in your HyperLynx software schematic.
7. Repeat the simulations and circuit adjustments until you are satisfied with the results. When the operation of the net meets your design requirements, implement changes to your I/O assignments in the Quartus II software and/or adjust your board routing constraints, component values, and placement to match the simulation.

 For more information about HyperLynx software, including schematic creation, simulation setup, model usage, product support, licensing, and training, refer to HyperLynx Help or the Mentor Graphics website at www.mentor.com.

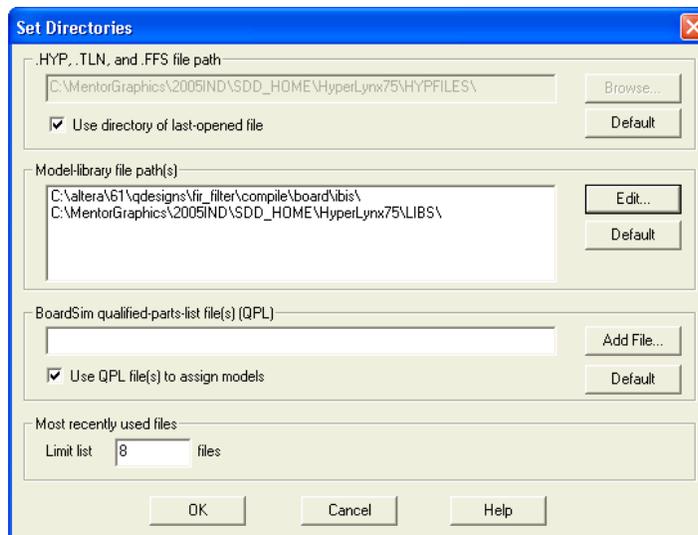
Configuring LineSim to Use Altera IBIS Models

You must configure LineSim to find and use the downloaded or generated IBIS models for your design. To do this, add the location of your **.ibs** file or files to the LineSim Model Library search path. Then you apply a selected model to a buffer in your schematic.

To add the Quartus II software's default IBIS model location, *<project directory>/board/ibis*, to the HyperLynx LineSim model library search path, perform the following steps in LineSim:

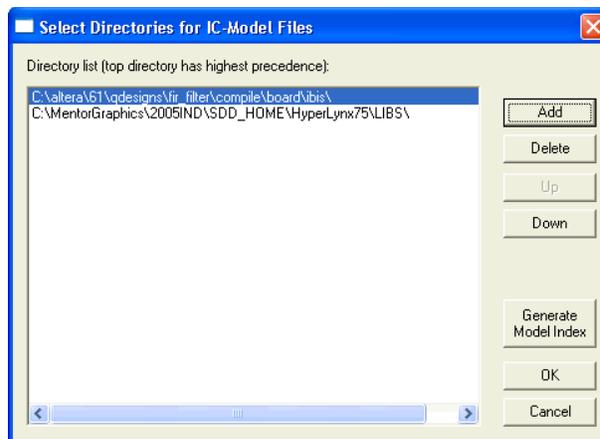
1. From the Options menu, click **Directories**. The **Set Directories** dialog box appears (Figure 7-4). The **Model-library file path(s)** list displays the order in which LineSim searches file directories for model files.

Figure 7-4. LineSim Set Directories Dialog Box



2. Click **Edit**. A dialog box appears where you can add directories and adjust the order in which LineSim searches them (Figure 7-5).

Figure 7-5. LineSim Select Directories Dialog Box



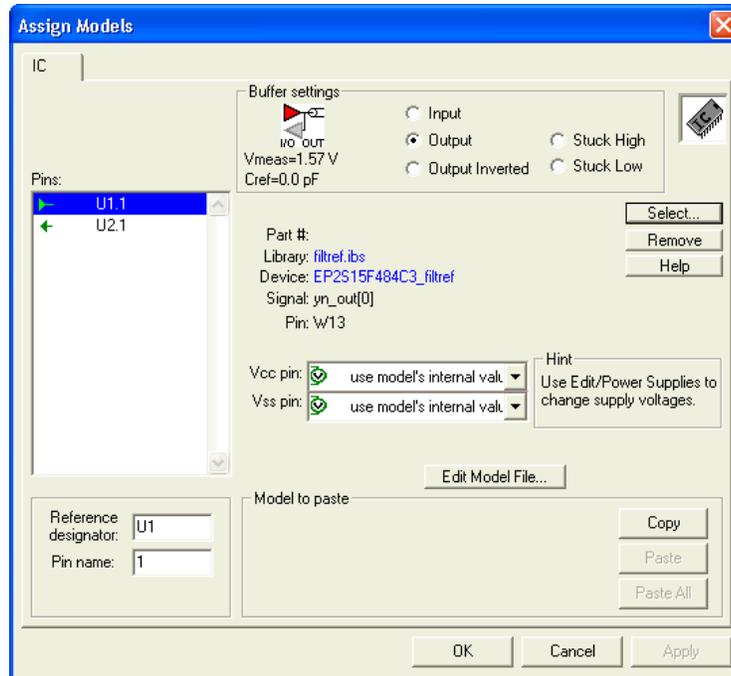
3. Click **Add**
4. Browse to the default IBIS model location, *<project directory>/board/ibis*. Click **OK**.
5. Click **Up** to move the IBIS model directory to the top of the list. Click **Generate Model Index** to update LineSim's model database with the models found in the added directory.
6. Click **OK**. The IBIS model directory for your project is added to the top of the Model-library file path(s) list.
7. To close the **Set Directories** dialog box, click **OK**.

Integrating Altera IBIS Models into LineSim Simulations

When the location for IBIS files has been set, you can assign the downloaded or generated IBIS models to the buffers in your schematic. To do this, perform the following steps:

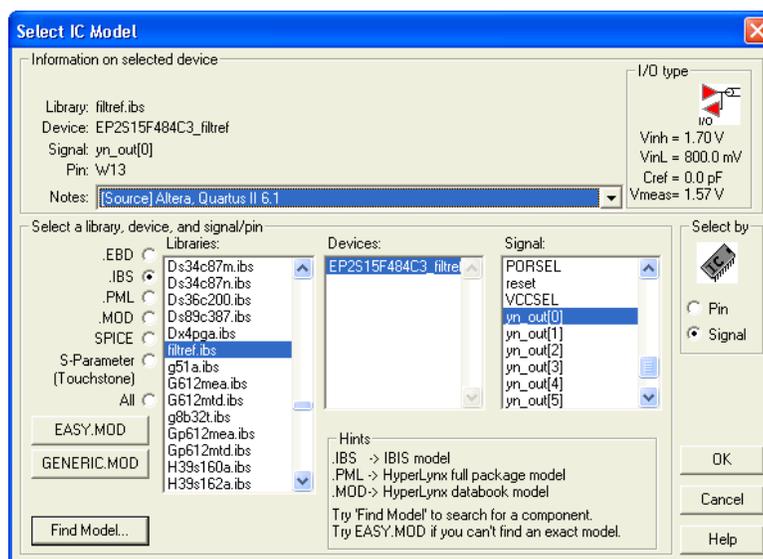
1. Double-click a buffer symbol in your schematic to open the **Assign Models** dialog box (Figure 7-6). You can also click **Assign Models** from the buffer symbol's right-click menu.

Figure 7-6. LineSim Assign Model Dialog Box



2. The pin of the buffer symbol you selected should be highlighted in the **Pins** list. If you want to assign a model to a different symbol or pin, select it from the list.
3. Click **Select**. The **Select IC Model** dialog box appears (Figure 7-7).

Figure 7-7. LineSim Select IC Model Dialog Box

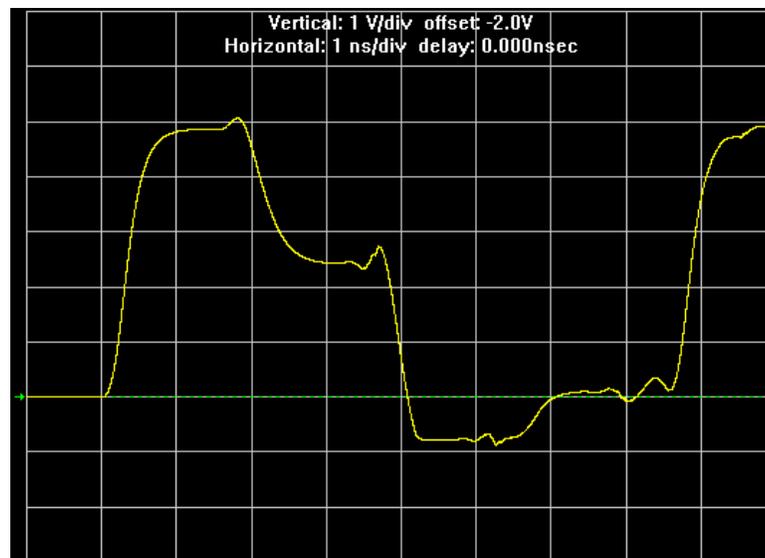


4. To filter the list of available libraries to display only IBIS models, select **.IBS**. Scroll through the **Libraries** list, and click the name of the library for your design. By default, this is *<project name>.ibs*.
5. The device for your design should be selected as the only item in the **Devices** list. If not, select your device from the list.
6. From the **Signal** list, select the name of the signal you want to simulate. You can also choose to select by device pin number.
7. Click **OK**. The **Assign Models** dialog box displays the selected **.ibs** file and signal.
8. If applicable to the signal you chose, adjust the buffer settings as required for the simulation.
9. Select and configure other buffer pins from the **Pins** list in the same manner.
10. Click **OK** when all I/O models are assigned.

Running and Interpreting LineSim Simulations

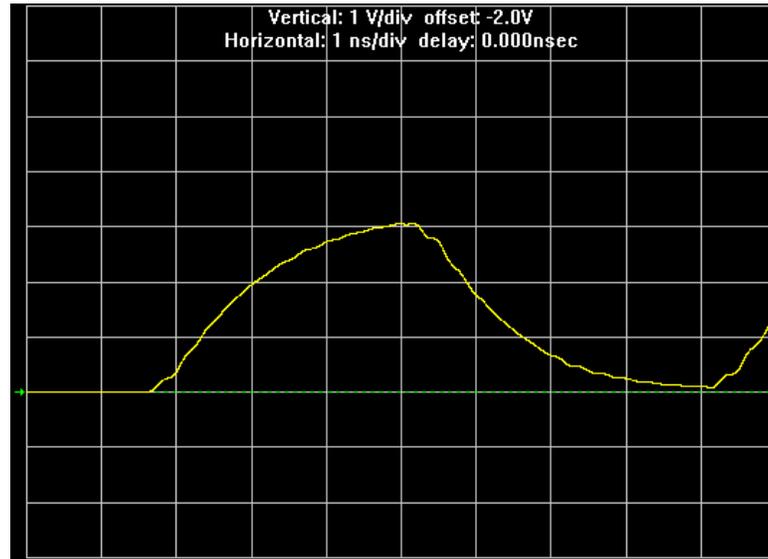
You can now run any desired simulations and make adjustments to the I/O assignments or simulation parameters as required. For example, if you see too much overshoot in the simulated signal at the destination buffer after running a simulation (as shown in [Figure 7-8](#)), you could adjust the drive strength I/O assignment setting to a lower value. Regenerate the **.ibs** file, and run the simulation again to verify whether the change fixed the problem.

Figure 7-8. Example of Overshoot in HyperLynx with IBIS Models



If you see a discontinuity or other anomalies at the destination, such as slow rise and fall times (as shown in Figure 7-9), adjust the termination scheme or termination component values. After making these changes, rerun the simulation to check whether your adjustments solved the problem. In this case, it is not necessary to regenerate the .ibs file.

Figure 7-9. Example of Signal Integrity Anomaly in HyperLynx with IBIS Models



 For more information about board-level signal integrity and to learn about ways to improve it with simple changes to your design, visit the [Altera Signal Integrity Center](#).

Simulation with HSPICE Models

HSPICE decks are used to perform highly accurate simulations by describing the physical properties of all aspects of a circuit precisely. HSPICE decks describe I/O buffers, board components, and all of the connections between them, as well as defining the parameters of the simulation to be run. By their nature, to be effective, HSPICE decks are highly customizable and require a detailed description of the circuit under simulation. For devices that support advanced I/O timing, when **Enable Advanced I/O Timing** is turned on, the HSPICE decks generated by the Quartus II HSPICE Writer automatically include board components and topology defined in the Board Trace Model. Configure the board components and topology in the Pin Planner or in the **Board Trace Model** tab of the **Device and Pin Options** dialog box. All HSPICE decks generated by the Quartus II software include compensation for the double count problem. For more information about the double count problem, refer to [“The Double Counting Problem in HSPICE Simulations”](#) on page 7-17. You can simulate with the default simulation parameters built in to the generated HSPICE decks or make adjustments to customize your simulation.

Supported Devices and Signaling

Beginning with Quartus II software version 6.1 and later, the HSPICE Writer supports the devices and signaling defined in [Table 7-2](#). Only Stratix III, Stratix II, and Cyclone III devices support the creation of a board trace model in the Quartus II software for automatic inclusion in an HSPICE deck. Other devices require the board description to be manually added to the HSPICE file.

Table 7-2. HSPICE Writer Device and Signaling Support

Device	Input	Output	Single-Ended	Differential	Automatic Board Trace Model Description
Stratix III	✓	✓	✓	✓	✓
Stratix II GX (non-HSSI pins)	✓	✓	✓	✓	—
Stratix II	✓	✓	✓	✓	✓
HardCopy® II	✓	✓	✓	✓	—
Cyclone III	✓	✓	✓	✓	✓

If you are using a Stratix II device for your design, you can turn on **Enable Advanced I/O Timing** and configure the board trace model for each I/O standard used in your design. Newer families have this feature turned on by default and it cannot be turned off. The HSPICE files include the board trace description you create in the Board Trace Model view in the Pin Planner or the **Board Trace Model** tab in the **Device and Pin Options** dialog box.



For more information about the **Enable Advanced I/O Timing** option and configuring board trace models for the I/O standards in your design, refer to the [I/O Management](#) chapter in volume 2 of the *Quartus II Handbook*.

Accessing HSPICE Simulation Kits

You can access the available HSPICE models at the [SPICE Models for Altera Devices](#) web page and also with the Quartus II software's HSPICE Writer tool. The Quartus II software HSPICE Writer tool removes many common sources of user error from the I/O simulation process. The HSPICE Writer tool automatically creates preconfigured I/O simulation spice decks that only require the addition of a user board model. All the difficult tasks required to configure the I/O modes and interpret the timing results are handled automatically by the HSPICE Writer tool.

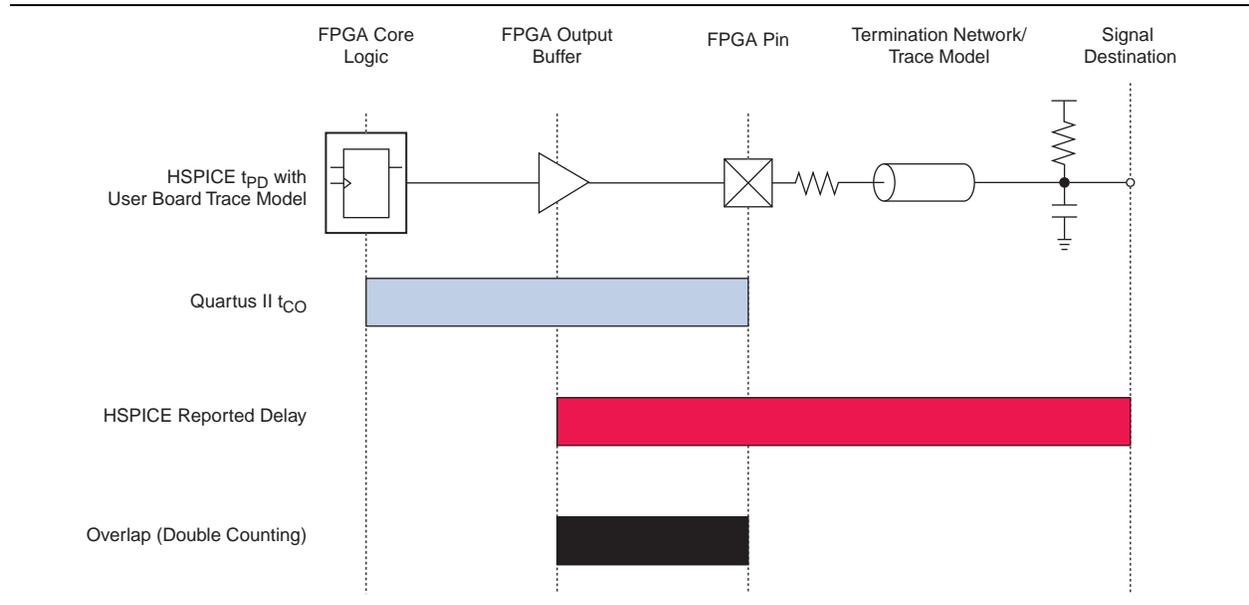
The Double Counting Problem in HSPICE Simulations

Simulating I/Os using accurate models is extremely helpful for finding and fixing FPGA I/O timing and board signal integrity issues before any boards are built. However, the usefulness of such simulations is directly related to the accuracy of the models used and whether the simulations are set up and performed correctly. To ensure accuracy in models and simulations created for FPGA output signals, the timing hand-off between t_{CO} timing in the Quartus II software and simulation-based board delay must be taken into account. If this hand-off is not handled correctly, the calculated delay could either count some of the delay twice or even miss counting some of the delay entirely.

Defining the Double Counting Problem

The double counting problem is inherent to the method output timing is analyzed versus the method used for HSPICE models. The timing analyzer tools in the Quartus II software measure delay timing for an output signal from the core logic of the FPGA design through the output buffer ending at the FPGA pin with a default capacitive load or a specified value for the selected I/O standard. This measurement is the t_{CO} timing variable as shown in Figure 7-10.

Figure 7-10. Double Counting Problem



HSPICE models for board simulation measure t_{PD} (propagation delay) from an arbitrary reference point in the output buffer, through the device pin, out along the board routing, and ending at the signal destination.

It is apparent immediately that if these two delays were simply added together, the delay between the output buffer and the device pin would be counted twice in the calculation. A model or simulation that does not account for this double count would create overly pessimistic simulation results, because the double-counted delay can limit I/O performance artificially. To fix the problem, it might seem that simply subtracting the overlap between t_{CO} and t_{PD} would account for the double count. However, this adjustment would not be accurate because each measurement is based on a different load.

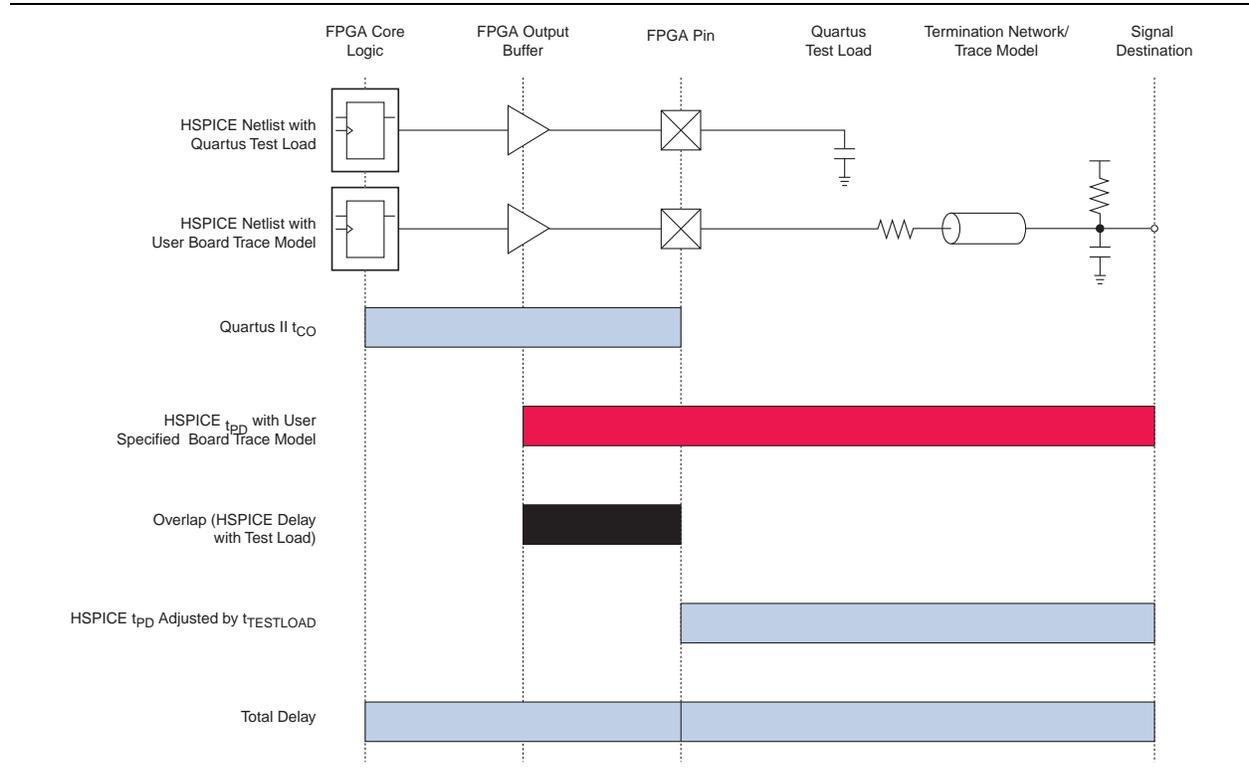


Input signals do not exhibit this problem because the HSPICE models for inputs stop at the FPGA pin instead of at the input buffer. In this case, simply adding the delays together produces an accurate measurement of delay timing.

The Solution to Double Counting

To adjust the measurements to account for the double-counting, the delay between the arbitrary point in the output buffer selected by the HSPICE model and the FPGA pin must be subtracted from either t_{CO} or t_{PD} before adding the results together. The subtracted delay must also be based on a common load between the two measurements. This is done by repeating the HSPICE model measurement, but with the same load used by the Quartus II software for the t_{CO} measurement. This second measurement, called $t_{TESTLOAD}$, is illustrated with the top circuit in Figure 7-11.

Figure 7-11. Common Test Loads Used for Output Timing



With $t_{TESTLOAD}$ known, the total delay for the output signal from the FPGA logic to the signal destination on the board, accounting for the double count, is calculated as shown in Equation 7-1.

Equation 7-1.

$$t_{\text{delay}} = t_{CO} + (t_{PD} - t_{TESTLOAD})$$

The preconfigured simulation files generated by the HSPICE Writer in the Quartus II software are designed to account for the double-counting problem based on this calculation automatically. Performing accurate timing simulations is easy without having to make adjustments for double counting manually.

HSPICE Writer Tool Flow

This section includes information to help you get started using the Quartus II software HSPICE Writer tool. The information in this section assumes you have a basic knowledge of the standard Quartus II software design flow, such as project and assignment creation, compilation, and timing analysis.

 For additional information about standard design flows, refer to the appropriate sections of the *Quartus II Handbook*.

Applying I/O Assignments

The first step in the HSPICE Writer tool flow is to configure the I/O standards and modes for each of the pins in your design properly. In the Quartus II software, these settings are represented by assignments that map I/O settings, such as pin selection, and I/O standard and drive strength, to corresponding signals in your design.

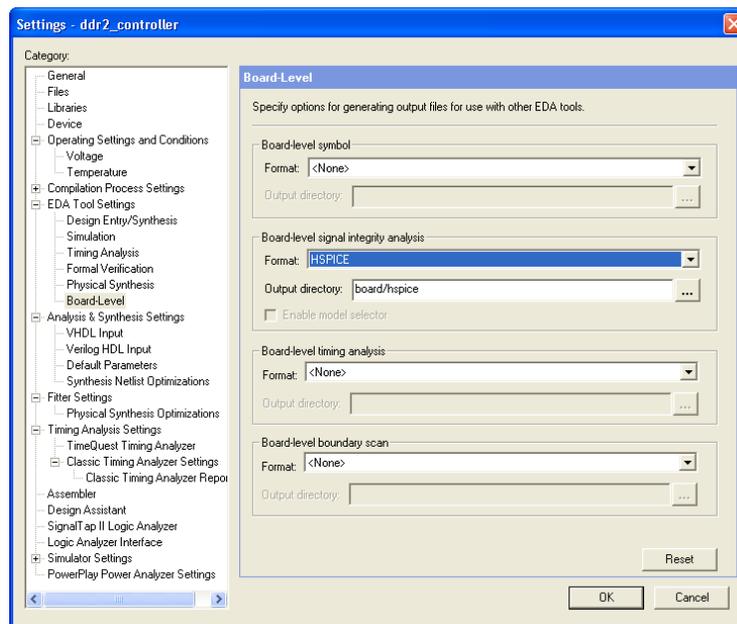
The Quartus II software provides multiple methods for creating these assignments:

- Using the Pin Planner
- Using the assignment editor
- Manually editing the .qsf file
- By making assignments in a scripted Quartus II flow using Tcl

Enabling HSPICE Writer

You must enable the HSPICE Writer in the **Settings** dialog box of the Quartus II software (Figure 7-12) to generate the HSPICE decks from the Quartus II software.

Figure 7-12. EDA Tool Settings: Board Level Options Dialog Box



Enabling HSPICE Writer Using Assignments

You can also use HSPICE Writer in conjunction with a scripted Tcl flow. To enable HSPICE Writer during a full compile, include the lines shown in [Example 7-1](#) in your Tcl script.

Example 7-1. Enable HSPICE Writer

```
set_global_assignment -name EDA_BOARD_DESIGN_SIGNAL_INTEGRITY_TOOL \  
"HSPICE (Signal Integrity)"  
  
set_global_assignment -name EDA_OUTPUT_DATA_FORMAT HSPICE \  
-section_id eda_board_design_signal_integrity  
  
set_global_assignment -name EDA_NETLIST_WRITER_OUTPUT_DIR <output_directory> \  
-section_id eda_board_design_signal_integrity
```

As with command-line invocation, specifying the output directory is optional. If not specified, the output directory defaults to **board/hspice**.

Naming Conventions for HSPICE Files

HSPICE Writer automatically generates simulation files and names them using the following naming convention:

<device>_<pin #>_<pin_name>_<in/out>.sp

For bidirectional pins, two spice decks are produced; one with the I/O buffer configured as an input, and the other with the I/O buffer configured as an output.

The Quartus II software supports alphanumeric pin names that contain the underscore (_) and dash (-) characters. Any illegal characters used in file names are converted automatically to underscores.

The contents of the HSPICE files are described in detail in [“Sample Output for I/O HSPICE Simulation Deck”](#) on page 7-33 and [“Sample Input for I/O HSPICE Simulation Deck”](#) on page 7-28.

Invoking HSPICE Writer

After HSPICE Writer is enabled, the HSPICE simulation files are generated automatically each time the project is completely compiled. The Quartus II software also provides an option to generate a new set of simulation files without having to recompile manually. In the Processing menu, click **Start EDA Netlist Writer** to generate new simulation files automatically.



You must perform both Analysis & Synthesis and Fitting on a design before invoking the HSPICE Writer tool.

Invoking HSPICE Writer from the Command Line

If you use a script-based flow to compile your project, you can create HSPICE model files by including the commands shown in [Example 7-2](#) in your Tcl script (.tcl file).

Example 7-2. Create HSPICE Model Files

```
set_global_assignment -name EDA_BOARD_DESIGN_SIGNAL_INTEGRITY_TOOL \
"HSPICE (Signal Integrity)"

set_global_assignment -name EDA_OUTPUT_DATA_FORMAT HSPICE \
-section_ideda_board_design_signal_integrity

set_global_assignment -name EDA_NETLIST_WRITER_OUTPUT_DIR <output_directory> \
-section_id eda_board_design_signal_integrity
```

The *<output_directory>* option specifies the location where HSPICE model files are saved. By default, the *<project_directory>/board/hspice* directory is used.

To invoke the HSPICE Writer tool through the command line, type the syntax shown in [Example 7-3](#).

Example 7-3. Invoke HSPICE Writer

```
quartus_eda.exe <project_name> --board_signal_integrity=on --format=HSPICE \
--output_directory=<output_directory>
```

<output_directory> specifies the location where the generated spice decks will be written (relative to the design directory). This is an optional parameter and defaults to *board/hspice*.

Customizing Automatically Generated HSPICE Decks

HSPICE models generated by the HSPICE Writer can be used for simulation as generated. A default board description is included, and a default simulation is set up to measure rise and fall delays for both input and output simulations, which compensates for the double counting problem. However, Altera recommends that you customize the board description to more accurately represent your routing and termination scheme.

The sample board trace loading in the generated HSPICE model files must be replaced by your actual trace model before you can run a correct simulation. To do this, open the generated HSPICE model files for all pins you want to simulate and locate the section shown in [Example 7-4](#).

Example 7-4. Sample Board Trace Section

```
* I/O Board Trace and Termination Description
* - Replace this with your board trace and termination description
```

You must replace the example load with a load that matches the design of your PCB board. This includes a trace model, termination resistors, and, for output simulations, a receiver model. The spice circuit node that represents the pin of the FPGA package is called **pin**. The node that represents the far pin of the external device is called **load-in** (for output SPICE decks) and **source-in** (for input SPICE decks).

For an input simulation, you must also modify the stimulus portion of the spice file. The section of the file that must be modified is indicated in the comment block shown in [Example 7-5](#).

Example 7-5. Sample Source Stimulus Section

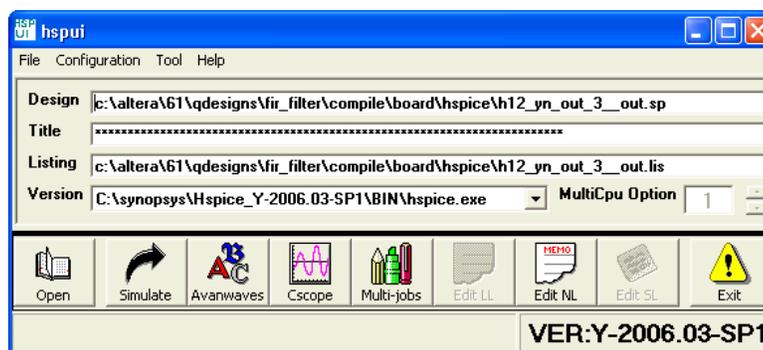
```
* Sample source stimulus placeholder
* - Replace this with your I/O driver model
```

Replace the sample stimulus model with a model for the device that will drive the FPGA.

Running an HSPICE Simulation

Because simulation parameters are configured directly in the HSPICE model files, running a simulation requires only that you open an HSPICE file in the HSPICE user interface and start the simulation. The HSPICE user interface window is shown in [Figure 7-13](#).

Figure 7-13. HSPICE User Interface Window



Click **Open** and browse to the location of the HSPICE model files generated by the Quartus II HSPICE Writer. The default location for HSPICE model files is *<project directory>/board/hspice*. Select the **.sp** file generated by the HSPICE Writer for the signal you want to simulate. Click **OK**.

To run the simulation, click **Simulate**. The status of the simulation is displayed in the window and saved in an **.lis** file with the same name as the **.sp** file when the simulation is complete. Check the **.lis** file if an error occurs during the simulation requiring a change in the **.sp** file to fix.

Interpreting the Results of an Output Simulation

By default, the automatically generated output simulation spice decks are set up to measure three delays for both rising and falling transitions. Two of the measurements, **tpd_rise** and **tpd_fall**, measure the double-counting corrected delay from the FPGA pin to the load pin. To determine the complete clock-edge to load-pin delay, add these numbers to the Quartus II software reported default loading t_{CO} delay.

The remaining four measurements, `tpd_uncomp_rise`, `tpd_uncomp_fall`, `t_dblcnt_rise`, and `t_dblcnt_fall`, are required for the double-counting compensation process and are not required for further timing usage. Refer to “[Simulation Analysis](#)” on page 7-33 for a description of these measurements.

Interpreting the Results of an Input Simulation

By default, the automatically generated input simulation SPICE decks are set up to measure delays from the source’s driver pin to the FPGA’s input pin for both rising and falling transitions. The propagation delay is reported by HSPICE measure statements as `tpd_rise` and `tpd_fall`. To determine the complete source driver pin-to-FPGA register delay, add these numbers to the Quartus II software reported T_H and T_{SU} input timing numbers.

Viewing and Interpreting Tabular Simulation Results

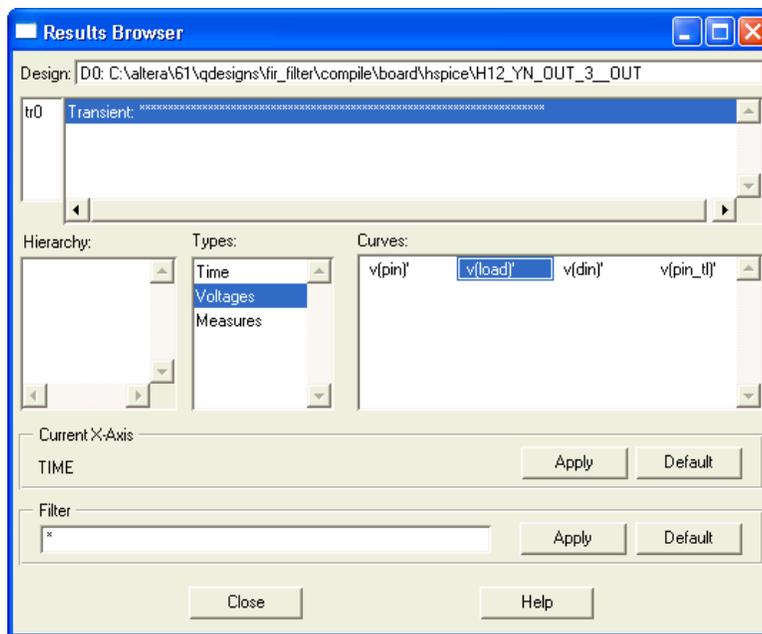
The `.lis` file stores the collected simulation data in tabular form. The default simulation configured by the HSPICE Writer produces delay measurements for rising and falling transitions on both input and output simulations. These measurements are found in the `.lis` file and named `tpd_rise` and `tpd_fall`. For output simulations, these values are already adjusted for the double count. To determine the complete delay from the FPGA logic to the load pin, add either of these measurements to the Quartus II t_{CO} delay. For input simulations, add either of these measurements to the Quartus II t_{SU} and t_H delay values to calculate the complete delay from the far end stimulus to the FPGA logic. Other values found in the `.lis` file, such as `tpd_uncomp_rise`, `tpd_uncomp_fall`, `t_dblcnt_rise`, and `t_dblcnt_fall`, are parts of the double count compensation calculation. These values are not necessary for further analysis.

Viewing Graphical Simulation Results

You can view the results of the simulation quickly as a graphical waveform display using the AvanWaves viewer included with HSPICE. With the default simulation configured by the HSPICE Writer, you can view the simulated waveforms at both the source and destination in input and output simulations.

To see the waveforms for the simulation, in the HSPICE user interface window, click **AvanWaves**. The AvanWaves viewer opens and displays the **Results Browser** as shown in Figure 7-14.

Figure 7-14. HSPICE AvanWaves Results Browser



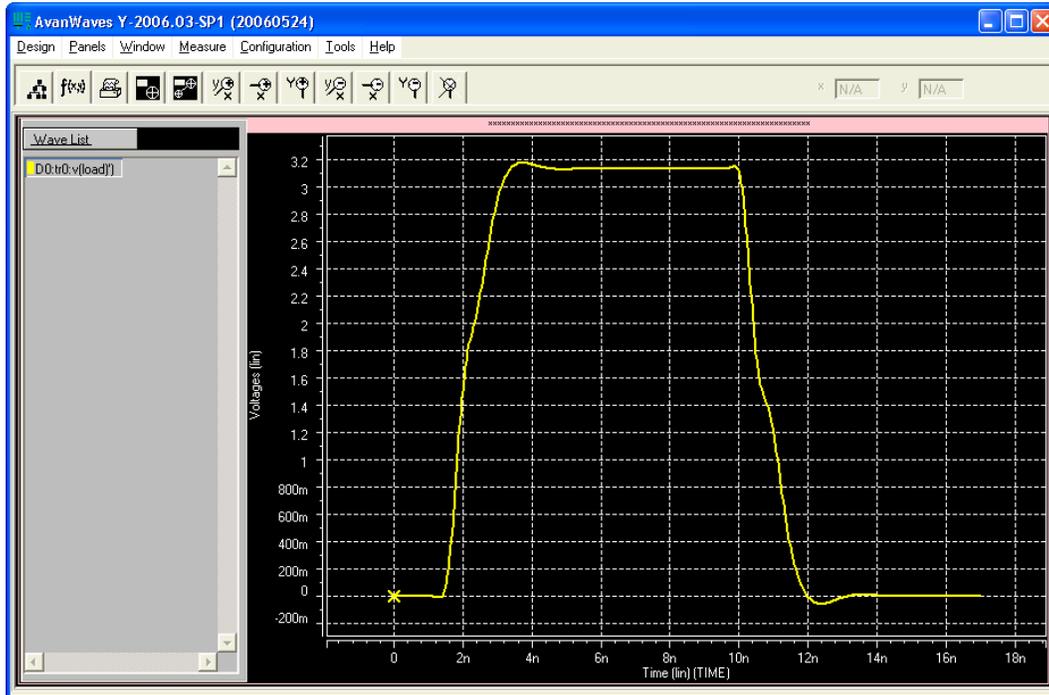
The **Results Browser** lets you select which waveform to view quickly in the main viewing window. If multiple simulations are run on the same signal, the list at the top of the **Results Browser** displays the results of each simulation. Click the simulation description to select which simulation to view. By default, the descriptions are derived from the first line of the HSPICE file, so the description might appear as a line of asterisks.

Select the type of waveform to view, by performing the following steps:

1. To see the source and destination waveforms with the default simulation, from the **Types** list, select **Voltages**.
2. On the **Curves** list, double-click the waveform you want to view. The waveform appears in the main viewing window.

You can zoom in and out and adjust the view as desired (Figure 7-15).

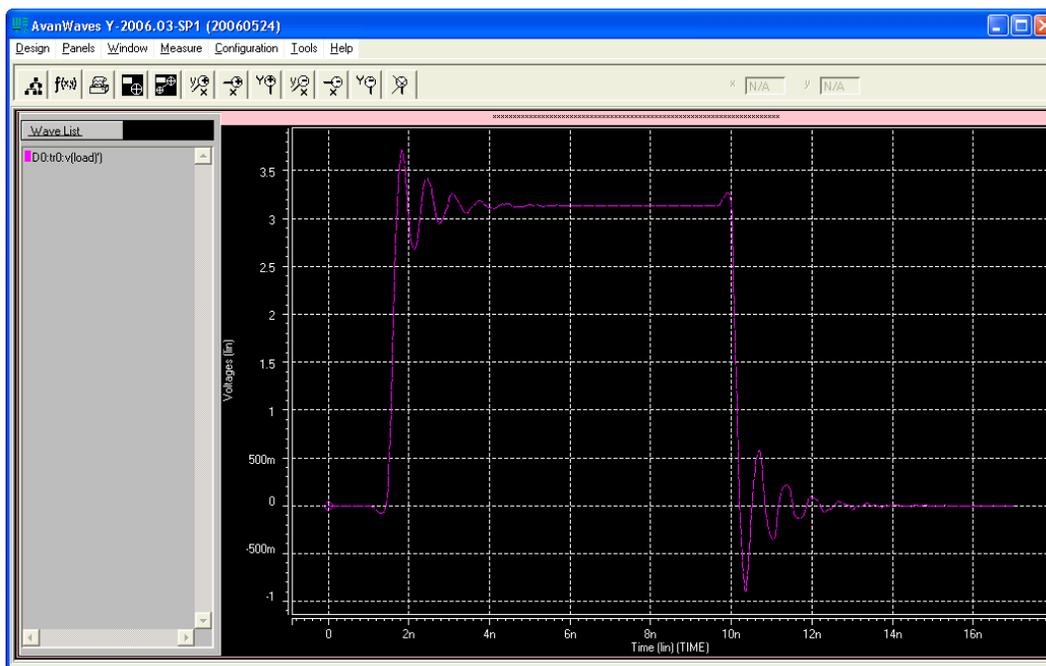
Figure 7-15. AvanWaves Waveform Viewer



Making Design Adjustments Based on HSPICE Simulations

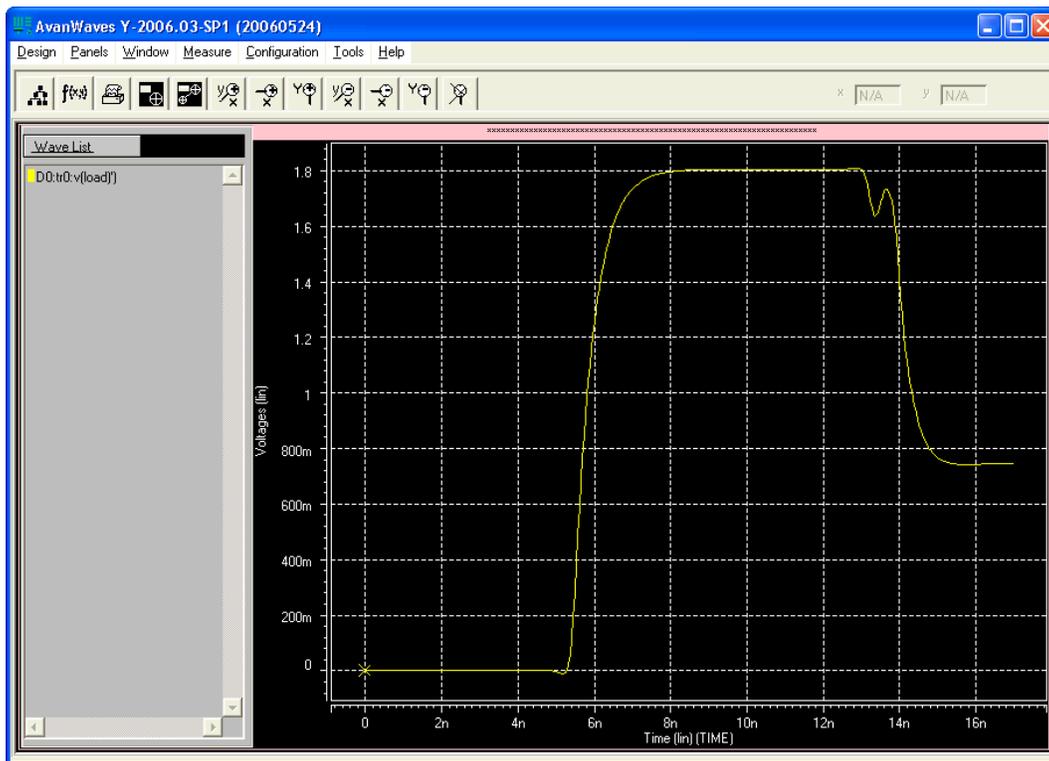
Based on the results of your simulations, you can make adjustments to the I/O assignments or simulation parameters if required. For example, after you run a simulation and see overshoot or ringing in the simulated signal at the destination buffer as shown in the example in Figure 7-16, you can adjust the drive strength I/O assignment setting to a lower value. Regenerate the HSPICE deck, and run the simulation again to verify that the change fixed the problem.

Figure 7-16. Example of Overshoot in the AvanWaves Waveform Viewer



If there is a discontinuity or any other anomalies at the destination as shown in the example in [Figure 7-17](#), adjust the board description in the Quartus II Board Trace Model (for Stratix II, Stratix III, or Cyclone III devices) or in the generated HSPICE model files to change the termination scheme or adjust termination component values. After making these changes, regenerate the HSPICE files if necessary, and rerun the simulation to verify whether your adjustments solved the problem.

Figure 7-17. Example of Signal Integrity Anomaly in the AvanWaves Waveform Viewer



For more information about board-level signal integrity and to learn about ways to improve it with simple changes to your FPGA design, refer to the [Altera Signal Integrity Center](#).

Sample Input for I/O HSPICE Simulation Deck

The following sections examine a typical HSPICE simulation spice deck for an I/O of type input. Each section presents the simulation file one block at a time.

Header Comment

The first block of an input simulation spice deck is the header comment. The purpose of this block is to provide an easily readable summary of how the simulation file has been automatically configured by the Quartus II software.

This block has two main components: The first component summarizes the I/O configuration relevant information such as device, speed grade, and so on. The second component specifies the exact test condition that the Quartus II software assumes for the given I/O standard. [Example 7-6](#) shows a header comment block.

Example 7-6. Header Comment Block

```

* Quartus II HSPICE Writer I/O Simulation Deck*

* This spice simulation deck was automatically generated by
* Quartus for the following IO settings:
*
* Device:          EP2S60F1020C3
* Speed Grade:    C3
* Pin:            AA4 (out96)
* Bank:           IO Bank 6 (Row I/O)
* I/O Standard:   LVTTL, 12mA
* OCT:            Off
*
* Quartus II's default I/O timing delays assume the following slow
* corner simulation conditions.
*
* Specified Test Conditions For Quartus II Tco
* Temperature:    85C (Slowest Temperature Corner)
* Transistor Model: TT (Typical Transistor Corner)
* Vccn:           3.135V (Vccn_min = Nominal - 5%)
* Vccpd:          2.97V (Vccpd_min = Nominal - 10%)
* Load:          No Load
* Vtt:            1.5675V (Voltage reference is Vccn/2)
*
* Note: The I/O transistors are specified to operate at least as
* fast as the TT transistor corner, actual production
* devices can be as fast as the FF corner. Any simulations
* for hold times should be conducted using the fast process
* corner with the following simulation conditions.
*   Temperature:    0C (Fastest Commercial Temperature Corner **)
*   Transistor Model: FF (Fastest Transistor Corner)
*   Vccn:           1.98V (Vccn_hold = Nominal + 10%)
*   Vccpd:          3.63V (Vccpd_hold = Nominal + 10%)
*   Vtt:            0.95V (Vtt_hold = Vccn/2 - 40mV)
*   Vcc:            1.25V (Vcc_hold = Maximum Recommended)
*   Package Model:  Short-circuit from pad to pin (no parasitics)
*
* Warnings:

```

Simulation Conditions

The simulation conditions block loads the appropriate process corner models for the transistors. This condition is automatically set up for the slow timing corner and is modified only if other simulation corners are desired. [Example 7-7](#) shows a simulation conditions block.

Example 7-7. Simulation Conditions Block

```

* Process Settings

.options brief
.inc `sii_tt.inc' * TT process corner

```

Simulation Options

The simulation options block configures the simulation temperature and configures HSPICE with typical simulation options. [Example 7-8](#) shows a simulation options block.

 For a detailed description of these options, consult your *HSPICE* manual.

Example 7-8. Simulation Options Block

```
* Simulation Options

.options brief=0
.options badchr co=132 scale=1e-6 acct ingold=2 nomod dv=1.0
+      dcstep=1 absv=1e-3 absi=1e-8 probe csdf=2 accurate=1
+      converge=1
.temp 85
```

Constant Definition

The constant definition block of the simulation file instantiates the voltage sources that controls the configuration modes of the I/O buffer. [Example 7-9](#) shows a constant definition block.

Example 7-9. Constant Definition Block

```
* Constant Definition

voeb      oeb      0      vc * Set to 0 to enable buffer output
vopdrain  opdrain  0      0  * Set to vc to enable open drain
vrambh    rambh    0      0  * Set to vc to enable bus hold
vrpullup  rpullup  0      0  * Set to vc to enable weak pullup
vpcdp5    rpcdp5   0      rp5 * Set the IO standard
vpcdp4    rpcdp4   0      rp4
vpcdp3    rpcdp3   0      rp3
vpcdp2    rpcdp2   0      rp2
vpcdp1    rpcdp1   0      rp1
vpcdp0    rpcdp0   0      rp0
vpcdn4    rpcdn4   0      rn4
vpcdn3    rpcdn3   0      rn3
vpcdn2    rpcdn2   0      rn2
vpcdn1    rpcdn1   0      rn1
vpcdn0    rpcdn0   0      rn0
vdin din      0      0
```

Where:

- Voltage source `voeb` controls the output enable of the buffer and is set to *disabled* for inputs.
- `vopdrain` controls the open drain mode for the I/O.
- `vrambh` controls the bus hold circuitry in the I/O.
- `vrpullup` controls the weak pullup.
- The next 11 voltages sources control the I/O standard of the buffer and are configured through a later library call.
- `vdin` is not used on input pins because it is the data pin for the output buffer.

Buffer Netlist

The buffer netlist block (Example 7-10) of the simulation spice deck loads all the load models required for the corresponding input pin.

Example 7-10. Buffer Netlist Block

```
* IO Buffer Netlist  
  
.include 'vio_buffer.inc'
```

Drive Strength

The drive strength block (Example 7-11) of the simulation SPICE deck loads the configuration bits necessary to configure the I/O into the proper I/O standard and drive strengths. Although these settings are not relevant to an input buffer, they are provided to allow the SPICE deck to be modifiable to support bidirectional simulations.

Example 7-11. Drive Strength Block

```
* Drive Strength Settings  
  
.lib 'drive_select_hio.lib' 3p3ttl_12ma
```

I/O Buffer Instantiation

The I/O buffer instantiation block of the simulation SPICE deck instantiates the necessary power supplies and I/O model components that are necessary to simulate the given I/O.

Example 7-12 shows I/O buffer instantiation.

Example 7-12. I/O Buffer Instantiation

```
I/O Buffer Instantiation

* Supply Voltages Settings
.param vcn=3.135
.param vpd=2.97
.param vc=1.15

* Instantiate Power Supplies|
vcc      vcc      0      vc      * FPGA core voltage
vss      vss      0      0      * FPGA core ground
vccn     vccn     0      vcn     * IO supply voltage
vssn     vssn     0      0      * IO ground
vccpd    vccpd    0      vpd     * Pre-drive supply voltage

* Instantiate I/O Buffer
xvio_buf din oeb opdrain die rambh
+ rpcdn4 rpcdn3 rpcdn2 rpcdn1 rpcdn0
+ rpcdp5 rpcdp4 rpcdp3 rpcdp2 rpcdp1 rpcdp0
+ rpullup vccn vccpd vcpad0 vio_buf

* Internal Loading on Pad
* - No loading on this pad due to differential buffer/support
*   circuitry

* I/O Buffer Package Model
* - Single-ended I/O standard on a Row I/O
.lib 'lib/package.lib' hio
xpkg die pin hio_pkg
```

Board Trace and Termination

The board trace and termination block of the simulation SPICE deck is provided only as an example (shown in Example 7-13). Replace this block with your own board trace and termination models.

Example 7-13. Board Trace and Termination Block

```
* I/O Board Trace and Termination Description
* - Replace this with your board trace and termination description

wtline pin vssn load vssn N=1 L=1 RLGCMODEL=tlinemodel
.MODEL tlinemodel W MODELTYPE=RLGC N=1 Lo=7.13n Co=2.85p
Rterm2 load vssn 1x
```

Stimulus Model

The stimulus model block of the simulation spice deck is provided only as a placeholder example (shown in Example 7-14). Replace this block with your own stimulus model. Options for this include an IBIS or HSPICE model, among others.

Example 7-14. Stimulus Model Block

```
* Sample source stimulus placeholder
* - Replace this with your I/O driver model

Vsource source 0 pulse(0 vcn 0s 0.4ns 0.4ns 8.5ns 17.4ns)
```

Simulation Analysis

The simulation analysis block (Example 7–15) of the simulation file is configured to measure the propagation delay from the source to the FPGA pin. Both the source and end point of the delay are referenced against the 50% V_{CCN} crossing point of the waveform.

Example 7–15. Simulation Analysis Block

```
* Simulation Analysis Setup

* Print out the voltage waveform at both the source and the pin
.print tran v(source) v(pin)
.tran 0.020ns 17ns

* Measure the propagation delay from the source pin to the pin
* referenced against the 50% voltage threshold crossing point

.measure TRAN tpd_rise TRIG v(source) val='vcn*0.5' rise=1
+ TARG v(pin) val = 'vcn*0.5' rise=1
.measure TRAN tpd_fall TRIG v(source) val='vcn*0.5' fall=1
+ TARG v(pin) val = 'vcn*0.5' fall=1
```

Sample Output for I/O HSPICE Simulation Deck

The following sections examine a typical HSPICE simulation SPICE deck for an I/O-type output. Each section presents the simulation file one block at a time.

Header Comment

The first block of an output simulation SPICE deck is the header comment, as shown in Example 7–16. The purpose of this block is to provide a readable summary of how the simulation file has been automatically configured by the Quartus II software.

This block has two main components:

- The first component summarizes the I/O configuration relevant information such as device, speed grade, and so on.
- The second component specifies the exact test condition that the Quartus II software assumes when generating t_{CO} delay numbers. This information is used as part of the double-counting correction circuitry contained in the simulation file.

The SPICE decks are preconfigured to calculate the slow process corner delay but can also be used to simulate the fast process corner as well. The fast corner conditions are listed in the header under the notes section.

The final section of the header comment lists any warning messages that you must consider when you use the SPICE decks.

Example 7-16. Header Comment Block

```
* Quartus II HSPICE Writer I/O Simulation Deck
*
* This spice simulation deck was automatically generated by
* Quartus II for the following IO settings:
*
* Device:          EP2S60F1020C3
* Speed Grade:    C3
* Pin:            AA4 (out96)
* Bank:           IO Bank 6 (Row I/O)
* I/O Standard:   LVTTL, 12mA
* OCT:            Off
*
* Quartus' default I/O timing delays assume the following slow
* corner simulation conditions.
* Specified Test Conditions For Quartus II Tco
* Temperature:    85C (Slowest Temperature Corner)
* Transistor Model: TT (Typical Transistor Corner)
* Vccn:           3.135V (Vccn_min = Nominal - 5%)
* Vccpd:          2.97V (Vccpd_min = Nominal - 10%)
* Load:          No Load
* Vtt:            1.5675V (Voltage reference is Vccn/2)
* For C3 devices, the TT transistor corner provides an
* approximation for worst case timing. However, for functionality
* simulations, it is recommended that the SS corner be simulated
* as well.
*
* Note: The I/O transistors are specified to operate at least as
* fast as the TT transistor corner, actual production
* devices can be as fast as the FF corner. Any simulations
* for hold times should be conducted using the fast process
* corner with the following simulation conditions.
* Temperature:    0C (Fastest Commercial Temperature Corner
** )
* Transistor Model: FF (Fastest Transistor Corner)
* Vccn:           1.98V (Vccn_hold = Nominal + 10%)
* Vccpd:          3.63V (Vccpd_hold = Nominal + 10%)
* Vtt:            0.95V (Vtt_hold = Vccn/2 - 40mV)
* Vcc:            1.25V (Vcc_hold = Maximum Recommended)
* Package Model:  Short-circuit from pad to pin
* Warnings:
```

Simulation Conditions

The simulation conditions block (Example 7-17) loads the appropriate process corner models for the transistors. This condition is automatically set up for the slow timing corner and must be modified only if other simulation corners are desired.



Two separate corners cannot be simulated at the same time. Instead, simulate the base case using the Quartus corner as one simulation and then perform a second simulation using the desired customer corner. The results of the two simulations can be manually added together.

Example 7-17. Simulation Conditions Block

```
* Process Settings

.options brief
.inc 'sii_tt.inc' * typical-typical process corner
```

Simulation Options

The simulation options block (Example 7-18) configures the simulation temperature and configures HSPICE with typical simulation options.



For a detailed description of these options, consult your *HSPICE* manual.

Example 7-18. Simulation Options Block

```
* Simulation Options
.options brief=0
.options badchr co=132 scale=1e-6 acct ingold=2 nomod dv=1.0
+      dcstep=1 absv=1e-3 absi=1e-8 probe csdf=2 accurate=1
+      converge=1
.temp 85
```

Constraint Definition

The constant definition block (Example 7-19) of the output simulation SPICE deck instantiates the voltage sources that controls the configuration modes of the I/O buffer.

Example 7-19. Constant Definition Block

```
* Constant Definition
voeb      oeb      0      0 * Set to 0 to enable buffer output
vopdrain  opdrain  0      0 * Set to vc to enable open drain
vrambh    rambh    0      0 * Set to vc to enable bus hold
vrpullup  rpullup  0      0 * Set to vc to enable weak pullup
vpci      rpci     0      0 * Set to vc to enable pci mode
vpcdp4    rpcdp4   0      rp4 * These control bits set the IO standard
vpcdp3    rpcdp3   0      rp3
vpcdp2    rpcdp2   0      rp2
vpcdp1    rpcdp1   0      rp1
vpcdp0    rpcdp0   0      rp0
vpcdn4    rpcdn4   0      rn4
vpcdn3    rpcdn3   0      rn3
vpcdn2    rpcdn2   0      rn2
vpcdn1    rpcdn1   0      rn1
vpcdn0    rpcdn0   0      rn0
vdin      din      0      pulse(0 vc 0s 0.2ns 0.2ns 8.5ns 17.4ns)
```

Where:

- Voltage source voeb controls the output enable of the buffer.
- vopdrain controls the open drain mode for the I/O.
- vrambh controls the bus hold circuitry in the I/O.
- vrpullup controls the weak pullup.
- vpci controls the PCI clamp.
- The next ten voltage sources control the I/O standard of the buffer and are configured through a later library call. Stratix III and Cyclone III devices have more bits and so might have more voltage sources listed in the constant definition block. They also have slew rate and delay chain settings.
- vdn is connected to the data input of the I/O buffer.
- The edge rate of the input stimulus is automatically set to the correct value by the Quartus II software.

I/O Buffer Netlist

The I/O buffer netlist block (Example 7-20) loads all of the models required for the corresponding pin. These include a model for the I/O output buffer, as well as any loads that might be present on the pin.

Example 7-20. I/O Buffer Netlist Block

```
*IO Buffer Netlist
.include 'hio_buffer.inc'
.include 'lvds_input_load.inc'
.include 'lvds_oct_load.inc'
```

Drive Strength

The drive strength block (Example 7-21) of the simulation spice deck loads the configuration bits for configuring the I/O to the proper I/O standard and drive strength. These options are set by the HSPICE Writer tool and are not changed for expected use.

Example 7-21. Drive Strength Block

```
* Drive Strength Settings  
  
.lib 'drive_select_hio.lib' 3p3ttl_12ma
```

Slew Rate and Delay Chain

Stratix III and Cyclone III devices have sections for configuring the slew rate and delay chain settings (Example 7-22).

Example 7-22. Slew Rate and Delay Chain Settings

```
* Programmable Output Delay Control Settings  
  
.lib 'lib/output_delay_control.lib' no_delay  
  
* Programmable Slew Rate Control Settings  
  
.lib 'lib/slew_rate_control.lib' slow_slow
```

I/O Buffer Instantiation

The I/O buffer instantiation block (Example 7–23) of the output simulation spice deck instantiates the necessary power supplies and I/O model components that are necessary to simulate the given I/O.

Example 7–23. I/O Buffer Instantiation Block

```

* I/O Buffer Instantiation

* Supply Voltages Settings
.param vcn=3.135
.param vpd=2.97
.param vc=1.15

* Instantiate Power Supplies
vvcv      vcc      0      vc      * FPGA core voltage
vvss      vss      0      0      * FPGA core ground
vvcn      vccn     0      vcn     * IO supply voltage
vvsn      vssn     0      0      * IO ground
vvcvcpd   vvcvcpd  0      vpd     * Pre-drive supply voltage

* Instantiate I/O Buffer
xhio_buf din oeb opdrain die rambh
+ rpcdn4 rpcdn3 rpcdn2 rpcdn1 rpcdn0
+ rpcdp4 rpcdp3 rpcdp2 rpcdp1 rpcdp0
+ rpullup vccn vvcvcpd vcpad0 hio_buf

* Internal Loading on Pad
* - This pad has an LVDS input buffer connected to it, along
*   with differential OCT circuitry. Both are disabled but
*   introduce loading on the pad that is modeled below.
xlvs_input_load die vss vccn lvds_input_load
xlvs_oct_load die vss vvcvcpd vccn vcpad0 vccn lvds_oct_load

* I/O Buffer Package Model
* - Single-ended I/O standard on a Row I/O
.lib 'lib/package.lib' hio
xpkg die pin hio_pkg

```

Board and Trace Termination

The board trace and termination block (Example 7–24) of the simulation SPICE deck is provided only as an example. Replace this block with your specific board loading models.

Example 7–24. Board Trace and Termination Block

```

* I/O Board Trace And Termination Description
* - Replace this with your board trace and termination description

wtline pin vssn load vssn N=1 L=1 RLGCMODEL=tlinemodel
.MODEL tlinemodel W MODELTYPE=RLGC N=1 Lo=7.13n Co=2.85p
Rterm2 load vssn 1x

```

Double-Counting Compensation Circuitry

The double-counting compensation circuitry block (Example 7-25) of the simulation SPICE deck instantiates a second I/O buffer that is used to measure double-counting. The buffer is configured identically to the user I/O buffer but is connected to the Quartus II software test load. The simulated delay of this second buffer can be interpreted as the amount of double-counting between the Quartus II software and HSPICE Writer simulated results.

As the amount of double-counting is constant for a given I/O standard on a given pin, consider separating the double-counting circuitry from the simulation file. In doing so, you can perform any number of I/O simulations while referencing the delay only once. For more information about the double-counting problem, refer to “The Double Counting Problem in HSPICE Simulations” on page 7-17.

Example 7-25. (Part 1 of 2) Double-Counting Compensation Circuitry Block

```
* Double Counting Compensation Circuitry
*
* The following circuit is designed to calculate the amount of
* double counting between Quartus II and the HSPICE models. If
* you have not changed the default simulation temperature or
* transistor corner the double counting will be automatically
* compensated by this spice deck. In the event you wish to
* simulate an IO at a different temperature or transistor corner
* you will need to remove this section of code and manually
* account for double counting. A description of Altera's
* recommended procedure for this process can be found in the
* Quartus II HSPICE Writer AppNote.
*

* Supply Voltages Settings
.param vcn_tl=3.135
.param vpd_tl=2.97

* Test Load Constant Definition
vopdrain_tl  opdrain_tl  0    0
vrambh_tl   rambh_tl   0    0
vrpullup_tl rpullup_tl 0    0

* Instantiate Power Supplies
vccn_tl     vccn_tl     0    vcn_tl
vvssn_tl    vssn_tl    0    0
vccpd_tl    vccpd_tl    0    vpd_tl

* Instantiate I/O Buffer
xhio_testload din oeb opdrain_tl die_tl rambh_tl
+ rpcdn4 rpcdn3 rpcdn2 rpcdn1 rpcdn0
+ rpcdp4 rpcdp3 rpcdp2 rpcdp1 rpcdp0
+ rpullup_tl vccn_tl vccpd_tl vcpad0_tl hio_buf

* Internal Loading on Pad
xlvds_input_testload die_tl vss vccn_tl lvds_input_load
xlvds_oct_testload die_tl vss vccpd_tl vccn_tl vcpad0_tl vccn_tl
lvds_oct_load
```

Example 7-25. (Part 2 of 2) Double-Counting Compensation Circuitry Block

```
* I/O Buffer Package Model
* - Single-ended I/O standard on a Row I/O
.lib 'lib/package.lib' hio
xpkg die pin hio_pkg

* Default Altera Test Load
* - 3.3V LVTTTL default test condition is an open load
```

Simulation Analysis

The simulation analysis block (Example 7-26) is set up to measure double-counting corrected delays. This is accomplished by measuring the uncompensated delay of the I/O buffer when connected to the user load, and when subtracting the simulated amount of double-counting from the test load I/O buffer.

Example 7-26. Simulation Analysis Block

```
*Simulation Analysis Setup

* Print out the voltage waveform at both the pin and far end load
.print tran v(pin) v(load)
.tran 0.020ns 17ns

* Measure the propagation delay to the load pin. This value will
* include some double counting with Quartus II's Tco
.measure TRAN tpd_uncomp_rise TRIG v(din) val='vc*0.5' rise=1
+ TARG v(load) val='vcn*0.5' rise=1
.measure TRAN tpd_uncomp_fall TRIG v(din) val='vc*0.5' fall=1
+ TARG v(load) val='vcn*0.5' fall=1

* The test load buffer can calculate the amount of double counting
.measure TRAN t_dbcnt_rise TRIG v(din) val='vc*0.5' rise=1
+ TARG v(pin_t1) val='vcn_t1*0.5' rise=1
.measure TRAN t_dbcnt_fall TRIG v(din) val='vc*0.5' fall=1
+ TARG v(pin_t1) val='vcn_t1*0.5' fall=1

* Calculate the true propagation delay by subtraction
.measure TRAN tpd_rise PARAM='tpd_uncomp_rise-t_dbcnt_rise'
.measure TRAN tpd_fall PARAM='tpd_uncomp_fall-t_dbcnt_fall'
```

Advanced Topics

The information in this section describes some of the more advanced topics and methods employed when setting up and running HSPICE simulation files.

PVT Simulations

The automatically generated HSPICE simulation files are set up to simulate the slow process corner using low voltage, high temperature, and slow transistors. To ensure a fully robust link, Altera recommends that you run simulations over all process corners.

To perform process, voltage, and temperature (PVT) simulations, manually modify the spice decks in a two step process:

1. Remove the double-counting compensation circuitry from the simulation file. This is required as the amount of double-counting is dependant upon how the Quartus II software calculates delays and is not based on which PVT corner is being simulated. By default, the Quartus II software provides timing numbers using the slow process corner.
2. Select the proper corner for the PVT simulation by setting the correct HSPICE temperature, changing the supply voltage sources, and loading the correct transistor models.

A more detailed description of HSPICE process corners can be found in the family-specific HSPICE model documentation. This document is available online with the HSPICE models as described in [“Accessing HSPICE Simulation Kits” on page 7–17](#).

Hold Time Analysis

Altera recommends performing worst-case hold time analysis using the fast corner models, which use fast transistors, high voltage, and low temperature. This involves modifying the SPICE decks to select the correct temperature option, change the supply voltage sources, and load the correct fast transistor models. The values of these parameters are located in the header comment section of the corresponding simulation deck files.

For a truly worst-case analysis, combine the HSPICE Writer hold time analysis results with the Quartus II software fast timing model. This requires that you change the double-counting compensation circuitry in the simulations files to also simulate the fast process corners, as this is what the Quartus II software uses for the fast timing model.



This method of hold time analysis is recommended only for globally synchronous buses. Do not apply this method of hold-time analysis to source synchronous buses. This is because the source synchronous clocking scheme is designed to cancel out some of the PVT timing effects. If this is not taken into account, the timing results will not be accurate. Proper source synchronous timing analysis is beyond the scope of this document.

I/O Voltage Variations

Use each of the FPGA family datasheets to verify the recommended operating conditions for supply voltages. For current FPGA families, the maximum recommended voltage corresponds to the fast corner, while the minimum recommended voltage corresponds to the slow corner. These voltage recommendations are specified at the power pins of the FPGA and are not necessarily the same voltage that are seen by the I/O buffers due to package IR drops.

The automatically generated HSPICE simulation files model this IR effect pessimistically by including a 50-mV IR drop on the V_{CCPD} supply when a high drive strength standard is being used.

Correlation Report

Correlation reports for the HSPICE I/O models are located in the family-specific HSPICE I/O buffer simulation kits. Refer to [“Accessing HSPICE Simulation Kits” on page 7–17](#) for additional information.

Conclusion

As FPGA devices are used in more high-speed applications, it becomes increasingly necessary to perform board-level signal integrity analysis simulations. You must run such simulations to ensure good signal integrity between the FPGA and any connected devices. The Quartus II software helps to simplify this process with the ability to automatically generate I/O buffer description models easily with the IBIS and HSPICE Writers. IBIS models can be integrated into a third-party signal integrity analysis workflow using a tool such as Mentor Graphics HyperLynx software, generating quick and accurate simulation results. HSPICE decks include preconfigured simulations and only require descriptions of board routing and stimulus models to create highly accurate simulation results using Synopsys HSPICE. Either type of simulation helps prevent unnecessary board spins, increasing your productivity and decreasing your costs.

Document Revision History

Table 7-3 shows the revision history for this chapter.

Table 7-3. Document Revision History

Date	Version	Changes
December 2010	10.0.1	Template update.
July 2010	10.0.0	Updated device support.
November 2009	9.1.0	No change to content.
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Was volume 3, chapter 12 in the 8.1.0 release. ■ No change to content.
November 2008	8.1.0	<ul style="list-style-type: none"> ■ Changed to 8-1/2 x 11 page size. ■ Added information for Stratix III devices. ■ Input signals for Cyclone III devices are supported.
May 2008	8.0.0	<ul style="list-style-type: none"> ■ Updated “Introduction” on page 12-1. ■ Updated Figure 12-1. ■ Updated Figure 12-3. ■ Updated Figure 12-13. ■ Updated “Output File Generation” on page 12-6. ■ Updated “Simulation with HSPICE Models” on page 12-17. ■ Updated “Invoking HSPICE Writer from the Command Line” on page 12-22. ■ Added “Sample Input for I/O HSPICE Simulation Deck” on page 12-29. ■ Added “Sample Output for I/O HSPICE Simulation Deck” on page 12-33. ■ Updated “Correlation Report” on page 12-41. ■ Added hyperlinks to referenced documents and websites throughout the chapter. ■ Made minor editorial updates.

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter discusses how the Quartus® II software interacts with the Mentor Graphics® I/O Designer software and the DxDesigner software to provide a complete FPGA-to-board design workflow.

With today's large, high-pin-count and high-speed FPGA devices, good and correct PCB design practices are essential to ensure correct system operation. The PCB design takes place concurrently with the design and programming of the FPGA. The FPGA or ASIC designer initially creates signal and pin assignments, and the board designer must correctly transfer these assignments to the symbols in their system circuit schematics and board layout. As the board design progresses, Altera recommends reassigning pins to optimize the PCB layout. Ensure that you inform the FPGA designer of the pin reassignments so that the new assignments are included in an updated placement and routing of the design.

The Mentor Graphics I/O Designer software allows you to take advantage of the full FPGA symbol design, creation, editing, and back-annotation flow supported by the Mentor Graphics tools.

This chapter covers the following topics:

- Performing design flow between the Quartus II software, the Mentor Graphics I/O Designer software, and the DxDesigner software
- Setting up the Quartus II software to create the design flow files
- Creating an I/O Designer database project to incorporate the Quartus II software signal and pin assignment data
- Updating signal and pin assignment changes between the I/O Designer software and the Quartus II software
- Generating symbols in the I/O Designer software
- Creating symbols in the DxDesigner software from the Quartus II software output files without the use of the I/O Designer software

This chapter is intended for board design and layout engineers who want to start the FPGA board integration while the FPGA is still in the design phase. Alternatively, the board designer can plan the FPGA pin-out and routing requirements in the Mentor Graphics tools and pass the information back to the Quartus II software for placement and routing. Part librarians can also benefit from this chapter by learning how to use output from the Quartus II software to create new library parts and symbols.

The procedures in this chapter require the following software:

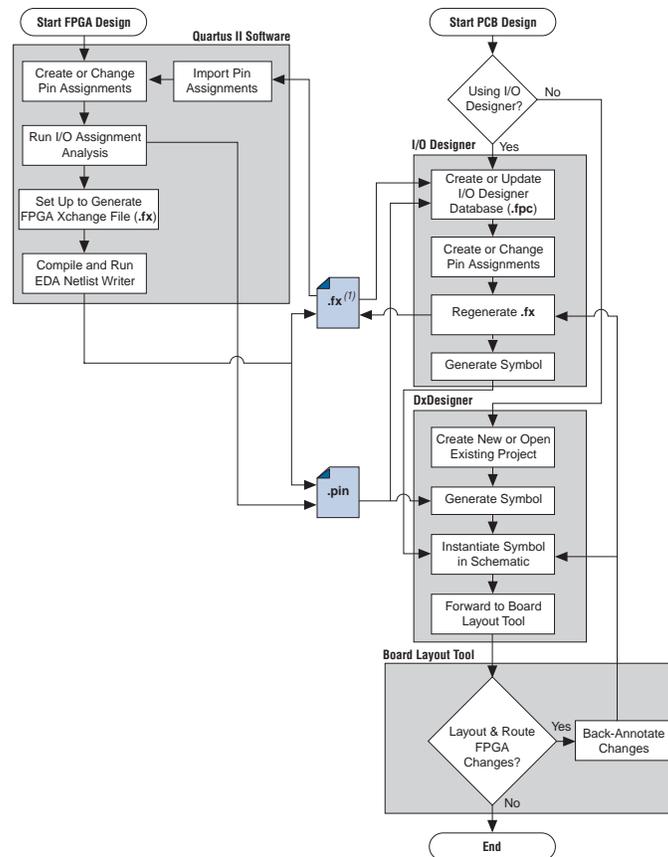
- The Quartus II software version 5.1 or later
- DxDesigner software version 2004 or later
- Mentor Graphics I/O Designer software (optional)

To obtain and license the Mentor Graphics tools and for product information, support, and training, refer to the Mentor Graphics website (www.mentor.com).

FPGA-to-PCB Design Flow

You can create a design flow integrating an Altera® FPGA design from the Quartus II software, and a circuit schematic in the DxDesigner software. Figure 8-1 shows the design flow with and without the I/O Designer software.

Figure 8-1. Design Flow with and Without the I/O Designer Software



Note to Figure 8-1:

- (1) The Quartus II software generates the **.fx** in the output directory you specify in the **Board-Level** page of the **Settings** dialog box. However, the Quartus II software and the I/O Designer software can import pin assignments from an **.fx** located in any directory. Altera recommends working with a backup **.fx** to prevent overwriting existing assignments or importing invalid assignments.

To perform the design flow shown in Figure 8-1, follow these steps:

1. In the Quartus II software, set up the board-level assignment settings to generate an **.fx** for symbol generation.
2. Compile your design to generate the **.fx** and Pin-Out File (**.pin**). You can locate the generated **.fx** and **.pin** files in the Quartus II project directory.

3. Create a board design with the DxDesigner software and the I/O Designer software by performing the following steps:
 - a. Create a new I/O Designer database based on the **.fx** and the **.pin** files.
 - b. In the I/O Designer software, make adjustments to signal and pin assignments.
 - c. Regenerate the **.fx** in the I/O Designer software to export the I/O Designer software changes to the Quartus II software.
 - d. Generate a single or fractured symbol for use in the DxDesigner software.
 - e. Add the symbol to the **sym** directory of a DxDesigner project, or specify a new DxDesigner project with the new symbol.
 - f. Instantiate the symbol in your DxDesigner schematic and export the design to the board layout tool.
 - g. Back-annotate pin changes created in the board layout tool to the DxDesigner software and back to the I/O Designer software and the Quartus II software.
4. Create a board design with the DxDesigner software without the I/O Designer software by performing the following steps:
 - a. Create a new DxBoardLink symbol with the **Symbol** wizard and reference the **.pin** from the Quartus II software in an existing DxDesigner project.
 - b. Instantiate the symbol in your DxDesigner schematic and export the design to a board layout tool.



You can update these symbols with design changes with or without the I/O Designer software. If you use the Mentor Graphics I/O Designer software and you change symbols with the DxDesigner software, you must reimport the symbols into I/O Designer to avoid overwriting your symbol changes.

Performing Simultaneous Switching Noise (SSN) Analysis of Your FPGA

With the Quartus II software, you can extract pin assignment data and perform SSN analysis of your design for designs targeting the Stratix III device family. You can perform SSN analysis early in the board layout stage as part of your overall pin planning process; however, you do not have to perform SSN analysis to generate pin assignment data from the Quartus II software. You can use the SSN Analyzer tool in the Quartus II software to optimize the pin assignments for better SSN performance.



For more information about the SSN Analyzer, refer to the *Simultaneous Switching Noise (SSN) Analysis and Optimizations* chapter in volume 2 of the *Quartus II Handbook* and *About the SSN Analyzer* in Quartus II Help.

Setting Up the Quartus II Software

You can transfer pin and signal assignments from the Quartus II software to the Mentor Graphics tools by generating `.pin` and `.fx` files (refer to [Figure 8-2](#)). The `.pin` is an output file generated by the Quartus II Fitter that contains pin assignment information. You can use the Quartus II Pin Planner to set and change the assignments contained in the `.pin` and then transfer the assignments to the Mentor Graphics tools. You cannot, however, import pin assignment changes from the Mentor Graphics tools into the Quartus II software with the `.pin`.

The `.pin` lists all used and unused pins on your selected Altera device. It also provides the following basic information fields for each assigned pin on the device:

- Pin signal name and usage
- Pin number
- Signal direction
- I/O standard
- Voltage
- I/O bank
- User or Fitter-assigned

The `.fx` is an input/output file generated by the Quartus II software and the I/O Designer software that can be imported and exported from both programs. The `.fx` generated by the Quartus II software lists only assigned pins and provides the following advanced information fields for each pin on a device:

- Pin number
- I/O bank
- Signal name
- Signal direction
- I/O standard
- Drive strength (mA)
- Termination enabling
- Slew rate
- IOB delay
- Swap group
- Differential pair type

The `.fx` generated by the I/O Designer software lists all pins, including unused pins. In addition to the advanced information fields listed above, the `.fx` generated by the Mentor Graphics I/O Designer software also includes the following information fields:

- Device pin name
- Pin set
- Pin set position

- Pin set group
- Super pin set group
- Super pin set position

 For more information about `.fx` files and the information fields added by the Mentor Graphics software, refer to *FPGA Xchange-Format File (.fx) Definition* in Quartus II Help and Mentor Graphics website (www.mentor.com) respectively.

The I/O Designer software can also read from or update a Quartus II Settings File (`.qsf`). The design flow uses the `.qsf` in a similar manner to the `.fx`, but does not transfer pin swap group information between the I/O Designer software and the Quartus II software.

 Because the `.qsf` contains additional information about your project that the Mentor Graphics I/O Designer software does not use, Altera recommends using the `.fx` instead of the `.qsf`.

 For more information about the `.qsf`, refer to *Quartus II Settings File (.qsf) Definition* in Quartus II Help.

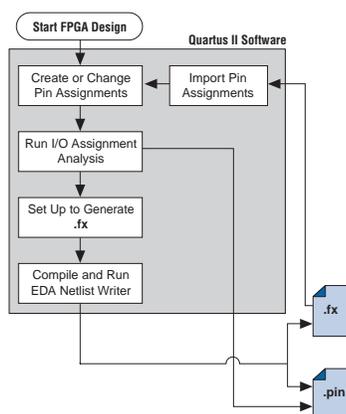
Generating a .pin File

The Quartus II software automatically generates the `.pin` after compiling your FPGA design or during I/O assignment analysis.

To start I/O assignment analysis, on the Processing menu, point to **Start** and then click **Start I/O Assignment Analysis**. The Quartus II Fitter generates the `.pin` and places the file in your Quartus II design directory with the name `<project name>.pin`. The Quartus II software cannot import assignments from an existing `.pin`.

Figure 8-2 shows how to generate `.pin` and `.fx` files.

Figure 8-2. Generating .pin and .fx Files (Note 1)



Note to Figure 8-2:

- (1) For more information about the full design flow, which includes the I/O Designer software, the DxDesigner software, and the board layout tool flowchart details, refer to Figure 8-1.

-  For more information about pin and signal assignment transfer and the files that the Quartus II software can import and export, refer to the *I/O Management* chapter in volume 2 of the *Quartus II Handbook*.

Generating an .fx File

You can generate an .fx in the Quartus II software for symbol generation in the Mentor Graphics I/O Designer software.

-  For more information about generating an .fx, refer to *Generating FPGA Xchange-Format Files for Use with Other EDA Tools* in Quartus II Help.

Creating a Backup .qsf

To create a backup .qsf of your current pin assignments, follow these steps:

1. On the Assignments menu, click **Import Assignments**. The **Import Assignments** dialog box appears.
2. In the **Import Assignments** dialog box, browse to your project and turn on **Copy existing assignments into <project name>.qsf.bak**.
3. Click **OK**.

-  For more information about pin and signal assignment transfer, and files the Quartus II software can import and export, refer to the *I/O Management* chapter in volume 2 of the *Quartus II Handbook*.

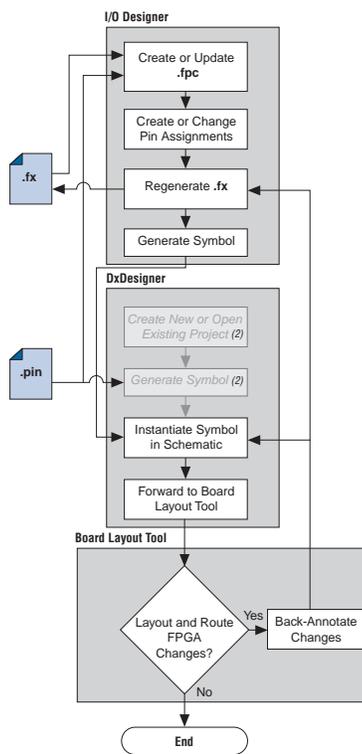
FPGA-to-Board Integration with the I/O Designer Software

The Mentor Graphics I/O Designer software allows you to integrate your FPGA and PCB designs. Pin and signal assignment changes can be made anywhere in the design flow with either the Quartus II Pin Planner or the I/O Designer software. The I/O Designer software facilitates moving these changes, as well as synthesis, placement, and routing changes, between the Quartus II software, an external synthesis tool (if used), and a schematic capture tool such as the DxDesigner software.

This section describes how to use the I/O Designer software to transfer pin and signal assignment information to and from the Quartus II software with an .fx, and how to create symbols for the DxDesigner software.

Figure 8-3 shows the design flow using the I/O Designer software.

Figure 8-3. Design Flow Using the I/O Designer Software (Note 1)



Notes to Figure 8-3:

- (1) For more information about the full design flow including the Quartus II software flowchart details, refer to Figure 8-1 on page 8-2.
- (2) These are DxDesigner software-specific steps in the design flow and are not part of the I/O Designer flow.

For more information about the I/O Designer software, and to obtain usage, support, and product updates, use the Help menu in the I/O Designer software or refer to the Mentor Graphics website (www.mentor.com).

I/O Designer Database Wizard

An .fpc file stores all I/O Designer project information. You can create a new database incorporating information for the .fx and .pin files generated by the Quartus II software using the I/O Designer **Database Wizard**. You can also create a new, empty database and manually add the assignment information. If there is no signal or pin assignment information currently available, you can create an empty database containing only a selection of the target device. This action is useful if you know the signals in your design and the pins you want to assign. You can transfer this information at a later time to the Quartus II software for placement and routing.

You can create an I/O Designer database with only a **.pin** or an **.fx**. However, if you are only using a **.pin**, you cannot import any I/O assignment changes made in the I/O Designer software back into the Quartus II software without first generating an **.fx**. If an **.fx** creates the I/O Designer database, the database may not contain all the available I/O assignment information. The **.fx** generated by the Quartus II software only lists pins with assigned signals. Because the **.pin** lists all device pins—whether signals are assigned to them or not—its use, along with the **.fx**, produces the most complete set of information for creating the I/O Designer database.

If you skip a step in the following process, you can complete the skipped step later. To return to a skipped step, on the Properties menu, click **File**. To create a new I/O Designer database using the **Database** wizard, follow these steps:

1. Start the I/O Designer software. The **Welcome to I/O Designer** dialog box appears. Select **Wizard to create new database** and click **OK**.



If the **Welcome to I/O Designer** dialog box does not appear, you can access the wizard through the menu. To access the wizard, on the File menu, click **Database Wizard**.

2. Click **Next**. The **Define HDL source file** page appears.



If no HDL files are available, or if the **.fx** contains your signal and pin assignments, you can skip Step 3 and proceed to Step 4.



For more information about creating and using HDL files in the Quartus II software, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*, or refer to the I/O Designer Help.

3. If you have created a Verilog HDL or VHDL file in your Quartus II software design, you can add a top-level Verilog HDL or VHDL file in the I/O Designer software. Adding a file allows you to create functional blocks or get signal names from your design. You must create all physical pin assignments in I/O Designer if you are not using an **.fx** or a **.pin**. Click **Next**. The **Database Name** page appears.
4. In the **Database Name** page, type your database file name. Click **Next**. The Database Location window appears.
5. Add a path to the new or an existing database in the **Location** field, or browse to a database location. Click **Next**. The **FPGA flow** page appears.
6. In the Vendor menu, click **Altera**.
7. In the Tool/Library menu, click **Quartus II 5.0**, or a later version of the Quartus II software.
8. Select the appropriate device family, device, package, and speed (if applicable), from the corresponding menus. Click **Next**. The **Place and route** page appears.



The Quartus II software version selections in the Tool/Library menu may not reflect the version of the Quartus II software currently installed in your system even if you are using the latest version of the I/O Designer software. The I/O Designer software uses the version number selection in this window to identify available or obsolete devices in that particular version of the Quartus II software. If you are unsure of the version to select,

use the latest version listed in the menu. If the device you are targeting does not appear in the device menu after making this selection, the device may be new and not yet added to the I/O Designer software. For I/O Designer software updates, contact Mentor Graphics or refer to their website (www.mentor.com).

9. In the **FPGAX file name** field, type or browse to the backup copy of the **.fx** generated by the Quartus II software.
10. In the **Pin report file name** field, type or browse to the **.pin** generated by the Quartus II software. Click **Next**.

You can also select a **.qsf** for update. The I/O Designer software can update the pin assignment information in the **.qsf** without affecting any other information in the file.

 You can select a **.pin** without selecting an **.fx** for import. The I/O Designer software does not generate a **.pin**. To transfer assignment information to the Quartus II software, select an additional file and file type. Altera recommends selecting an **.fx** in addition to a **.pin** for transferring all the assignment information in the **.fx** and **.pin** files.

 In some versions of the I/O Designer software, the standard file picker may incorrectly look for a **.pin** instead of an **.fx**. In this case, select **All Files (*.*)** from the **Save as type** list and select the file from the list.

11. The **Synthesis** page appears. On the **Synthesis** page, you can specify an external synthesis tool and a synthesis constraints file for use with the tool. If you do not use an external synthesis tool, click **Next**.

 For more information about third-party synthesis tools, refer to *Volume 3: Verification* of the *Quartus II Handbook*.

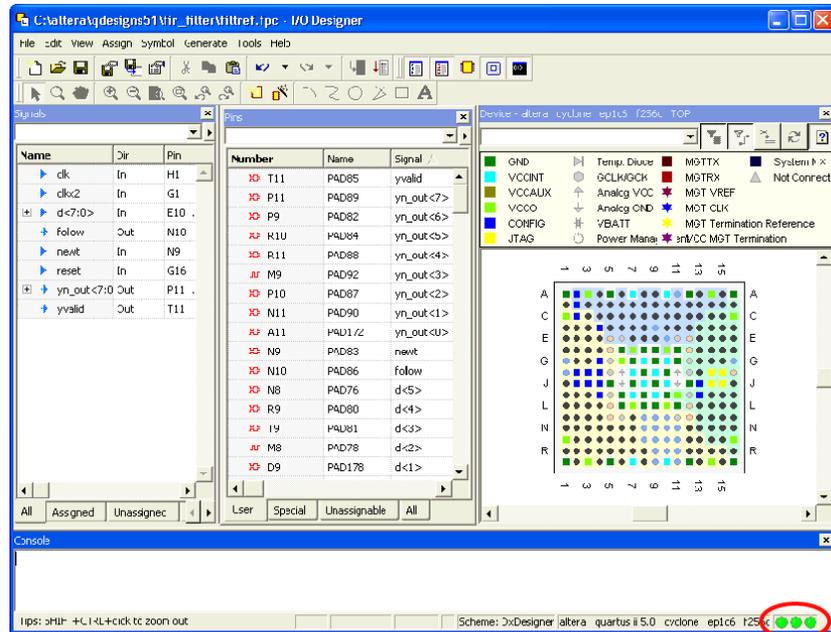
12. On the **PCB Flow** page, you can select an existing schematic project or create a new project as a symbol information destination.
 - To select an existing project, select **Choose existing project** and click **Browse** after the **Project Path** field. The **Select project** dialog box appears. Select the project.
 - To create a new project, in the **Select project** dialog box, select **Create new empty project**. Type the project file name in the **Name** field and browse to the location where you want to save the file. Click **OK**.

If you have not specified a design tool to which you can send symbol information in the I/O Designer software, click **Advanced** in the **PCB Flow** page and select your design tool. If you select the DxDesigner software, you have the option to specify a Hierarchical Occurrence Attributes (**.oat**) file to import into the I/O Designer software. Click **Next** and then click **Finish** to create the database.

 In I/O Designer version 2005 or later, the **Update Wizard** dialog box (refer to [Figure 8-7 on page 8-13](#)) appears if you are creating the database with the **Database** wizard. Use the **Update Wizard** dialog box to confirm creation of the I/O Designer database using the selected **.fx** and **.pin** files.

Use the I/O Designer software and your newly created database to make pin assignment changes, create pin swap groups, or adjust signal and pin properties in the I/O Designer GUI (Figure 8-4).

Figure 8-4. Mentor Graphics I/O Designer Main Window



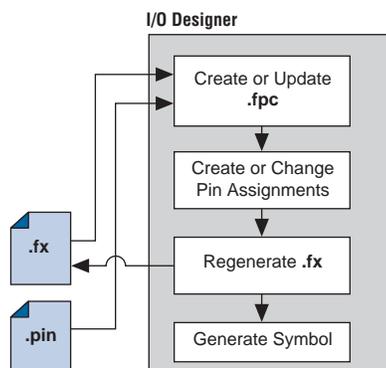
Database Update Indicator

For more information about using the I/O Designer software and the DxDesigner software, refer to the Mentor Graphics website (www.mentor.com) or refer to the I/O Designer software or the DxDesigner Help.

Updating Pin Assignments from the Quartus II Software

As the design process continues, the FPGA designer must make changes to the logic design in the Quartus II software that places signals on different pins after recompiling the design, or manually with the Quartus II Pin Planner. These types of changes must be carried forward to the circuit schematic and board layout tools to ensure that signals connect to the correct pins on the FPGA. Updating the `.fx` and the `.pin` files in the Quartus II software facilitates this flow (Figure 8-5).

Figure 8-5. Updating the I/O Designer Pin Assignments in the Design Flow (Note 1)



Note to Figure 8-5:

- (1) For more information about the full design flow, which includes the Quartus II software, the DxDesigner software, and the board layout tool flowchart details, refer to Figure 8-1 on page 8-2.

To update the `.fx` in your selected output directory and the `.pin` in your project directory after making changes to the design, perform one of the following tasks:

- compile, or
- start EDA Netlist Writer.

You must rerun the I/O Assignment Analyzer whenever you make I/O changes in the Quartus II software. To rerun the I/O Assignment Analyzer, perform one of the following tasks:

- on the Processing menu, click **Start Compilation**, or
- on the Processing menu, click **Start I/O Assignment Analysis**.

 For more information about setting up the `.fx` and running the I/O Assignment Analyzer, refer to the *I/O Management* chapter in volume 2 of the *Quartus II Handbook*.

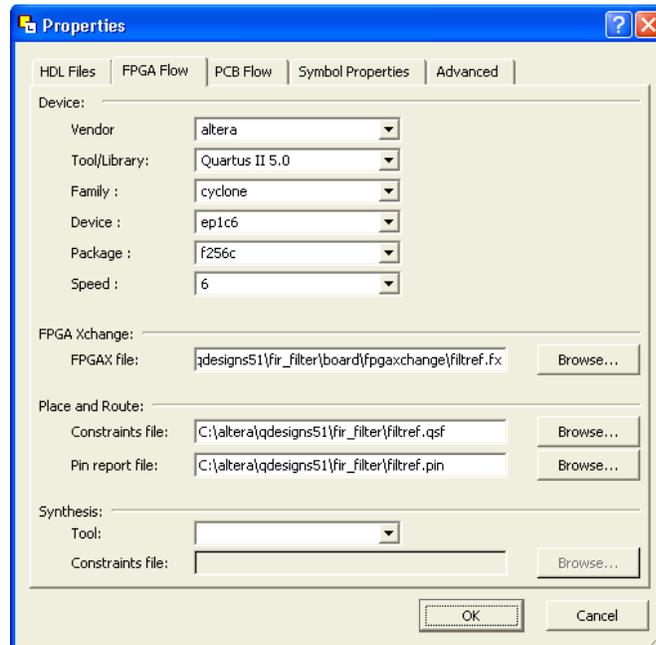


If your I/O Designer database points to the `.fx` generated by the Quartus II software instead of a backup copy of the file, updating the file in the Quartus II software overwrites any changes made to the file by the I/O Designer software. If there are I/O Designer assignments in the `.fx` that you want to preserve, create a backup copy of the file before updating it in the Quartus II software, and verify that your I/O Designer database points to the backup copy. To point to the backup copy, perform the steps in the following section.

Whenever you update the **.fx** or the **.pin** in the Quartus II software, the I/O Designer database imports the changes. You must set up the locations for the files in the I/O Designer software.

1. To set up the file locations, on the File menu, click **Properties**. The project **Properties** dialog box appears (Figure 8-6).

Figure 8-6. Project Properties Dialog Box

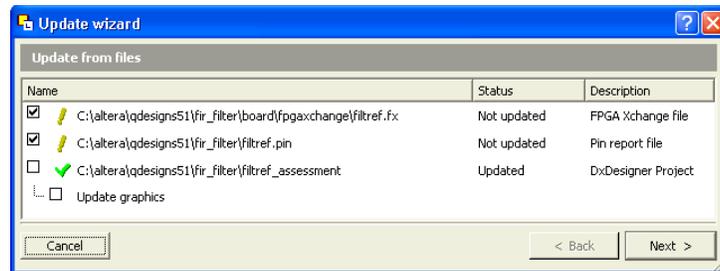


2. Under **FPGA Xchange**, click **Browse** to select the **.fx** file name and location.
3. To specify a Pin report file, under **Place and Route**, click **Browse** to select the **.pin** file name and location.

After you have set up these file locations, the I/O Designer software monitors these files for changes. If the **.fx** or **.pin** changes during the design flow, three indicators flash red in the lower right corner of the I/O Designer GUI (refer to Figure 8-4 on page 8-10). You can continue working or click on the indicators to open the I/O Designer **Update Wizard** dialog box. If you have made changes to your design in the Quartus II software that result in an updated **.fx** or **.pin** and the update indicators do not flash or you have previously canceled an indicated update, manually open the **Update Wizard** dialog box. To open the **Update Wizard** dialog box, on the File menu, click **Update**.

The I/O Designer **Update Wizard** dialog box lists the updated files associated with the database (Figure 8-7).

Figure 8-7. Update Wizard Dialog Box



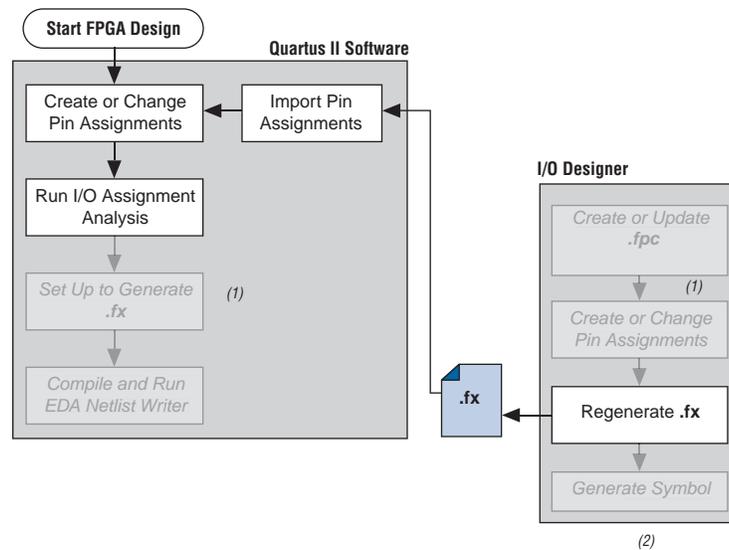
The paths to the updated files have yellow exclamation points and the **Status** column shows **Not updated**, indicating that the database has not yet been updated with the newer information contained in the files. A checkmark to the left of any updated file indicates that the file updates the database. Turn on any files you want to use to update the I/O Designer database, and click **Next**. If you are not satisfied with the database update, on the Edit menu, click **Undo**.

 You can update the I/O Designer database using the **.fx** and the **.pin** files simultaneously. Turning on the **.fx** and the **.pin** files for update causes the **Update Wizard** dialog box to provide options for using assignments from one file or the other exclusively or merging the assignments contained in both files into the I/O Designer database. Versions of the I/O Designer software older than version 2005 merge assignments contained in multiple files.

Sending Pin Assignment Changes to the Quartus II Software

In the same way that the FPGA designer can make adjustments that affect the PCB design, the board designer can make changes to optimize signal routing and layout that must be applied to the FPGA. The FPGA designer can take these required changes back into the Quartus II software to refit the logic to match the adjustments to the pin-out. The I/O Designer software accommodates this reverse flow as shown in Figure 8-8.

Figure 8-8. Updating the Quartus II Pin Assignments in the Reverse Design Flow



Notes to Figure 8-8:

- (1) These are software-specific steps in the design flow and are not necessary for the reverse flow steps of the design.
- (2) For more information about the full design flow, which includes the complete I/O Designer software, the DxDesigner software, and the board layout tool flowchart details, refer to Figure 8-1 on page 8-2.

You can make pin assignment changes directly in the I/O Designer software, or the software can automatically update changes made in a board layout tool that are back-annotated to a schematic entry program such as the DxDesigner software. You must update the .fx to reflect these updates in the Quartus II software. To perform this update in the I/O Designer software, on the Generate menu, click **FPGA Xchange File**.



If your I/O Designer database points to the .fx generated by the Quartus II software instead of a backup copy, updating the file from the I/O Designer software overwrites any changes made to the file by the Quartus II software. If there are assignments from the Quartus II software in the file that you want to preserve, create a backup copy of the file before updating it in the I/O Designer software, and verify that your I/O Designer database points to the backup copy. To point to the backup copy, perform the steps in [“Updating Pin Assignments from the Quartus II Software”](#) on page 8-11.

You must import the updated .fx into the Quartus II software. To import the file, follow these steps:

1. Start the Quartus II software and open your project.
2. On the Assignments menu, click **Import Assignments**.
3. In the File name box, click **Browse** and from the **Files of type** list, select **FPGA Xchange Files (*.fx)**.
4. Select the .fx and click **Open**.
5. Click **OK**.

Protecting Assignments in the Quartus II Software

To protect assignments in the Quartus II software, follow these steps:

1. Start the Quartus II software.
2. On the Assignments menu, click **Import Assignments**. The **Import Assignments** dialog box appears.
3. Turn on **Copy existing assignments into <project name>.qsf.bak before importing** before importing the .fx. This action automatically creates a backup copy of the Quartus II constraints file that contains all your current pin assignments.

Generating Symbols for the DxDesigner Software

Along with circuit simulation, circuit board schematic creation is one of the first tasks required in the design of a new PCB. Schematics must understand how the PCB works, and to generate a netlist for a board layout tool for board design and routing. The I/O Designer software allows you to create schematic symbols based on the FPGA design exported from the Quartus II software.

Most FPGA devices contain hundreds of pins, requiring large schematic symbols that may not fit on a single schematic page. Symbol designs in the I/O Designer software can be split or fractured into various functional blocks, allowing multiple part fractures on the same schematic page or across multiple pages. In the DxDesigner software, these part fractures join together with the use of the HETERO attribute.

The I/O Designer software can generate symbols for use in various Mentor Graphics schematic entry tools, and can import changes back-annotated by board layout tools to update the database and feed updates back to the Quartus II software with the .fx. This section discusses symbol creation specifically for the DxDesigner software.

You can create schematic symbols with the I/O Designer software in the following ways:

- Manually
- Using the I/O Designer **Symbol** wizard
- Importing previously created symbols from the DxDesigner software

The I/O Designer **Symbol** wizard can be used as a design base that allows you to quickly create a symbol for manual editing at a later time. If you have created symbols in a DxDesigner project and want to apply a different FPGA design to them, you can manually import these symbols from the DxDesigner project. To import the symbols, start the I/O Designer software, and on the File menu, click **Import Symbol**.

- For more information about importing symbols from the DxDesigner software into an I/O Designer database, refer to the I/O Designer Help.

Symbols created in the I/O Designer software are either functional, physical (PCB), or both. Signals imported into the database, usually from Verilog HDL or VHDL files, are the basis of a functional symbol. No physical device pins must be associated with the signals to generate a functional symbol. This section focuses on board-level PCB symbols with signals directly mapped to physical device pins through assignments in either the Quartus II Pin Planner or in the I/O Designer database.

- For more information about manually creating, importing, and editing symbols in the I/O Designer software, as well as the different types of symbols the software can generate, refer to the I/O Designer Help.

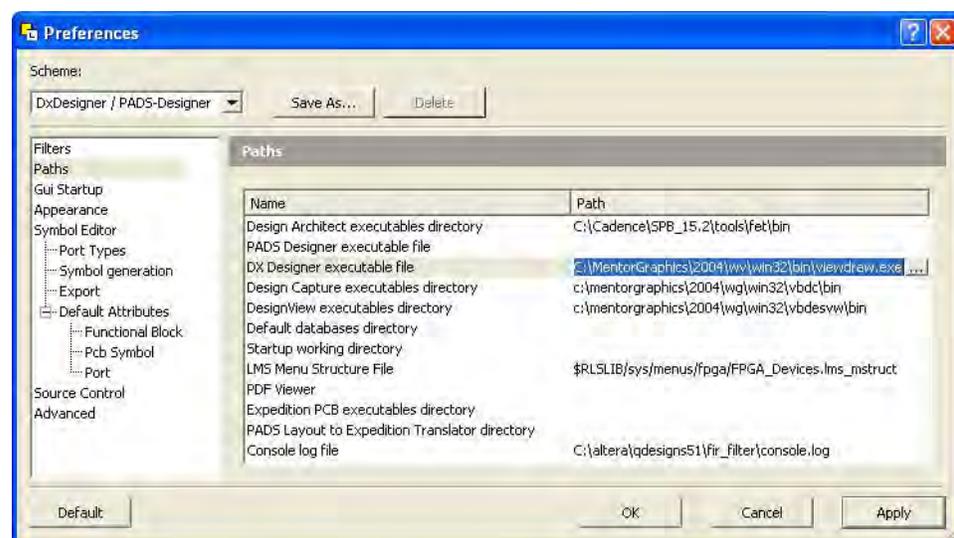
Setting Up the I/O Designer Software to Work with the DxDesigner Software

To verify if you are set up to export symbols to a DxDesigner project, or to manually set up the I/O Designer software to work with the DxDesigner software, you must set the path to the DxDesigner executable, set the export type to DxDesigner, and set the path to a DxDesigner project directory.

To set these options, follow these steps:

1. Start the I/O Designer software.
2. On the Tools menu, click **Preferences**. The **Preferences** dialog box appears.
3. Click **Paths**, double-click on the **DxDesigner executable file** path field, and click **Browse** to select the location of the DxDesigner application (Figure 8-9).
4. Click **Apply**.

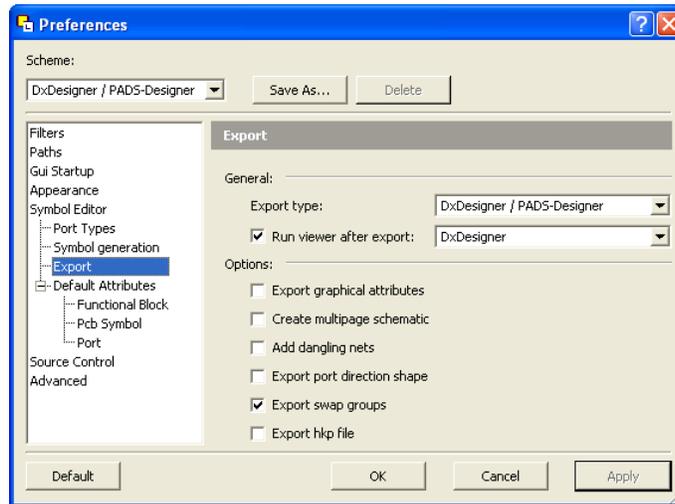
Figure 8-9. Path Preferences Dialog Box



5. Click **Symbol Editor** and click **Export**. In the Export type menu, under **General**, select **DxDesigner/PADS-Designer** (Figure 8-10).

6. Click **Apply** and click **OK**.

Figure 8-10. Symbol Editor Export Preferences



7. On the File menu, click **Properties**. The **Properties** dialog box appears.
8. Click the **PCB Flow** tab and click **Path to a DxDesigner project directory**.
9. Click **OK**.

If you do not have a new DxDesigner project in the **Database** wizard and a DxDesigner project, you must create a new database with the DxDesigner software, and point the I/O Designer software to this new project.

 For more information about creating and working with DxDesigner projects, refer to the DxDesigner Help.

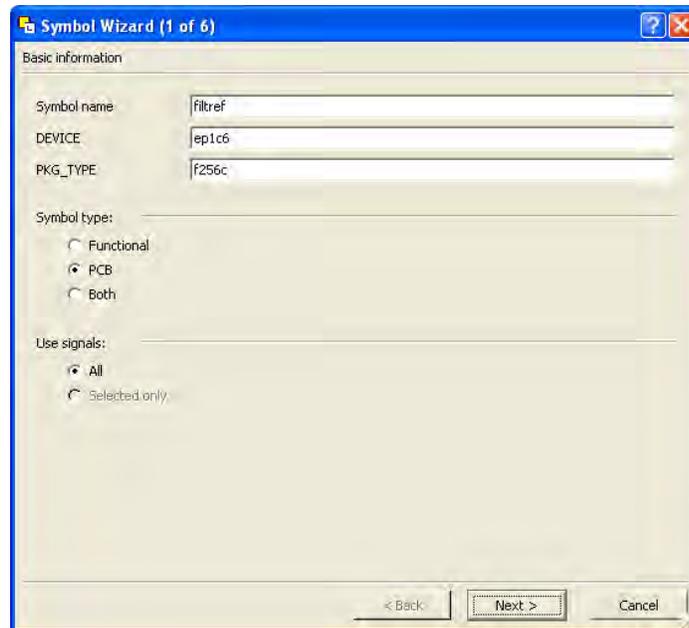
Creating Symbols with the Symbol Wizard

You can create, fracture, and edit FPGA symbols based on Altera devices with the I/O Designer **Symbol** wizard. To create a symbol based on a selected Altera FPGA device, follow these steps:

1. Start the I/O Designer software.

- Click **Symbol Wizard** in the toolbar, or on the Symbol menu, click **Symbol Wizard**. The **Symbol Wizard (1 of 6)** page appears (Figure 8-11).

Figure 8-11. Symbol Wizard



- On page 1 of the **Symbol Wizard** page, in the **Symbol name** field, type the symbol name. The **DEVICE** and **PKG_TYPE** fields are automatically populated with the device and package information. Under **Symbol type**, click **PCB**. Under **Use signals**, click **All**.
- Click **Next**. The **Symbol Wizard (2 of 6)** page appears.

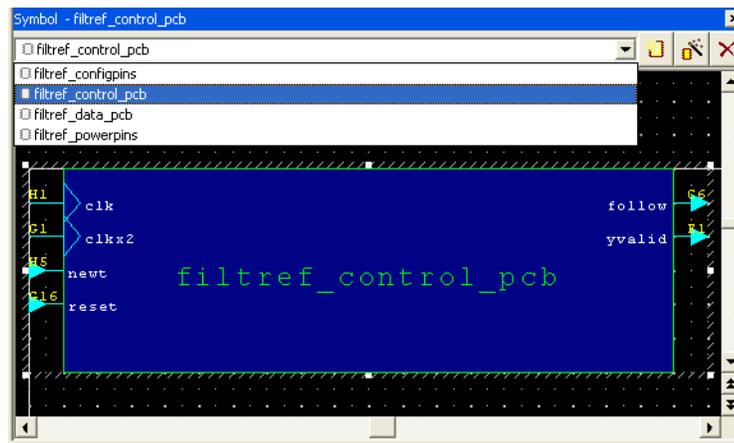
 If the **DEVICE** and **PKG_TYPE** fields are blank or incorrect, cancel the **Symbol wizard** and select the correct device information. On the File menu, click **Properties**. In the Properties window, click the **FPGA Flow** tab and enter the correct device information.

- On page 2 of the **Symbol Wizard** page, select fracturing options for your symbol. If you are using the **Symbol wizard** to edit a previously created fractured symbol, you must turn on **Reuse existing fractures** to preserve your current fractures. Select other options on this page as appropriate for your symbol.
- Click **Next**. The **Symbol Wizard (3 of 6)** page appears.
- Additional fracturing options are available on page 3 of the **Symbol Wizard** page. After selecting the necessary options, click **Next**. The **Symbol Wizard (4 of 6)** page appears.
- On page 4 of the **Symbol Wizard** page, select the options for the appearance of the symbols. Select the necessary options and click **Next**. The **Symbol Wizard (5 of 6)** page appears.

9. On page 5 of the **Symbol Wizard** page, define what information you want to label for the entire symbol and for individual pins. Select the necessary options and click **Next**. The **Symbol Wizard (6 of 6)** page appears.
10. On the final page of the **Symbol Wizard** page, add additional signals and pins that have not been placed in the symbol. Click **Finish** when you complete your selections.

You can view your symbol and any fractures you created with the Symbol Editor (Figure 8-12). You can edit parts of the symbol, delete fractures, or rerun the **Symbol** wizard.

Figure 8-12. The I/O Designer Symbol Editor



If assignments in the I/O Designer database are updated, the symbols created in the I/O Designer software automatically reflect these changes. Assignment changes can be made in the I/O Designer software, with an updated **.fx** from the Quartus II software, or from a back-annotated change in your board layout tool.

Exporting Symbols to the DxDesigner Software

After you have completed your symbols, export the symbols to your DxDesigner project. To generate all the fractures of a symbol, on the Generate menu, click **All Symbols**. To generate a symbol for the currently displayed symbol in Symbol Editor, click **Current Symbol Only**. The **/sym** directory in your DxDesigner project saves each symbol in the database as a separate file. The symbols can be instantiated in your DxDesigner schematics.

 For more information about working with DxDesigner projects, refer to the DxDesigner Help.

Scripting Support

The I/O Designer software features a command line Tcl interpreter. All commands issued through the GUI in the I/O Designer software translate into Tcl commands run by the tool. You can view the generated Tcl commands and run scripts, or type individual commands in the I/O Designer Console window.

This scripting support section includes commands that perform some of the operations described in this chapter.

If you want to change the `.fx` from which the I/O Designer software updates assignments, type the following command at an I/O Designer Tcl prompt:

```
set_fpga_xchange_file <file name>
```

You can type the following command to update the I/O Designer database with assignment updates made in the Quartus II software after specifying the `.fx`:

```
update_from_fpga_xchange_file
```

You can type the following command to update the `.fx` with changes made to the assignments in the I/O Designer software for transfer back into the Quartus II software:

```
generate_fpga_xchange_file
```

You can type the following command if you want to import assignment data from a `.pin` created by the Quartus II software:

```
set_pin_report_file -quartus_pin <file name>
```

You can run the I/O Designer **Symbol** wizard with the following command:

```
symbolwizard
```

You can set the DxDesigner project directory path where symbols are saved with the following command:

```
set_dx_designer_project -path <path>
```



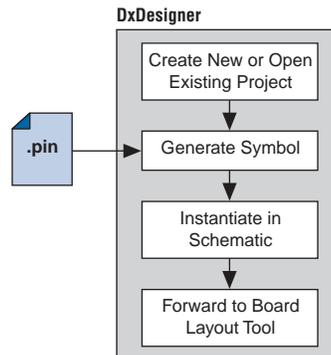
For more information about Tcl scripting and Tcl scripting with the Quartus II software, refer to the *Tcl Scripting* chapter in volume 2 of the *Quartus II Handbook*. For more information about the Tcl scripting capabilities of the I/O Designer software as well as a list of available commands, refer to the I/O Designer Help.

FPGA-to-Board Integration with the DxDesigner Software

The Mentor Graphics DxDesigner software is a design entry tool for schematic capture. You can use it to create flat circuit schematics for all the PCB design types. You can also use the DxDesigner software to create hierarchical schematics that facilitate design reuse and a team-based design. You can use the DxDesigner software in the design flow alone or in conjunction with the I/O Designer software. However, if you use the DxDesigner software without the I/O Designer software, the design flow is one-way, using only the `.pin` generated by the Quartus II software.

You can only make signal and pin assignment changes in the Quartus II software and these changes reflect as updated symbols in a DxDesigner schematic. You cannot back-annotate changes made in a board layout tool or in a DxDesigner symbol to the Quartus II software. Figure 8-13 shows the design flow without the I/O Designer software.

Figure 8-13. Design Flow Without the I/O Designer Software (Note 1)



Note to Figure 8-13:

(1) For more information about the full design flow, which includes the Quartus II software, the I/O Designer software, and the board layout tool flowchart details, refer to Figure 8-1 on page 8-2.

For more information about the DxDesigner software, including usage, support, training, and product updates, refer to the Mentor Graphics website (www.mentor.com), or choose Schematic Design Help Topics in the DxDesigner Help.

DxDesigner Project Settings

New projects in the DxDesigner software are set up to create FPGA symbols by default. However, if you are using the I/O Designer software with the DxDesigner software, you must enable the DxBoardLink Flow options for complete support and compatibility with the I/O Designer software.

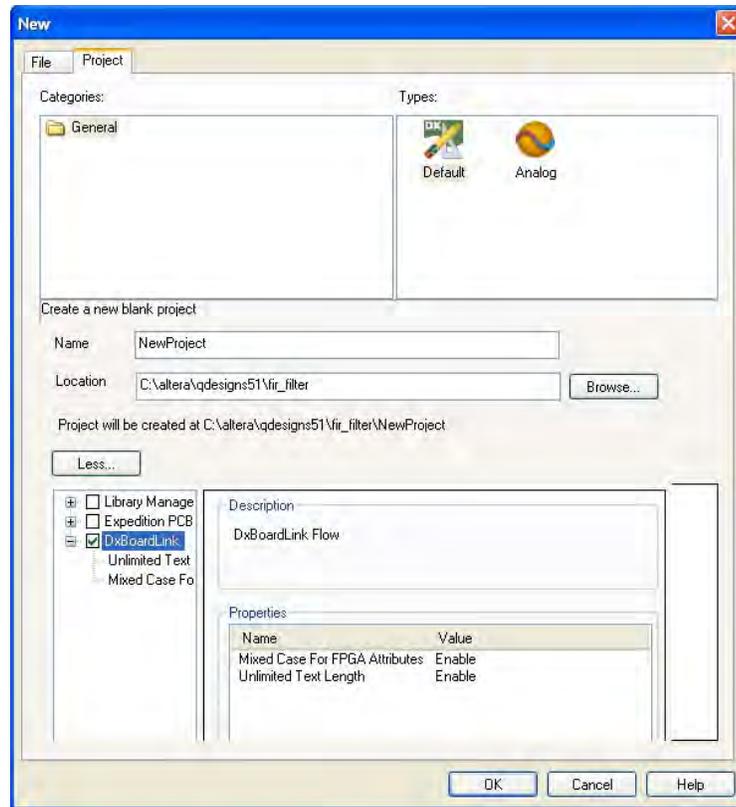
You can enable the DxBoardLink flow design configuration during or after creating a new DxDesigner project.

To enable the DxBoardLink flow design configuration when creating a new DxDesigner project, follow these steps:

1. Start the DxDesigner software.

- On the File menu, click **New** and click the **Project** tab. The **New** dialog box appears (Figure 8-14).

Figure 8-14. New Project Dialog Box

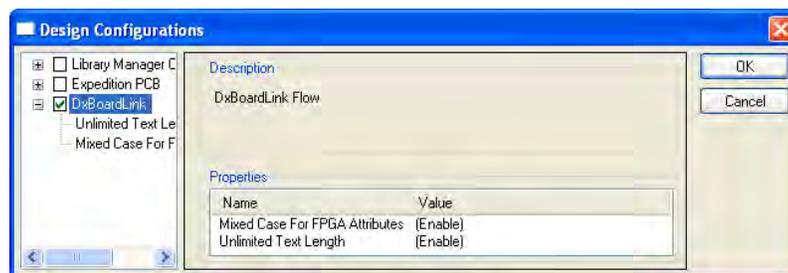


- Click **More**. Turn on **DxBoardLink** (Figure 8-14).



To enable the DxBoardLink Flow design configuration in an existing project, click **Design Configurations** in the Design Configuration toolbar and turn on **DxBoardLink** (Figure 8-15).

Figure 8-15. DxBoardLink Design Configuration



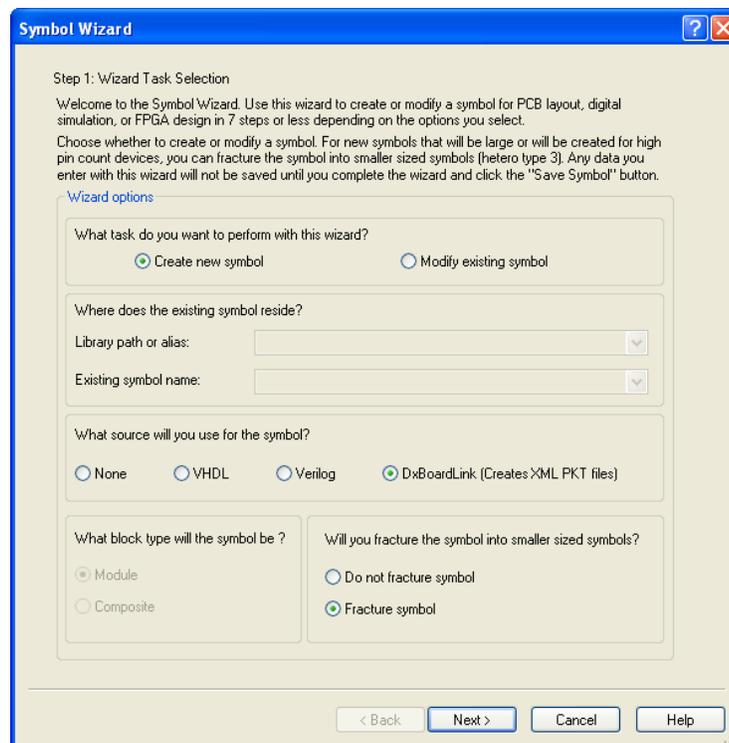
DxDesigner Symbol Wizard

You can create schematic symbols in the DxDesigner software manually or with the **Symbol** wizard. The DxDesigner **Symbol** wizard is similar to the I/O Designer **Symbol** wizard, but with fewer fracturing options.

FPGA symbols based on Altera devices can be created, fractured, and edited with the DxDesigner **Symbol** wizard. To start the **Symbol** wizard, follow these steps:

1. Start the DxDesigner software.
2. Click **Symbol Wizard** in the toolbar, or on the File menu, click **New**. The **New** window appears. Click the **File** tab and create a new file of type **Symbol Wizard**.
3. Type the new symbol name in the name field and click **OK**. The **Symbol Wizard** page appears (Figure 8-16).

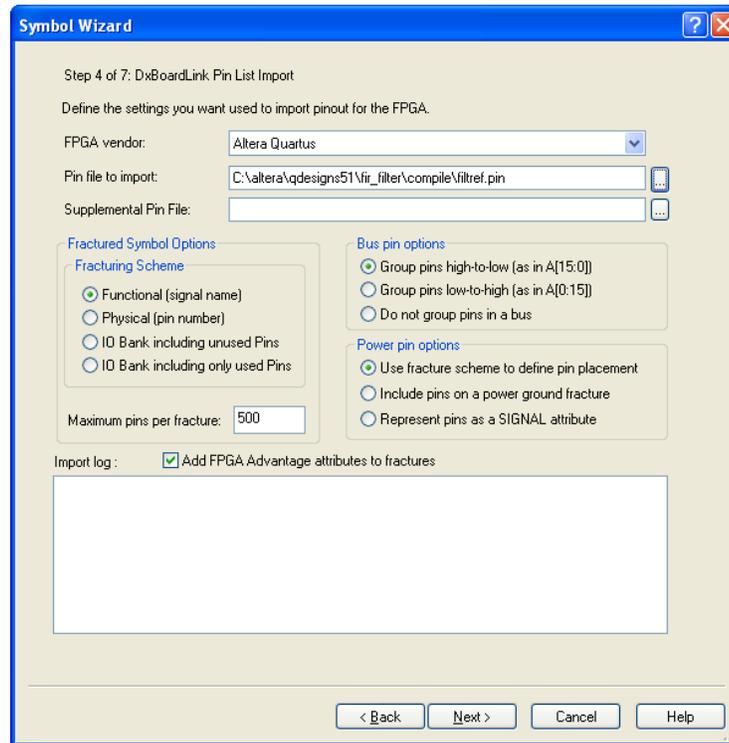
Figure 8-16. Wizard Task Selection



4. On the **Wizard Task Selection** page, choose to create a new symbol or modify an existing symbol. If you are modifying an existing symbol, specify the library path or alias, and select the existing symbol. If you are creating a new symbol, select **DxBoardLink** for the symbol source. The DxDesigner block type defaults to **Module** because the FPGA design does not have an underlying DxDesigner schematic. Choose whether or not to fracture the symbol. After making your selections, click **Next**. The **New Symbol and Library Name** page appears.
5. On the **New Symbol and Library Name** page, type a name for the symbol, an overall part name for all the symbol fractures, and a library name for the new library created for this symbol. By default, the part and library names are the same as the symbol name. Click **Next**. The **Symbol Parameters** page appears.

- On the **Symbol Parameters** page, specify the appearance of the generated symbol and how it matches up with the grid you have set in your DxDesigner project schematic. After making your selections, click **Next**. The **DxBoardLink Pin List Import** page appears (Figure 8-17).

Figure 8-17. DxBoardLink Pin List Import



- On the **DxBoardLink Pin List Import** page, in the **FPGA vendor** list, select **Altera Quartus**. In the **Pin-Out file to import** field, browse to and select the **.pin** from your Quartus II design project directory. You can also select choices from the Fracturing Scheme, Bus pin, and Power pin options. After making your selections, click **Next**. The **Symbol Attributes** page appears.
- On the **Symbol Attributes** page, select to create or modify symbol attributes for use in the DxDesigner software. After making your selections, click **Next**. The **Pin Settings** page appears.
- On the **Pin Settings** page, make any final adjustments to pin and label location and information. Each tabbed spreadsheet represents a fracture of your symbol. After making your selections, click **Save Symbol**.

After creating the symbol, you can examine and place any fracture of the symbol in your schematic. You can locate separate files of all the fractures you created in the library you specified or created in the **/sym** directory in your DxDesigner project. You can add the symbols to your schematics or you can manually edit the symbols or with the **Symbol** wizard.

-  Symbols created in the DxDesigner software can be edited and updated with newer versions of the .pin generated by the Quartus II software. However, you cannot fracture a symbol again because symbol fracturing is permanent. To create new fractures for your design, create a new symbol in the **Symbol** wizard, and perform the steps in “DxDesigner Symbol Wizard” on page 8-23.
-  For more information about creating, editing, and instantiating component symbols in DxDesigner, choose Schematic Design Help Topics from the Help menu in the DxDesigner software.

Conclusion

Transferring a complex, high-pin-count FPGA design to a PCB for prototyping or manufacturing is a daunting process that can lead to errors in the PCB netlist or design, especially when multiple engineers are working on different parts of the project. The design workflow available when using the Quartus II software with the Mentor Graphics toolset assists the FPGA designer and the board designer in preventing errors and focusing their attention on the design.

Document Revision History

Table 8-1 shows the revision history for this chapter.

Table 8-1. Document Revision History

Date	Version	Changes
December 2010	10.0.1	Template update.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Removed Reference Document section. ■ General style editing. ■ Added a link to Help in “Performing Simultaneous Switching Noise (SSN) Analysis of Your FPGA.” ■ Removed Figure 8-4 on page 8-9 and Figure 8-5 on page 8-11. ■ Updated “Generating an .fx File.”
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Added minor information about simultaneous switching noise (SSN) analysis on “Performing Simultaneous Switching Noise (SSN) Analysis of Your FPGA.” ■ General style editing.
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Was chapter 6 in the 8.1.0 release. ■ Removed Figures that were numbered 6-4, 6-6, 6-7, and 6-8 in v8.1.0.
November 2008	8.1.0	Changed to 8½” × 11” page size. No change to content.
May 2008	8.0.0	Updated references.

-  For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).
-  Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter addresses how the Quartus® II software interacts with the Cadence Allegro Design Entry HDL software and the Cadence Allegro Design Entry CIS (Component Information System) software (also known as OrCAD Capture CIS) to provide a complete FPGA-to-board integration design workflow.

With today's large, high-pin-count and high-speed FPGA devices, good PCB design practices are important to ensure the correct operation of your system. The PCB design takes place concurrently with the design and programming of the FPGA. An FPGA or ASIC designer initially creates the signal and pin assignments and the board designer must transfer these assignments to the symbols used in their system circuit schematics and board layout correctly. As the board design progresses, you must perform pin reassignments to optimize the layout. You must communicate pin reassignments to the FPGA designer to ensure the new assignments are processed through the FPGA with updated placement and routing.

This chapter is intended for board design and layout engineers who want to begin the FPGA board integration process while the FPGA is still in the design phase. Part librarians can also benefit from this chapter by learning the method to use output from the Quartus II software to create new library parts and symbols.

This chapter discusses the following topics:

- Cadence tool description, history, and comparison.
- The general design flow between the Quartus II software and the Cadence Allegro Design Entry HDL software and the Cadence Allegro Design Entry CIS software.
- Generating schematic symbols from your FPGA design for use in the Cadence Allegro Design Entry HDL software.
- Updating Design Entry HDL symbols when making signal and pin assignment changes in the Quartus II software.
- Creating schematic symbols in the Cadence Allegro Design Entry CIS software from your FPGA design.
- Updating symbols in the Cadence Allegro Design Entry CIS software when making signal and pin assignment changes in the Quartus II software.
- Using Altera-provided device libraries in the Cadence Allegro Design Entry CIS software.

The procedures in this chapter require the following software:

- The Quartus II software version 5.1 or later
- The Cadence Allegro Design Entry HDL software or the Cadence Allegro Design Entry CIS software version 15.2 or later
- The OrCAD Capture software with the optional CIS option version 10.3 or later (optional)

 These programs are very similar because the Cadence Allegro Design Entry CIS software is based on the OrCAD Capture software. This chapter refers to the Cadence Allegro Design Entry CIS software; however, any procedural information can also apply to the OrCAD Capture software unless otherwise noted.

 For more information about obtaining and licensing the Cadence tools and for product information, support, and training, refer to the Cadence website (www.cadence.com). For more information about the OrCAD Capture software and the CIS option, refer to the Cadence website (www.cadence.com). For more information about Cadence and OrCAD support and training, refer to the EMA Design Automation website (www.ema-eda.com).

Product Comparison

The Cadence and OrCAD design tools are different in their function and location of product information. Table 9-1 lists the Cadence and OrCAD products described in this chapter and provides information about changes, product information, and support.

Table 9-1. Cadence and OrCAD Product Comparison

Description	Cadence Allegro Design Entry HDL	Cadence Allegro Design Entry CIS	OrCAD Capture CIS
Former Name	Concept HDL Expert	Capture CIS Studio	—
History	More commonly known by its former name, Cadence renamed all board design tools in 2004 under the Allegro name.	Based directly on OrCAD Capture CIS, the Cadence Allegro Design Entry CIS software is still developed by OrCAD but sold and marketed by Cadence. EMA provides support and training.	The basis for Design Entry CIS is still developed by OrCAD for continued use by existing OrCAD customers. EMA provides support and training for all OrCAD products.
Vendor Design Flow	Cadence Allegro 600 series, formerly known as the Expert Series, for high-end, high-speed design.	Cadence Allegro 200 series, formerly known as the Studio Series, for small- to medium-level design.	—
Information and Support	www.cadence.com www.ema-eda.com	www.cadence.com www.ema-eda.com	www.cadence.com www.ema-eda.com

FPGA-to-PCB Design Flow

You can create a design flow integrating an Altera FPGA design from the Quartus II software through a circuit schematic in the Cadence Allegro Design Entry HDL software or the Cadence Allegro Design Entry CIS software. Figure 9-1 shows the design flow with the Cadence Allegro Design Entry HDL software. Figure 9-2 shows the design flow with the Cadence Allegro Design Entry CIS software.

Figure 9-1. Design Flow with the Cadence Allegro Design Entry HDL Software

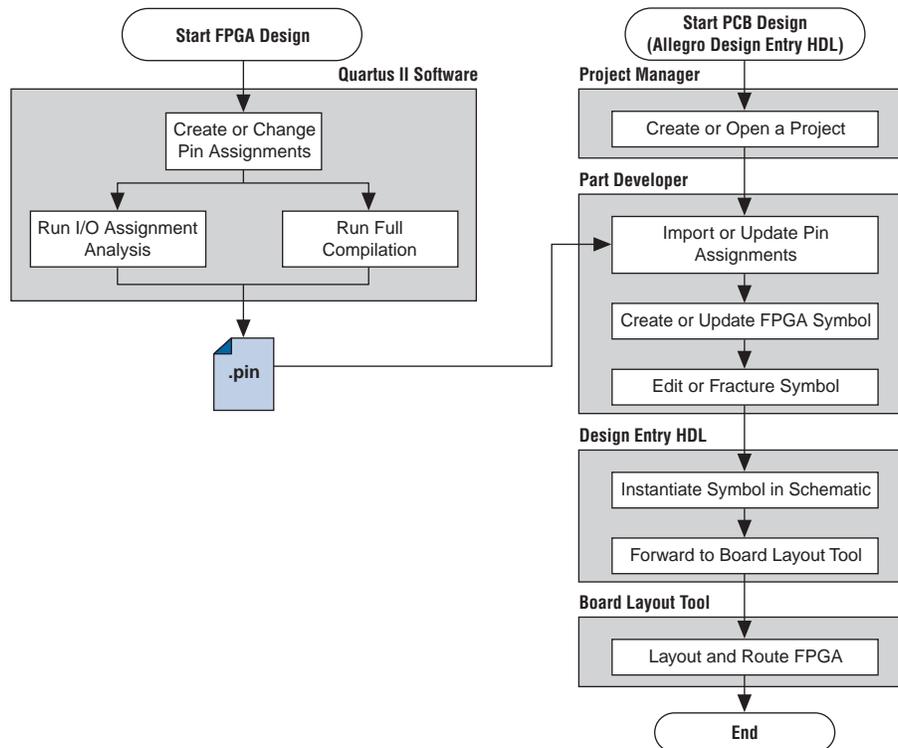


Figure 9-2. Design Flow with the Cadence Allegro Design Entry CIS Software

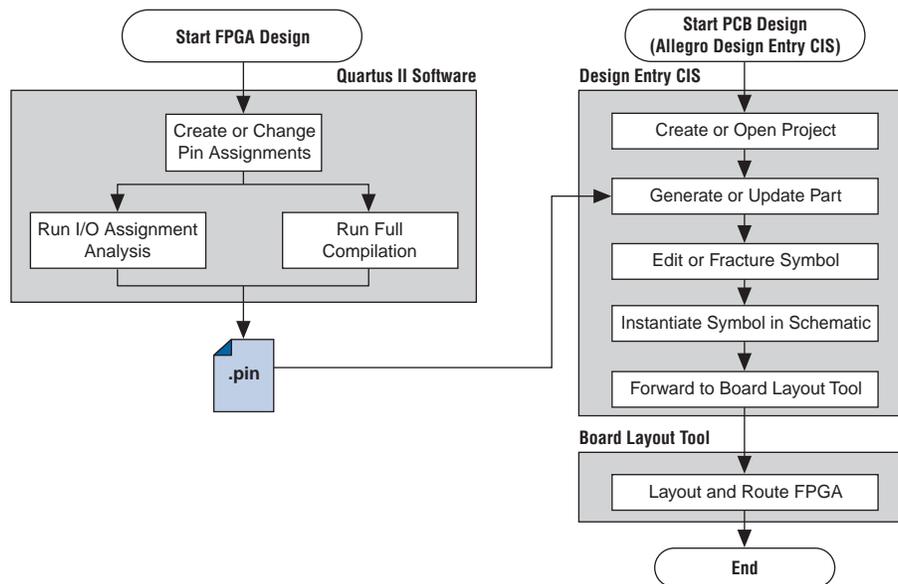


Figure 9-1 and Figure 9-2 show the possible design flows, depending on your tool choice. To create FPGA symbols using the Cadence Allegro PCB Librarian Part Developer tool, you must obtain the Cadence PCB Librarian Expert license. You can update symbols with changes made to the FPGA design using any of these tools.

To integrate an Altera FPGA design starting in the Quartus II software through to a circuit schematic in the Cadence Allegro Design Entry HDL software or the Cadence Allegro Design Entry CIS software, follow these steps:

1. In the Quartus II software, compile your design to generate a Pin-Out File (.pin) to transfer the assignments to the Cadence software.
2. If you are using the Cadence Allegro Design Entry HDL software for your schematic design, follow these steps:
 - a. Open an existing project or create a new project in the Cadence Allegro Project Manager tool.
 - b. Construct a new symbol or update an existing symbol using the Cadence Allegro PCB Librarian Part Developer tool.
 - c. With the Cadence Allegro PCB Librarian Part Developer tool, edit your symbol or fracture it into smaller parts (optional).
 - d. Instantiate the symbol in your Cadence Allegro Design Entry HDL software schematic and transfer the design to your board layout tool.

or

If you are using the Cadence Allegro Design Entry CIS software for your schematic design, follow these steps:

- a. Generate a new part in a new or existing Cadence Allegro Design Entry CIS project, referencing the .pin output file from the Quartus II software. You can also update an existing symbol with a new .pin.

- b. Split the symbol into smaller parts as necessary.
- c. Instantiate the symbol in your Cadence Allegro Design Entry CIS schematic and transfer the design to your board layout tool.

Performing Simultaneous Switching Noise (SSN) Analysis of Your FPGA

With the Quartus II software, you can extract pin assignment data and perform SSN analysis of your FPGA design for designs targeting the Stratix III device family. You can analyze SSN in your device early in the board layout stage as part of your overall pin planning process; however, you do not have to perform SSN analysis to generate pin assignment data from the Quartus II software. You can use the SSN Analyzer tool to optimize the pin assignments for better SSN performance of your device.

 For more information about the SSN Analyzer, refer to *About the SSN Analyzer* in Quartus II Help and the *Simultaneous Switching Noise (SSN) Analysis and Optimizations* chapter in volume 2 of the *Quartus II Handbook*.

Setting Up the Quartus II Software

You can transfer pin and signal assignments from the Quartus II software to the Cadence design tools by generating the Quartus II project **.pin**. The **.pin** is an output file generated by the Quartus II Fitter containing pin assignment information. You can use the Quartus II Pin Planner to set and change the assignments in the **.pin** and then transfer the assignments to the Cadence design tools. You cannot, however, import pin assignment changes from the Cadence design tools into the Quartus II software with the **.pin**.

The **.pin** lists all used and unused pins on your selected Altera device. The **.pin** also provides the following basic information fields for each assigned pin on the device:

- Pin signal name and usage
- Pin number
- Signal direction
- I/O standard
- Voltage
- I/O bank
- User or Fitter-assigned

 For more information about using the Quartus II Pin Planner to create or change pin assignment details, refer to the *I/O Management* chapter in volume 2 of the *Quartus II Handbook*.

Generating a .pin File

The Quartus II Fitter generates a **.pin** during a full compilation of your FPGA design, or when performing I/O assignment analysis on your design. You can locate the **.pin** in your Quartus II project directory with the name *<project name>.pin*.

-  For more information about pin and signal assignment transfer and the files that the Quartus II software can import and export, refer to the *I/O Management* chapter in volume 2 of the *Quartus II Handbook*.

FPGA-to-Board Integration with the Cadence Allegro Design Entry HDL Software

The Cadence Allegro Design Entry HDL software is a schematic capture tool and is part of the Cadence 600 series design flow. Use the Cadence Allegro Design Entry HDL software to create flat circuit schematics for all types of PCB design. The Cadence Allegro Design Entry HDL software can also create hierarchical schematics to facilitate design reuse and team-based design. With the Cadence Allegro Design Entry HDL software, the design flow from FPGA-to-board is one-way, using only the `.pin` generated by the Quartus II software. You can only make signal and pin assignment changes in the Quartus II software and these changes reflect as updated symbols in a Cadence Allegro Design Entry HDL project. For more information about the design flow with the Cadence Allegro Design Entry HDL software, refer to [Figure 9-1 on page 9-3](#).

-  Routing or pin assignment changes made in a board layout tool or a Cadence Allegro Design Entry HDL software symbol cannot be back-annotated to the Quartus II software.
-  For more information about the Cadence Allegro Design Entry HDL software and the Cadence Allegro PCB Librarian Part Developer tool, including licensing, support, usage, training, and product updates, refer to the Help in the software or to the Cadence website (www.cadence.com).

Creating Symbols

In addition to circuit simulation, circuit board schematic creation is one of the first tasks required when designing a new PCB. Schematics must understand how the PCB works, and to generate a netlist for a board layout tool for board design and routing. The Cadence Allegro PCB Librarian Part Developer tool allows you to create schematic symbols based on FPGA designs exported from the Quartus II software.

You can create symbols for the Cadence Allegro Design Entry HDL project with the Cadence Allegro PCB Librarian Part Developer tool, which is available in the Cadence Allegro Project Manager tool. Altera recommends using the Cadence Allegro PCB Librarian Part Developer tool to import FPGA designs into the Cadence Allegro Design Entry HDL software.

You must obtain a PCB Librarian Expert license from Cadence to run the Cadence Allegro PCB Librarian Part Developer tool. The Cadence Allegro PCB Librarian Part Developer tool provides a GUI with many options for creating, editing, fracturing, and updating symbols. If you do not use the Cadence Allegro PCB Librarian Part Developer tool, you must create and edit symbols manually in the Symbol Schematic View in the Cadence Allegro Design Entry HDL software.

 If you do not have a PCB Librarian Expert license, you can automatically create FPGA symbols using the programmable IC (PIC) design flow found in the Cadence Allegro Project Manager tool. For more information about using the PIC design flow, refer to the Help in the Cadence design tools, or go to the Cadence website (www.cadence.com).

Before creating a symbol from an FPGA design, you must open a Cadence Allegro Design Entry HDL project with the Cadence Allegro Project Manager tool. If you do not have an existing Cadence Allegro Design Entry HDL project, you can create one with the Cadence Allegro Design Entry HDL software. The Cadence Allegro Design Entry HDL project directory with the name *<project name>.cpm* contains your Cadence Allegro Design Entry HDL projects.

While the Cadence Allegro PCB Librarian Part Developer tool refers to symbol fractures as slots, the other tools described in this chapter use different names to refer to symbol fractures. Table 9-2 lists the symbol fracture naming conventions for each of the tools addressed in this chapter.

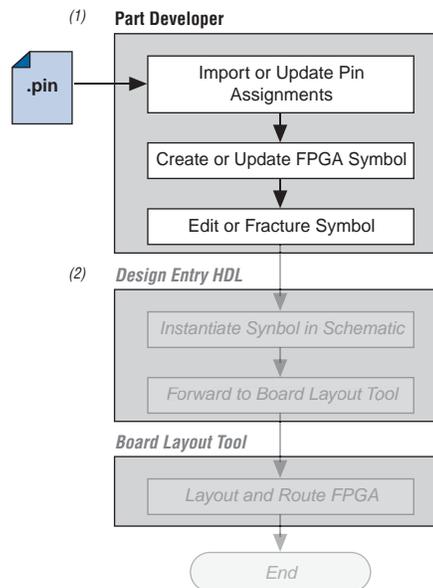
Table 9-2. Symbol Fracture Naming

	Cadence Allegro PCB Librarian Part Developer Tool	Cadence Allegro Design Entry HDL Software	Cadence Allegro Design Entry CIS Software
During symbol generation	Slots	—	Sections
During symbol schematic instantiation	—	Versions	Parts

Cadence Allegro PCB Librarian Part Developer Tool

You can create, fracture, and edit schematic symbols for your designs using the Cadence Allegro PCB Librarian Part Developer tool. Symbols designed in the Cadence Allegro PCB Librarian Part Developer tool can be split or fractured into several functional blocks called slots, allowing multiple smaller part fractures to exist on the same schematic page or across multiple pages. Figure 9-3 shows how the Cadence Allegro PCB Librarian Part Developer tool fits into the design flow.

Figure 9-3. Cadence Allegro PCB Librarian Part Developer Tool in the Design Flow



Notes to Figure 9-3:

- (1) For more information about the full design flow flowchart, refer to Figure 9-1 on page 9-3.
- (2) Grayed out steps are not part of the FPGA symbol creation or update process.

To run the Cadence Allegro PCB Librarian Part Developer tool, you must open a Cadence Allegro Design Entry HDL project in the Cadence Allegro Project Manager tool. To open the Cadence Allegro PCB Librarian Part Developer tool, on the Flows menu, click **Library Management**, and then click **Part Developer**.

Import and Export Wizard

After starting the Cadence Allegro PCB Librarian Part Developer tool, use the **Import and Export** wizard to import your pin assignments from the Quartus II software.



Altera recommends using your PCB Librarian Expert license file. To point to your PCB Librarian Expert license file, on the File menu, click **Change Product** and then select the correct product license.

To access the Import and Export wizard, follow these steps:

1. On the File menu, click **Import and Export**.
2. Select **Import ECO-FPGA**, and then click **Next**.

3. In the **Select Source** page of the **Import and Export** wizard, specify the following settings:
 - a. In the **Vendor** list, select **Altera**.
 - b. In the **PnR Tool** list, select **quartusII**.
 - c. In the **PR File** box, browse to select the **.pin** in your Quartus II project directory.
 - d. Click **Simulation Options** to select simulation input files.
 - e. Click **Next**.
4. In the **Select Destination** dialog box, specify the following settings:
 - a. Under **Select Component**, click **Generate Custom Component** to create a new component in a library,
or
Click **Use standard component** to base your symbol on an existing component.



Altera recommends creating a new component if you previously created a generic component for an FPGA device. Generic components can cause some problems with your design. When you create a new component, you can place your pin and signal assignments from the Quartus II software on this component and reuse the component as a base when you have a new FPGA design.

- b. In the **Library** list, select an existing library. You can select from the cells in the selected library. Each cell represents all the symbol versions and part fractures for a particular part. In the **Cell** list, select the existing cell to use as a base for your part.
- c. In the **Destination Library** list, select a destination library for the component. Click **Next**.
- d. Review and edit the assignments you import into the Cadence Allegro PCB Librarian Part Developer tool based on the data in the **.pin** and then click **Finish**. The location of each pin is not included in the **Preview of Import Data** page of the **Import and Export** wizard, but input pins are on the left side of the created symbol, output pins on the right, power pins on the top, and ground pins on the bottom.

Editing and Fracturing Symbol

After creating your new symbol in the Cadence Allegro PCB Librarian Part Developer tool, you can edit the symbol graphics, fracture the symbol into multiple slots, and add or change package or symbol properties.

The Part Developer Symbol Editor contains many graphical tools to edit the graphics of a particular symbol. To edit the symbol graphics, select the symbol in the cell hierarchy. The **Symbol Pins** tab appears. You can edit the preview graphic of the symbol in the **Symbol Pins** tab.

8. Select whether to create a new symbol or modify an existing symbol in each slot. Click **OK**.

The newly generated or modified slot symbols appear as separate symbols in the cell hierarchy. Each of these symbols can be edited individually.



The Cadence Allegro PCB Librarian Part Developer tool allows you to remap pin assignments in the **Package Pin** tab of the main Cadence Allegro PCB Librarian Part Developer window. If signals remap to different pins in the Cadence Allegro PCB Librarian Part Developer tool, the changes reflect only in regenerated symbols for use in your schematics. You cannot transfer pin assignment changes to the Quartus II software from the Cadence Allegro PCB Librarian Part Developer tool, which creates a potential mismatch of the schematic symbols and assignments in the FPGA design. If pin assignment changes are necessary, make the changes in the Quartus II Pin Planner instead of the Cadence Allegro PCB Librarian Part Developer tool, and update the symbol as described in the following sections.

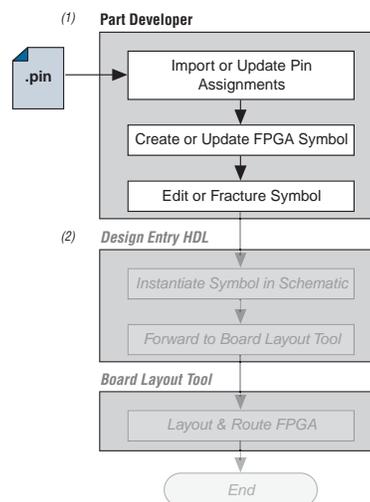


For more information about creating, editing, and organizing component symbols with the Cadence Allegro PCB Librarian Part Developer tool, refer to the Part Developer Help.

Updating FPGA Symbols

As the design process continues, you must make logic changes in the Quartus II software, placing signals on different pins after recompiling the design, or use the Quartus II Pin Planner to make changes manually. The board designer can request such changes to improve the board routing and layout. To ensure signals connect to the correct pins on the FPGA, you must carry forward these types of changes to the circuit schematic and board layout tools. Updating the **.pin** in the Quartus II software facilitates this flow. [Figure 9-5](#) shows this part of the design flow.

Figure 9-5. Updating the FPGA Symbol in the Design Flow



Notes to [Figure 9-5](#):

- (1) For more information about the full design flow flowchart, refer to [Figure 9-1 on page 9-3](#).
- (2) Grayed out steps are not part of the FPGA symbol update process.

To update the symbol using the Cadence Allegro PCB Librarian Part Developer tool after updating the **.pin**, follow these steps:

1. On the File menu, click **Import and Export**. The Import and Export wizard appears.
2. In the list of actions to perform, select **Import ECO - FPGA**. Click **Next**. The **Select Source** dialog box appears.
3. Select the updated source of the FPGA assignment information. In the **Vendor** list, select **Altera**. In the **PnR Tool** list, select **quartusII**. In the **PR File** field, click **browse** to specify the updated **.pin** in your Quartus II project directory. Click **Next**. The Select Destination window appears.
4. Select the source component and a destination cell for the updated symbol. To create a new component based on the updated pin assignment data, select **Generate Custom Component**. Selecting **Generate Custom Component** replaces the cell listed under the **Specify Library and Cell** name header with a new, nonfractured cell. You can preserve these edits by selecting **Use standard component and select the existing library and cell**. Select the destination library for the component and click **Next**. The **Preview of Import Data** dialog box appears.
5. Make any additional changes to your symbol. Click **Next**. A list of ECO messages appears summarizing the changes made to the cell. To accept the changes and update the cell, click **Finish**.
6. The main Cadence Allegro PCB Librarian Part Developer window appears. You can edit, fracture, and generate the updated symbols as usual from the main Cadence Allegro PCB Librarian Part Developer window.



If the Cadence Allegro PCB Librarian Part Developer tool is not set up to point to your PCB Librarian Expert license file, an error message appears in red at the bottom of the message text window of the Part Developer when you select the **Import and Export** command. To point to your PCB Librarian Expert license, on the File menu, click **Change Product**, and select the correct product license.

Instantiating the Symbol in the Cadence Allegro Design Entry HDL Software

To instantiate the symbol in your Cadence Allegro Design Entry HDL schematic after saving the new symbol in the Cadence Allegro PCB Librarian Part Developer tool, follow these steps:

1. In the Cadence Allegro Project Manager tool, switch to the board design flow.
2. On the Flows menu, click **Board Design**.
3. To start the Cadence Allegro Design Entry HDL software, click **Design Entry**.
4. To add the newly created symbol to your schematic, on the Component menu, click **Add**. The **Add Component** dialog box appears.
5. Select the new symbol library location, and select the name of the cell you created from the list of cells.

The symbol attaches to your cursor for placement in the schematic. To fracture the symbol into slots, right-click the symbol and choose **Version** to select one of the slots for placement in the schematic.

-  For more information about the Cadence Allegro Design Entry HDL software, including licensing, support, usage, training, and product updates, refer to the Help in the software or go to the Cadence website (www.cadence.com).

FPGA-to-Board Integration with Cadence Allegro Design Entry CIS Software

The Cadence Allegro Design Entry CIS software is a schematic capture tool (part of the Cadence 200 series design flow based on OrCAD Capture CIS). Use the Cadence Allegro Design Entry CIS software to create flat circuit schematics for all types of PCB design. You can also create hierarchical schematics to facilitate design reuse and team-based design using the Cadence Allegro Design Entry CIS software. With the Cadence Allegro Design Entry CIS software, the design flow from FPGA-to-board is unidirectional using only the **.pin** generated by the Quartus II software. You can only make signal and pin assignment changes in the Quartus II software. These changes reflect as updated symbols in a Cadence Allegro Design Entry CIS schematic project. [Figure 9-2 on page 9-4](#) shows the design flow with the Cadence Allegro Design Entry CIS software.

-  Routing or pin assignment changes made in a board layout tool or a Cadence Allegro Design Entry CIS symbol cannot be back-annotated to the Quartus II software.

-  For more information about the Cadence Allegro Design Entry CIS software, including licensing, support, usage, training, and product updates, refer to the Help in the software, go to the Cadence (www.cadence.com) or go to the EMA Design Automation website (www.ema-eda.com).

Creating a Cadence Allegro Design Entry CIS Project

The Cadence Allegro Design Entry CIS software has built-in support for creating schematic symbols using pin assignment information imported from the Quartus II software.

To create a new project in the Cadence Allegro Design Entry CIS software, follow these steps:

1. On the File menu, point to **New** and click **Project**. The New Project wizard starts.
When you create a new project, you can select the PC Board wizard, the Programmable Logic wizard, or a blank schematic.
2. Select the PC Board wizard to create a project where you can select which part libraries to use, or select a blank schematic.

The Programmable Logic wizard only builds an FPGA logic design in the Cadence Allegro Design Entry CIS software.

Your new project is in the specified location and consists of the following files:

- OrCAD Capture Project File (**.opj**)
- Schematic Design File (**.dsn**)

Generating a Part

After you create a new project or open an existing project in the Cadence Allegro Design Entry CIS software, you can generate a new schematic symbol based on your Quartus II FPGA design. You can also update an existing symbol. The Cadence Allegro Design Entry CIS software stores component symbols in OrCAD Library File (.olb). When you place a symbol in a library attached to a project, it is immediately available for instantiation in the project schematic.

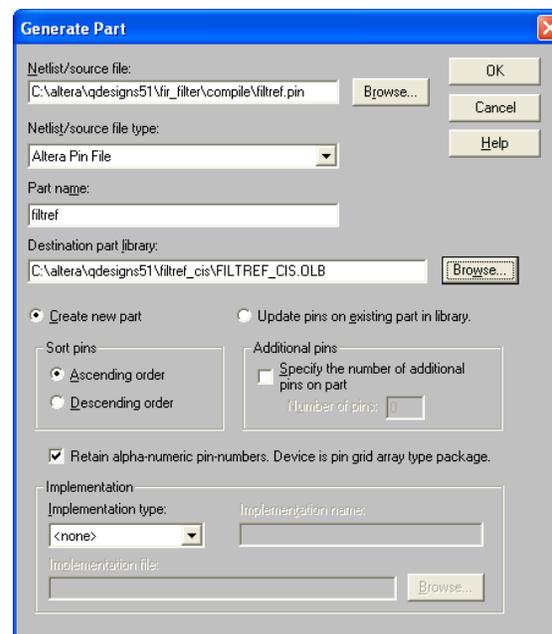
You can add symbols to an existing library or you can create a new library specifically for the symbols generated from your FPGA designs. To create a new library, follow these steps:

1. On the File menu, point to **New** and click **Library** in the Cadence Allegro Design Entry CIS software to create a default library named **library1.olb**. This library appears in the **Library** folder in the Project Manager window of the Cadence Allegro Design Entry CIS software.
2. To specify a desired name and location for the library, right-click the new library and select **Save As**. Saving the new library creates the library file.

You can now create a new symbol to represent your FPGA design in your schematic. To generate a schematic symbol, follow these steps:

1. Start the Cadence Allegro Design Entry CIS software.
2. On the Tools menu, click **Generate Part**. The **Generate Part** dialog box appears (Figure 9-6).

Figure 9-6. Generate Part Dialog Box

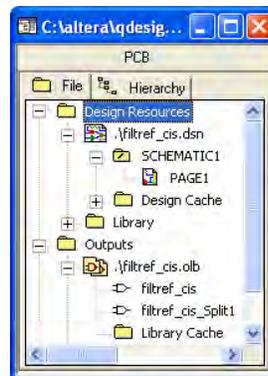


3. To specify the .pin from your Quartus II design, in the **Netlist/source file type** field, click **Browse**.
4. In the **Netlist/source file type** list, select **Altera Pin File**.

5. Type the new part name.
6. Specify the **Destination part library** for the symbol. Failing to select an existing library for the part creates a new library with a default name that matches the name of your Cadence Allegro Design Entry CIS project.
7. To create a new symbol for this design, select **Create new part**. If you updated your **.pin** in the Quartus II software and want to transfer any assignment changes to an existing symbol, select **Update pins on existing part in library**.
8. Select any other desired options and set **Implementation type** to **<none>**. The symbol is for a primitive library part based only on the **.pin** and does not require special implementation. Click **OK**.
9. Review the Undo warning and click **Yes** to complete the symbol generation.

You can locate the generated symbol in the selected library or in a new library found in the **Outputs** folder of the design in the Project Manager window (Figure 9-7). Double-click the name of the new symbol to see its graphical representation and edit it manually using the tools available in the Cadence Allegro Design Entry CIS software.

Figure 9-7. Project Manager Window

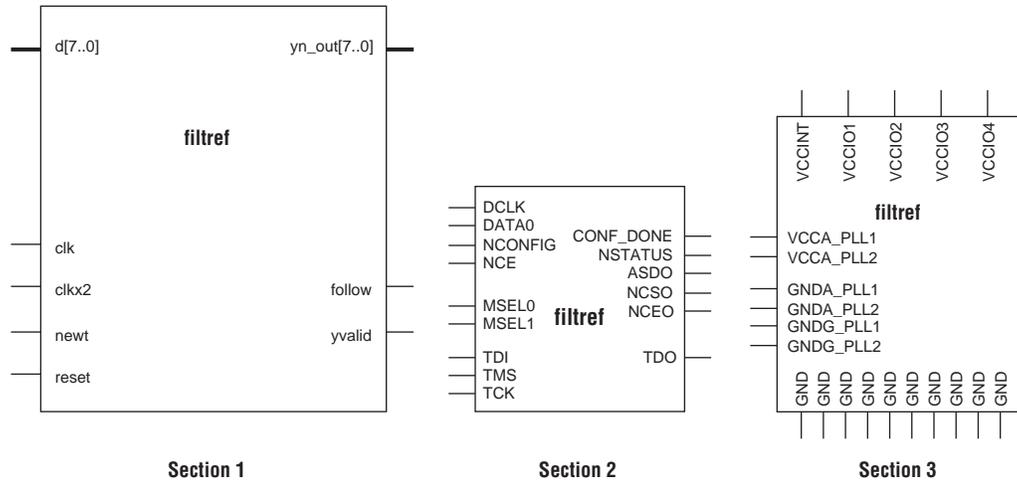


 For more information about creating and editing symbols in the Cadence Allegro Design Entry CIS software, refer to the Help in the software.

Splitting a Part

After saving a new symbol in a project library, you can fracture the symbol into multiple parts called sections. Fracturing a part into separate sections is useful for FPGA designs. A single symbol for most FPGA packages might be too large for a single schematic page. Splitting the part into separate sections allows you to organize parts of the symbol by function, creating cleaner circuit schematics. For example, you can create one slot for an I/O symbol, a second slot for a JTAG symbol, and a third slot for a power/ground symbol. Figure 9-8 shows a part fractured into separate sections.

Figure 9-8. Splitting a Symbol into Multiple Sections (Notes 1), (2)



Notes to Figure 9-8:

- (1) Figure 9-8 represents a Cyclone device with JTAG or passive serial (PS) mode configuration option settings. Symbols created for other devices or other configuration modes might have different sets of configuration pins, but can be fractured in a similar manner.
- (2) The power/ground section shows only a representation of power and ground pins because the device contains a high number of power and ground pins.



Although symbol generation in the Design Entry CIS software refers to symbol fractures as sections, the other tools described in this chapter use different names to refer to symbol fractures.

To split a part into sections, select the part in its library in the Project Manager window of the Cadence Allegro Design Entry CIS software. On the Tools menu, click **Split Part** or right-click the part and choose **Split Part**. The **Split Part Section Input Spreadsheet** appears (Figure 9-9).

Figure 9-9. Split Part Section Input Spreadsheet

Number	Name	Type	Order	Length	User Assig	I/O Bank	Voltage	I/O Standard	Location	Section
1	H1	clk	Input	0	Line				Left	1
2	G1	clkx2	Input	1	Line				Left	1
3	K13	CONF_DONE	Passive	2	Line				Left	2
4	C7	d[0]	Input	3	Line				Left	1
5	A6	d[1]	Input	4	Line				Left	1
6	D7	d[2]	Input	5	Line				Left	1
7	B7	d[3]	Input	6	Line				Left	1
8	B8	d[4]	Input	7	Line				Left	1
9	M7	d[5]	Input	8	Line				Left	1
10	A8	d[6]	Input	9	Line				Left	1
11	B6	d[7]	Input	10	Line				Left	1
12	H2	DATA0	Input	11	Line				Left	2
13	K4	DCLK	Bidirectional	12	Line				Left	2
14	C6	follow	Output	13	Line				Right	1
15	J3	MSEL0	Passive	14	Line				Left	2
16	J2	MSEL1	Passive	15	Line				Left	2
17	J4	nCE	Passive	16	Line				Left	2
18	H4	nCEO	Passive	17	Line				Left	2
19	H3	nCONFIG	Passive	18	Line				Left	2
20	H5	newt	Input	19	Line				Left	1
21	J13	nSTATUS	Passive	20	Line				Left	2
22	G16	reset	Input	21	Line				Left	1

Each row in the spreadsheet represents a pin in the symbol. The **Section** column indicates the section of the symbol to which each pin is assigned. You can locate all pins in a new symbol in section 1. You can change the values in the **Section** column to assign pins to various sections of the symbol. You can also specify the side of a section on the location of the pin by changing the values in the **Location** column. When you are ready, click **Split**. A new symbol appears in the same library as the original with the name *<original part name>_Split1*.

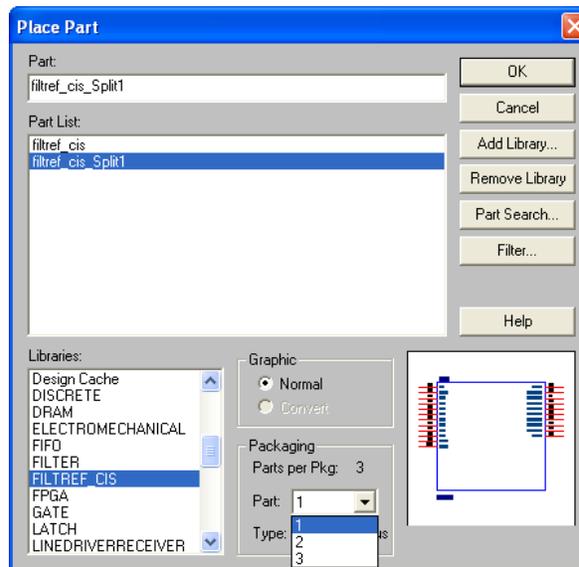
View and edit each section individually. To view the new sections of the part, double-click the part. The Part Symbol Editor window appears and the first section of the part displays for editing. On the View menu, click **Package** to view thumbnails of all the part sections. To edit the section of the symbol, double-click the thumbnail.

 For more information about splitting parts into sections and editing symbol sections in the Cadence Allegro Design Entry CIS software, refer to the Help in the software.

Instantiating a Symbol in a Design Entry CIS Schematic

After saving a new symbol in a library in your Cadence Allegro Design Entry CIS project, you can instantiate the new symbol on a page in your schematic. Open a schematic page in the Project Manager window of the Cadence Allegro Design Entry CIS software. To add the newly created symbol to your schematic on the schematic page, on the Place menu, click **Part**. The **Place Part** dialog box appears (Figure 9-10).

Figure 9-10. Place Part Dialog Box



Select the new symbol library location and the newly created part name. If you select a part that is split into sections, you can select the section to place from the **Part** pop-up menu. Click **OK**. The symbol attaches to your cursor for placement in the schematic. To place the symbol, click on the schematic page.



For more information about using the Cadence Allegro Design Entry CIS software, refer to the Help in the software.

Altera Libraries for the Cadence Allegro Design Entry CIS Software

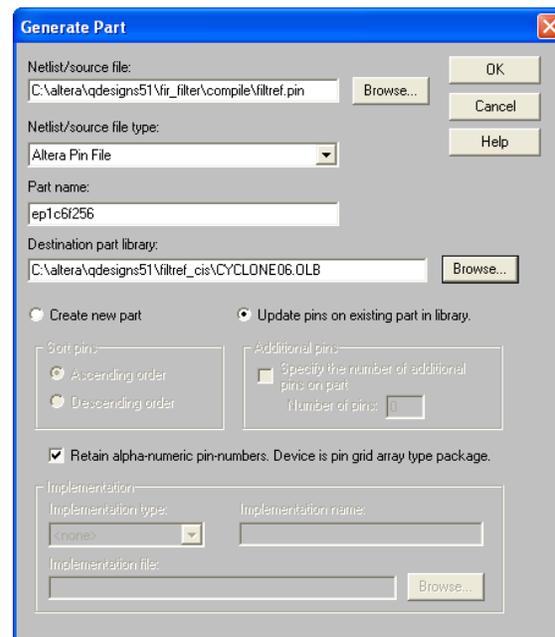
Altera provides downloadable **.olb** for many of its device packages. You can add these libraries to your Cadence Allegro Design Entry CIS project and update the symbols with the pin assignments contained in the **.pin** generated by the Quartus II software. You can use the downloaded library symbols as a base for creating custom schematic symbols with your pin assignments that you can edit or fracture. This method increases productivity by reducing the amount of time it takes to create and edit a new symbol.

To use the Altera-provided libraries with your Cadence Allegro Design Entry CIS project, follow these steps:

1. Download the library of your target device from the Download Center page found through the Support page on the Altera website (www.altera.com).

2. Create a copy of the appropriate **.olb** to maintain the original symbols. Place the copy in a convenient location, such as your Cadence Allegro Design Entry CIS project directory.
3. In the Project Manager window of the Cadence Allegro Design Entry CIS software, click once on the **Library** folder to select it. On the Edit menu, click **Project** or right-click the **Library** folder and choose **Add File** to select the copy of the downloaded **.olb** and add it to your project. You can locate the new library in the list of part libraries for your project.
4. On the Tools menu, click **Generate Part**. The **Generate Part** dialog box appears (Figure 9-11).

Figure 9-11. Generate Part Dialog Box



5. In the **Netlist/source file** field, click **Browse** to specify the **.pin** in your Quartus II design.
6. From the **Netlist/source file type** list, select **Altera Pin File**.
7. For **Part name**, type the name of the target device the same as it appears in the downloaded library file. For example, if you are using a device from the **CYCLONE06.OLB** library, type the part name to match one of the devices in this library such as **ep1c6f256**. You can rename the symbol in the Project Manager window after updating the part.
8. Set the **Destination part library** to the copy of the downloaded library you added to the project.
9. Select **Update pins on existing part in library**. Click **OK**.
10. Click **Yes**.

The symbol is updated with your pin assignments. Double-click the symbol in the Project Manager window to view and edit the symbol. On the View menu, click **Package** if you want to view and edit other sections of the symbol. If the symbol in the downloaded library is fractured into sections, you can edit each section but you cannot further fracture the part. You can generate a new part without using the downloaded part library if you require additional sections.

 For more information about creating, editing, and fracturing symbols in the Cadence Allegro Design Entry CIS software, refer to the Help in the software.

Conclusion

Transferring a complex, high-pin-count FPGA design to a PCB for prototyping or manufacturing is a daunting process and can lead to errors in the PCB netlist or design, especially when different engineers are working on different parts of the project. The design workflow available when the Quartus II software is used with tools from Cadence assists the FPGA designer and the board designer in preventing such errors and focusing all attention on the design.

Document Revision History

Table 9-3 shows the revision history for this chapter.

Table 9-3. Document Revision History

Date	Version	Changes
December 2010	10.0.1	Template update.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ General style editing. ■ Removed Referenced Document Section. ■ Added a link to Help in “Performing Simultaneous Switching Noise (SSN) Analysis of Your FPGA” on page 9-5.
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Added “Performing Simultaneous Switching Noise (SSN) Analysis of Your FPGA” on page 9-5. ■ General style editing. ■ Edited Figure 9-4 on page 9-10 and Figure 9-8 on page 9-16.
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Chapter 9 was previously Chapter 7 in the 8.1 software release. ■ No change to content.
November 2008	8.1.0	Changed to 8-1/2 x 11 page size.
May 2008	8.0.0	Updated references.

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter provides guidelines for reviewing printed circuit board (PCB) schematics with the Quartus® II software. Altera FPGAs and CPLDs offer a multitude of configurable options to allow you to implement a custom application-specific circuit on your PCB.

Your Quartus II project provides important information specific to your programmable logic design, which you can use in conjunction with the device literature available on Altera's website to ensure that you implement the correct board-level connections in your schematic.

This chapter highlights the important options in the Quartus II software, including **Settings** dialog box options, the Fitter report, and Messages window to which you should refer when creating and reviewing your PCB schematic. The Quartus II software also provides useful tools, such as the Pin Planner and the SSN Analyzer, to assist you during your PCB schematic review process.

The “[Reviewing Quartus II Software Settings](#)” section provides information about the settings you can make in the Quartus II software to help you review your PCB schematic. After verifying options in the Quartus II software, you can compile your design and use the data generated in the Fitter report, which is described in “[Reviewing Device Pin-Out Information in the Fitter Report](#)” on page 10–4 to verify settings in your PCB schematic. You should also ensure that you carefully review error and warning messages, as described in “[Reviewing Compilation Error and Warning Messages](#)” on page 10–5.

In addition to verifying your settings in the **Settings** dialog box and Fitter report, and checking messages, you can turn on additional settings, as described in “[Using Additional Quartus II Software Features](#)” on page 10–6 and “[Running the HardCopy Design Readiness Check](#)” on page 10–6.

Finally, Quartus II software tools, such as the Pin Planner and the SSN Analyzer, described in “[Using Additional Quartus II Software Tools](#)” on page 10–6, help you to verify proper I/O placement.

You should use this chapter in conjunction with Altera's device family-specific literature.



For more information, refer to the [Schematic Review Worksheets](#) and the [Pin Connection Guidelines](#) pages of the Altera.com website.

Reviewing Quartus II Software Settings

The **Device** dialog box in the Quartus II software allows you to specify device-specific assignments and settings. You can use the **Device** dialog box to specify general project-wide options, including specific device and pin options, which help you to implement correct board-level connections in your PCB schematic.

The **Device** dialog box provides project-specific device information, including the target device and any migration devices you specify. Using migration devices can impact the number of available user I/O pins and internal resources, as well as require connection of some user I/O pins to power/ground pins to support migration.

If you want to use vertical migration, which allows you to use different devices with the same package, you can specify your list of migration devices in the **Migration Devices** dialog box. The Fitter places the pins in your design based on your targeted migration devices, and allows you to use only I/O pins that are common to all of the migration devices.

- ④ For more information about the **Migration Devices** dialog box in the Quartus II software, refer to *Migration Devices Dialog Box* in Quartus II Help.

If a migration device has pins that are power or ground, but the pins are also user I/O pins on a different device in the migration path, the Fitter ensures that these pins are not used as user I/O pins. You must ensure that these pins are connected to the appropriate plane on the PCB.

If you are migrating from a smaller device with NC (no-connect) pins to a larger device with power or ground pins in the same package, you can safely connect the NC pins to power or ground pins to facilitate successful migration.

Device and Pins Options Dialog Box Settings

You can verify important design-specific data in the **Device and Pin Options** dialog box when reviewing your PCB schematic, including options found on the **Configuration**, **Unused Pin**, **Dual-Purpose Pins**, and **Voltage** pages.

Configuration Page Settings

The **Configuration** page of the **Device and Pin Options** dialog box specifies the configuration scheme and configuration device for the target device. Use the **Configuration** page settings to verify the configuration scheme with the MSEL pin settings used on your PCB schematic and the I/O voltage of the configuration scheme.

Your specific configuration settings may impact the availability of some dual-purpose I/O pins in user mode. Refer to “*Dual-Purpose Pins Page Settings*” on page 10-3 for more information.

Unused Pin Page Settings

The **Unused Pin** page specifies the behavior of all unused pins in your design. Use the **Unused Pin** page to ensure that unused pin settings are compatible with your PCB. For example, if you reserve all unused pins as outputs driving ground, you must ensure that you do not connect unused I/O pins to VCC pins on your PCB. Connecting unused I/O pins to VCC pins may result in contention that could lead to higher than expected current draw and possible device overstress.

The **Reserve all unused pins** list shows available unused pin state options for the target device. The default state for each pin is the recommended setting for each device family.

When you reserve a pin as output driving ground, the Fitter connects a ground signal to the output pin internally. You should connect the output pin to the ground plane on your PCB, although you are not required to do so. Connecting the output driving ground to the ground plane is known as creating a virtual ground pin, which helps to minimize simultaneous switching noise (SSN) and ground bounce effects.

Dual-Purpose Pins Page Settings

The **Dual-Purpose Pins** page specifies how configuration pins should be used after device configuration completes. You can set the function of the dual-purpose pins by selecting a value for a specific pin in the **Dual-purpose pins** list. Pin functions should match your PCB schematic. The available options on the **Dual-Purpose Pins** page may differ depending on the selected configuration mode.

Voltage Page Settings

The **Voltage** page specifies the default VCCIO I/O bank voltage and the default I/O bank voltage for the pins on the target device. VCCIO I/O bank voltage settings made in the **Voltage** page are overridden by I/O standard assignments made on I/O pins in their respective banks. Refer to the “[Reviewing Device Pin-Out Information in the Fitter Report](#)” on page 10-4 for more details about the I/O bank voltages for your design.

Error Detection CRC Page Settings

The **Error Detection CRC** page specifies error detection cyclic redundancy check (CRC) use for the target device. When **Enable error detection CRC** is turned on, the device checks the validity of the programming data in the devices. Any changes made in the data while the device is in operation generates an error.

Turning on the **Enable open drain on CRC error pin** option allows the CRC ERROR pin to be set as an open-drain pin in some devices, which decouples the voltage level of the CRC ERROR pin from VCCIO voltage. You must connect a pull-up resistor to the CRC ERROR pin on your PCB if you turn on this option.

In addition to settings in the **Device** dialog box, you should verify settings in the **Voltage** page of the **Settings** dialog box.

- For more information about the **Device and Pins Options** dialog box in the Quartus II software, refer to *Device and Pin Options Dialog Box* in Quartus II Help.

Voltage Page Settings

The **Voltage** page, under **Operating Settings and Conditions** in the **Settings** dialog box, allows you to specify voltage operating conditions for timing and power analyses. Ensure that the settings in the **Voltage** page match the settings in your PCB schematic, especially if the target device includes transceivers.

The **Voltage** page settings requirements differ depending on the settings of the transceiver instances in the design. Refer to the Fitter report for the required settings, and verify that the voltage settings are correctly set up for your PCB schematic.

- For more information about voltage settings, refer to the [Pin Connection Guidelines](#) page of the Altera.com website.

Once you verify your settings in the **Device** and **Settings** dialog boxes, you can verify your device pin-out with the Fitter report.

Reviewing Device Pin-Out Information in the Fitter Report

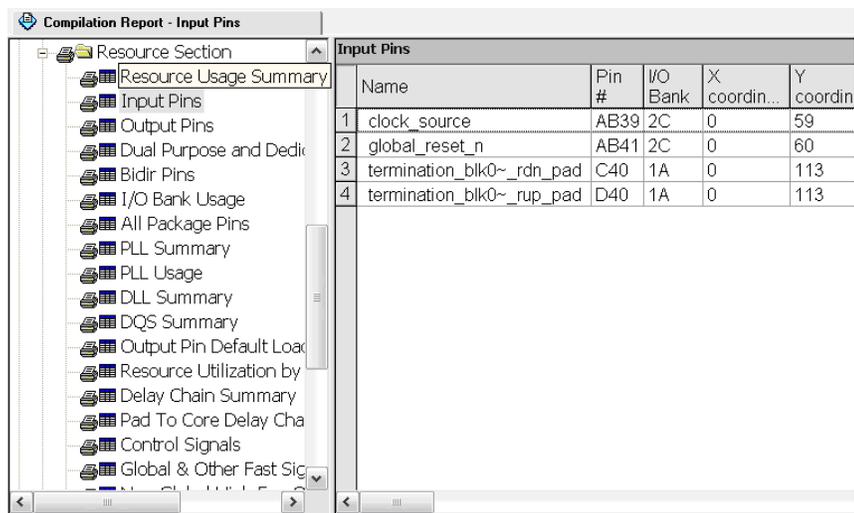
After you compile your design, you can use the reports in the Resource section of the Fitter report to check your device pin-out in detail.

The Input Pins, Output Pins, and Bidirectional Pins reports identify all the user I/O pins in your design and the features enabled for each I/O pin. For example, you can find use of weak internal pull-ups, PCI clamp diodes, and on-chip termination (OCT) pin assignments in these sections of the Fitter report. You can check the pin assignments reported in the Input Pins, Output Pins, and Bidirectional Pins reports against your PCB schematic to determine whether your PCB requires external components.

These reports also identify whether you made pin assignments or if the Fitter automatically placed the pins. If the Fitter changed your pin assignments, you should make these changes user assignments because the location of pin assignments made by the Fitter may change with subsequent compilations.

Figure 10-1 shows the pins the Fitter chose for the OCT external calibration resistor connections (RUP/RDN) and the name of the associated termination block in the Input Pins report. You should make these types of assignments user assignments.

Figure 10-1. Resource Section Report



Compilation Report - Input Pins					
Resource Section					
Input Pins					
	Name	Pin #	I/O Bank	X coordin...	Y coordin
1	clock_source	AB39	2C	0	59
2	global_reset_n	AB41	2C	0	60
3	termination_blk0~_rdn_pad	C40	1A	0	113
4	termination_blk0~_rup_pad	D40	1A	0	113

The I/O Bank Usage report provides a high-level overview of the VCCIO and VREF requirements for your design, based on your I/O assignments. Verify that the requirements in this report match the settings in your PCB schematic. All unused I/O banks, and all banks with I/O pins with undefined I/O standards, default the VCCIO voltage to the voltage defined in the **Voltage** page of the **Device and Pin Options** dialog box.

The All Package Pins report lists all the pins on your device, including unused pins, dedicated pins and power/ground pins. You can use this report to verify pin characteristics, such as the location, name, usage, direction, I/O standard and voltage for each pin with the pin information in your PCB schematic. In particular, you should verify the recommended voltage levels at which you connect unused dedicated inputs and I/O and power pins, especially if you selected a migration device. Use the All Package Pins report to verify that you connected all the device voltage rails to the voltages reported.

Errors commonly reported include connecting the incorrect voltage to the predriver supply (VCCPD) pin in a specific bank, or leaving dedicated clock input pins floating. Unused input pins that should be connected to ground are designated as **GND+** in the **Pin Name/Usage** column in the All Package Pins report.

You can also use the All Package Pins report to check transceiver-specific pin connections and verify that they match the PCB schematic. Unused transceiver pins have the following requirements, based on the pin designation in the Fitter report:

- **GXB_GND***—Unused GXB receiver or dedicated reference clock pin. This pin must be connected to GXB_GND through a 10k Ohm resistor.
- **GXB_NC**—Unused GXB transmitter or dedicated clock output pin. This pin must be disconnected.

Some transceiver power supply rails have dual voltage capabilities, such as VCCA_L/R and VCCH_L/R, that depend on the settings you created for the ALTGX MegaWizard Plug-In Manager. Because these user-defined settings overwrite the default settings, you should use the All Package Pins report to verify that these power pins on the device symbol in the PCB schematics are connected to the voltage required by the transceiver. An incorrect connection may cause the transceiver to function not as expected.

If your design includes a memory interface, the DQS Summary report provides an overview of each DQ pin group. You can use this report to quickly confirm that the correct DQ/DQS pins are grouped together. This section also provides information on DLL usage.

Finally, the Fitter Device Options report summarizes some of the settings made in the **Device and Pin Options** dialog box. Verify that these settings match your PCB schematics.

Reviewing Compilation Error and Warning Messages

If your project does not compile without error or warning messages, you should resolve the issues identified by the Compiler before signing off on your pin-out or PCB schematic. Error messages often indicate illegal or unsupported use of the device resources and IP.

Additionally, you should cross-reference fitting and timing analysis warnings with the design implementation. Timing may be constrained due to nonideal pin placement. You should investigate if you can reassign pins to different locations to prevent fitting and timing analysis warnings. Ensure that you review each warning and consider its potential impact on the design.

Running the HardCopy Design Readiness Check

You should run the HardCopy Design Readiness Check for designs including a HardCopy device in the migration path. This tool checks for issues that must be addressed prior to handing off the design to the Altera HardCopy Design Center for the HardCopy back-end process, including possible issues with device resource and I/O usage. Verify all warning messages generated during the HardCopy Design Readiness Check.

- For more information about the HardCopy Design Readiness Check and designing for HardCopy devices in the Quartus II software, refer to the *Quartus II Support for HardCopy Series Devices* chapter in volume 1 of the *Quartus II Handbook*.

Using Additional Quartus II Software Features

You can generate IBIS files, which contain models specific to your design and selected I/O standards and options, with the Quartus II software.

Because board-level simulation is important to verify, you should check for potential signal integrity issues. You can turn on the **Board-Level Signal Integrity** feature in the **EDA Tool Settings** page of the **Settings** dialog box.

- For more information about signal integrity analysis in the Quartus II software, refer to the *Signal Integrity Analysis with Third-Party Tools* chapter in volume 3 of the *Quartus II Handbook*.

Additionally, using advanced I/O timing allows you to enter physical PCB information to accurately model the load seen by an output pin. This feature facilitates accurate I/O timing analysis.

- For more information about advanced I/O timing, refer to the *I/O Management* chapter in volume 2 of the *Quartus II Handbook*.

Using Additional Quartus II Software Tools

This section describes additional tools found in the Quartus II software, specifically the Pin Planner and the SSN Analyzer, and how you can use these tools to assist you with reviewing your PCB schematics.

Pin Planner

The Pin Planner provides a graphical representation of the target device and is a useful visual aid that shows the device package and I/O assignments. The Pin Planner allows you to view the I/O banks, VREF groups, edges and DQ/DQS pin groups, which enables you to verify the expected placement of pin groups.

You can use the Pin Planner to verify the location of clock inputs, and whether they have been placed on dedicated clock input pins, which is recommended when your design uses PLLs.

You can also use the Pin Planner to verify the placement of dedicated SERDES pins. SERDES receiver inputs can be placed only on DIFFIO_RX pins, while SERDES transmitter outputs can be placed only on DIFFIO_TX pins.

The Pin Planner gives a visual indication of signal-to-signal proximity in the Pad View window, and also provides information about differential pin pair placement, such as the placement of pseudo-differential signals.

- For more information about the Pin Planner, refer to the *I/O Management* chapter in volume 2 of the *Quartus II Handbook*.

SSN Analyzer

The SSN Analyzer supports pin planning by estimating the voltage noise caused by the simultaneous switching of output pins on the device. Because of the importance of the potential SSN performance for a specific I/O placement, you can use the SSN Analyzer to analyze the effects of aggressor I/O signals on a victim I/O pin.

- For more information about the SSN Analyzer, refer to the *Simultaneous Switching Noise (SSN) Analysis and Optimizations* chapter in volume 2 of the *Quartus II Handbook*.

Conclusion

This chapter describes guidelines and descriptions of settings to verify when reviewing your PCB schematic with the Quartus II software. You can use settings in the **Settings** dialog box; information in the Fitter report and Messages window; and the Pin Planner and SSN Analyzer during the PCB schematic review process.

Document Revision History

Table 10-1 shows the revision history for this chapter.

Table 10-1. Document Revision History

Date	Version	Changes
December 2010	10.0.1	Changed to new document template. No change to content.
July 2010	10.0.0	Initial release.

- For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).
- Take an [online survey](#) to provide feedback about this handbook chapter.

This section introduces features in the Quartus® II software that you can use to optimize area, timing, power, and compilation time when you design for programmable logic devices (PLDs).

This section includes the following chapters:

- [Chapter 11, Design Optimization Overview](#)

This chapter summarizes features in the Quartus II software that you can use to achieve the highest design performance when you design for PLDs, especially high density FPGAs.

- [Chapter 12, Reducing Compilation Time](#)

This chapter describes techniques for reducing the amount of time it takes to compile and recompile your design, accelerating your design process.

- [Chapter 13, Area and Timing Optimization](#)

This chapter describes a broad spectrum of Quartus II software features and design techniques to reduce resource usage and improve timing performance when designing for Altera® devices. This chapter also explains how and when to use some of the features described in other chapters of the *Quartus II Handbook*.

- [Chapter 14, Power Optimization](#)

This chapter describes the power-driven compilation feature and flow in detail, as well as low power design techniques that can further reduce power consumption in your design.

- [Chapter 15, Analyzing and Optimizing the Design Floorplan](#)

You can use the Chip Planner to perform design analysis and create a design floorplan. This chapter discusses how to analyze and optimize the design floorplan with the Chip Planner.

- [Chapter 16, Netlist Optimizations and Physical Synthesis](#)

This chapter explains how the physical synthesis optimizations in the Quartus II software can improve your quality of results. This chapter also provides information about preserving and writing out a new netlist, and provides guidelines for applying the various options.

This chapter introduces features in Altera's Quartus® II software that you can use to achieve the highest design performance when you design for programmable logic devices (PLDs), especially high density FPGAs.

Introduction

Physical implementation can be an intimidating and challenging phase of the design process. The Quartus II software provides a comprehensive environment for FPGA designs, delivering unmatched performance, efficiency, and ease-of-use.

In a typical design flow, you must synthesize your design with Quartus II integrated synthesis or a third-party tool, place and route your design with the Fitter, and use the TimeQuest timing analyzer to ensure your design meets the timing requirements. With the PowerPlay Power Analyzer, you ensure the design's power consumption is within limits. .

Physical Implementation

Most optimization issues involve preserving previous results, reducing area, reducing critical path delay, reducing power consumption, and reducing runtime. The Quartus II software includes advisors to address each of these issues and helps you optimize your design. Run these advisors during physical implementation for advice about your specific design.

You can reduce the time spent on design iterations by following the recommended design practices for designing with Altera® devices. Design planning is critical for successful design timing implementation and closure.



For more information, refer to the *Design Planning with the Quartus II Software* chapter in volume 1 of the *Quartus II Handbook*.

Trade-Offs and Limitations

Many optimization goals can conflict with one another, so you might need to make trade-offs between different goals. For example, one major trade-off during physical implementation is between resource usage and critical path timing, because certain techniques (such as logic duplication) can improve timing performance at the cost of increased area. Similarly, a change in power requirements can result in area and timing trade-offs, such as if you reduce the number of high-speed tiles available, or if you attempt to shorten high-power nets at the expense of critical path nets.

In addition, system cost and time-to-market considerations can affect the choice of device. For example, a device with a higher speed grade or more clock networks can facilitate timing closure at the expense of higher power consumption and system cost.

Finally, not all designs can be realized in a hardware circuit with limited resources and given constraints. If you encounter resource limitations, timing constraints, or power constraints that cannot be resolved by the Fitter, consider rewriting parts of the HDL code.

 For more information, refer to the *Area and Timing Optimization* chapter in volume 2 of the *Quartus II Handbook*.

Preserving Results and Enabling Teamwork

For some Quartus II Fitter algorithms, small changes to the design can have a large impact on the final result. For example, a critical path delay can change by 10% or more because of seemingly insignificant changes. If you are close to meeting your timing objectives, you can use the Fitter algorithm to your advantage by changing the fitter seed, which changes the pseudo-random result of the Fitter.

Conversely, if you cannot meet timing on a portion of your design, you can partition that portion and prevent it from recompiling if an unrelated part of the design is changed. This feature, known as incremental compilation, can reduce the Fitter runtimes by up to 70% if the design is partitioned, such that only small portions require recompilation at any one time.

When you use incremental compilation, you can apply design optimization options to individual design partitions and preserve performance in other partitions by leaving them untouched. Many optimization techniques often result in longer compilation times, but by applying them only on specific partitions, you can reduce this impact and complete iterations more quickly.

In addition, by physically floorplanning your partitions with LogicLock regions, you can enable team-based flows and allow multiple people to work on different portions of the design.

 For more information, refer to *Quartus II Incremental Compilation for Hierarchical and Team-Based Designs* in volume 1 of the *Quartus II Handbook* and *About Incremental Compilation* in Quartus II Help.

Reducing Area

By default, the Quartus II Fitter might physically spread a design over the entire device to meet the set timing constraints. If you prefer to optimize your design to use the smallest area, you can change this behavior. If you require reduced area, you can enable certain physical synthesis options to modify your netlist to create a more area-efficient implementation, but at the cost of increased runtime and decreased performance.

 For more information, refer to the *Area and Timing Optimization* and *Netlist Optimizations and Physical Synthesis* chapters in volume 2 and the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*.

Reducing Critical Path Delay

To meet complex timing requirements involving multiple clocks, routing resources, and area constraints, the Quartus II software offers a close interaction between synthesis, timing analysis, floorplan editing, and place-and-route processes.

By default, the Quartus II Fitter tries to meet the specified timing requirements and stops trying when the requirements are met. Therefore, using realistic constraints is important to successfully close timing. If you under-constrain your design, you may get sub-optimal results. By contrast, if you over-constrain your design, the Fitter might over-optimize non-critical paths at the expense of true critical paths. In addition, you might incur an increased area penalty. Compilation time may also increase because of excessively tight constraints.

If your resource usage is very high, the Quartus II Fitter might have trouble finding a legal placement. In such circumstances, the Fitter automatically modifies some of its settings to try to trade off performance for area.

The Quartus II Fitter offers a number of advanced options that can help you improve the performance of your design when you properly set constraints. Use the Timing Optimization Advisor to determine which options are best suited for your design.

If you use incremental compilation, you can help resolve inter-partition timing requirements by locking down the results one partition at a time or by guiding the placement of the partitions with LogicLock regions. You might be able to improve the timing on such paths by placing the partitions optimally to reduce the length of critical paths. Once your inter-partition timing requirements are met, use incremental compilation to preserve the results and work on partitions that have not met timing requirements.

In high-density FPGAs, routing accounts for a major part of critical path timing. Because of this, duplicating or retiming logic can allow the Fitter to reduce delay on critical paths. The Quartus II software offers push-button netlist optimizations and physical synthesis options that can improve design performance at the expense of considerable increases of compilation time and area. Turn on only those options that help you keep reasonable compilation times and resource usage. Alternately, you can modify your HDL to manually duplicate or retime logic.

Reducing Power Consumption

The Quartus II software has features that help reduce design power consumption. The PowerPlay power optimization options control the power-driven compilation settings for Synthesis and the Fitter.



For more information, refer to the *Power Optimization* chapter in volume 2 of the *Quartus II Handbook*.

Reducing Runtime

Many Fitter settings influence compilation time. Most of the default settings in the Quartus II software are set for reduced compilation time. You can modify these settings based on your project requirements.

The Quartus II software supports parallel compilation in computers with multiple processors. This can reduce compilation times by up to 15% while giving the identical result as serial compilation.

You can also reduce compilation time with your iterations by using incremental compilation. Use incremental compilation when you want to change parts of your design, while keeping most of the remaining logic unchanged.

Using Quartus II Tools

The following sections describe several Quartus II tools that you can use to help optimize your design.

Design Analysis

The Quartus II software provides tools that help with a visual representation of your design. You can use the RTL Viewer to see a schematic representation of your design before synthesis and place-and-route. The Technology Map Viewer provides a schematic representation of the design implementation in the selected device architecture after synthesis and place-and-route. It can also include timing information.

With incremental compilation, the Design Partition Planner and the Chip Planner allow you to partition and layout your design at a higher level. In addition, you can perform many different tasks with the Chip Planner, including: making floorplan assignments, implementing engineering change orders (ECOs), and performing power analysis. Also, you can analyze your design and achieve a faster timing closure with the Chip Planner. The Chip Planner provides physical timing estimates, critical path display, and routing congestion view to help guide placement for optimal performance.



For more information, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Designs* and *Best Practices for Incremental Compilation Partitions and Floorplan Assignments* chapters in volume 1 and the *Engineering Change Management with the Chip Planner* chapter in volume 2 of the *Quartus II Handbook*.

Advisors

The Quartus II software includes several advisors to help you optimize your design and reduce compilation time. You can complete your design faster by following the recommendations in the Compilation Time Advisor, Incremental Compilation Advisor, Timing Optimization Advisor, Area Optimization Advisor, Resource Optimization Advisor, and Power Optimization Advisor. These advisors give recommendations based on your project settings and your design constraints.



For more information about advisors, refer to Quartus II Help.

Design Space Explorer

Use the Design Space Explorer (DSE) to find optimal settings in the Quartus II software. DSE automatically tries different combinations of netlist optimizations and advanced Quartus II software compiler settings, and reports the best settings for your design, based on your chosen primary optimization goal. You can try different seeds with the DSE if you are fairly close to meeting your timing or area requirements and find one seed that meets timing or area requirements. Finally, the DSE can run the different compilations on multiple computers in parallel, which shortens the timing closure process.

- For more information, refer to [About Design Space Explorer](#) in Quartus II Help.

Conclusion

The Quartus II software includes a number of features and tools that you can use to optimize area, timing, power, and compilation time when you design for programmable logic devices (PLDs).

Document Revision History

Table 11-1 shows the revision history for this chapter.

Table 11-1. Document Revision History

Date	Version	Changes
December 2010	10.0.2	Changed to new document template. No change to content.
August 2010	10.0.1	Corrected link
July 2010	10.0.0	Initial release. Chapter based on topics and text in Section III of volume 2.

- For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).
- Take an [online survey](#) to provide feedback about this handbook chapter.

The Quartus® II software offers a number of features and techniques to help reduce compilation time.

This chapter describes techniques to reduce compilation time when designing for Altera® devices, and includes the following topics:

- “Compilation Time Optimization Techniques”
- “Compilation Time Advisor” on page 12-2
- “Strategies to Reduce the Overall Compilation Time” on page 12-2
- “Reducing Synthesis Time and Synthesis Netlist Optimization Time” on page 12-5
- “Reducing Placement Time” on page 12-7
- “Reducing Routing Time” on page 12-8
- “Reducing Static Timing Analysis Time” on page 12-9
- “Setting Process Priority” on page 12-10

Compilation Time Optimization Techniques

In long compilations, most of the time is spent in the Analysis and Synthesis and Fitter modules. Analysis and Synthesis includes physical synthesis optimizations done during synthesis, if you have turned on physical synthesis optimizations. The Fitter includes two steps, placement and routing, and also includes physical synthesis if you turned on the physical synthesis option with **Normal** or **Extra** effort levels. The **Flow Elapsed Time** section of the Compilation Report shows how much time is spent running the Analysis and Synthesis and Fitter modules. The Fitter Messages report in the **Fitter** section of the Compilation Report shows the time that was spent in placement and the time that was spent in routing.

Placement is the process of finding optimum locations for the logic in your design. Placement includes Quartus II pre-Fitter operations, which place dedicated logic such as clocks, PLLs, and transceiver blocks. Routing is the process of connecting the nets between the logic in your design. There are many possible placements for the logic in a design, and finding better placements typically uses more compilation time. Good logic placement allows you to more easily meet your timing requirements and makes the design easier to route.



The applicable messages are indicated as shown in the following example, with each time component in two-digit format, and days shown only if applicable:

```
Info: Fitter placement operations ending: elapsed time =
<days:hours:mins:secs>
Info: Fitter routing operations ending: elapsed time =
<days:hours:mins:secs>
```

While the Fitter is running (including Placement and Routing), messages similar to the following message are displayed every hour to indicate Fitter operations are progressing normally.

```
Info: Placement optimizations have been running for 4 hour(s)
```

Compilation Time Advisor

A Compilation Time Advisor is available in the Quartus II software, which helps you to reduce compilation time. Run the Compilation Time Advisor on the Tools menu by pointing to **Advisors** and clicking **Compilation Time Advisor**. You can find all the compilation time optimizing techniques described in this section in the Compilation Time Advisor as well.

Strategies to Reduce the Overall Compilation Time

This section discusses strategies you can use to reduce overall compilation time, including the following topics:

- [“Using Parallel Compilation with Multiple Processors”](#)
- [“Using Incremental Compilation”](#) on page 12-4
- [“Using the Smart Compilation Setting”](#) on page 12-4
- [“Using Rapid Recompile”](#) on page 12-4

Using Parallel Compilation with Multiple Processors

The Quartus II software can run some algorithms in parallel to take advantage of multiple processors and reduce compilation time when more than one processor is available. Parallel compilation is turned on by default in the Quartus II software and the software can detect if multiple processors are available. You can also specify the maximum number of processors that the software can use if you want to reserve some of the available processors for other tasks. The Quartus II software supports up to 16 processors. The software does not necessarily use all the processors that you specify during a given compilation, but it never uses more than the specified number of processors. This allows you to work on other tasks on your computer without it becoming slow or less responsive.

If you have partitioned your design and enabled parallel compilation, the Quartus II software can use different processors to compile those partitions simultaneously during the Analysis and Synthesis stage. This may result in high peak memory usage during Analysis and Synthesis.

By allowing the Quartus II software to use two processors, you can reduce the compilation time by up to 10% on systems with two processing cores and by up to 20% on systems with four cores. With certain design flows in which timing analysis runs alone, using multiple processors can reduce the time required for timing analysis by an average of 10% when using two processors. This reduction can reach an average of 15% when using four processors.

The actual reduction in compilation time depends on the design and on the specific settings used for compilation. For example, compilations with multi-corner optimization turned on benefit more from using multiple processors than do compilations that do not use multi-corner optimization. The runtime requirement is not reduced for some other compilation goals, such as Analysis and Synthesis. The Fitter (`quartus_fit`) and the Quartus II TimeQuest Timing Analyzer (`quartus_sta`) stages in the compilation can, in certain cases, benefit from the use of multiple processors. The average number of processors used for these stages is shown in the Compilation Report, on the **Flow Elapsed Time** panel. A more detailed breakdown of processor usage is also shown in the Parallel Compilation panel of the appropriate report, such as the Fit report. This panel is only displayed if parallel compilation is enabled.

This feature is available for Arria® series, Cyclone®, HardCopy III, HardCopy IV, MAX® II, MAX V (limited support), and Stratix® series devices.



Do not consider processors with Intel Hyper-Threading to be more than one processor. If you have a single processor with Intel Hyper-Threading enabled, you should set the number of processors to one. Altera recommends that you do not use the Intel Hyper-Threading feature for Quartus II compilations, as it can increase runtimes.

The Quartus II software can detect the number of processors available on a computer and use available processors to reduce compilation time. You can also control the number of processors used during a compilation on a per user basis.

- ② For more information, refer to *Processing Page (Options Dialog Box)* in Quartus II Help.
- ② For information about how to control the number of processors used during compilation for a specific project, refer to *Compilation Process Settings Page (Settings Dialog Box)* in Quartus II Help.

Using multiple processors does not affect the quality of the fit. For a given Fitter seed on a specific design, the fit is exactly the same, regardless of whether the Quartus II software uses one processor or multiple processors. The only difference between compilations using a different number of processors is the compilation time.

You can also set the number of processors available for Quartus II compilation using the following Tcl command in your script.

```
set_global_assignment -name NUM_PARALLEL_PROCESSORS <value> ←
```

In this case, *<value>* is an integer from 1 to 16.

If you want the Quartus II software to detect the number of processors and use all of them for running the compilation, use the following Tcl command in your script:

```
set_global_assignment -name NUM_PARALLEL_PROCESSORS ALL ←
```

Using Incremental Compilation

The incremental compilation feature can speed up design iteration time by up to 70% for small design changes, and helps you reach design timing closure more efficiently. Design iterations can be made faster by recompiling only a particular design partition and merging results with previous compilation results from other partitions. You can also use physical synthesis optimization techniques for specific design partitions while leaving other parts of the design untouched to preserve performance.

If you are using a third-party synthesis tool, you can create separate atom netlist files for parts of your design that you already have synthesized and optimized so that you update only the parts of the design that change.

In the standard incremental compilation design flow, the top-level design is divided into partitions, which can be compiled and optimized together in the top-level Quartus II project. You can preserve fitting results and performance for completed partitions while other parts of the design are changing, which reduces the compilation time for each design iteration because neither synthesis nor fitting is performed for unchanged partitions in the design.

The incremental compilation feature also facilitates team-based incremental compilation design flows that allow designers to create and optimize design blocks independently, when necessary, and support third-party IP integration.



For information about the full incremental compilation flow in the Quartus II software, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*. For information about creating multiple netlist files in third-party tools for use with incremental compilation, refer to the appropriate chapter in *Section III. Synthesis* in volume 1 of the *Quartus II Handbook*.



For additional information about incremental compilation, refer to *About Incremental Compilation* in Quartus II Help.

Using the Smart Compilation Setting

Smart compilation can reduce compilation time by skipping compiler stages that are not required to recompile the design. This is especially useful when you perform multiple compilation iterations during the optimization phase of the design process. However, smart compilation uses more disk space. To turn on smart compilation, on the Assignments menu, click **Settings**. In the **Category** list, select **Compilation Process Settings** and turn on **Use smart compilation**.



Smart compilation skips entire compiler stages (such as Analysis and Synthesis) when they are not required. This feature is different from incremental compilation, which you can use to compile parts of your design while preserving results for unchanged parts.

Using Rapid Recompile

The Rapid Recompile feature maximizes designer productivity when making small engineering change order (ECO)-style design changes after a full compilation, reducing compilation times by an average of 50%. Rapid Recompile also significantly improves designer productivity during timing closure by preserving critical timing during late design changes.

The Rapid Recompile option can be used stand-alone or along with standard incremental flow for compatible nodes in your design. A compatible node is a node that can be matched to a node from previous compilation results. Rapid Recompile allows the Quartus II software to reuse placement and routing resources of compatible nodes from previous results with a high degree of confidence.

If you enable the Rapid Recompile feature, you will see the compile time reduction after a full compilation. Turn on the Rapid Recompile feature in later compilations to see further reductions. The **Incremental Compilation Preservation Summary** section in the Fitter Report provides details about the placement and routing preservation for your design.

The performance of Rapid Recompile is largely dependent on the nature of the design change. If the Quartus II software determines that full optimization is necessary for design performance, you may not see much compilation time reduction. For example, if the total time taken by the Fitter is dominated by the time taken for fitter preparation operations, using this feature may not save you a lot of compilation time. Applying extensive, global optimizations to a small user change may be required to obtain optimal performance. Be sure to select the right flow to achieve your end goals.



If you see the message Fitter has failed to locate previous placement information during the compilation of your design, Rapid Recompile will not provide any compile time reduction.



For more information about this feature, refer to *Incremental Compilation Page (Settings Dialog Box)* in Quartus II Help.

Reducing Synthesis Time and Synthesis Netlist Optimization Time

You can reduce synthesis time by reducing your use of netlist optimizations and by using incremental compilation (with Netlist Type set to **Post-Synthesis**) without affecting the Fitter time. For tips for reducing synthesis time when using third-party EDA synthesis tools, refer to your synthesis software's documentation.

Settings to Reduce Synthesis Time and Synthesis Netlist Optimization Time

You can use Quartus II integrated synthesis to synthesize and optimize HDL designs, and you can use synthesis netlist optimizations to optimize netlists that were synthesized by third-party EDA software. When using Quartus II Integrated Synthesis, you can also enable specific Physical Synthesis Optimizations during Analysis and Synthesis. Using these netlist optimizations can cause the Analysis and Synthesis module to take much longer to run. Read the Analysis and Synthesis messages to find out how much time these optimizations take. The compilation time spent in Analysis and Synthesis is usually small compared to the compilation time spent in the Fitter.

If your design meets your performance requirements without synthesis netlist optimizations, turn off the optimizations to save time. If you require synthesis netlist optimizations to meet performance, you can optimize parts of your design hierarchy separately to reduce the overall time spent in analysis and synthesis.

Turn off settings that are not useful. In general, if you carry over compilation settings from a previous project, evaluate all settings and keep only those that you need.

Use Appropriate Coding Style to Reduce Synthesis Time

The way you code your design in HDL can effect the synthesis time. For example, if you want to infer RAM blocks from your code, you must follow the guidelines for inferring RAMs. If not, those blocks are implemented as registers, and if you are trying to infer a large memory, it will use a large amount of resources within the FPGA. This will cause routing congestion and increase compilation time drastically. In general, if you see very high routing utilizations in some blocks, it is a good idea to review the code for such blocks.



For more information about coding guidelines, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*.

Using Early Timing Estimation

The Quartus II software provides an Early Timing Estimation feature that estimates your design's timing results before the software performs full placement and routing. On the Processing menu, point to **Start**, and click **Start Early Timing Estimate** to generate initial compilation results after you have run analysis and synthesis. When you want a quick estimate of a design's performance before proceeding with further design or synthesis tasks, this command can save significant compilation time. Using this feature provides a timing estimate 2.5× faster (on average) than running a full compilation (8.5× faster in best case), although the fit is not fully optimized or routed. Therefore, the timing report is only an estimate. On average, the estimated delays are within 15% of the final timing results as achieved by a full compilation.

You can specify the type of delay estimates to use with Early Timing Estimation. On the Assignments menu, click **Settings**. In the **Category** list, select **Compilation Process Settings**, and select **Early Timing Estimate**. On the **Early Timing Estimate** page, the following options are available:

- The **Realistic** option, which is the default, generates delay estimates that are likely to be close to the results of a full compilation.
- The **Optimistic** option uses delay estimates that are lower than those likely to be achieved by a full compilation, which results in an optimistic performance estimate.
- The **Pessimistic** option uses delay estimates that are higher than those likely to be achieved by a full compilation, which results in a pessimistic performance estimate.

All three options offer the same reduction in compilation time.

You can view the critical paths in the design by locating these paths in the Chip Planner from the **TimeQuest Timing Report** panel. Then, if necessary, you can add or modify floorplan constraints such as LogicLock regions, or make other changes to the design. You can then rerun the **Early Timing Estimate** to quickly assess the impact of any floorplan assignments or logic changes, enabling you to try different design variations and find the best solution.

Reducing Placement Time

The time required to place a design depends on two factors: the number of ways the logic in the design can be placed in the device and the settings that control how hard the placer works to find a good placement. You can reduce the placement time in two ways:

- Change the settings for the placement algorithm
- Use incremental compilation to preserve the placement for parts of the design

Sometimes there is a trade-off between placement time and routing time. Routing time can increase if the placer does not run long enough to find a good placement. When you reduce placement time, make sure that it does not increase routing time and negate the overall time reduction.

Fitter Effort Setting

The highest fitter effort setting, **Standard Fit**, takes the most runtime, but does not always yield a better result than using the default **Auto Fit**. For designs with very tight timing requirements, both **Auto Fit** and **Standard Fit** use the maximum effort during optimization. Using **Auto Fit** is recommended for reducing compilation time. If you are certain that your design has only easy-to-meet timing constraints, you can select **Fast Fit** for an even greater runtime savings.

Placement Effort Multiplier Settings

You can control the amount of time the Fitter spends in placement by reducing one aspect of placement effort with the **Placement Effort Multiplier** option. On the Assignments menu, click **Settings**. Select **Fitter Settings**, and click **More Settings**. Under **Existing Option Settings**, select **Placement Effort Multiplier**. The default is 1.0. Legal values must be greater than 0 and can be non-integer values. Numbers between 0 and 1 can reduce fitting time, but also can reduce placement quality and design performance. Numbers higher than 1 increase placement time and placement quality, but can reduce routing time for designs with routing congestion. For example, a value of 4 increases placement time by approximately 2 to 4 times, but might result in better placement, which can result in reduced routing time.

Final Placement Optimization Levels

The **Final Placement Optimization Level** option specifies whether the Fitter performs final placement optimizations. This can be set to **Always**, **Never**, and **Automatically**. Performing optimizations can improve register-to-register timing and fitting, but might require longer compilation times. The default setting of **Automatically** can be used with the **Auto Fit** Fitter Effort Level (also the default) to let the Fitter decide whether these optimizations should run based on the routability and timing requirements of the design.

Setting the **Final Placement Optimization Level** to **Never** often reduces your compilation time, but typically affects routability negatively and reduces timing performance.

To change the **Final Placement Optimization Level**, on the Assignments menu, click **Settings**. The **Settings** dialog box appears. From the **Category** list, select **Fitter Settings**, and then click the **More Settings** button. Select **Final Placement Optimization Level**, and then from the list, select the required setting.

Physical Synthesis Effort Settings

You can use the physical synthesis options to optimize your post-synthesis netlist and improve your timing performance. These options, which affect placement, can significantly increase compilation time.

If your design meets your performance requirements without physical synthesis options, turn them off to save time. You also can use the **Physical synthesis effort** setting on the **Physical Synthesis Optimizations** page under **Compilation Process Settings** in the **Category** list to reduce the amount of extra compilation time that these optimizations use. The **Fast** setting directs the Quartus II software to use a lower level of physical synthesis optimization that, compared to the **Normal** physical synthesis effort level, can cause a smaller increase in compilation time. However, the lower level of optimization can result in a smaller increase in design performance.

Limit to One Fitting Attempt

This option causes the software to quit after one fitting attempt option, instead of repeating placement and routing with increased effort. For hard to fit designs, consider increasing the **Placement Effort Multiplier** setting and the **Limit to One Fitting Attempt** setting. This saves you time, because if the design is hard to fit and does not result in a valid fit, the compilation stops after the first attempt..

From the Assignments menu, select **Settings**. On the **Fitter Settings** page, turn on **Limit to one fitting attempt**.

 For more details about this option, refer to “Limit to One Fitting Attempt” in the *Area and Timing Optimization* chapter in volume 2 of the Quartus II Handbook.

Preserving Placement, Incremental Compilation, and LogicLock Regions

Preserving information about previous placements can make future placements faster. The incremental compilation feature provides an easy-to-use methodology for preserving placement results. For more information, refer to “Using Incremental Compilation” on page 12-4.

Reducing Routing Time

The time required to route a design depends on three factors: the device architecture, the placement of the design in the device, and the connectivity between different parts of the design. Typically, the routing time is not a significant amount of the compilation time. If your design takes a long time to route, perform one or more of the following actions:

- Check for routing congestion
- Let the placer run longer to find a more routable placement
- Use incremental compilation to preserve routing information for parts of your design

Identifying Routing Congestion in the Chip Planner

To identify areas of routing congestion in your design, open the Chip Planner. On the Tools menu, click **Chip Planner**. To view the routing congestion in the Chip Planner, click the Layers icon located next to the Task menu. Under **Background Color Map**, select **Routing Utilization**. Routing resource usage above 90% indicates routing congestion. You can change the connections in your design to reduce routing congestion. If the area with routing congestion is in a LogicLock region or between LogicLock regions, change or remove the LogicLock regions and recompile the design. If the routing time remains the same, the time is a characteristic of the design and the placement. If the routing time decreases, consider changing the size, location, or contents of LogicLock regions to reduce congestion and decrease routing time.

In some cases, routing congestion may be a result of the HDL coding style used in the design. Once you identify areas that are congested using the Chip Planner, review the HDL code for the blocks placed in those areas to see if interconnect usage can be reduced by code changes.

The Quartus II compilation messages contain information about average and peak interconnect usage. In general, peak interconnect usage over 75%, or average interconnect usage over 60%, could be an indication that it might be difficult to fit the design. Similarly, peak interconnect usage over 90%, or average interconnect usage over 75%, are likely to have increased chances of not getting a valid fit.



For information about identifying areas of congested routing using the Chip Planner, refer to the “Viewing Routing Congestion” subsection in the *Analyzing and Optimizing the Design Floorplan* chapter in volume 2 of the *Quartus II Handbook*.

Placement Effort Multiplier Setting

Some designs might be time consuming and difficult to route because the placement is not optimal. In such cases, you can increase the Placement Effort Multiplier to get a better placement. Doing so might increase the placement time, but it can reduce the routing time, and even overall compilation time in some cases.

Preserving Routing with Incremental Compilation

Preserving the previous routing results for part of the design can reduce future routing time. Incremental compilation provides an easy-to-use methodology that preserves placement and routing results. For more information, refer to “Using Incremental Compilation” on page 12-4 and the references listed in the section.

Reducing Static Timing Analysis Time

If you are performing timing-driven synthesis, the Quartus II software runs the TimeQuest analyzer during Analysis and Synthesis. The Quartus II Fitter also runs the TimeQuest analyzer during placement and routing. If there are incorrect constraints in the SDC file, the Quartus II software may spend time processing constraints unnecessarily several times. If false paths and multicycle paths in the design are not specified completely, the TimeQuest analyzer may spend time analyzing paths that are not relevant to your design. Also, if constraints are redefined in the SDC files, the TimeQuest analyzer may spend additional time processing them. In the compilation messages, look for indications about SDC constraints being redefined, and update the constraints file to avoid this situation. Also, make sure to

provide the correct timing constraints to your design, because the software cannot assume design intent, such as which paths should be considered as false paths or multicycle paths. When you specify these assignments correctly, the TimeQuest analyzer skips analysis for those paths, and the Fitter does not spend additional time optimizing those paths.

Setting Process Priority

It might be necessary to reduce the computing resources allocated to the compilation at the expense of increased compilation time. It can be convenient to reduce the resource allocation to the compilation with single processor machines if you also have to run other tasks at the same time.

- ❓ For more information about setting process priority, refer to *Processing Page (Options Dialog Box)* in Quartus II Help.

Conclusion

The Quartus II software provides many features to reduce compilation time and achieve optimal results. Using the recommended techniques described in this chapter can help you reduce compilation time.

Document Revision History

Table 12-1 shows the revision history for this chapter.

Table 12-1. Document Revision History

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Template update. ■ Added details about peak and average interconnect usage. ■ Added new section “Reducing Static Timing Analysis Time” on page 12-9. ■ Minor changes throughout chapter.
July 2010	10.0.0	Initial release.

-  For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

-  Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter describes techniques to reduce resource usage and improve timing performance when designing for Altera® devices.

Good optimization techniques are essential for achieving the best results when designing for programmable logic devices (PLDs). The optimization features available in the Quartus® II software allow you to meet design requirements by applying these techniques at multiple points in the design process.

This chapter also explains how and when to use some of the features described in other chapters of the *Quartus II Handbook*.

This chapter includes the following topics:

- “Optimizing Your Design”
- “Design Analysis” on page 13–10
- “Resource Utilization Optimization Techniques (LUT-Based Devices)” on page 13–15
- “Timing Optimization Techniques (LUT-Based Devices)” on page 13–27
- “Resource Utilization Optimization Techniques (Macrocell-Based CPLDs)” on page 13–43
- “Timing Optimization Techniques (Macrocell-Based CPLDs)” on page 13–49
- “Scripting Support” on page 13–53

The application of these techniques varies from design to design. Applying each technique does not always improve results. Settings and options in the Quartus II software have default values that generally provide the best trade-off between compilation time, resource utilization, and timing performance. You can adjust these settings to determine whether other settings provide better results for your design.

When using advanced optimization settings and tools, benchmark their effect on your results and use them only if they improve results for your design.

You can use the optimization flow described in this chapter to explore various compiler settings and determine the techniques that provide the best results.

Optimizing Your Design

The first stage in the optimization process is to perform an initial compilation on your design. “Initial Compilation: Required Settings” on page 13–2 provides guidelines for some of the settings and assignments that are recommended for your initial compilation. “Initial Compilation: Optional Fitter Settings” on page 13–6 describes settings that you might turn on based on your design requirements. “Design Analysis” on page 13–10 explains how to analyze the compilation results.

-  You can use incremental compilation in the optimization process. Incremental compilation can preserve timing to aid in timing closure, as well as compilation time reduction; however, it can cause a slight increase in resource utilization.
-  For more details about Quartus II incremental compilation flow, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.
-  To view information about timing analysis results, refer to *Viewing Timing Analysis Results (TimeQuest Timing Analyzer)* in Quartus II Help.

After you have analyzed the compilation results, perform the optimization stages in the recommended order, as described in this chapter.

For LUT-based devices (FPGAs, MAX[®] II series devices), perform optimizations in the following order:

1. If your design does not fit, refer to “[Resource Utilization Optimization Techniques \(LUT-Based Devices\)](#)” on page 13–15 before trying to optimize I/O timing or register-to-register timing.
2. If your design does not meet the required I/O timing performance, refer to “[I/O Timing Optimization Techniques \(LUT-Based Devices\)](#)” on page 13–56 before trying to optimize register-to-register timing.
3. If your design does not meet the required slack on any of the clock domains in the design, refer to “[Register-to-Register Timing Optimization Techniques \(LUT-Based Devices\)](#)” on page 13–56.

For macrocell-based devices (MAX 7000 and MAX 3000 CPLDs), perform optimizations in the following order:

1. If your design does not fit, refer to “[Resource Utilization Optimization Techniques \(Macrocell-Based CPLDs\)](#)” on page 13–43 before trying to optimize I/O timing or register-to-register timing.
2. If your timing performance requirements are not met, refer to “[Timing Optimization Techniques \(Macrocell-Based CPLDs\)](#)” on page 13–49.

-  For device-independent techniques to reduce compilation time, refer to the “[Compilation-Time Optimization Techniques](#)” section in the *Reducing Compilation Time* chapter in volume 2 of the *Quartus II Handbook*.

You can use these techniques in the GUI or with Tcl commands. For more information about scripting techniques, refer to “[Scripting Support](#)” on page 13–53.

Initial Compilation: Required Settings

This section describes the basic assignments and settings for your initial compilation. Check the following compilation assignments before compiling the design in the Quartus II software. Significantly varied compilation results can occur depending on the assignments you set.

Verify the following settings:

- “[Device Settings](#)” on page 13–3

- [“I/O Assignments”](#)
- [“Timing Requirement Settings”](#)
- [“Device Migration Settings”](#) on page 13-5
- [“Partitions and Floorplan Assignments for Incremental Compilation”](#) on page 13-5

Device Settings

Specific device assignments determine the timing model that the Quartus II software uses during compilation. Choose the correct speed grade to obtain accurate results and the best optimization. The device size and the package determine the device pin-out and the number of resources available in the device.

To select the target device, on the Assignments menu, click **Device**.

In a Tcl script, use the following command to set the device:

```
set_global_assignment -name DEVICE <device> ←
```

I/O Assignments

The I/O standards and drive strengths specified for a design affect I/O timing. Specify I/O assignments so that the Quartus II software uses accurate I/O timing delays in timing analysis and Fitter optimizations.

The Quartus II software can select pin locations automatically. If your pin locations are not fixed due to PCB layout requirements, leave pin locations unconstrained. If your pin locations are already fixed, make pin assignments to constrain the compilation appropriately. [“Resource Utilization Optimization Techniques \(Macrocell-Based CPLDs\)”](#) on page 13-43 includes recommendations for making pin assignments that can have a large effect on your results in smaller macrocell-based architectures.

Use the Assignment Editor and Pin Planner to assign I/O standards and pin locations.



For more information about I/O standards and pin constraints, refer to the appropriate device handbook. For information about planning and checking I/O assignments, refer to the [I/O Management](#) chapter in volume 2 of the *Quartus II Handbook*.



For information about using the Assignment Editor, refer to [About the Assignment Editor](#) in Quartus II Help.

Timing Requirement Settings

You must use comprehensive timing requirement settings to achieve the best results for the following reasons:

- Correct timing assignments allow the software to work hardest to optimize the performance of the timing-critical parts of the design and make trade-offs for performance. This optimization can also save area or power utilization in non-critical parts of the design.

- The Quartus II software performs physical synthesis optimizations based on timing requirements (refer to “[Physical Synthesis Optimizations](#)” on page 13-35 for more information).
- Depending on the **Fitter Effort** setting, the Quartus II Fitter can reduce runtime considerably if your timing requirements are being met. For a description of the different effort levels, refer to “[Fitter Effort Setting](#)” on page 13-8

Use your real requirements to get the best results. If you apply more demanding timing requirements than you actually need, increased resource usage, higher power utilization, increased compilation time, or all of these may result.

The Quartus II TimeQuest Timing Analyzer checks your design against the timing constraints. The Compilation Report and timing analysis reporting commands show whether timing requirements are met and provide detailed timing information about paths that violate timing requirements.

To create timing constraints for the TimeQuest analyzer, create a Synopsys Design Constraints File (.sdc). You can also enter constraints in the TimeQuest GUI. Use the `write_sdc` command, or, on the Constraints menu in the TimeQuest analyzer, click **Write SDC File** to write your constraints to an .sdc file. You can add an .sdc file to your project on the **Quartus II Settings** page under **Timing Analysis Settings**.

 If you already have an .sdc file in your project, using the `write_sdc` command from the command line or using the **Write SDC File** option from the TimeQuest GUI enables you to create a new .sdc file, combining the constraints from your current .sdc file and any new constraints added through the GUI or command window, or overwrite the existing .sdc file with your newly applied constraints.

Ensure that every clock signal has an accurate clock setting constraint. If clocks arrive from a common oscillator, they can be considered related. Ensure that all related or derived clocks are set up correctly in the constraints. All I/O pins that require I/O timing optimization must be constrained. Specify both minimum and maximum timing constraints as applicable. If there is more than one clock or there are different I/O requirements for different pins, make multiple clock settings and individual I/O assignments instead of using a global constraint.

Make any complex timing assignments required in the design, including false path and multicycle path assignments. Common situations for these types of assignments include reset or static control signals, cases in which it is not important how long it takes a signal to reach a destination, and paths that can operate in more than one clock cycle. These assignments allow the Quartus II software to make appropriate trade-offs between timing paths and can enable the Compiler to improve timing performance in other parts of the design.

 For more information about timing assignments and timing analysis, refer to *The Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook* and the *Quartus II TimeQuest Timing Analyzer Cookbook*. For more information about how to specify multicycle exceptions in the TimeQuest analyzer, refer to *AN 481: Applying Multicycle Exceptions in the TimeQuest Timing Analyzer*.

 To ensure that constraints or assignments have been applied to all design nodes, you can report all unconstrained paths in your design.

While using the Quartus II TimeQuest analyzer, you can report all the unconstrained paths in your design with the **Report Unconstrained Paths** command in the **Task** pane or the `report_ucp` Tcl command.

Device Migration Settings

If you anticipate a change to the target device later in the design cycle, either because of changes in the design or other considerations, plan for it at the beginning of your design cycle. Whenever you select a target device in the **Settings** dialog box, you can also list any other compatible devices you can migrate to by clicking on the **Migration Devices** button on the **Device** dialog box. If you plan to move your design to a HardCopy® device, make sure to select the device from the **HardCopy** list under **Companion device** on the **Device** dialog box.

Selecting the migration device and companion device early in the design cycle helps to minimize changes to the design at a later stage.

Partitions and Floorplan Assignments for Incremental Compilation

The Quartus II incremental compilation feature enables hierarchical and team-based design flows in which you can compile parts of your design while other parts of the design remain unchanged, or import parts of your design from separate Quartus II projects.

Using incremental compilation for your design with good design partitioning methodology can often help to achieve timing closure. Creating LogicLock™ regions and using incremental compilation can help you achieve timing closure block by block, and preserve the timing performance between iterations, which helps achieve timing closure for the entire design.

Using incremental compilation may also help reduce compilation times.

 For more information, refer to the “Incremental Compilation” section in the *Reducing Compilation Time* chapter in volume 2 of the *Quartus II Handbook*.

If you want to take advantage of incremental compilation for a team-based design flow to reduce your compilation times, or to improve the timing performance of your design during iterative compilation runs, make meaningful design partitions and create a floorplan for your design partitions. Good assignments can improve your results. Assignments can negatively affect a design’s results if you do not follow Altera’s recommendations.

 If you plan to use incremental compilation, you must create a floorplan for your design. If you are not using incremental compilation, this step is optional.

 For guidelines about how to create partition and floorplan assignments for your design, refer to the *Best Practices for Incremental Compilation Partitions and Floorplan Assignments* chapter in volume 1 of the *Quartus II Handbook*.

Initial Compilation: Optional Fitter Settings

This section describes optional Fitter settings that can help to compile your design. You can selectively set all the optional settings that help to improve performance. These settings vary between designs and there is no standard set that applies to all designs. Significantly different compilation results can occur depending on the assignments you have set.

The following settings are optional:

- “Optimize Hold Timing”
- “Limit to One Fitting Attempt” on page 13-7
- “Optimize Multi-Corner Timing” on page 13-7
- “Fitter Effort Setting” on page 13-8

To turn on these settings, follow these steps:

1. On the Assignments menu, click **Settings**.
2. In the **Category** list, select **Fitter Settings**. The **Fitter Settings** page appears.
3. Turn on the appropriate options.

Optimize Hold Timing

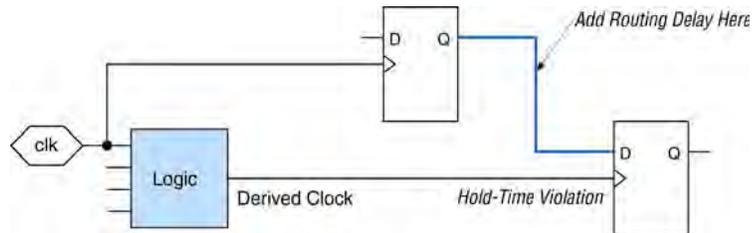
The **Optimize Hold Timing** option directs the Quartus II software to optimize minimum delay timing constraints. This option is available for all Altera device families except MAX 3000 and MAX 7000 series devices. By default, the Quartus II software optimizes hold timing for all paths for designs using devices newer than Arria GX, Stratix III, and Cyclone III. By default, the Quartus II software optimizes hold timing only for I/O paths and minimum TPD paths for older devices.

When you turn on **Optimize Hold Timing**, the Quartus II software adds delay to paths to guarantee that the minimum delay requirements are satisfied. In the **Fitter Settings** pane, if you select **I/O Paths and Minimum TPD Paths** (the default choice for older devices such as Cyclone II and Stratix II family of devices if you turn on **Optimize Hold Timing**), the Fitter works to meet the following criteria:

- Hold times (t_H) from device input pins to registers
- Minimum delays from I/O pins to I/O registers or from I/O registers to I/O pins
- Minimum clock-to-out time (t_{CO}) from registers to output pins

If you select **All Paths**, the Fitter also works to meet hold requirements from registers to registers, as in [Figure 13-1](#), where a derived clock generated with logic causes a hold time problem on another register. However, if your design has internal hold time violations between registers, Altera recommends that you correct the problems by making changes to your design, such as using a clock enable signal instead of a derived or gated clock.

Figure 13-1. Optimize Hold Timing Option Fixing an Internal Hold Time Violation



For design practices that can help eliminate internal hold time violations, refer to the [Recommended Design Practices](#) chapter in volume 1 of the *Quartus II Handbook*.

Limit to One Fitting Attempt

A design might fail to fit for several reasons, such as logic overuse or illegal assignments. For most failures, the Quartus II software informs you of the problem. However, if the design uses too much routing, the Quartus II software makes up to two additional attempts to fit your design. Each of these fit attempts takes significantly longer than the previous attempt.

For large designs, you might not want to wait for all three fitting attempts to be completed. To have the Quartus II software issue an error message after the first failed attempt, turn on **Limit to one fitting attempt** on the **Fitter Settings** page.

For instructions about how to lower the design's routing utilization, so your design can be made to fit into the target device if it fails to fit due to the lack of routing resources, refer to ["Routing" on page 13-24](#)

Optimize Multi-Corner Timing

Historically, FPGA timing analysis has been performed using only worst-case delays, which are described in the slow corner timing model. However, due to process variation and changes in the operating conditions, delays on some paths can be significantly smaller than those in the slow corner timing model. This can result in hold time violations on those paths, and in rare cases, additional setup time violations.

Also, because of the small process geometries of the Cyclone III, Stratix III, and newer device families, the slowest circuit performance of designs targeting these devices does not necessarily occur at the highest operating temperature. The temperature at which the circuit is slowest depends on the selected device, the design, and the Quartus II compilation results. Therefore, the Quartus II software provides the Cyclone III series, Stratix III, and newer device families with three different timing corners—Slow 85°C corner, Slow 0°C corner, and Fast 0°C corner. For other device families, two timing corners are available—Fast 0°C and Slow 85°C corner.

By default, the Fitter optimizes constraints using only the slow corner timing model. You can turn on the **Optimize multi-corner timing** option to instruct the Fitter to also optimize constraints considering all available timing corners, at the cost of a slight increase in runtime. By optimizing for all timing corners, you can create a design implementation that is more robust across process, temperature, and voltage variations. While optimizing for multi-corner timing, the Fitter chooses one of the two slow corners that is known to have more critical timing (depending on the chosen device), along with the fast corner. This option is available only for Arria, Cyclone, HardCopy, MAX II, MAX V, and Stratix series devices.

Using the different timing models can be important to account for process, voltage, and temperature variations for each device. Turning this option on increases compilation time by approximately 10%.

For designs with external memory interfaces such as DDR and QDR, Altera recommends that you turn on the **Optimize multi-corner timing** setting.

Fitter Effort Setting

Fitter effort refers to the amount of effort the Quartus II software uses to fit your design. To set the Fitter effort, on the Assignments menu, click **Settings**. In the **Category** list, select **Fitter Settings**. The **Fitter effort** settings are **Auto Fit**, **Standard Fit**, and **Fast Fit**. The default setting depends on the device family specified.

Auto Fit

The **Auto Fit** option (available for Arria, Cyclone, HardCopy, MAX II, MAX V, and Stratix series devices) focuses the full Fitter effort only on those aspects of the design that require further optimization. **Auto Fit** can significantly reduce compilation time relative to **Standard Fit** if your design has easy-to-meet timing requirements, low routing resource utilization, or both. However, those designs that require full optimization generally receive the same effort as is achieved by selecting **Standard Fit**. **Auto Fit** is the default Fitter effort setting for all devices for which this option is available.

If you want the Fitter to attempt to exceed the timing requirements by a certain margin instead of simply meeting them, specify a minimum slack in the **Desired worst case slack** box.



Specifying a minimum slack does not guarantee that the Fitter achieves the slack requirement; it only guarantees that the Fitter applies full optimization unless the target slack is exceeded.

In some designs with multiple clocks, it might be possible to improve the timing performance on one clock domain while reducing the performance on other clock domains by over-constraining the most important clock. If you use this technique, perform a sweep over multiple seeds to ensure that any performance improvements that you see are real gains. For more information, refer to [“Fitter Seed” on page 13-39](#).

Over-constraining the clock for which you require maximum slack, while using the **Auto Fit** option, increases the chances that the Fitter is able to meet this requirement.

The **Auto Fit** option also causes the Quartus II Fitter to optimize for shorter compilation times instead of maximum possible performance if the design includes no timing assignments.

If your design has aggressive timing requirements or is hard to route, the placement does not stop early and the compilation time is the same as using the **Standard Fit** option.

The **Auto Fit** option might increase the number of routing wires used. This can lead to an increase in the dynamic power when compared to using the **Standard Fit** option, unless the **Extra effort** option in the **PowerPlay power optimization** list is also enabled. When you turn on **Extra effort**, **Auto Fit** continues to optimize for reduction of wire usage even after meeting the register-to-register requirement, and there is no adverse effect on the dynamic power consumption relative to using **Standard Fit**. If dynamic power consumption is a concern, select **Extra effort** in both the **Analysis & Synthesis Settings** and the **Fitter Settings** pages.



For more details, refer to the “Power Driven Compilation” section in the *Power Optimization* chapter in volume 2 of the *Quartus II Handbook*.

Standard Fit

Use the **Standard Fit** option to exceed specified timing requirements and achieve the best possible timing results and lowest routing resource utilization for your design. The **Standard Fit** setting usually increases compilation time relative to **Auto Fit**, because it applies full optimization, regardless of the design requirement. In designs with no timing assignments, on average, using the **Standard Fit** option results in a f_{MAX} about 10% higher than that achieved using the **Auto Fit** option. In designs where timing requirements can be easily met, using the **Standard Fit** option can result in considerably longer compilation times than using the **Auto Fit** option.

Fast Fit

The **Fast Fit** option reduces the amount of optimization effort for each algorithm employed during fitting. This option reduces the compilation time by about 50%, resulting in a fit that has, on average, 10% lower f_{MAX} than that achieved using the **Standard Fit** setting.

Design Assistant

You can run the Design Assistant to analyze the post-fitting results of your design during a full compilation. The Design Assistant checks rules related to gated clocks, reset signals, asynchronous design practices, and signal race conditions. This is especially useful during the early stages of your design, so that you can work on any areas of concern in your design before proceeding with design optimization.

- For more information about the Design Assistant, refer to *About the Design Assistant* and *Analyzing Designs with the Design Assistant* in Quartus II Help.

Design Analysis

The initial compilation establishes whether the design achieves a successful fit and meets the specified timing requirements. This section describes how to analyze your design results in the Quartus II software.

Error and Warning Messages

After compiling your design, evaluate all error and warning messages to see if any design or setting changes are required. If changes are required, make these changes and recompile the design before proceeding with design optimization.

To suppress messages that you have already evaluated and do not want to see again, right-click on the message in the Messages window and click **Suppress**.

 For more information about message suppression, refer to the “Message Suppression” section in the *Managing Quartus II Projects* chapter in volume 2 of the *Quartus II Handbook*.

Ignored Timing Constraints

The Quartus II software ignores illegal, obsolete, and conflicting constraints.

You can view a list of ignored constraints by clicking **Report Ignored Constraints** in the Reports menu in the TimeQuest GUI or by typing the following command to generate a list of ignored timing constraints:

```
report_sdc -ignored -panel_name "Ignored Constraints" ←
```

If any constraints were ignored, analyze why they were ignored. If necessary, correct the constraints and recompile the design before proceeding with design optimization.

 For more information about the `report_sdc` command and its options, refer to the *Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

Resource Utilization

Determining device utilization is important regardless of whether a successful fit is achieved. If your compilation results in a no-fit error, resource utilization information is important for analyzing the fitting problems in your design. If your fitting is successful, review the resource utilization information to determine whether the future addition of extra logic or other design changes might introduce fitting difficulties. Also, review the resource utilization information to determine if it is impacting timing performance.

To determine resource usage, refer to the **Flow Summary** section of the Compilation Report. This section reports how many resources are used, including pins, memory bits, digital signal processing, and phase-locked loops (PLLs). The **Flow Summary** indicates whether the design exceeds the available device resources. More detailed information is available by viewing the reports under **Resource Section** in the **Fitter** section of the Compilation Report.

 For Arria series, Stratix II, Stratix III, Stratix IV, and Stratix V devices, a device with low utilization does not necessarily have the lowest adaptive logic module (ALM) utilization possible. For these devices, the Fitter uses adaptive look-up tables (ALUTs) in different ALMs—even when the logic can be placed within one ALM—to achieve the best timing and routing results. In achieving these results, the Fitter can spread logic throughout the device. As the device fills up, the Fitter automatically searches for logic functions with common inputs to place in one ALM. The number of partnered ALUTs and packed registers also increases. Therefore, a design that is reported as close to 100% full might still have space for extra logic if logic and registers can be packed together more aggressively.

If resource usage is reported as less than 100% and a successful fit cannot be achieved, either there are not enough routing resources or some assignments are illegal. In either case, a message appears in the **Processing** tab of the Messages window describing the problem.

If the Fitter finishes faster than the Fitter runs on similar designs, a resource might be over-utilized or there might be an illegal assignment. If the Quartus II software seems to run for an excessively long time compared to runs on similar designs, a legal placement or route probably cannot be found. In the Compilation Report, look for errors and warnings that indicate these types of problems.

For more information about how to get a quick error message on hard-to-fit designs, refer to [“Limit to One Fitting Attempt” on page 13-7](#).

You can use the Chip Planner to find areas of the device that have routing congestion. If you find areas with very high congestion, analyze the cause of the congestion. Issues such as high fan-out nets not using global resources, an improperly chosen optimization goal (speed versus area), very restrictive floorplan assignments, or the coding style can cause routing congestion. After you identify the cause, modify the source or settings to reduce routing congestion.

 For details about using the Chip Planner tool, refer to the [Analyzing and Optimizing the Design Floorplan](#) chapter in volume 2 of the *Quartus II Handbook* and [About the Chip Planner](#) in Quartus II Help.

I/O Timing (Including t_{PD})

TimeQuest analyzer supports the Synopsys Design Constraints (SDC) format for constraining your design. When using the TimeQuest analyzer for timing analysis, use the `set_input_delay` constraint to specify the data arrival time at an input port with respect to a given clock. For output ports, use the `set_output_delay` command to specify the data arrival time at an output port’s receiver with respect to a given clock. You can use the `report_timing` Tcl command to generate the I/O timing reports.

The I/O paths that do not meet the required timing performance are reported as having negative slack and are highlighted in red in the TimeQuest analyzer **Report** pane. In cases where you do not apply an explicit I/O timing constraint to an I/O pin, the Quartus II timing analysis software still reports the **Actual** number, which is the timing number that must be met for that timing parameter when the device runs in your system.

 For more information about how timing numbers are calculated, refer to the *Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

Register-to-Register Timing

This section contains the following sections:

- “Timing Analysis with the TimeQuest Timing Analyzer”
- “Tips for Analyzing Failing Paths” on page 13-14
- “Tips for Analyzing Failing Clock Paths that Cross Clock Domains” on page 13-14

Timing Analysis with the TimeQuest Timing Analyzer

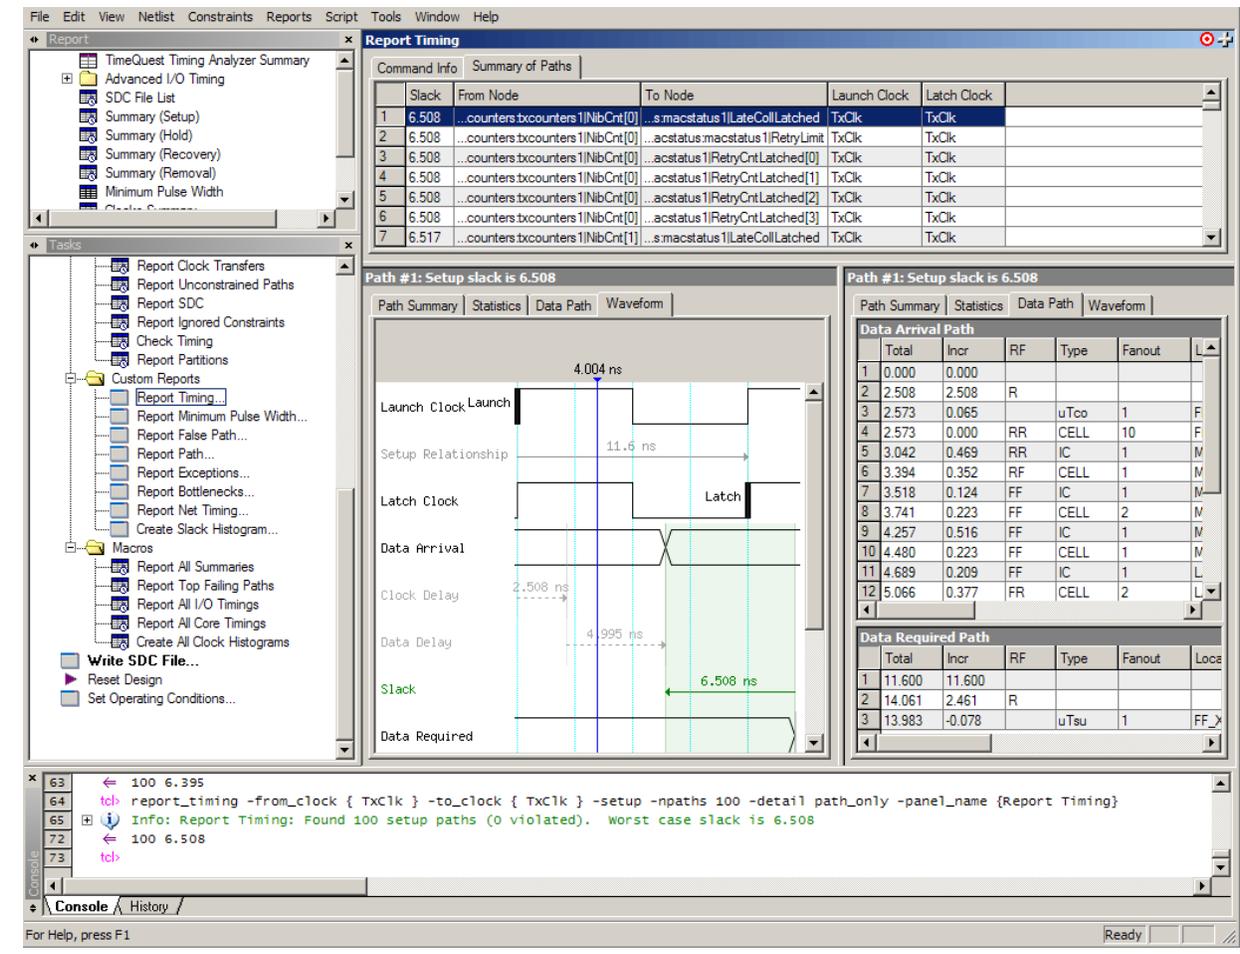
If you are using the TimeQuest analyzer, analyze all valid register-to-register paths by using appropriate constraints. Use the `report_timing` command to generate the required timing reports for any register-to-register path. Your design meets timing requirements when you do not have negative slack on any register-to-register path on any of the clock domains.

When you select a path listed in the TimeQuest **Report Timing** pane, the tabs in the corresponding path detail pane show a path summary of source and destination registers and their timing, statistics about the path delay, detailed information about the complete data path with all nodes in the path and the waveforms of the relevant signals (Figure 13-2). To locate a selected path in the Chip Planner or the Technology Map Viewer by using the shortcut menu, right-click on a path, point to **Locate**, and click **Locate in Chip Planner**. The Chip Planner appears with the path highlighted. Similarly, if you know that a path is not a valid path, you can set it to be a false path using the shortcut menu.

To see the path details of any selected path, click on the **Data Path** tab in the path details pane. This displays the details of the Data Arrival Path, as well as the Data Required Path. For a graphical view of the information, click on the **Waveform** tab.

You can locate critical paths in the Chip Planner from the TimeQuest timing analysis report panel.

Figure 13-2. TimeQuest Analyzer GUI



For more information about how timing analysis results are calculated, refer to the *Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

You also can see the logic in a particular path by locating the logic in the RTL Viewer or Technology Map Viewer. These viewers allow you to see a gate-level or technology-mapped representation of your design netlist. To locate a timing path in one of the viewers, right-click on a path in the report, point to **Locate**, and click **Locate in RTL Viewer** or **Locate in Technology Map Viewer**. When you locate a timing path in the Technology Map Viewer, the annotated schematic displays the same delay information that is shown when you use the **List Paths** command.

For more information about netlist viewers, refer to the *Analyzing Designs with Quartus II Netlist Viewers* chapter in volume 1 of the *Quartus II Handbook*.

Tips for Analyzing Failing Paths

When you are analyzing clock path failures, examine reports and waveforms to determine if the correct constraints are being applied, and add multicycle or false paths as appropriate.

Focus on improving the paths that show the worst slack. The Fitter works hardest on paths with the worst slack. If you fix these paths, the Fitter might be able to improve the other failing timing paths in the design.

Check for particular nodes that appear in many failing paths. Look for paths that have common source registers, destination registers, or common intermediate combinational nodes. In some cases, the registers might not be identical, but are part of the same bus. In the timing analysis report panels, clicking on the **From** or **To** column headings can be helpful to sort the paths by the source or destination registers. Clicking first on **From**, then on **To**, uses the registers in the **To** column as the primary sort and **From** as the secondary sort. If you see common nodes, these nodes indicate areas of your design that might be improved through source code changes or Quartus II optimization settings. Constraining the placement for just one of the paths might decrease the timing performance for other paths by moving the common node further away in the device.

Tips for Analyzing Failing Clock Paths that Cross Clock Domains

When analyzing clock path failures, check whether these paths cross between two clock domains. This is the case if the **From Clock** and **To Clock** in the timing analysis report are different. There can also be paths that involve a different clock in the middle of the path, even if the source and destination register clock are the same. To analyze these paths in more detail, right-click on the entry in the report and click **List Paths**.

Expand the **List Paths** entry in the Messages window and analyze the largest register-to-register requirement. Evaluate the setup relationship between the source and destination (launch edge and latch edge) to determine if that is reducing the available setup time. For example, the path can start at a rising edge and end at a falling edge, which reduces the setup relationship by one half clock cycle.

Check to see if the PLL phase shift is reducing the setup requirement. You might be able to adjust this using PLL parameters and settings.

Paths that cross clock domains are generally protected with synchronization logic (for example, FIFOs or double-data synchronization registers) to allow asynchronous interaction between the two clock domains. In such cases, you can ignore the timing paths between registers in the two clock domains while running timing analysis, even if the clocks are related.

The Fitter attempts to optimize all failing timing paths. If there are paths that can be ignored for optimization and timing analysis, but the paths do not have constraints that instruct the Fitter to ignore them, the Fitter tries to optimize those paths as well. In some cases, optimizing unnecessary paths can prevent the Fitter from meeting the timing requirements on timing paths that are critical to the design. It is beneficial to specify all paths that can be ignored, so that the Fitter can put more effort into the paths that must meet their timing requirements instead of optimizing paths that can be ignored.

- For more details about how to ignore timing paths that cross clock domains, refer to the *Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

Evaluate the clock skew between the source clock and the destination clock to determine if that is reducing the available setup time. You can check the shortest and longest clock path reports to see what is causing the clock skew. Avoid using combinational logic in clock paths because it contributes to clock skew. Differences in the logic or in its routing between the source and destination can cause clock skew problems and result in warnings during compilation.

Global Routing Resources

Global routing resources are designed to distribute high-fan-out, low-skew signals (such as clocks) without consuming regular routing resources. Depending on the device, these resources can span the entire chip, or some smaller portion, such as a quadrant. The Quartus II software attempts to assign signals to global routing resources automatically, but you might be able to make more suitable assignments manually.

- Refer to the relevant device handbook for details about the number and types of global routing resources available.

Check the global signal utilization in your design to ensure that appropriate signals have been placed on global routing resources. In the Compilation Report, open the Fitter report and click the **Resource Section**. Analyze the Global & Other Fast Signals and Non-Global High Fan-out Signals reports to determine whether any changes are required.

You might be able to reduce clock skew for high fan-out signals by placing them on global routing resources. Conversely, you can reduce the insertion delay of low fan-out signals by removing them from global routing resources. Doing so can improve clock enable timing and control signal recovery/removal timing, but increases clock skew. Use the **Global Signal** setting in the Assignment Editor to control global routing resources.

Resource Utilization Optimization Techniques (LUT-Based Devices)

After design analysis, the next stage of design optimization is to improve resource utilization. Complete this stage before proceeding to I/O timing optimization or register-to-register timing optimization. Ensure that you have already set the basic constraints described in “[Initial Compilation: Required Settings](#)” on page 13-2 before proceeding with the resource utilization optimizations discussed in this section. If a design does not fit into a specified device, use the techniques in this section to achieve a successful fit. After you optimize resource utilization and your design fits in the desired target device, optimize I/O timing as described in “[I/O Timing Optimization Techniques \(LUT-Based Devices\)](#)” on page 13-56. These tips are valid for all FPGA families and the MAX II family of CPLDs.

Using the Resource Optimization Advisor

The Resource Optimization Advisor provides guidance in determining settings that optimize the resource usage. To run the Resource Optimization Advisor, on the Tools menu, point to **Advisors**, and click **Resource Optimization Advisor**.

The Resource Optimization Advisor provides step-by-step advice about how to optimize the resource usage (logic element, memory block, DSP block, I/O, and routing) of your design. Some of the recommendations in these categories might conflict with each other. Altera recommends evaluating the options and choosing the settings that best suit your requirements.

Resolving Resource Utilization Issues Summary

Resource utilization issues can be divided into the following three categories:

- Issues relating to I/O pin utilization or placement, including dedicated I/O blocks such as PLLs or LVDS transceivers (refer to [“I/O Pin Utilization or Placement”](#)).
- Issues relating to logic utilization or placement, including logic cells containing registers and look-up tables as well as dedicated logic, such as memory blocks and DSP blocks (refer to [“Logic Utilization or Placement”](#) on page 13-17).
- Issues relating to routing (refer to [“Routing”](#) on page 13-24).

I/O Pin Utilization or Placement

Use the suggestions in the following sections to help you resolve I/O resource problems.

Use I/O Assignment Analysis

On the Processing menu, point to **Start** and click **Start I/O Assignment Analysis** to help with pin placement. The **Start I/O Assignment Analysis** command allows you to check your I/O assignments early in the design process. You can use this command to check the legality of pin assignments before, during, or after compilation of your design. If design files are available, you can use this command to accomplish more thorough legality checks on your design’s I/O pins and surrounding logic. These checks include proper reference voltage pin usage, valid pin location assignments, and acceptable mixed I/O standards.

Common issues with I/O placement relate to the fact that differential standards have specific pin pairings, and certain I/O standards might be supported only on certain I/O banks.

If your compilation or I/O assignment analysis results in specific errors relating to I/O pins, follow the recommendations in the error message. Right-click on the message in the Messages window and click **Help** to open the Quartus II Help topic for this message.

Modify Pin Assignments or Choose a Larger Package

If a design that has pin assignments fails to fit, compile the design without the pin assignments to determine whether a fit is possible for the design in the specified device and package. You can use this approach if a Quartus II error message indicates fitting problems due to pin assignments.

If the design fits when all pin assignments are ignored or when several pin assignments are ignored or moved, you might have to modify the pin assignments for the design or select a larger package.

If the design fails to fit because insufficient I/Os are available, a successful fit can often be obtained by using a larger device package (which can be the same device density) that has more available user I/O pins.

 For more information about I/O assignment analysis, refer to the *I/O Management* chapter in volume 2 of the *Quartus II Handbook*.

Logic Utilization or Placement

Use the suggestions in the following subsections to help you resolve logic resource problems, including logic cells containing registers and lookup tables (LUTs), as well as dedicated logic such as memory blocks and DSP blocks.

Optimize Synthesis for Area, Not Speed

If your design fails to fit because it uses too much logic, resynthesize the design to improve the area utilization. First, ensure that you have set your device and timing constraints correctly in your synthesis tool. Particularly when area utilization of the design is a concern, ensure that you do not over-constrain the timing requirements for the design. Synthesis tools generally try to meet the specified requirements, which can result in higher device resource usage if the constraints are too aggressive.

If resource utilization is an important concern, some synthesis tools offer an easy way to optimize for area instead of speed. If you are using Quartus II integrated synthesis, select **Balanced** or **Area** for the **Optimization Technique**. You can also specify this logic option for specific modules in your design with the Assignment Editor in cases where you want to reduce area using the **Area** setting (potentially at the expense of register-to-register timing performance) while leaving the default **Optimization Technique** setting at **Balanced** (for the best trade-off between area and speed for certain device families) or **Speed**. You can also use the **Speed Optimization Technique for Clock Domains** logic option to specify that all combinational logic in or between the specified clock domain(s) is optimized for speed.

In some synthesis tools, not specifying an f_{MAX} requirement can result in less resource utilization.

 In the Quartus II software, the **Balanced** setting typically produces utilization results that are very similar to those produced by the **Area** setting, with better performance results. The **Area** setting can give better results in some cases.

 For information about setting timing requirements and synthesis options in Quartus II integrated synthesis and other synthesis tools, refer to the appropriate chapter in *Section III. Synthesis* in volume 1 of the *Quartus II Handbook*, or your synthesis software's documentation.

The Quartus II software provides additional attributes and options that can help improve the quality of your synthesis results.

Restructure Multiplexers

Multiplexers form a large portion of the logic utilization in many FPGA designs. By optimizing your multiplexed logic, you can achieve a more efficient implementation in your Altera device.

❓ For more information about this option, refer to *Restructure Multiplexers logic option* in Quartus II Help.

👉 For design guidelines to achieve optimal resource utilization for multiplexer designs, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*.

Perform WYSIWYG Primitive Resynthesis with Balanced or Area Setting

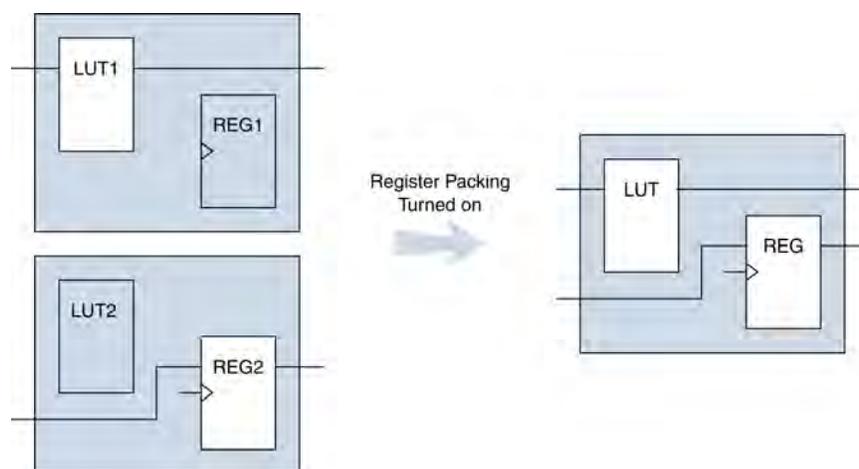
❓ For information about this logic option, refer to *Perform WYSIWYG Primitive Resynthesis logic option* in Quartus II Help.

👉 The **Balanced** setting typically produces utilization results that are very similar to the **Area** setting with better performance results. The **Area** setting can give better results in some cases. Performing WYSIWYG resynthesis for area in this way typically reduces register-to-register timing performance.

Use Register Packing

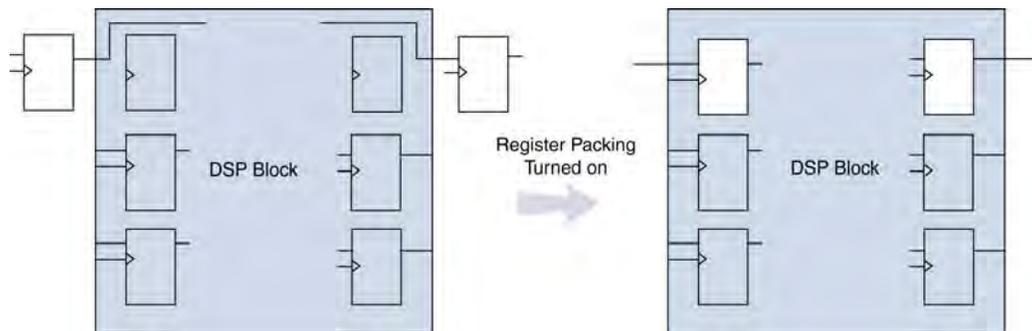
The **Auto Packed Registers** option implements the functions of two cells into one logic cell by combining the register of one cell in which only the register is used with the LUT of another cell in which only the LUT is used. [Figure 13-3](#) shows register packing and the gain of one logic cell in the design.

Figure 13-3. Register Packing



Registers can also be packed into DSP blocks (Figure 13-4).

Figure 13-4. Register Packing in DSP Blocks



The following list shows the most common cases in which register packing helps to optimize a design:

- A LUT can be implemented in the same cell as an unrelated register with a single data input
- A LUT can be implemented in the same cell as the register that is fed by the LUT
- A LUT can be implemented in the same cell as the register that feeds the LUT
- A register can be packed into a RAM block
- A register can be packed into a DSP block
- A register can be packed into an I/O Element (IOE)

The following options are available for register packing (for supported device families):

- **Auto**—This is the default setting for register packing. This setting tells the Fitter to attempt to achieve the best performance while maintaining a fit for the design in the specified device. The Fitter combines all combinational (LUT) and sequential (register) functions that benefit circuit speed. In addition, more aggressive combinations of unrelated combinational and sequential functions are performed to the extent required to reduce the area of the design to achieve a fit in the specified device. This option is available only for Arria series, Cyclone series, and Stratix series devices.
- **Minimize Area**—Aggressively packs registers to reduce area, even at the cost of design performance.
- **Minimize Area with Chains**—Aggressively packs registers to reduce area. This option packs registers with carry chains. It also converts registers into register cascade chains and packs them with other logic to reduce area. This option is available only for Arria series, Cyclone series, MAX II, and Stratix series devices.
- **Normal**—Packs registers when this is not expected to adversely affect timing results.
- **Off**—Does not pack registers.

- **Sparse**—In this mode, the combinational (LUT) and sequential (register) functions are combined such that the combined logic has either a combinational output or a sequential output, but not both. This mode is available only for Arria series, Cyclone II, Cyclone III, Stratix II, Stratix III, Stratix IV, and Stratix V devices. This option results in a higher logic array block (LAB) usage, but might give you better timing performance because of reduced routing congestion.
- **Sparse Auto**—In this mode, the Quartus II Fitter starts with sparse mode packing, and then attempts to achieve best performance while maintaining a fit for the specified device. Later optimizations are carried out in a way similar to the **Auto** mode. This mode is available only for Arria series, Cyclone II, Cyclone III, Stratix II, Stratix III, Stratix IV, and Stratix V devices.

❓ For more information, refer to *Auto Packed Registers logic option* in Quartus Help.

Remove Fitter Constraints

A design with conflicting constraints or constraints that are difficult to meet may not fit in the targeted device. This can occur when the location or LogicLock assignments are too strict and not enough routing resources are available on the device.

In this case, use the **Routing Congestion** task in the Chip Planner to locate routing problems in the floorplan, then remove any location or LogicLock region assignments in that area. If your design still does not fit, the design is over-constrained. To correct the problem, remove all location and LogicLock assignments and run successive compilations, incrementally constraining the design before each compilation. You can delete specific location assignments in the Assignment Editor or the Chip Planner. To remove LogicLock assignments in the Chip Planner, in the LogicLock Regions Window, or on the Assignments menu, click **Remove Assignments**. Turn on the assignment categories you want to remove from the design in the **Available assignment categories** list.

 For more information about the **Routing Congestion** task in the Chip Planner, refer to *Analyzing and Optimizing the Design Floorplan* in volume 2 of the *Quartus II Handbook*.

Change State Machine Encoding

State machines can be encoded using various techniques. Using binary or gray code encoding typically results in fewer state registers than one-hot encoding, which requires one register for every state bit. If your design contains state machines, changing the state machine encoding to one that uses the minimal number of registers may reduce resource utilization. The effect of state machine encoding varies depending on the way your design is structured.

If your design does not manually encode the state bits, you can specify the state machine encoding in your synthesis tool. When using Quartus II integrated synthesis, turn on the **Minimal Bits** setting for the **State Machine Processing** option.

❓ For more information, refer to *State Machine Processing logic option* in Quartus II Help.

You can also specify this logic option for specific modules or state machines in your design with the Assignment Editor.

You can also use the following Tcl command in scripts to modify the state machine encoding.

```
set_global_assignment -name state_machine_processing <value>
```

In this case, <value> can be AUTO, ONE-HOT, MINIMAL BITS, or USER-ENCODE.

Flatten the Hierarchy During Synthesis

Synthesis tools typically provide the option of preserving hierarchical boundaries, which can be useful for verification or other purposes. However, optimizing across hierarchical boundaries allows the synthesis tool to perform the most logic minimization, which can reduce area. Therefore, to achieve the best results, flatten your design hierarchy whenever possible.

If you are using Quartus II incremental compilation, you cannot flatten your design across design partitions. Incremental compilation always preserves the hierarchical boundaries between design partitions. Follow Altera's recommendations for design partitioning, such as registering partition boundaries to reduce the effect of cross-boundary optimizations.

-  For more information about using incremental compilation and recommendations for design partitioning, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.

Retarget Memory Blocks

If the design fails to fit because it runs out of device memory resources, your design may require a certain type of memory the device does not have. For example, a design that requires two M-RAM blocks cannot be targeted to a Stratix EP1S10 device, which has only one M-RAM block. You might be able to obtain a fit by building one of the memories with a different size memory block, such as an M4K memory block.

If the memory block was created with the MegaWizard™ Plug-In Manager, open the MegaWizard Plug-In Manager and edit the RAM block type so it targets a new memory block size.

ROM and RAM memory blocks can also be inferred from your HDL code, and your synthesis software can place large shift registers into memory blocks by inferring the ALTSHIFT_TAPS megafunction. This inference can be turned off in your synthesis tool to cause the memory to be placed in logic instead of in memory blocks.

-  For more information, refer to *Auto RAM Replacement logic option*, *Auto ROM Replacement logic option*, and *Auto Shift Register Replacement logic option* in Quartus II Help.

Depending on your synthesis tool, you can also set the RAM block type for inferred memory blocks. In Quartus II integrated synthesis, set the **ramstyle** attribute to the desired memory type for the inferred RAM blocks, or set the option to **logic**, to implement the memory block in standard logic instead of a memory block.

Consider the resource utilization by hierarchy in the report file, and determine whether there is an unusually high register count in any of the modules. Some coding styles can prevent the Quartus II software from inferring RAM blocks from the source code because of their architectural implementation, and forces the software to implement the logic in flipflops. As an example, a function such as an asynchronous reset on a register bank might make it incompatible with the RAM blocks in the device architecture, so that the register bank is implemented in flipflops. It is often possible to move a large register bank into RAM by slight modification of associated logic.

-  For more information about memory inference control in other synthesis tools, refer to the appropriate chapter in *Section III. Synthesis* in volume 1 of the *Quartus II Handbook*, or your synthesis software's documentation. For more information about coding styles and HDL examples that ensure memory inference, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*.

Use Physical Synthesis Options to Reduce Area

The physical synthesis options for fitting can help you decrease the resource usage. When you enable these settings for physical synthesis for fitting, the Quartus II software makes placement-specific changes to the netlist that reduce resource utilization for a specific Altera device.

-  The compilation time might increase considerably when you use physical synthesis options.

With the Quartus II software, you can apply physical synthesis options to specific instances, which can reduce the impact on compilation time. Physical synthesis instance assignments allow you to enable physical synthesis algorithms for specific portions of their design.

The following physical synthesis optimizations for fitting are available:

- Physical synthesis for combinational logic
- Map logic into memory

-  For more information, refer to *Physical Synthesis Optimizations Page (Settings Dialog Box)* in Quartus II Help.

Retarget or Balance DSP Blocks

A design might not fit because it requires too many DSP blocks. All DSP block functions can be implemented with logic cells, so you can retarget some of the DSP blocks to logic to obtain a fit.

If the DSP function was created with the MegaWizard Plug-In Manager, open the MegaWizard Plug-In Manager and edit the function so it targets logic cells instead of DSP blocks. The Quartus II software uses the `DEDICATED_MULTIPLIER_CIRCUITRY` megafunction parameter to control the implementation.

DSP blocks also can be inferred from your HDL code for multipliers, multiply-adders, and multiply-accumulators. This inference can be turned off in your synthesis tool. When you are using Quartus II integrated synthesis, you can disable inference by turning off the **Auto DSP Block Replacement** logic option for your entire project. On the Assignments menu, click **Settings**. In the **Category** list, select **Analysis & Synthesis Settings**, click **More Settings**, and turn off **Auto DSP Block Replacement**. Alternatively, you can disable the option for a specific block with the Assignment Editor.

 For more information about disabling DSP block inference in other synthesis tools, refer to the appropriate chapter in *Section III. Synthesis* in volume 1 of the *Quartus II Handbook*, or your synthesis software's documentation.

The Quartus II software also offers the **DSP Block Balancing** logic option, which implements DSP block elements in logic cells or in different DSP block modes. The default **Auto** setting allows DSP block balancing to convert the DSP block slices automatically as appropriate to minimize the area and maximize the speed of the design. You can use other settings for a specific node or entity, or on a project-wide basis, to control how the Quartus II software converts DSP functions into logic cells and DSP blocks. Using any value other than **Auto** or **Off** overrides the `DEDICATED_MULTIPLIER_CIRCUITRY` parameter used in megafunction variations.

 For more details about the Quartus II logic options described in this section, refer to *Auto DSP Block Replacement* and *DSP Block Balancing* in Quartus II Help.

Optimize Source Code

If your design does not fit because of logic utilization, and the methods described in the preceding sections do not sufficiently improve the resource utilization of the design, modify the design at the source to achieve the desired results. You can often improve logic significantly by making design-specific changes to your source code. This is typically the most effective technique for improving the quality of your results.

If your design does not fit into available LEs or ALMs, but you have unused memory or DSP blocks, check to see if you have code blocks in your design that describe memory or DSP functions that are not being inferred and placed in dedicated logic. You might be able to modify your source code to allow these functions to be placed into dedicated memory or DSP resources in the target device.

Ensure that your state machines are recognized as state machine logic and optimized appropriately in your synthesis tool. State machines that are recognized are generally optimized better than if the synthesis tool treats them as generic logic. In the Quartus II software, you can check for the State Machine report under **Analysis & Synthesis** in the Compilation Report. This report provides details, including the state encoding for each state machine that was recognized during compilation. If your state machine is not being recognized, you might have to change your source code to enable it to be recognized.

- For coding style guidelines, including examples of HDL code for inferring memory and DSP functions, refer to the “Instantiating Altera Megafunctions” and the “Inferring Multiplier and DSP Functions from HDL Code” sections of the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*. For guidelines and sample HDL code for state machines, refer to the “General Coding Guidelines” section of the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*.
- For additional HDL coding examples, refer to *AN 584: Timing Closure Methodology for Advanced FPGAs*.

Use a Larger Device

If a successful fit cannot be achieved because of a shortage of LEs or ALMs, memory, or DSP blocks, you might require a larger device.

Routing

Use the suggestions in the following subsections to help you resolve routing resource problems.

Set Auto Packed Registers to Sparse or Sparse Auto

This option is useful for reducing LE or ALM count in a design. This option is available for all Altera devices supported by the Quartus II software.

This option can be set in the Assignment Editor, or you can set this option by clicking **More Settings** on the **Fitter Settings** page in the **Settings** dialog box

- For more information, refer to *Auto Packed Registers* in Quartus II Help.

Set Fitter Aggressive Routability Optimizations to Always

Use this option if your design does not fit due to excessive routing wire utilization.

- For more information, refer to *Fitter Aggressive Routability Optimizations logic option* in Quartus II Help.

If there is a significant imbalance between placement and routing time (during the first fitting attempt), it might be because of high wire utilization. By turning on this option, you might be able to reduce your compilation time.

On average, this option can save up to 6% wire utilization, but can also reduce performance by up to 4%, depending on the device.

These optimizations are used automatically when the Fitter performs more than one fitting attempt, but turning the option on increases the optimization effort on the first fitting attempt. This option also ensures that the Quartus II software uses maximum optimization to reduce routability, even if the Fitter Effort is set to **Auto Fit**.

Increase Placement Effort Multiplier

Increasing the placement effort can improve the routability of the design, allowing the software to route a design that otherwise requires too many routing resources.

- ❓ For more information refer to *Placement Effort Multiplier logic option* in Quartus II Help.

Increased effort is used automatically when the Fitter performs more than one fitting attempt. Setting a multiplier higher than one (before compilation) increases the optimization effort on the first fitting attempt. The second and third fitting loops increase the Placement Effort Multiplier to 4 and then to 16. These loops result in increased compilation times, with possible improvement in the quality of placement.

You can modify the Placement Effort Multiplier using the following Tcl command:

```
set_global_assignment -name PLACEMENT_EFFORT_MULTIPLIER <value> ←
```

<value> can be any positive, non-zero number.

Increasing placement effort is likely to reduce congestion during routing, and help fit hard-to-route designs.

Increase Router Effort Multiplier

The Router Effort Multiplier controls how quickly the router tries to find a valid solution. The default value is 1.0 and legal values must be greater than 0. Numbers higher than 1 (values up to 3 are generally considered reasonable) help designs that are difficult to route by increasing the routing effort, but do not necessarily resolve routing congestion. Numbers closer to 0 (for example, 0.1) can reduce router runtime, but usually reduce routing quality slightly. Experimental evidence shows that a multiplier of 3.0 reduces overall wire usage by about 2%. There is usually no gain in performance beyond a multiplier value of 3.

You can set the Router Effort Multiplier to a value higher than the default value for difficult-to-route designs. To set the Router Effort Multiplier, on the Assignments menu, click **Settings**, and then click **Fitter Settings**. Click the **More Settings** button. From the options available, select **Router Effort Multiplier** and edit the value in the dialog box that appears.

You can modify the Router Effort Multiplier by entering the following Tcl command:

```
set_global_assignment -name ROUTER_EFFORT_MULTIPLIER <value> ←
```

<value> can be any positive, non-zero number.

Remove Fitter Constraints

A design with conflicting constraints or constraints that are difficult to meet may not fit the targeted device. This can occur when location or LogicLock assignments are too strict and there are not enough routing resources.

In this case, use the **Routing Congestion** task in the Chip Planner to locate routing problems in the floorplan, then remove all location and LogicLock region assignments from that area. If the local constraints are removed, and the design still does not fit, the design is over-constrained. To correct the problem, remove all location and LogicLock assignments and run successive compilations, incrementally constraining the design before each compilation. You can delete specific location assignments in the Assignment Editor or the Chip Planner. Remove LogicLock assignments in the Chip Planner, in the LogicLock Regions Window, or on the Assignments menu, click **Remove Assignments**. Turn on the assignment categories you want to remove from the design in the **Available assignment categories** list.

 For more information about the **Routing Congestion** task in the Chip Planner, refer to the *Analyzing and Optimizing the Design Floorplan* chapter in volume 2 of the *Quartus II Handbook*. You can also refer to the Quartus II Help.

Optimize Synthesis for Area, Not Speed

In some cases, resynthesizing the design to improve the area utilization can also improve the routability of the design. First, ensure that you have set your device and timing constraints correctly in your synthesis tool. Ensure that you do not over-constrain the timing requirements for the design, particularly when the area utilization of the design is a concern. Synthesis tools generally try to meet the specified requirements, which can result in higher device resource usage if the constraints are too aggressive.

If resource utilization is important to improving the routing results in your design, some synthesis tools offer an easy way to optimize for area instead of speed. If you are using Quartus II integrated synthesis, on the Assignments menu, click **Settings**. In the **Category** list, select **Analysis & Synthesis Settings**, and select **Balanced** or **Area** under **Optimization Technique**.

You can also specify this logic option for specific modules in your design with the Assignment Editor in cases where you want to reduce area using the **Area** setting (potentially at the expense of register-to-register timing performance). You can apply the setting to specific modules while leaving the default **Optimization Technique** setting at **Balanced** (for the best trade-off between area and speed for certain device families) or **Speed**. You can also use the **Speed Optimization Technique for Clock Domains** logic option to specify that all combinational logic in or between the specified clock domain(s) is optimized for speed.

 In the Quartus II software, the **Balanced** setting typically produces utilization results that are very similar to those obtained with the **Area** setting, with better performance results. The **Area** setting can yield better results in some unusual cases.

In some synthesis tools, not specifying an f_{MAX} requirement can result in less resource utilization, which can improve routability.

 For information about setting timing requirements and synthesis options in Quartus II integrated synthesis and other synthesis tools, refer to the appropriate chapter in *Section III. Synthesis* in volume 1 of the *Quartus II Handbook*, or your synthesis software's documentation.

Optimize Source Code

If your design does not fit because of routing problems and the methods described in the preceding sections do not sufficiently improve the routability of the design, modify the design at the source to achieve the desired results. You can often improve results significantly by making design-specific changes to your source code, such as duplicating logic or changing the connections between blocks that require significant routing resources.

Use a Larger Device

If a successful fit cannot be achieved because of a shortage of routing resources, you might require a larger device.

Timing Optimization Techniques (LUT-Based Devices)

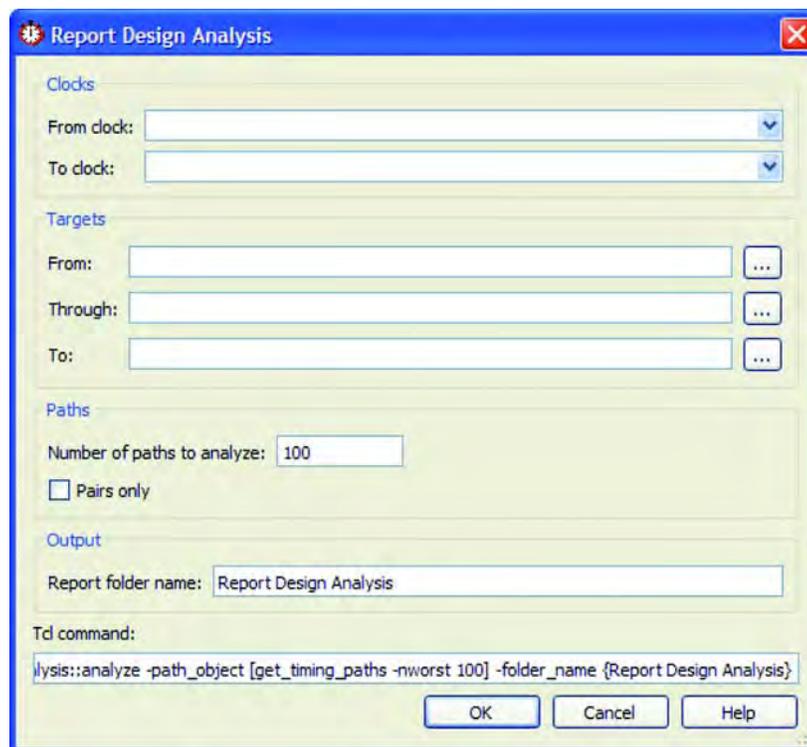
This section contains guidelines that might help you if your design does not meet its timing requirements.

Debugging Timing Failures in the TimeQuest Analyzer

Beginning with the Quartus II software version 10.1, a new **Report Timing Closure Recommendations** task is available in the Custom Reports section of the **Tasks** pane of the TimeQuest analyzer. Use this report to get more information and help on the failing paths in your design. This feature is available for Arria II GX, Cyclone III, Cyclone IV, Stratix III, Stratix IV, and Stratix V device families.

Selecting the **Report Timing Closure Recommendations** task opens the **Report Design Analysis** dialog box (Figure 13-5).

Figure 13-5. Report Design Analysis Dialog Box



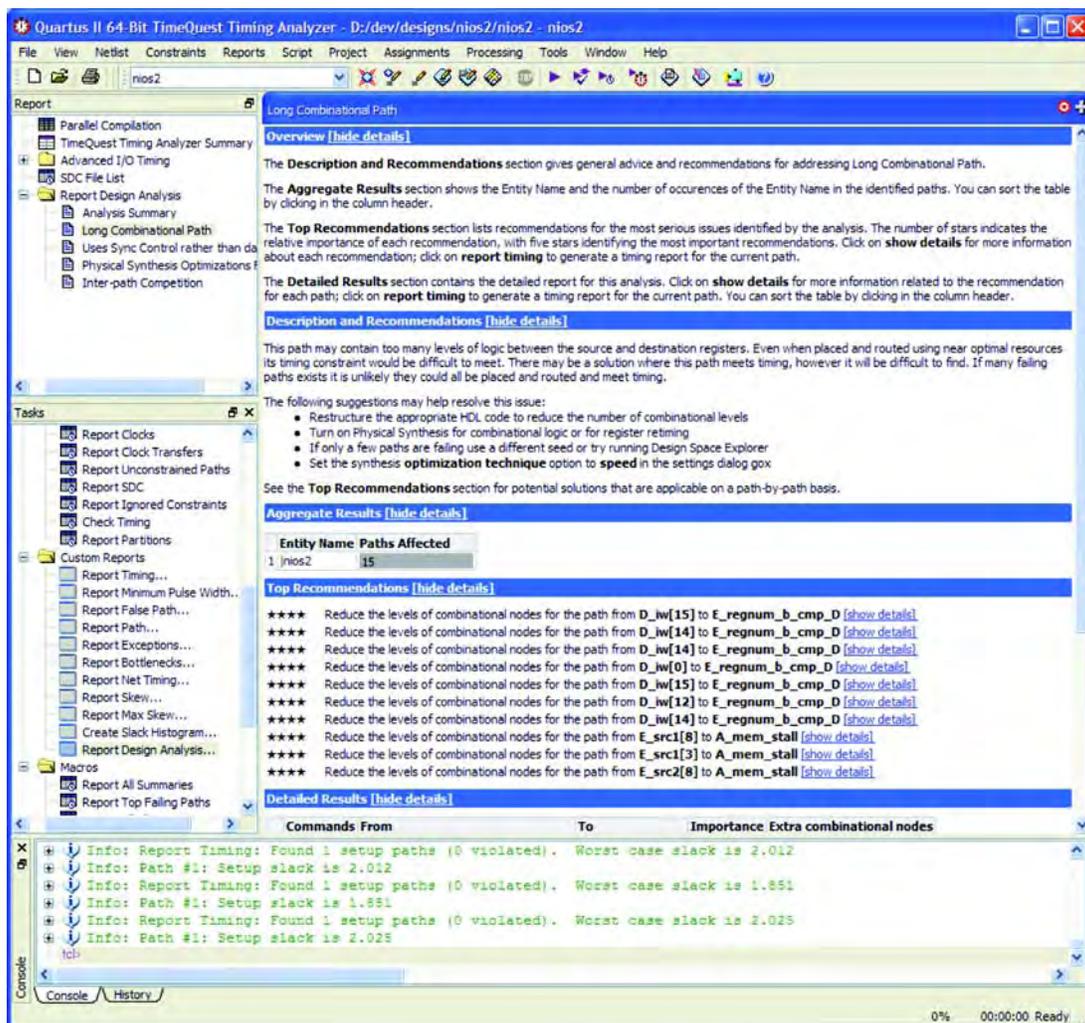
When you run the **Report Timing Closure Recommendations** task, you get specific recommendations about failing paths in your design and changes that you can make to potentially fix the failing paths.

From the dialog box (Figure 13-5), you can select paths based on the clock domain, filter by nodes on path, and choose the number of paths to analyze.

After running this command in the TimeQuest analyzer, examine the reports in the Report Timing Closure Recommendations folder in the **Report** pane of the TimeQuest analyzer GUI. Each recommendation has star symbols (*) associated with it. Recommendations with more stars are more likely to help you close timing on your design.

Figure 13-6 shows an example report.

Figure 13-6. Example Report



The reports give you the most probable causes of failure for each path being analyzed. The reports are organized into sections, depending on the type of issues found in the design, such as large clock skew, restricted optimizations, unbalanced logic, skipped optimizations, coding style that has too many levels of logic between registers, or region or partition constraints specific to your project.

You will see recommendations that may help you fix the failing paths. For a detailed analysis of the critical paths, run the **report_timing** command on specified paths. In the **Extra Fitter Information** tab of the Path report panel, you will also see detailed Fitter-related information that may help you visualize the issue and take appropriate action if user constraints cause a specific placement.

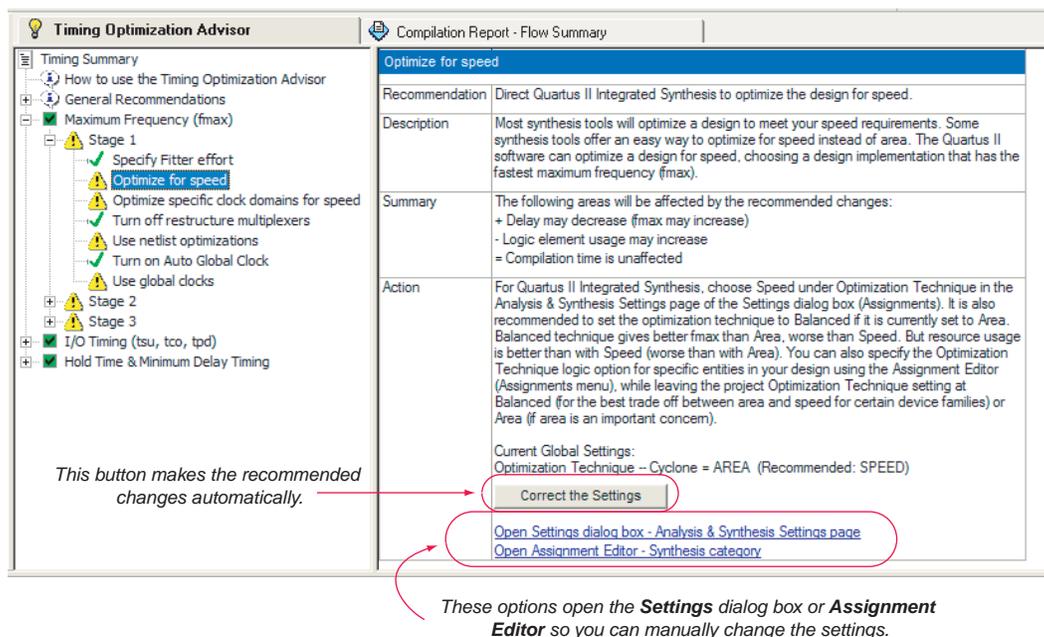
Timing Optimization Advisor

The Timing Optimization Advisor guides you in making settings that optimize your design to meet your timing requirements. To run the Timing Optimization Advisor, on the Tools menu, point to **Advisors**, and click on **Timing Optimization Advisor**. This advisor describes many of the suggestions made in this section.

When you open the Timing Optimization Advisor after compilation, you find recommendations to improve the timing performance of your design. Some of the recommendations in these advisors can contradict each other. Altera recommends evaluating these options and choosing the settings that best suit the given requirements.

The example in [Figure 13-7](#) shows the Timing Optimization Advisor after compiling a design that meets its frequency requirements, but requires setting changes to improve the timing.

Figure 13-7. Timing Optimization Advisor



When you expand one of the categories in the Advisor, such as **Maximum Frequency (fmax)** or **I/O Timing (tsu, tco, tpd)**, the recommendations are divided into stages. The stages show the order in which to apply the recommended settings. The first stage contains the options that are easiest to change, make the least drastic changes to your design optimization, and have the least effect on compilation time. Icons indicate whether each recommended setting has been made in the current project. In [Figure 13-7](#), the checkmark icons in the list of recommendations for Stage 1 indicate

recommendations that are already implemented. The warning icons indicate recommendations that are not followed for this compilation. The information icons indicate general suggestions. For these entries, the advisor does not report whether these recommendations were followed, but instead explains how you can achieve better performance. Refer to the “How to use” page in the Advisor for a legend that provides more information for each icon.

There is a link from each recommendation to the appropriate location in the Quartus II user interface where you can change the settings. For example, consider the **Synthesis Netlist Optimizations** page of the **Settings** dialog box or the **Global Signals** category in the Assignment Editor. This approach provides the most control over which settings are made and helps you learn about the settings in the software. In some cases, you can also use the **Correct the Settings** button to automatically make the suggested change to global settings.

For some entries in the advisor, a button appears that allows you to further analyze your design and gives you more information. The advisor provides a table with the clocks in the design and indicates whether they have been assigned a timing constraint.

I/O Timing Optimization

The next stage of design optimization focuses on I/O timing. Ensure that you have made the appropriate assignments as described in “**Initial Compilation: Required Settings**” on page 13-2, and that the resource utilization is satisfactory before proceeding with I/O timing optimization. The suggestions provided in this section are applicable to all Altera FPGA families and to the MAX II family of CPLDs.

Because changes to the I/O paths affect the internal register-to-register timing, complete this stage before proceeding to the register-to-register timing optimization stage as described in the “**Register-to-Register Timing Optimization Techniques (LUT-Based Devices)**” on page 13-34.

The options presented in this section address how to improve I/O timing, including the setup delay (t_{SU}), hold time (t_H), and clock-to-output (t_{CO}) parameters.

Improving Setup and Clock-to-Output Times Summary

Table 13-1 shows the recommended order in which to use techniques to reduce t_{SU} and t_{CO} times. Checkmarks indicate which timing parameters are affected by each technique. Reducing t_{SU} times increases hold (t_H) times.

Table 13-1. Improving Setup and Clock-to-Output Times (Note 1) (Part 1 of 2)

Technique	Affects t_{SU}	Affects t_{CO}
Ensure that the appropriate constraints are set for the failing I/Os (page 13-3)	✓	✓
Use timing-driven compilation for I/O (page 13-31)	✓	✓
Use fast input register (page 13-32)	✓	—
Use fast output register, fast output enable register, and fast OCT register (page 13-32)	—	✓
Decrease the value of Input Delay from Pin to Input Register or set Decrease Input Delay to Input Register = ON	✓	—
Decrease the value of Input Delay from Pin to Internal Cells , or set Decrease Input Delay to Internal Cells = ON	✓	—

Table 13–1. Improving Setup and Clock-to-Output Times (Note 1) (Part 2 of 2)

Technique	Affects t_{SU}	Affects t_{CO}
Decrease the value of Delay from Output Register to Output Pin , or set Increase Delay to Output Pin = OFF (page 13–33)	—	✓
Increase the value of Input Delay from Dual-Purpose Clock Pin to Fan-Out Destinations (page 13–33)	✓	—
Use PLLs to shift clock edges (page 13–33)	✓	✓
Use the Fast Regional Clock (page 13–34)	—	✓
For MAX II series devices, set Guarantee I/O paths to zero, Hold Time at Fast Timing Corner to OFF , or when t_{SU} and t_{PD} constraints permit (page 13–34)	✓	—
Increase the value of Delay to output enable pin or set Increase delay to output enable pin (page 13–33)	—	✓

Note to Table 13–1:

(1) These options may not apply to all device families.

Timing-Driven Compilation

This option moves registers into I/O elements if required to meet t_{SU} or t_{CO} assignments, duplicating the register if necessary (as in the case in which a register fans out to multiple output locations). This option is turned on by default and is a global setting. The option does not apply to MAX II series devices because they do not contain I/O registers.

The **Optimize IOC Register Placement for Timing** option affects only pins that have a t_{SU} or t_{CO} requirement. Using the I/O register is possible only if the register directly feeds a pin or is fed directly by a pin. This setting does not affect registers with any of the following characteristics:

- Have combinational logic between the register and the pin
- Are part of a carry or cascade chain
- Have an overriding location assignment
- Use the asynchronous load port and the value is not 1 (in device families where the port is available)

Registers with the characteristics listed are optimized using the regular Quartus II Fitter optimizations.

❓ For more information, refer to *Optimize IOC Register Placement for Timing logic option* in Quartus II Help.

Fast Input, Output, and Output Enable Registers

You can place individual registers in I/O cells manually by making fast I/O assignments with the Assignment Editor. For an input register, use the **Fast Input Register** option; for an output register, use the **Fast Output Register** option; and for an output enable register, use the **Fast Output Enable Register** option. Stratix II devices also support the **Fast OCT (on-chip termination) Register** option. In MAX II series devices, which have no I/O registers, these assignments lock the register into the LAB adjacent to the I/O pin if there is a pin location assignment for that I/O pin.

If the fast I/O setting is on, the register is always placed in the I/O element. If the fast I/O setting is off, the register is never placed in the I/O element. This is true even if the **Optimize IOC Register Placement for Timing** option is turned on. If there is no fast I/O assignment, the Quartus II software determines whether to place registers in I/O elements if the **Optimize IOC Register Placement for Timing** option is turned on.

The four fast I/O options (**Fast Input Register**, **Fast Output Register**, **Fast Output Enable Register**, and **Fast OCT Register**) can also be used to override the location of a register that is in a LogicLock region, and force it into an I/O cell. If this assignment is applied to a register that feeds multiple pins, the register is duplicated and placed in all relevant I/O elements. In MAX II series devices, the register is duplicated and placed in each distinct LAB location that is next to an I/O pin with a pin location assignment.

Programmable Delays

Various programmable delay options can be used to minimize the t_{SU} and t_{CO} times. For Arria, Cyclone, MAX II, MAX V, and Stratix series devices, the Quartus II software automatically adjusts the applicable programmable delays to help meet timing requirements. Programmable delays are advanced options to use only after you compile a project, check the I/O timing, and determine that the timing is unsatisfactory. For detailed information about the effect of these options, refer to the device family handbook or data sheet.

After you have made a programmable delay assignment and compiled the design, you can view the implemented delay values for every delay chain for every I/O pin in the **Delay Chain Summary** section of the Compilation Report.

You can assign programmable delay options to supported nodes with the Assignment Editor. You can also view and modify the delay chain setting for the target device with the Chip Planner and Resource Property Editor. When you use the Resource Property Editor to make changes after performing a full compilation, recompiling the entire design is not necessary; you can save changes directly to the netlist. Because these changes are made directly to the netlist, the changes are not made again automatically when you recompile the design. The change management features allow you to reapply the changes on subsequent compilations.

Although the programmable delays in newer devices are user-controllable, Altera recommends their use for advanced users only. However, the Quartus II software might use the programmable delays internally during the Fitter phase.

 For more details about Stratix III programmable delays, refer to the *Stratix III Device Handbook* and *AN 474: Implementing Stratix III Programmable I/O Delay Settings in the Quartus II Software*. For more information about using the Chip Planner and Resource Property Editor, refer to the *Engineering Change Management with the Chip Planner* chapter in volume 2 of the *Quartus II Handbook*.

 For details about the programmable delay logic options available for Altera devices, refer to the following Quartus II Help topics:

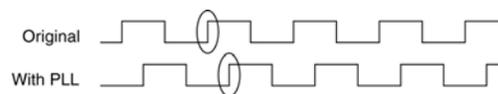
- *Decrease Input Delay to Input Register*
- *Input Delay from Pin to Input Register*

- *Decrease Input Delay to Internal Cells*
- *Input Delay from Pin to Internal Cells*
- *Decrease Input Delay to Output Register*
- *Increase Delay to Output Enable Pin*
- *Output Enable Pin Delay*
- *Increase Delay to Output Pin*
- *Delay from Output Register to Output Pin*
- *Increase Input Clock Enable Delay*
- *Input Delay from Dual-Purpose Clock Pin to Fan-Out Destinations*
- *Increase Output Clock Enable Delay*
- *Increase Output Enable Clock Enable Delay*
- *Increase t_{zx} Delay to Output Pin*

Use PLLs to Shift Clock Edges

Using a PLL typically improves I/O timing automatically. If the timing requirements are still not met, most devices allow the PLL output to be phase shifted to change the I/O timing. Shifting the clock backwards gives a better t_H at the expense of t_{SU} , while shifting it forward gives a better t_{SU} at the expense of t_H (refer to [Figure 13-8](#)). This technique can be used only in devices that offer PLLs with the phase shift option.

Figure 13-8. Shift Clock Edges Forward to Improve t_{SU} at the Expense of t_H



You can achieve the same type of effect in certain devices by using the programmable delay called **Input Delay from Dual Purpose Clock Pin to Fan-Out Destinations**.

- ❓ For more information, refer to *Input Delay from Dual-Purpose Clock Pin to Fan-Out Destinations* in Quartus II Help.

Use Fast Regional Clock Networks and Regional Clocks Networks

Altera devices have a variety of hierarchical clock structures. These include dedicated global clock networks (GCLKs), regional clock networks (RCLKs), fast regional clock networks (FCLK) and periphery clock networks (PCLKs). The available resources differ between various Altera device families.

- 🔗 Refer to the appropriate device handbook to get the number of various clocking resources available in your target device.

In general, fast regional clocks have less delay to I/O elements than regional and global clocks, and are used for high fan-out control signals. Regional clocks provide the lowest clock delay and skew for logic contained in a single quadrant. Placing clocks on these low-skew and low-delay clock nets provides better t_{CO} performance.

Change How Hold Times are Optimized for MAX II Devices

For MAX II series devices, you can use the **Guarantee I/O paths have zero hold time at Fast Timing Corner** option to control how hold time is optimized by the Quartus II software.

- For details, refer to *Guarantee I/O Paths Have Zero Hold Time at Fast Corner logic option* in Quartus II Help.

Register-to-Register Timing Optimization Techniques (LUT-Based Devices)

The next stage of design optimization is to improve register-to-register (f_{MAX}) timing. The following sections provide available options if the performance requirements are not achieved after compilation.

Coding style affects the performance of your design to a greater extent than other changes in settings. Always evaluate your code and make sure to use synchronous design practices.

- For more details about synchronous design practices and coding styles, refer to the *Recommended Design Practices* chapter in volume 1 of the *Quartus II Handbook*.

- When using the TimeQuest analyzer, register-to-register timing optimization is the same as maximizing the slack on the clock domains in your design. You can use the techniques described in this section to improve the slack on different timing paths in your design.

Before optimizing your design, understand the structure of your design as well as the type of logic affected by each optimization. An optimization can decrease performance if the optimization does not benefit your logic structure.

Improving Register-to-Register Timing Summary

The choice of options and settings to improve the timing margin (slack) or to improve register-to-register timing depends on the failing paths in the design. To achieve the results that best approximate your performance requirements, apply the following techniques and compile the design after each step:

1. Ensure that your timing assignments are complete and correct. For details, refer to *“Timing Requirement Settings”* on page 13-3.
2. Ensure that you have reviewed all warning messages from your initial compilation and check for ignored timing assignments. Refer to *“Design Analysis”* on page 13-10 for details and fix any of these problems before proceeding with optimization.
3. Apply netlist synthesis optimization options.

Apply the following synthesis options to optimize for speed:

- “Optimize Synthesis for Speed, Not Area” on page 13-37
- “Flatten the Hierarchy During Synthesis” on page 13-38
- “Set the Synthesis Effort to High” on page 13-38
- “Change State Machine Encoding” on page 13-38
- “Prevent Shift Register Inference” on page 13-39
- “Use Other Synthesis Options Available in Your Synthesis Tool” on page 13-39

4. Apply the following options for physical synthesis optimization:
 - Perform physical synthesis for combinational logic
 - Perform automatic asynchronous signal pipelining
 - Perform register duplication
 - Perform register retiming
 - Perform logic to memory mapping
5. Try different Fitter seeds (page 13-39). You can omit this step if a large number of critical paths are failing, or if the paths are failing badly.
6. Make LogicLock assignments (page 13-41) to control placement.
7. Make design source code modifications to fix areas of the design that are still failing timing requirements by significant amounts (page 13-40).
8. Make location assignments, or as a last resort, perform manual placement by back-annotating the design (page 13-42).

You can use the Design Space Explorer (DSE) to automate the process of running several different compilations with different settings.

- ② For more information, refer to *About Design Space Explorer* in Quartus II Help.

If these techniques do not achieve performance requirements, additional design source code modifications might be required (page 13-40).

Physical Synthesis Optimizations

The Quartus II software offers physical synthesis optimizations that can help improve the performance of many designs regardless of the synthesis tool used. Physical synthesis optimizations can be applied both during synthesis and during fitting.

Physical synthesis optimizations that occur during the synthesis stage of the Quartus II compilation operate either on the output from another EDA synthesis tool or as an intermediate step in Quartus II integrated synthesis. These optimizations make changes to the synthesis netlist to improve either area or speed, depending on your selected optimization technique and effort level.

To view and modify the synthesis netlist optimization options, on the Assignments menu, click **Settings**. In the **Category** list, expand **Compilation Process Settings** and select **Physical Synthesis Optimizations**.

If you use a third-party EDA synthesis tool and want to determine if the Quartus II software can remap the circuit to improve performance, you can use the **Perform WYSIWYG Primitive Resynthesis** option. This option directs the Quartus II software to unmap the LEs in an atom netlist to logic gates and then map the gates back to Altera-specific primitives. Using Altera-specific primitives enables the Fitter to remap the circuits using architecture-specific techniques.

- ② For more information, refer to *Perform WYSIWYG Primitive Resynthesis logic option* in Quartus II Help.

The Quartus II technology mapper optimizes the design for **Speed**, **Area**, or **Balanced**, according to the setting of the **Optimization Technique** option. Set this option to **Speed** or **Balanced**.

- ② For more information, refer to *Optimization Technique logic option* in Quartus II Help.

The physical synthesis optimizations occur during the Fitter stage of the Quartus II compilation. Physical synthesis optimizations make placement-specific changes to the netlist that improve speed performance results for a specific Altera device.

The following physical synthesis optimizations are available during the Fitter stage for improving performance:

- Physical synthesis for combinational logic
- Automatic asynchronous signal pipelining
- Physical synthesis for registers
 - Register duplication
 - Register retiming

 You can apply physical synthesis options on specific instances if you want the performance gain from physical synthesis only on parts of your design.

- ② For more information, refer to *Physical Synthesis Optimizations Page (Settings Dialog Box)* in Quartus II Help.

To apply physical synthesis assignments for fitting on a per instance basis, use the Quartus II Assignment Editor. The following assignments are available as instance assignments:

- Perform physical synthesis for combinational logic
- Perform register duplication for performance
- Perform register retiming for performance
- Perform automatic asynchronous signal pipelining

Follow these steps:

1. In the Assignment Editor, indicate the module instance you want to apply to the specific physical synthesis setting in the **To** tab.
2. Select the required physical synthesis assignment in the **Assignment Name** tab.
3. In the **Value** tab, select **ON**.

4. In the **Enabled** tab, select **Yes**.

Turn Off Extra-Effort Power Optimization Settings

If PowerPlay power optimization settings are set to **Extra Effort**, your design performance can be affected. If improving timing performance is more important than reducing power use, set the PowerPlay power optimization setting to **Normal**.

-  For more information, refer to *PowerPlay Power Optimization logic option* in Quartus II Help.

-  For more information about reducing power use, refer to the *Power Optimization* chapter in volume 2 of the *Quartus II Handbook*.

Optimize Synthesis for Speed, Not Area

The manner in which the design is synthesized has a large impact on design performance. Design performance varies depending on the way the design is coded, the synthesis tool used, and the options specified when synthesizing. Change your synthesis options if a large number of paths are failing, or if specific paths are failing badly and have many levels of logic.

Set your device and timing constraints in your synthesis tool. Synthesis tools are timing-driven and optimized to meet specified timing requirements. If you do not specify target frequency, some synthesis tools optimize for area.

Some synthesis tools offer an easy way to instruct the tool to focus on speed instead of area.

-  For more information, refer to *Optimization Technique logic option* in Quartus II Help

You can also specify this logic option for specific modules in your design with the Assignment Editor while leaving the default **Optimization Technique** setting at **Balanced** (for the best trade-off between area and speed for certain device families) or **Area** (if area is an important concern). You can also use the **Speed Optimization Technique for Clock Domains** option in the Assignment Editor to specify that all combinational logic in or between the specified clock domain(s) is optimized for speed.

To achieve best performance with push-button compilation, follow the recommendations in the following sections for other synthesis settings. You can use the DSE to experiment with different Quartus II synthesis options to optimize your design for the best performance.

-  For information about setting timing requirements and synthesis options in Quartus II integrated synthesis and third-party synthesis tools, refer to the appropriate chapter in *Section III. Synthesis* in volume 1 of the *Quartus II Handbook*, or refer to your synthesis software documentation.

-  For more information about the Design Space Explorer, refer to *About Design Space Explorer* in Quartus II Help.

Flatten the Hierarchy During Synthesis

Synthesis tools typically let you preserve hierarchical boundaries, which can be useful for verification or other purposes. However, the best optimization results generally occur when the synthesis tool optimizes across hierarchical boundaries, because doing so often allows the synthesis tool to perform the most logic minimization, which can improve performance. Whenever possible, flatten your design hierarchy to achieve the best results. If you are using Quartus II incremental compilation, you cannot flatten your design across design partitions. Incremental compilation always preserves the hierarchical boundaries between design partitions. Follow Altera's recommendations for design partitioning, such as registering partition boundaries to reduce the effect of cross-boundary optimizations.



For more information about using incremental compilation and recommendations for design partitioning, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.

Set the Synthesis Effort to High

Some synthesis tools offer varying synthesis effort levels to trade off compilation time with synthesis results. Set the synthesis effort to high to achieve best results when applicable.

Change State Machine Encoding

State machines can be encoded using various techniques. One-hot encoding, which uses one register for every state bit, usually provides the best performance. If your design contains state machines, changing the state machine encoding to one-hot can improve performance at the cost of area.



For more information, refer to *State Machine Processing logic option* in Quartus II Help.

Duplicate Logic for Fan-Out Control

Duplicating logic or registers can help improve timing in cases where moving a register in a failing timing path to reduce routing delay creates other failing paths, or where there are timing problems due to the fan-out of the registers. Most often, timing failures occur not because of the high fan-out registers, but because of the location of those registers. Duplicating registers, where source and destination registers are physically close, can help improve slack on critical paths.

Many synthesis tools support options or attributes that specify the maximum fan-out of a register. When using Quartus II integrated synthesis, you can set the **Maximum Fan-Out** logic option in the Assignment Editor to control the number of destinations for a node so that the fan-out count does not exceed a specified value. You can also use the `maxfan` attribute in your HDL code. The software duplicates the node as required to achieve the specified maximum fan-out.



Logic duplication using **Maximum Fan-Out** assignments normally increases resource utilization and can potentially increase compilation time, depending on the placement and the total resource usage within the selected device. The improvement in timing performance that results because of **Maximum Fan-Out** assignments is very design-specific.

If you are using **Maximum Fan-Out** assignments, Altera recommends benchmarking your design with and without these assignments to evaluate whether they give the expected improvement in timing performance. Use the assignments only when you get improved results.

You can manually duplicate registers in the Quartus II software regardless of the synthesis tool used. To duplicate a register, apply the **Manual Logic Duplication** option to the register with the Assignment Editor.

-  For more information, refer to *Manual Logic Duplication logic option* in Quartus II Help.

Prevent Shift Register Inference

In some cases, turning off the inference of shift registers increases performance. Doing so forces the software to use logic cells to implement the shift register instead of implementing the registers in memory blocks using the ALTSHIFT_TAPS megafunction. If you implement shift registers in logic cells instead of memory, logic utilization is increased.

Use Other Synthesis Options Available in Your Synthesis Tool

With your synthesis tool, experiment with the following options if they are available:

- Turn on register balancing or retiming
- Turn on register pipelining
- Turn off resource sharing

These options can increase performance, but typically increase the resource utilization of your design.

Fitter Seed

The Fitter seed affects the initial placement configuration of the design. Changing the seed value changes the Fitter results, because the fitting results change whenever there is a change in the initial conditions. Each seed value results in a somewhat different fit, and you can experiment with several different seeds to attempt to obtain better fitting results and timing performance.

When there are changes in your design, there is some random variation in performance between compilations. This variation is inherent in placement and routing algorithms—there are too many possibilities to try them all and get the absolute best result, so the initial conditions change the compilation result.

-  Any design change that directly or indirectly affects the Fitter has the same type of random effect as changing the seed value. This includes any change in source files, **Analysis & Synthesis Settings**, **Fitter Settings**, or **Timing Analyzer Settings**. The same effect can appear if you use a different computer processor type or different operating system, because different systems can change the way floating point numbers are calculated in the Fitter.

If a change in optimization settings slightly affects the register-to-register timing or number of failing paths, you cannot always be certain that your change caused the improvement or degradation, or whether it could be due to random effects in the Fitter. If your design is still changing, running a seed sweep (compiling your design with multiple seeds) determines whether the average result has improved after an optimization change and whether a setting that increases compilation time has benefits worth the increased time (such as setting the **Physical Synthesis Effort** to **Extra**). The sweep also shows the amount of random variation to expect for your design.

If your design is finalized, you can compile your design with different seeds to obtain one optimal result. However, if you subsequently make any changes to your design, you might need to perform seed sweep again.

On the Assignments menu, select **Fitter Settings** to control the initial placement with the seed. You can use the DSE to perform a seed sweep easily.

You can use the following Tcl command from a script to specify a Fitter seed:

```
set_global_assignment -name SEED <value> ←
```

- ② For more information about compiling your design with different seeds using the Design Space Explorer (DSE seed sweep), refer to *About Design Space Explorer* in Quartus II Help.

Set Maximum Router Timing Optimization Level

To improve routability in designs where the router did not pick up the optimal routing lines, set the **Router Timing Optimization Level** to **Maximum**. This setting determines how aggressively the router tries to meet timing requirements. Setting this option to **Maximum** can increase design speed slightly at the cost of increased compilation time. Setting this option to **Minimum** can reduce compilation time at the cost of slightly reduced design speed. The default value is **Normal**.

- ② For more information, refer to *Router Timing Optimization Level logic option* in Quartus II Help.

Optimize Source Code

If the methods described in the preceding sections do not sufficiently improve timing of the design, modify your design files to achieve the desired results. Try restructuring the design to use pipelining or more efficient coding techniques. In many cases, optimizing the design's source code can have a very significant effect on your design performance. In fact, optimizing your source code is typically the most effective technique for improving the quality of your results, and is often a better choice than using LogicLock or location assignments.

If the critical path in your design involves memory or DSP functions, check whether you have code blocks in your design that describe memory or functions that are not being inferred and placed in dedicated logic. You might be able to modify your source code to cause these functions to be placed into high-performance dedicated memory or resources in the target device.

Ensure that your state machines are recognized as state machine logic and optimized appropriately in your synthesis tool. State machines that are recognized are generally optimized better than if the synthesis tool treats them as generic logic. In the Quartus II software, you can check for the State Machine report under **Analysis & Synthesis** in the Compilation Report. This report provides details, including the state encoding for each state machine that was recognized during compilation. If your state machine is not being recognized, you might have to change your source code to enable it to be recognized.

-  For coding style guidelines including examples of HDL code for inferring memory, functions, guidelines, and sample HDL code for state machines, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*.
-  For additional HDL coding examples, refer to *AN 584: Timing Closure Methodology for Advanced FPGAs*.

LogicLock Assignments

Using LogicLock assignments to improve timing performance is only recommended for older Altera devices, such as the MAX II family. For other device families, especially for larger devices such as Arria and Stratix series devices, using LogicLock assignments to improve timing performance is not recommended. For these devices, the LogicLock feature is intended to be used for performance preservation and to floorplan your design.

LogicLock assignments do not always improve the performance of the design. In many cases, you cannot improve upon results from the Fitter by making location assignments. If there are existing LogicLock assignments in your design, remove the assignments if your design methodology permits it. Recompile the design to see if the assignments are making the performance worse.

When making LogicLock assignments, it is important to consider how much flexibility to give the Fitter. LogicLock assignments provide more flexibility than hard location assignments. Assignments that are more flexible require higher Fitter effort, but reduce the chance of design over-constraint. The following types of LogicLock assignments are available, listed in the order of decreasing flexibility:

- Auto size, floating location regions
- Fixed size, floating location regions
- Fixed size, locked location regions

-  For more information about using LogicLock regions, refer to the *Analyzing and Optimizing the Design Floorplan* chapter in volume 2 of the *Quartus II Handbook*.

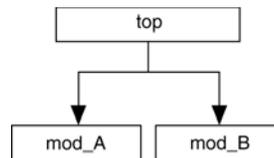
To determine what to put into a LogicLock region, refer to the timing analysis results and analyze the critical paths in the Chip Planner. The register-to-register timing paths in the Timing Analyzer section of the Compilation Report help you recognize patterns.

The following sections describe cases in which LogicLock regions can help to optimize a design.

Hierarchy Assignments

For a design with the hierarchy shown in [Figure 13-9](#), which has failing paths in the timing analysis results similar to those shown in [Table 13-2](#), `mod_A` is probably a problem module. In this case, a good strategy to fix the failing paths is to place the `mod_A` hierarchy block in a LogicLock region so that all the nodes are closer together in the floorplan.

Figure 13-9. Design Hierarchy



[Table 13-2](#) shows the failing paths connecting two regions together within `mod_A` listed in the timing analysis report.

Table 13-2. Failing Paths in a Module Listed in Timing Analysis

From	To
mod_A reg1	mod_A reg9
mod_A reg3	mod_A reg5
mod_A reg4	mod_A reg6
mod_A reg7	mod_A reg10
mod_A reg0	mod_A reg2

Hierarchical LogicLock regions are also important if you are using an incremental compilation flow. Place each design partition for incremental compilation in a separate LogicLock region to reduce conflicts and ensure good results as the design develops. You can use auto size and floating location regions to find a good design floorplan, but fix the size and placement to achieve the best results in future compilations.

 For more information about using incremental compilation and recommendations for creating a design floorplan using LogicLock regions, refer to the [Quartus II Incremental Compilation for Hierarchical and Team-Based Design](#) and [Best Practices for Incremental Compilation and Floorplan Assignments](#) chapters in volume 1 of the *Quartus II Handbook*, and [Analyzing and Optimizing the Design Floorplan](#) chapter in volume 2 of the *Quartus II Handbook*.

Location Assignments and Back-Annotation

If a small number of paths are failing to meet their timing requirements, you can use hard location assignments to optimize placement. Location assignments are less flexible for the Quartus II Fitter than LogicLock assignments. In some cases, when you are familiar with your design, you can enter location constraints in a way that produces better results.

 Improving fitting results, especially for larger devices, such as Arria and Stratix series devices, can be difficult. Location assignments do not always improve the performance of the design. In many cases, you cannot improve upon the results from the Fitter by making location assignments.

Metastability Analysis and Optimization Techniques

Metastability problems can occur when a signal is transferred between circuitry in unrelated or asynchronous clock domains, because the designer cannot guarantee that the signal will meet its setup and hold time requirements. The mean time between failure (MTBF) is an estimate of the average time between instances when metastability could cause a design failure.

 For more information about metastability and MTBF, refer to the *Understanding Metastability in FPGAs* white paper.

You can use the Quartus II software to analyze the average MTBF due to metastability when a design synchronizes asynchronous signals, and optimize the design to improve the MTBF. These metastability features are supported only for designs constrained with the TimeQuest analyzer, and for select device families.

If the MTBF of your design is low, refer to the Metastability Optimization section in the Timing Optimization Advisor, which suggests various settings that can help optimize your design in terms of metastability.

 For details about the metastability features in the Quartus II software, refer to the *Managing Metastability with the Quartus II Software* chapter in volume 1 of the *Quartus II Handbook*. This chapter describes how to enable metastability analysis and identify the register synchronization chains in your design, provides details about metastability reports, and provides additional guidelines for managing metastability.

Resource Utilization Optimization Techniques (Macrocell-Based CPLDs)

The following recommendations help you take advantage of the macrocell-based architecture in the MAX 7000 and MAX 3000 devices to yield maximum speed, reliability, and device resource utilization while minimizing fitting difficulties.

After design analysis, the first stage of design optimization is to improve resource utilization. Complete this stage before proceeding to timing optimization. First, ensure that you have set the basic constraints described in *“Initial Compilation: Required Settings”* on page 13-2. If your design is not fitting into a specified device, use the techniques in this section to achieve a successful fit.

Use Dedicated Inputs for Global Control Signals

MAX 7000 and MAX 3000 devices have four dedicated inputs that can be used for global register control. Because the global register control signals can bypass the logic cell array and directly feed registers, product terms can be preserved for primary logic. Also, because each signal has a dedicated path into the LAB, global signals also can bypass logic and data path interconnect resources.

Because the dedicated input pins are designed for high fan-out control signals and provide low skew, always assign global signals (such as clock, clear, and output enable) to the dedicated input pins.

You can use logic-generated control signals for global control signals instead of dedicated inputs. However, the following list shows the disadvantages of using logic-generated control signals:

- More resources are required (logic cells, interconnect).
- More data skew is introduced.
- If the logic-generated control signals have high fan-out, the design can be more difficult to fit.

By default, the Quartus II software uses dedicated inputs for global control signals automatically. You can assign control signals to dedicated input pins in one of the following ways:

- In the Assignment Editor, select one of the two following methods:
 - Assign pins to dedicated pin locations.
 - Assign a **Global Signal** setting to the pins.
- On the Assignments menu, click **Settings**. On the **Analysis & Synthesis Settings** page, click **More Settings**, and in the **Existing Option settings** section, select **Auto Global Register Control Signals**.
- Insert a GLOBAL primitive after the pins.
- If you have already assigned pins for the design in the MAX+PLUS® II software, on the Assignments menu, click **Import Assignments**.

Reserve Device Resources

Because pin and logic option assignments can be necessary for board layout and performance requirements, and because full utilization of the device resources can increase the difficulty of fitting the design, Altera recommends that you leave 10% of the logic cells and 5% of the I/O pins unused to accommodate future design modifications. Following the Altera-recommended device resource reservation guidelines for macrocell-based CPLDs increases the chance that the Quartus II software can fit the design during recompilation after changes or assignments have been made.

Pin Assignment Guidelines and Procedures

Sometimes user-specified pin assignments are necessary for board layout. This section discusses pin assignment guidelines and procedures.

To minimize fitting issues with pin assignments, follow these guidelines:

- Assign speed-critical control signals to dedicated inputs.
- Assign output enables to appropriate locations.
- Estimate fan-in to assign output pins to the appropriate LAB.
- Assign output pins that require parallel expanders to macrocells numbered 4 to 16.

-  Altera recommends that you allow the Quartus II software to select pin assignments automatically when possible. You can use the Quartus II Pin Advisor feature (accessible from the Tools menu) for pin connection guidelines.
-  For more information about the Pin Advisor, refer to *Pin Advisor Command* in Quartus II Help.

Control Signal Pin Assignments

Assign speed-critical control signals to dedicated input pins. Every MAX 7000 and MAX 3000 device has four dedicated input pins (GCLK1, OE2/GCLK2, OE1, and GCLRn). You can assign clocks to global clock dedicated inputs (GCLK1 and OE2/GCLK2), clear to the global clear dedicated input (GCLRn), and speed-critical output enable to global OE dedicated inputs (OE1 and OE2/GCLK2).

Output Enable Pin Assignments

Occasionally, because the total number of required output enable pins is more than the dedicated input pins, output enable signals must be assigned to I/O pins.

-  To minimize possible fitting errors when assigning the output enable pins for MAX 7000 and MAX 3000 devices, refer to *Pin-Out Files for Altera Devices* on the Altera website (www.altera.com).

Estimate Fan-In When Assigning Output Pins

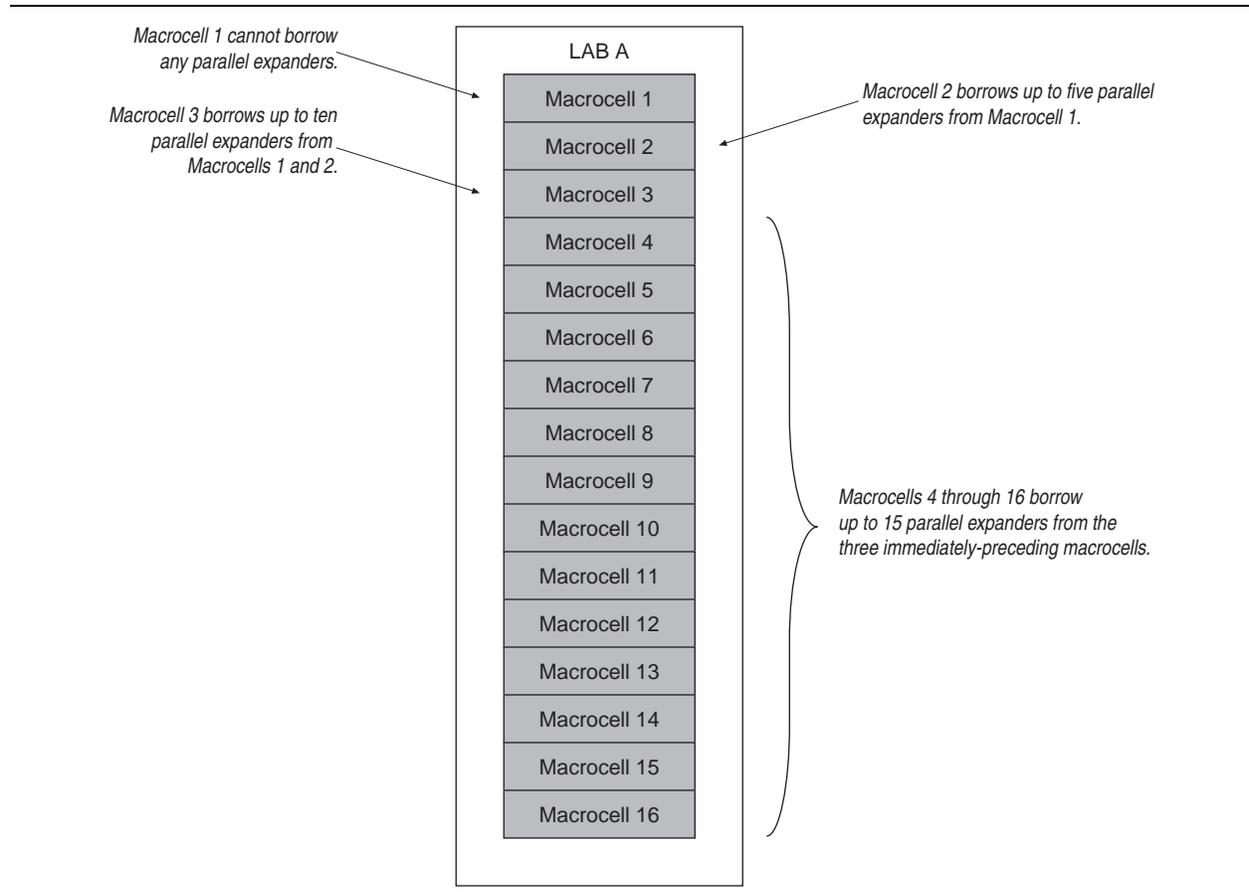
Macrocells with high fan-in can cause more placement problems for the Quartus II Fitter than those with low fan-in. The maximum fan-in per LAB should not exceed 36 in MAX 7000 and MAX 3000 devices. Therefore, estimate the fan-in of logic (such as an x -input AND gate) that feeds each output pin. If the total fan-in of logic that feeds each output pin in the same LAB exceeds 36, compilation can fail. To save resources and prevent compilation errors, avoid assigning pins that have high fan-in.

Outputs Using Parallel Expander Pin Assignments

[Figure 13-10](#) illustrates how parallel expanders are used within a LAB. MAX 7000 and MAX 3000 devices contain chains that can lend or borrow parallel expanders. The Quartus II Fitter places macrocells in a location that allows them to lend and borrow parallel expanders appropriately.

As shown in Figure 13-10, only macrocells 2 through 16 can borrow parallel expanders. Therefore, assign output pins that might require parallel expanders to pins adjacent to macrocells 4 through 16. Altera recommends using macrocells 4 through 16 because they can borrow the largest number of parallel expanders.

Figure 13-10. LAB Macrocells and Parallel Expander Associations



Resolving Resource Utilization Problems

Two common Quartus II compilation fitting issues cause errors: excessive macrocell usage and lack of routing resources. Macrocell usage errors occur when the total number of macrocells in the design exceed the available macrocells in the device. Routing errors occur when the available routing resources are insufficient to implement the design. Check the Message window for the compilation results.

 Messages in the Messages window are also copied in the Report Files. Right-click on a message and click **Help** for more information.

Resolving Macrocell Usage Issues

Occasionally, a design requires more macrocell resources than are available in the selected device, which results in the design not fitting. The following list provides tips for resolving macrocell usage issues as well as tips to minimize the number of macrocells used:

- On the Assignments menu, click **Settings**. In the **Category** list, select **Analysis & Synthesis Settings**, click **More Settings**, and turn off **Auto Parallel Expanders**. If the design's clock frequency (f_{MAX}) is not an important design requirement, turn off parallel expanders for all or part of the project. The design usually requires more macrocells if parallel expanders are turned on.
- Change **Optimization Technique** from **Speed** to **Area**. Selecting **Area** instructs the compiler to give preference to area utilization rather than speed (f_{MAX}). On the Assignments menu, click **Settings**. In the **Category** list, change the **Optimization Technique** option in the **Analysis & Synthesis Settings** page.
- Use D-type flipflops instead of latches. Altera recommends that you always use D-type flipflops instead of latches in your design because D-type flipflops can reduce the macrocell fan-in, and thus reduce macrocell usage. The Quartus II software uses extra logic to implement latches in MAX 7000 and MAX 3000 designs because MAX 7000 and MAX 3000 macrocells contain D-type flipflops instead of latches.
- Use asynchronous clear and preset instead of synchronous clear and preset. To reduce the product term usage, use asynchronous clear and preset in your design whenever possible. Using other control signals such as synchronous clear produces macrocells and pins with higher fan-out.



After following the suggestions in this section, if your project still does not fit the targeted device, consider using a larger device. When upgrading to a different density, the vertical package-migration feature of the MAX 7000 and MAX 3000 device families allows pin assignments to be maintained.

Resolving Routing Issues

Routing is another resource that can cause design fitting issues. For example, if the total fan-in into a LAB exceeds the maximum allowed, a no-fit error can occur during compilation. If your design does not fit the targeted device because of routing issues, consider the following suggestions:

- Use dedicated inputs/global signals for high fan-out signals. The dedicated inputs in MAX 7000 and MAX 3000 devices are designed for speed-critical and high fan-out signals. Always assign high fan-out signals to dedicated inputs/global signals.
- Change the **Optimization Technique** option from **Speed** to **Area**. This option can resolve routing resource and macrocell usage issues. Refer to [“Resolving Macrocell Usage Issues” on page 13-47](#).

- Reduce the fan-in per cell. If you are not limited by the number of macrocells used in the design, you can use the **Fan-in per cell (%)** option to reduce the fan-in per cell. The allowable values are 20–100%; the default value is 100%. Reducing the fan-in can reduce localized routing congestion but increase the macrocell count. You can set this logic option in the Assignment Editor or under **More Settings** in the **Analysis & Synthesis Settings** page of the **Settings** dialog box.
- On the Assignments menu, click **Settings**. In the **Category** list, select **Analysis & Synthesis Settings**, click **More Options**, and turn off **Auto Parallel Expanders**. By turning off the parallel expanders, you give the Quartus II software more fitting flexibility for each macrocell, allowing macrocells to be relocated. For example, each macrocell (previously grouped together in the same LAB) can be moved to a different LAB to reduce routing constraints.
- Insert logic cells. Inserting logic cells reduces fan-in and shared expanders used per macrocell, increasing routability. By default, the Quartus II software automatically inserts logic cells when necessary. Otherwise, **Auto Logic Cell** can be disabled as follows. On the Assignments menu, click **Settings**. In the **Category** list, select **Analysis & Synthesis Settings**. Under **More Settings**, turn off **Auto Logic Cell Insertion**. Refer to “[Using LCELL Buffers to Reduce Required Resources](#)” for more information.
- Change pin assignments. If you want to discard your pin assignments, you can let the Quartus II Fitter ignore some or all of the assignments.



If you prefer reassigning pins to increase routing efficiency, refer to “[Pin Assignment Guidelines and Procedures](#)” on page 13-44.

Using LCELL Buffers to Reduce Required Resources

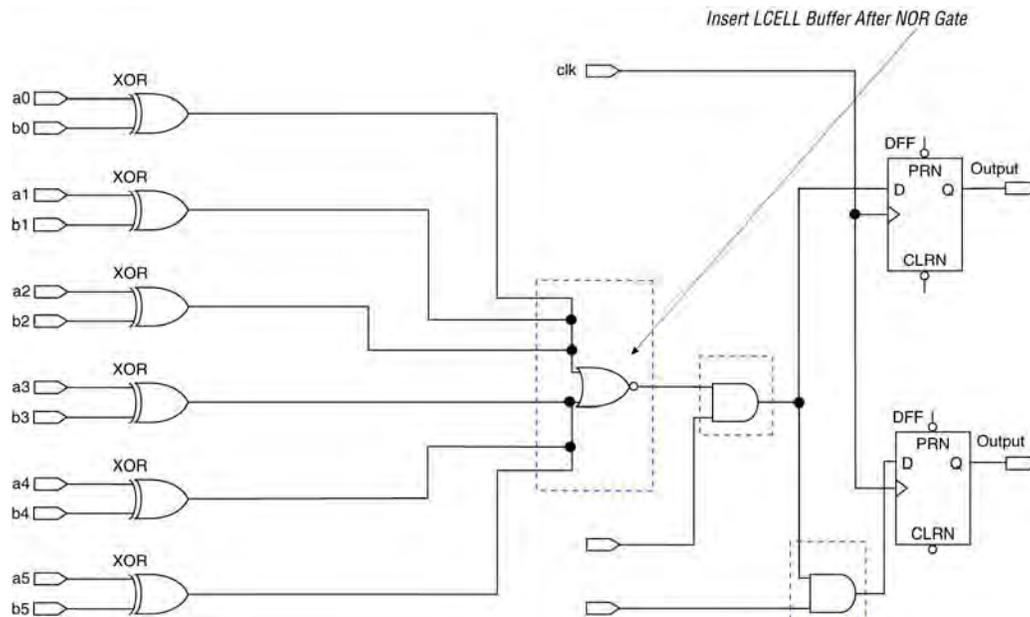
Complex logic, such as multilevel XOR gates, are often implemented with more than one macrocell. When this occurs, the Quartus II software automatically allocates shareable expanders—or additional macrocells (called synthesized logic cells)—to supplement the logic resources that are available in a single macrocell. You can also break down complex logic by inserting logic cells in the project to reduce the average fan-in and the total number of shareable expanders required. Manually inserting logic cells can provide greater control over speed-critical paths.

Instead of using the Quartus II software’s **Auto Logic Cell Insertion** option, you can manually insert logic cells. However, Altera recommends that you use the **Auto Logic Cell Insertion** option unless you know which part of the design is causing the congestion.

A good location to manually insert LCELL buffers is where a single complex logic expression feeds multiple destinations in your design. You can insert an LCELL buffer just after the complex expression; the Quartus II Fitter extracts this complex expression and places it in a separate logic cell. Rather than duplicate all the logic for each destination, the Quartus II software feeds the single output from the logic cell to all destinations.

To reduce fan-in and prevent no-fit compilations caused by routing resource issues, insert an LCELL buffer after a NOR gate (Figure 13-11). The design in Figure 13-11 was compiled for a MAX 7000AE device. Without the LCELL buffer, the design requires two macrocells and eight shareable expanders, and the average fan-in is 14.5 macrocells. However, with the LCELL buffer, the design requires three macrocells and eight shareable expanders, and the average fan-in is just 6.33 macrocells.

Figure 13-11. Reducing the Average Fan-In by Inserting LCELL Buffers



Timing Optimization Techniques (Macrocell-Based CPLDs)

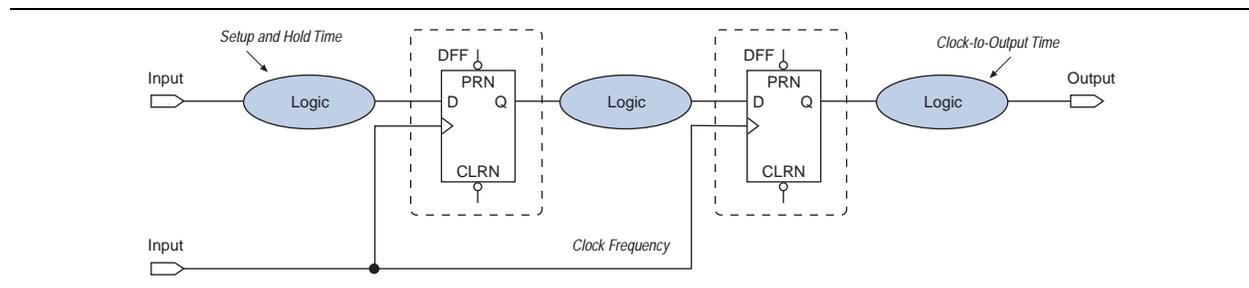
After resource optimization, design optimization focuses on timing. Ensure that you have made the appropriate assignments as described in “Initial Compilation: Required Settings” on page 13-2, and that the resource utilization is satisfactory before proceeding with timing optimization.

The following five timing parameters are primarily responsible for a design’s performance:

- Setup time (t_{SU})—the propagation time for input data signals
- Hold time (t_{H})—the propagation time for input data signals
- Clock-to-output time (t_{CO})—the propagation time for output signals
- Pin-to-pin delays (t_{PD})—the time required for a signal from an input pin to propagate through combinational logic and appear at an external output pin
- Maximum clock frequency (f_{MAX})—the internal register-to-register performance

This section provides guidelines to improve the timing if the timing requirements are not met. Figure 13-12 shows the parts of the design that determine the t_{SU} , t_H , t_{CO} , t_{PD} , and f_{MAX} timing parameters.

Figure 13-12. Main Timing Parameters that Determine the System's Performance



When you are analyzing a design to improve performance, be sure to consider the two major contributors to long delay paths:

- Excessive levels of logic
- Excessive loading (high fan-out)

When a MAX 7000 or MAX 3000 device signal drives more than one LAB, the programmable interconnect array (PIA) delay increases by 0.1 ns per additional LAB fan-out. Therefore, to minimize the added delay, concentrate the destination macrocells into fewer LABs, minimizing the number of LABs that are driven. The main cause of long delays in circuit design is excessive levels of logic.

Improving Setup Time

Sometimes the t_{SU} timing reported by the Quartus II Fitter does not meet your timing requirements. To improve the t_{SU} timing, refer to the following guidelines:

- Turn on the **Fast Input Register** option using the Assignment Editor. The **Fast Input Register** option allows input pins to directly drive macrocell registers via the fast-input path, thus minimizing the pin-to-register delay. This option is useful when a pin drives a D-type flipflop and there is no combinational logic between the pin and the register.
- Reduce the amount of logic between the input and the register. Excessive logic between the input pin and register causes more delays. To improve setup time, Altera recommends that you reduce the amount of logic between the input pin and the register whenever possible.
- Reduce fan-out. The delay from input pins to macrocell registers increases when the fan-out of the pins increases. To improve the setup time, minimize the fan-out.

Improving Clock-to-Output Time

To improve a design's clock-to-output time, minimize the register-to-output-pin delay. To improve the t_{CO} timing, refer to the following guidelines:

- Use the global clock. In addition to minimizing the delay from a register to an output pin, minimizing the delay from the clock pin to the register can also improve t_{CO} timing. Always use the global clock for low-skew and speed-critical signals.
- Reduce the amount of logic between the register and output pin. Excessive logic between the register and the output pin causes more delay. Always minimize the amount of logic between the register and output pin for faster clock-to-output time.

Table 13-3 shows the timing results for an EPM7064AETC100-4 device when a combination of the **Fast Input Register** option, global clock, and minimal logic is used. When the **Fast Input Register** option is turned on, the t_{SU} timing is improved (t_{SU} decreases from 1.6 ns to 1.3 ns and from 2.8 ns to 2.5 ns). The t_{CO} timing is improved when the global clock is used for low-skew and speed-critical signals (t_{CO} decreases from 4.3 ns to 3.1 ns). However, if there is additional logic used between the input pin and the register or the register and the output pin, the t_{SU} and t_{CO} delays increase.

Table 13-3. EPM7064AETC100-4 Device Timing Results

Number of Registers	t_{SU} (ns)	t_H (ns)	t_{CO} (ns)	Global Clock Used	Fast Input Register Option	D Input Location	Q Output Location	Additional Logic Between:	
								D Input Location & Register	Register & Q Output Location
1	1.3	1.2	4.3	—	On	LAB A	LAB A	—	—
1	1.6	0.3	4.3	—	Off	LAB A	LAB A	—	—
1	2.5	0	3.1	✓	On	LAB A	LAB A	—	—
1	2.8	0	3.1	✓	Off	LAB A	LAB A	—	—
1	3.6	0	3.1	✓	Off	LAB A	LAB A	✓	—
1	2.8	0	7.0	✓	Off	LAB D	LAB A	—	✓
16 with the same D and clock inputs	2.8	0	All 6.2	✓	Off	LAB D	LAB A, B	—	—
32 with the same D and clock inputs	2.8	0	All 6.4	✓	Off	LAB C	LAB A, B, C	—	—

Improving Propagation Delay (t_{PD})

Achieving fast propagation delay (t_{PD}) timing is required in many system designs. However, if there are long delay paths through complex logic, achieving fast propagation delays can be difficult. To improve your design's t_{PD} , refer to the following guidelines:

- On the Assignments menu, click **Settings**. In the **Category** list, select **Analysis & Synthesis Settings**, and turn on **Auto Parallel Expanders**. Turning on the parallel expanders for individual nodes or sub-designs can increase the performance of complex logic functions. However, if the project's pin or logic cell assignments use parallel expanders placed physically together with macrocells (which can reduce routability), parallel expanders can cause the Quartus II Fitter to have difficulties finding and optimizing a fit. Additionally, the number of macrocells required to implement the design increases and results in a no-fit error during compilation if the device resources are limited. For more information about turning on the **Auto Parallel Expanders** option, refer to ["Resolving Macrocell Usage Issues" on page 13-47](#).
- Set the **Optimization Technique** to **Speed**. By default, the Quartus II software sets the **Optimization Technique** option to **Speed** for MAX 7000 and MAX 3000 devices. Reset the **Optimization Technique** option to **Speed** only if you previously set it to **Area**. On the Assignments menu, click **Settings**. In the **Category** list, select **Analysis & Synthesis Settings**, and turn on **Speed** under **Optimization Technique**.

Improving Maximum Frequency (f_{MAX})

Maintaining the system clock at or above a certain frequency is a major goal in circuit design. For example, if you have a fully synchronous system that must run at 100 MHz, the longest delay path from the output of any register to the inputs of the registers it feeds must be less than 10 ns. Maintaining the system clock speed can be difficult if there are long delay paths through complex logic. Altera recommends that you perform the following guidelines to increase your design's clock speed (f_{MAX}):

- On the Assignments menu, click **Settings**. In the **Category** list, select **Analysis & Synthesis Settings**, click **More Settings**, and turn on **Auto Parallel Expanders**. Turning on the parallel expanders for individual nodes or subdesigns can increase the performance of complex logic functions. However, if the project's pin or logic cell assignments use parallel expanders placed physically together with macrocells (which can reduce routability), parallel expanders can cause the Quartus II compiler to have difficulties finding and optimizing a fit. Additionally, the number of macrocells required to implement the design also increases and can result in a no-fit error during compilation if the device's resources are limited. For more information about using the **Auto Parallel Expanders** option, refer to ["Resolving Macrocell Usage Issues" on page 13-47](#).
- Use global signals or dedicated inputs. Altera MAX 7000 and MAX 3000 devices have dedicated inputs that provide low skew and high speed for high fan-out signals. Minimize the number of control signals in the design and use the dedicated inputs to implement them.
- Set the **Optimization Technique** to **Speed**. By default, the Quartus II software sets the **Optimization Technique** option to **Speed** for MAX 7000 and MAX 3000 devices. Reset the **Optimization Technique** option to **Speed** only if you have previously set it to **Area**. You can reset the **Optimization Technique** option. In the **Category** list, select **Analysis & Synthesis Settings**, and turn on **Speed** under **Optimization Technique**.
- Pipeline the design. Pipelining, which increases clock frequency (f_{MAX}), refers to dividing large blocks of combinational logic by inserting registers.

Optimizing Source Code—Pipelining for Complex Register Logic

If the methods described in the preceding sections do not sufficiently improve your results, modify the design at the source to achieve the desired results. Using additional register stages (pipeline registers) consumes more device resources, but it also lowers the propagation delay between registers, allowing you to maintain high system clock speed.

-  Refer to the application note *AN 584: Timing Closure Methodology for Advanced FPGA Designs* for more information about pipelining registers and other examples of optimizing source code.

Other Optimization Resources

The Quartus II software has additional resources to help you optimize your design for resource, performance, compilation time, and power.

Design Space Explorer

The DSE automates the process of running multiple compilations with different settings. You can use the DSE to try the techniques described in this chapter. The DSE utility helps automate the process of finding the best set of options for your design. The DSE explores the design space by applying various optimization techniques and analyzing the results.

-  For more information, refer to *About Design Space Explorer* in Quartus II Help.

Other Optimization Advisors

The Power Optimization Advisor provides guidance for reducing power consumption. In addition, the Incremental Compilation Advisor provides suggestions to improve your results when partitioning your design for a hierarchical or team-based design flow using the Quartus II incremental compilation feature.

-  For more information about using the Power Optimization Advisor, refer to the *Power Optimization* chapter in volume 2 of the *Quartus II Handbook*. For more information about using the Incremental Compilation Advisor, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.

Scripting Support

You can run procedures and make settings described in this chapter in a Tcl script. You can also run some procedures at a command prompt. For detailed information about scripting command options, refer to the Quartus II command-line and Tcl API Help browser. To run the Help browser, type the following command at the command prompt:

```
quartus_sh --qhelp ←
```

 For more information about Tcl scripting, refer to the *Tcl Scripting* chapter in volume 2 of the *Quartus II Handbook*. Refer to the *Quartus II Settings File Manual* for information about all settings and constraints in the Quartus II software. For more information about command-line scripting, refer to the *Command-Line Scripting* chapter in volume 2 of the *Quartus II Handbook*.

You can specify many of the options described in this section either in an instance, or at a global level, or both.

Use the following Tcl command to make a global assignment:

```
set_global_assignment -name <.qsf variable name> <value> ←
```

Use the following Tcl command to make an instance assignment:

```
set_instance_assignment -name <.qsf variable name> <value> \  
-to <instance name> ←
```

 If the <value> field includes spaces (for example, “Standard Fit”), the value must be enclosed by straight double quotation marks.

Initial Compilation Settings

The Quartus II Settings File (.qsf) variable name is used in the Tcl assignment to make the setting along with the appropriate value. The **Type** column indicates whether the setting is supported as a global setting, an instance setting, or both.

Table 13-4 lists the .qsf file variable name and applicable values for the settings discussed in “Initial Compilation: Required Settings” on page 13-2. Table 13-5 shows the list of advanced compilation settings.

Table 13-4. Initial Compilation Settings

Setting Name	.qsf File Variable Name	Values	Type
Device Setting	DEVICE	<device part number>	Global
Use Smart Compilation	SPEED_DISK_USAGE_TRADEOFF	SMART, NORMAL	Global
Optimize IOC Register Placement For Timing	OPTIMIZE_IOC_REGISTER_PLACEMENT_FOR_TIMING	ON, OFF	Global
Optimize Hold Timing	OPTIMIZE_HOLD_TIMING	OFF, IO PATHS AND MINIMUM TPD PATHS, ALL PATHS	Global
Fitter Effort	FITTER_EFFORT	STANDARD FIT, FAST FIT, AUTO FIT	Global

Table 13-5. Advanced Compilation Settings

Setting Name	.qsf File Variable Name	Values	Type
Router Effort Multiplier	ROUTER_EFFORT_MULTIPLIER	Any positive, non-zero value	Global
Router Timing Optimization level	ROUTER_TIMING_OPTIMIZATION_LEVEL	NORMAL, MINIMUM, MAXIMUM	Global
Final Placement Optimization	FINAL_PLACEMENT_OPTIMIZATION	ALWAYS, AUTOMATICALLY, NEVER	Global

Resource Utilization Optimization Techniques (LUT-Based Devices)

Table 13-6 lists the .qsf file variable name and applicable values for the settings discussed in “Resource Utilization Optimization Techniques (LUT-Based Devices)” on page 13-15.

Table 13-6. Resource Utilization Optimization Settings

Setting Name	.qsf File Variable Name	Values	Type
Auto Packed Registers (1)	AUTO_PACKED_REGISTERS_<device family name>	OFF, NORMAL, MINIMIZE AREA, MINIMIZE AREA WITH CHAINS, AUTO	Global, Instance
Perform WYSIWYG Primitive Resynthesis	ADV_NETLIST_OPT_SYNTH_WYSIWYG_REMAP	ON, OFF	Global, Instance
Physical Synthesis for Combinational Logic for Reducing Area	PHYSICAL_SYNTHESIS_COMBO_LOGIC_FOR_AREA	ON, OFF	Global, Instance
Physical Synthesis for Mapping Logic to Memory	PHYSICAL_SYNTHESIS_MAP_LOGIC_TO_MEMORY_FOR AREA	ON, OFF	Global, Instance
Optimization Technique	<device family name>_OPTIMIZATION_TECHNIQUE	AREA, SPEED, BALANCED	Global, Instance
Speed Optimization Technique for Clock Domains	SYNTH_CRITICAL_CLOCK	ON, OFF	Instance
State Machine Encoding	STATE_MACHINE_PROCESSING	AUTO, ONE-HOT, MINIMAL BITS, USER-ENCODE	Global, Instance
Auto RAM Replacement	AUTO_RAM_RECOGNITION	ON, OFF	Global, Instance
Auto ROM Replacement	AUTO_ROM_RECOGNITION	ON, OFF	Global, Instance
Auto Shift Register Replacement	AUTO_SHIFT_REGISTER_RECOGNITION	ON, OFF	Global, Instance
Auto Block Replacement	AUTO_DSP_RECOGNITION	ON, OFF	Global, Instance
Number of Processors for Parallel Compilation	NUM_PARALLEL_PROCESSORS	Integer between 1 and 4 inclusive, or ALL	Global

Note to Table 13-6:

(1) Allowed values for this setting depend on the device family that is selected.

I/O Timing Optimization Techniques (LUT-Based Devices)

Table 13-7 lists the .qsf file variable name and applicable values for the I/O timing optimization settings.

Table 13-7. I/O Timing Optimization Settings

Setting Name	.qsf File Variable Name	Values	Type
Optimize IOC Register Placement For Timing	OPTIMIZE_IOC_REGISTER_PLACEMENT_FOR_TIMING	ON, OFF	Global
Fast Input Register	FAST_INPUT_REGISTER	ON, OFF	Instance
Fast Output Register	FAST_OUTPUT_REGISTER	ON, OFF	Instance
Fast Output Enable Register	FAST_OUTPUT_ENABLE_REGISTER	ON, OFF	Instance
Fast OCT Register	FAST_OCT_REGISTER	ON, OFF	Instance

Register-to-Register Timing Optimization Techniques (LUT-Based Devices)

Table 13-8 lists the .qsf file variable name and applicable values for the settings discussed in “Register-to-Register Timing Optimization Techniques (LUT-Based Devices)” on page 13-34.

Table 13-8. Register-to-Register Timing Optimization Settings

Setting Name	.qsf File Variable Name	Values	Type
Perform WYSIWYG Primitive Resynthesis	ADV_NETLIST_OPT_SYNTH_WYSIWYG_REMAP	ON, OFF	Global, Instance
Perform Physical Synthesis for Combinational Logic	PHYSICAL_SYNTHESIS_COMBO_LOGIC	ON, OFF	Global, Instance
Perform Register Duplication	PHYSICAL_SYNTHESIS_REGISTER_DUPLICATION	ON, OFF	Global, Instance
Perform Register Retiming	PHYSICAL_SYNTHESIS_REGISTER_RETIMING	ON, OFF	Global, Instance
Perform Automatic Asynchronous Signal Pipelining	PHYSICAL_SYNTHESIS_ASYNCHRONOUS_SIGNAL_PIPELINING	ON, OFF	Global, Instance
Physical Synthesis Effort	PHYSICAL_SYNTHESIS_EFFORT	NORMAL, EXTRA, FAST	Global
Fitter Seed	SEED	<integer>	Global
Maximum Fan-Out	MAX_FANOUT	<integer>	Instance
Manual Logic Duplication	DUPLICATE_ATOM	<node name>	Instance
Optimize Power during Synthesis	OPTIMIZE_POWER_DURING_SYNTHESIS	NORMAL, OFF, EXTRA_EFFORT	Global
Optimize Power during Fitting	OPTIMIZE_POWER_DURING_FITTING	NORMAL, OFF, EXTRA_EFFORT	Global

Duplicate Logic for Fan-Out Control

The manual logic duplication option accepts wildcards. This is an easy and powerful duplication technique that you can use without editing your source code. You can use this technique, for example, to make a duplicate of a large fan-out node for all of its destinations in a certain design hierarchy, such as `hierarchy_A`. To make such an assignment with Tcl, use a command similar to [Example 13-1](#).

Example 13-1. Duplication Technique

```
set_instance_assignment -name DUPLICATE_ATOM high_fanout_to_A -from \
high_fanout_node -to *hierarchy_A*
```

Conclusion

Using the recommended techniques described in this chapter can help you close timing quickly on complex designs, reduce iterations by providing more intelligent and better links between analysis and assignment tools, and balance multiple design constraints including multiple clocks, routing resources, and area constraints.

The Quartus II software provides many features to achieve optimal results. Follow the techniques presented in this chapter to efficiently optimize a design for area or timing performance, or to reduce compilation time.

Document Revision History

[Table 13-9](#) shows the revision history for this chapter.

Table 13-9. Document Revision History (Part 1 of 3)

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Added links to Help ■ Updated device support ■ Added “Debugging Timing Failures in the TimeQuest Analyzer” section ■ Removed Classic Timing Analyzer references ■ Other updates throughout chapter
August 2010	10.0.1	Corrected link
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Moved Compilation Time Optimization Techniques section to new <i>Reducing Compilation Time</i> chapter ■ Removed references to Timing Closure Floorplan ■ Moved Smart Compilation Setting and Early Timing Estimation sections to new <i>Reducing Compilation Time</i> chapter ■ Added Other Optimization Resources section ■ Removed outdated information ■ Changed references to DSE chapter to Help links ■ Linked to Help where appropriate ■ Removed Referenced Documents section

Table 13-9. Document Revision History (Part 2 of 3)

Date	Version	Changes
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Removed unsupported Timing Closure Floorplan references ■ Removed references to unsupported device families ■ Added several notes ■ Minor text edits
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Was chapter 8 in the 8.1.0 release. ■ Updated the following sections: <ul style="list-style-type: none"> ■ “Timing Analysis with the TimeQuest Timing Analyzer” on page 10-14 ■ “Perform WYSIWYG Resynthesis with Balanced or Area Setting” on page 10-22 ■ “Use Physical Synthesis Options to Reduce Area” on page 10-26 ■ “Metastability Analysis and Optimization Techniques” on page 10-32 ■ “Use Fast Regional Clock Networks and Regional Clocks Networks” on page 10-39 ■ “Register-to-Register Timing Optimization Techniques (LUT-Based Devices)” on page 10-40 ■ “Physical Synthesis Optimizations” on page 10-41 ■ “Duplicate Logic for Fan-Out Control” on page 10-45 ■ “LogicLock Assignments” on page 10-49 ■ “Enable Beneficial Skew Optimization” on page 10-48 ■ “Use Multiple Processors for Parallel Compilation” on page 10-65 ■ Removed “Analyze Your Design for Megastability” ■ Updated Table 10-11 and Table 10-9 ■ Removed Tables 8-1, 8-2, 8-3, 8-6, and 8-7 from version 8.1

Table 13–9. Document Revision History (Part 3 of 3)

Date	Version	Changes
November 2008	8.1.0	<ul style="list-style-type: none"> ■ Changed document to 8½" × 11" page size. ■ Updated the following sections: <ul style="list-style-type: none"> ■ “Optimizing Your Design” on page 10–2 ■ “Timing Requirement Settings” on page 10–4 ■ “Optimize Hold Timing” on page 10–8 ■ “Limit to One Fitting Attempt” on page 10–9 ■ “Auto Fit” on page 10–10 ■ “Fast Fit” on page 10–11 ■ “Ignored Timing Assignments” on page 10–12 ■ “I/O Timing (Including tPD)” on page 10–13 ■ “Register-to-Register Timing” on page 10–14 ■ “Timing Analysis with the TimeQuest Timing Analyzer” on page 10–14 ■ “Use I/O Assignment Analysis” on page 10–20 ■ “Flatten the Hierarchy During Synthesis” on page 10–25 ■ “Retarget Memory Blocks” on page 10–25 ■ “Use Physical Synthesis Options to Reduce Area” on page 10–26 ■ “Increase Placement Effort Multiplier” on page 10–30 ■ “Metastability Analysis and Optimization Techniques” on page 10–32 ■ “Synthesis Netlist Optimizations and Physical Synthesis Optimizations” on page 10–43 ■ “Incremental Compilation” on page 10–65 ■ “Use Multiple Processors for Parallel Compilation” on page 10–66 ■ Updated Table 10–9 on page 10–73 and Table 10–11 on page 10–75.
May 2008	8.0.0	<ul style="list-style-type: none"> ■ Updated links ■ Updated the following sections: <ul style="list-style-type: none"> ■ Other Optimization Resources] ■ Setting Process Priority ■ Location Assignment and Back-Annotation ■ Fitter Effort Setting ■ Synthesis Netlist Optimizations and Physical Synthesis Optimizations ■ Fast Fit ■ Added Metastability Analysis ■ Added Enable Beneficial Skew Optimization and Analyze Your Design for Metastability ■ Removed figures from “Optimizing Source Code—Pipelining for Complex Register Logic ■ Updated Table 8-5

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

The Quartus® II software offers power-driven compilation to fully optimize device power consumption. Power-driven compilation focuses on reducing your design's total power consumption using power-driven synthesis and power-driven place-and-route. This chapter describes the power-driven compilation feature and flow in detail, as well as low power design techniques that can further reduce power consumption in your design. The techniques primarily target Arria® GX, Stratix® and Cyclone® series of devices, and HardCopy® II devices. These devices utilize a low-k dielectric material that dramatically reduces dynamic power and improves performance. Arria series, Stratix II, Stratix III, Stratix IV, and Stratix V device families include efficient logic structures called adaptive logic modules (ALMs) that obtain maximum performance while minimizing power consumption. Cyclone device families offer the optimal blend of high performance and low power in a low-cost FPGA.

 For more information about a device-specific architecture, refer to the device handbook, available from the [Literature and Technical Documentation](#) page on the [Altera website](#).

Altera provides the Quartus II PowerPlay Power Analyzer to aid you during the design process by delivering fast and accurate estimations of power consumption. You can minimize power consumption, while taking advantage of the industry's leading FPGA performance, by using the tools and techniques described in this chapter.

 For more information about the PowerPlay Power Analyzer, refer to the [PowerPlay Power Analysis](#) chapter in volume 3 of the *Quartus II Handbook*.

Total FPGA power consumption is comprised of I/O power, core static power, and core dynamic power. This chapter focuses on design optimization options and techniques that help reduce core dynamic power and I/O power. In addition to these techniques, there are additional power optimization techniques available for Stratix III and Stratix IV devices. These techniques include:

- Selectable Core Voltage (available only for Stratix III devices)
- Programmable Power Technology
- Device Speed Grade Selection

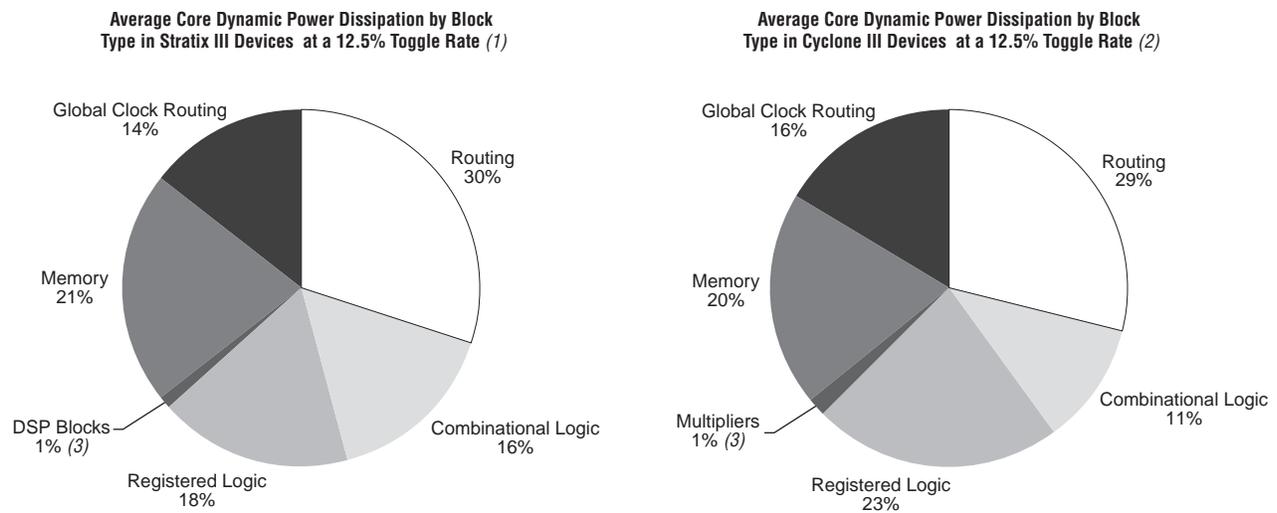
 For more information about power optimization techniques available for Stratix III devices, refer to [AN 437: Power Optimization in Stratix III FPGAs](#). For more information about power optimization techniques available for Stratix IV devices, refer to [AN 514: Power Optimization in Stratix IV FPGAs](#).

Power Dissipation

This section describes the sources of power dissipation in Stratix III and Cyclone III devices. You can refine techniques that reduce power consumption in your design by understanding the sources of power dissipation.

Figure 14-1 shows the power dissipation of Stratix III and Cyclone III devices in different designs. All designs were analyzed at a fixed clock rate of 100 MHz and exhibited varied logic resource utilization across available resources.

Figure 14-1. Average Core Dynamic Power Dissipation



Notes to Figure 14-1:

- (1) 103 different designs were used to obtain these results.
- (2) 96 different designs were used to obtain these results.
- (3) In designs using DSP blocks, DSPs consumed 5% of core dynamic power.

As shown in Figure 14-1, a significant amount of the total power is dissipated in routing for both Stratix III and Cyclone III devices, with the remaining power dissipated in logic, clock, and RAM blocks.

In Stratix and Cyclone device families, a series of column and row interconnect wires of varying lengths provide signal interconnections between logic array blocks (LABs), memory block structures, and digital signal processing (DSP) blocks or multiplier blocks. These interconnects dissipate the largest component of device power.

FPGA combinational logic is another source of power consumption. The basic building block of logic in the latest Stratix series devices is the ALM, and in Cyclone II, Cyclone III and Cyclone IV GX devices, it is the logic element (LE).



For more information about ALMs and LEs in Cyclone II, Cyclone III, Cyclone IV GX, Stratix II, Stratix III, Stratix IV, and Stratix V, devices, refer to the respective device handbook.

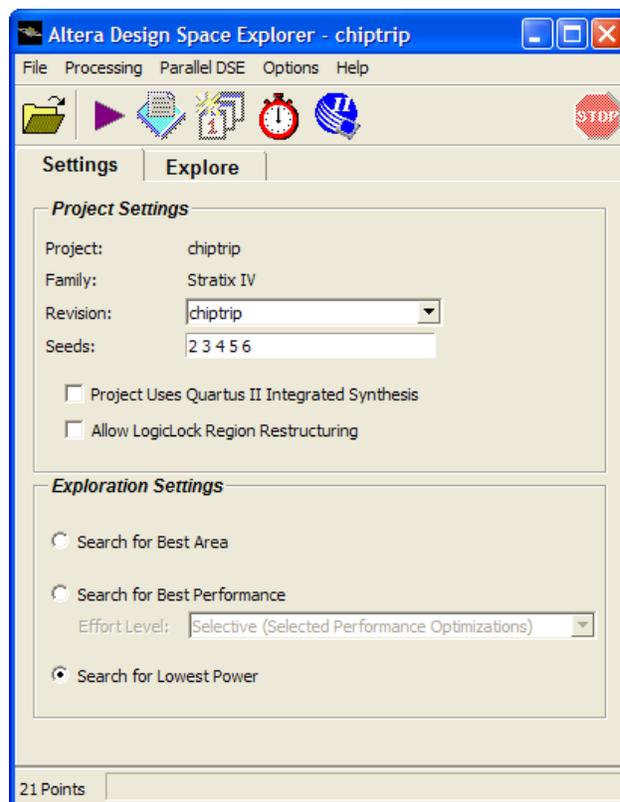
Memory and clock resources are other major consumers of power in FPGAs. Stratix II devices feature the TriMatrix memory architecture. TriMatrix memory includes 512-bit M512 blocks, 4-Kbit M4K blocks, and 512-Kbit M-RAM blocks, which are configurable to support many features. Stratix IV and Stratix III TriMatrix on-chip memory is an enhancement based upon the Stratix II FPGA TriMatrix memory and includes three sizes of memory blocks: MLAB blocks, M9K blocks, and M144K blocks. Stratix III, Stratix IV, and Stratix V devices feature Programmable Power Technology, an advanced architecture that enables a smooth trade-off between speed and power. The core of each Stratix III, Stratix IV, and Stratix V device is divided into tiles, each of which may be put into a high-speed or low-power mode. The primary benefit of Programmable Power Technology is to reduce static power, with a secondary benefit being a small reduction in dynamic power. Cyclone II devices have 4-Kbit M4K memory blocks, and Cyclone III and Cyclone IV GX devices have 9-Kbit M9K memory blocks.

Design Space Explorer

Design Space Explorer (DSE) is a simple, easy-to-use, design optimization utility that is included in the Quartus II software. DSE explores and reports optimal Quartus II software options for your design, targeting either power optimization, design performance, or area utilization improvements. You can use DSE to implement the techniques described in this chapter.

Figure 14-2 shows the DSE user interface. The **Settings** tab is divided into **Project Settings** and **Exploration Settings**.

Figure 14-2. Design Space Explorer User Interface



The **Search for Lowest Power** option, under **Exploration Settings**, uses a predefined exploration space that targets overall design power improvements. This setting focuses on applying different options that specifically reduce total design thermal power.

By default, the Quartus II PowerPlay Power Analyzer is run for every exploration performed by the DSE when the **Search for Lowest Power** option is selected. This helps you debug your design and determine trade-offs between power requirements and performance optimization.

- ④ For more information about the DSE, refer to *About Design Space Explorer* in Quartus II Help.

Power-Driven Compilation

The standard Quartus II compilation flow consists of Analysis and Synthesis, placement and routing, Assembly, and Timing Analysis. Power-driven compilation takes place at the Analysis and Synthesis and Place-and-Route stages.

Quartus II software settings that control power-driven compilation are located in the **PowerPlay power optimization** list on the **Analysis & Synthesis Settings** page, and the **PowerPlay power optimization** list on the **Fitter Settings** page. The following sections describes these power optimization options at the Analysis and Synthesis and Fitter levels.

Power-Driven Synthesis

Synthesis netlist optimization occurs during the synthesis stage of the compilation flow. The optimization technique makes changes to the synthesis netlist to optimize your design according to the selection of area, speed, or power optimization. This section describes power optimization techniques at the synthesis level.

The **Analysis & Synthesis Settings** page allows you to specify logic synthesis options. The **PowerPlay power optimization** option is available for all devices supported by the Quartus II software except MAX[®] 3000 and MAX 7000 devices. (Figure 14-3).

Figure 14-3. Analysis & Synthesis Settings Page

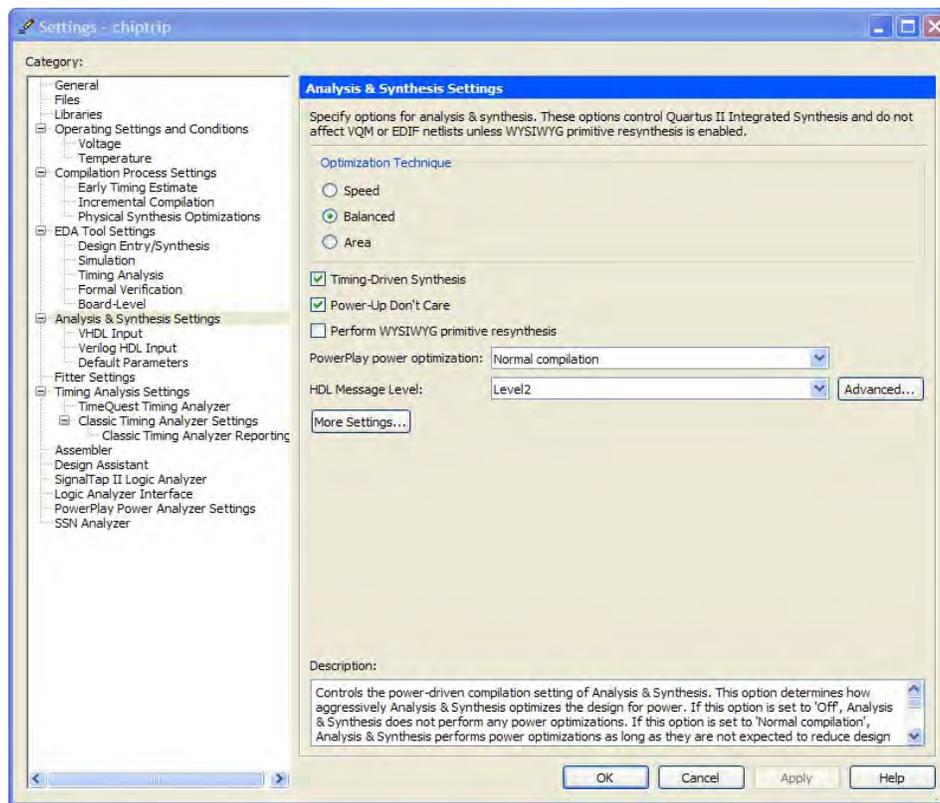


Table 14-1 shows the settings in the **PowerPlay power optimization** list. You can apply these settings on a project or entity level.

Table 14-1. Optimize Power During Synthesis Options

Settings	Description
Off	No netlist, placement, or routing optimizations are performed to minimize power.
Normal compilation (Default)	Low compute effort algorithms are applied to minimize power through netlist optimizations as long as they are not expected to reduce design performance.
Extra effort	High compute effort algorithms are applied to minimize power through netlist optimizations. Max performance might be impacted.

The **Normal compilation** setting is turned on by default. This setting performs memory optimization and power-aware logic mapping during synthesis.

Memory blocks can represent a large fraction of total design dynamic power as described in “Reducing Memory Power Consumption” on page 14-14. Minimizing the number of memory blocks accessed during each clock cycle can significantly reduce memory power. Memory optimization involves effective movement of user-defined read/write enable signals to associated read-and-write clock enable signals for all memory types (Figure 14-4).

Figure 14-4. Memory Transformation

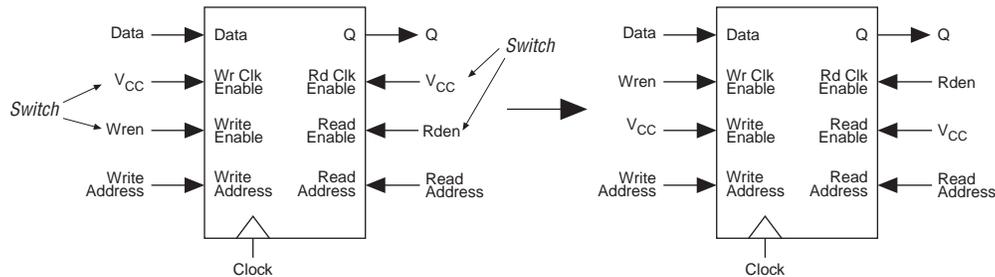


Figure 14-4 shows a default implementation of a simple dual-port memory block in which write-clock enable signals and read-clock enable signals are connected to V_{CC}, making both read and write memory ports active during each clock cycle. Memory transformation effectively moves the read-enable and write-enable signals to the respective read-clock enable and write-clock enable signals. By using this technique, memory ports are shut down when they are not accessed. This significantly reduces your design’s memory power consumption. For more information about clock enable signals, refer to “Reducing Memory Power Consumption” on page 14-14. For Stratix III, Stratix IV, and Stratix V devices, the memory transformation takes place at the Fitter level by selecting the **Normal compilation** settings for the power optimization option.

In Stratix III, Cyclone III, Cyclone IV GX, and Stratix III devices, the specified read-during-write behavior can significantly impact the power of single-port and bidirectional dual-port RAMs. It is best to set the read-during-write parameter to “Don’t care” (at the HDL level), as it allows an optimization whereby the read-enable signal can be set to the inversion of the existing write-enable signal (if one exists). This allows the core of the RAM to shut down (that is, not toggle), which saves a significant amount of power.

The other type of power optimization that takes place with the **Normal compilation** setting is power-aware logic mapping. The power-aware logic mapping reduces power by rearranging the logic during synthesis to eliminate nets with high toggle rates.

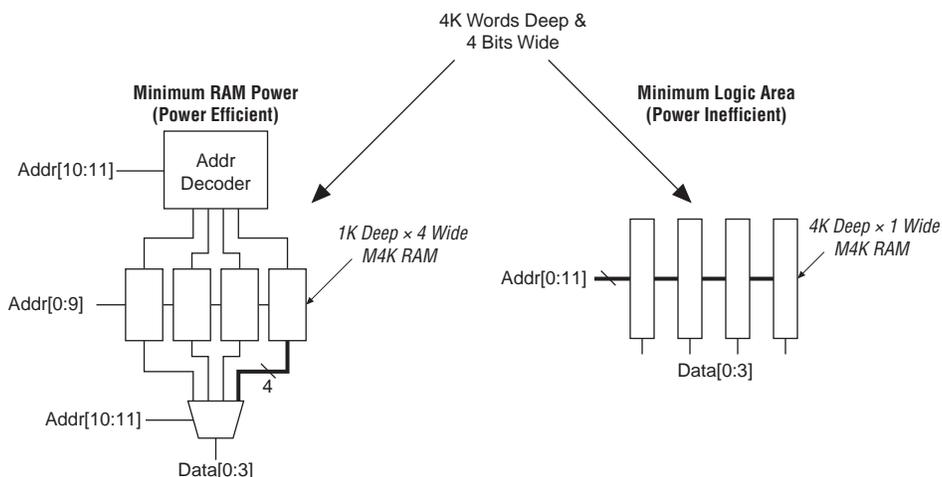
The **Extra effort** setting performs the functions of the **Normal compilation** setting and other memory optimizations to further reduce memory power by shutting down memory blocks that are not accessed. This level of memory optimization can require extra logic, which can reduce design performance.

The **Extra effort** setting also performs power-aware memory balancing. Power-aware memory balancing automatically chooses the best memory configuration for your memory implementation and provides optimal power saving by determining the number of memory blocks, decoder, and multiplexer circuits required. If you have not previously specified target-embedded memory blocks for your design’s memory functions, the power-aware balancer automatically selects them during memory implementation.

Figure 14-5 shows an example of a 4k × 4 (4k deep and 4 bits wide) memory implementation in two different configurations using M4K memory blocks available in Stratix II devices. The minimum logic area implementation uses M4K blocks configured as 4k × 1. This implementation is the default in the Quartus II software because it has the minimum logic area (0 logic cells) and the highest speed. However, all four M4K blocks are active on each memory access in this implementation, which increases RAM power. The minimum RAM power implementation is created by selecting **Extra effort** in the **PowerPlay power optimization** list. This implementation automatically uses four M4K blocks configured as 1k × 4 for optimal power saving. An address decoder is implemented by the RAM megafunction to select which of the four M4K blocks should be activated on a given cycle, based on the state of the top two user address bits. The RAM megafunction automatically implements a multiplexer to feed the downstream logic by choosing the appropriate M4K output. This implementation reduces RAM power because only one M4K block is active on any cycle, but it requires extra logic cells, costing logic area and potentially impacting design performance.

There is a trade-off between power saved by accessing fewer memories and power consumed by the extra decoder and multiplexor logic. The Quartus II software automatically balances the power savings against the costs to choose the lowest power configuration for each logical RAM. The benchmark data shows that the power-driven synthesis can reduce memory power consumption by as much as 60% in Stratix devices.

Figure 14-5. 4K × 4 Memory Implementation Using Multiple M4K Blocks



Memory optimization options can also be controlled by the `Low_Power_Mode` parameter in the **Default Parameters** page of the **Settings** dialog box. The settings for this parameter are **None**, **Auto**, and **ALL**. **None** corresponds to the **Off** setting in the **PowerPlay power optimization** list. **Auto** corresponds to the **Normal compilation** setting and **ALL** corresponds to the **Extra effort** setting, respectively. You can apply PowerPlay power optimization either on a compiler basis or on individual entities. The `Low_Power_Mode` parameter always takes precedence over the **Optimize Power for Synthesis** option for power optimization on memory.

You can also set the `MAXIMUM_DEPTH` parameter manually to configure the memory for low power optimization. This technique is the same as the power-aware memory balancer, but it is manual rather than automatic like the **Extra effort** setting in the **PowerPlay power optimization** list. You can set the `MAXIMUM_DEPTH` parameter for memory modules manually in the megafunction instantiation or in the MegaWizard™ Plug-In Manager for power optimization as described in *“Reducing Memory Power Consumption”* on page 14-14. The `MAXIMUM_DEPTH` parameter always takes precedence over the **Optimize Power for Synthesis** options for power optimization on memory optimization.

- ② For step-by-step instructions on how to perform power-driven synthesis, refer to *Running a Power-Optimized Compilation* in Quartus II Help.

Power-Driven Fitter

The **Fitter Settings** page enables you to specify options for fitting (Figure 14-6). The **PowerPlay power optimization** option is available for Arria GX, Arria II GX, Cyclone II, Cyclone III, Cyclone IV, HardCopy series, Stratix II, Stratix II GX, Stratix III, Stratix IV, and Stratix V devices.

Figure 14-6. Fitter Settings Page

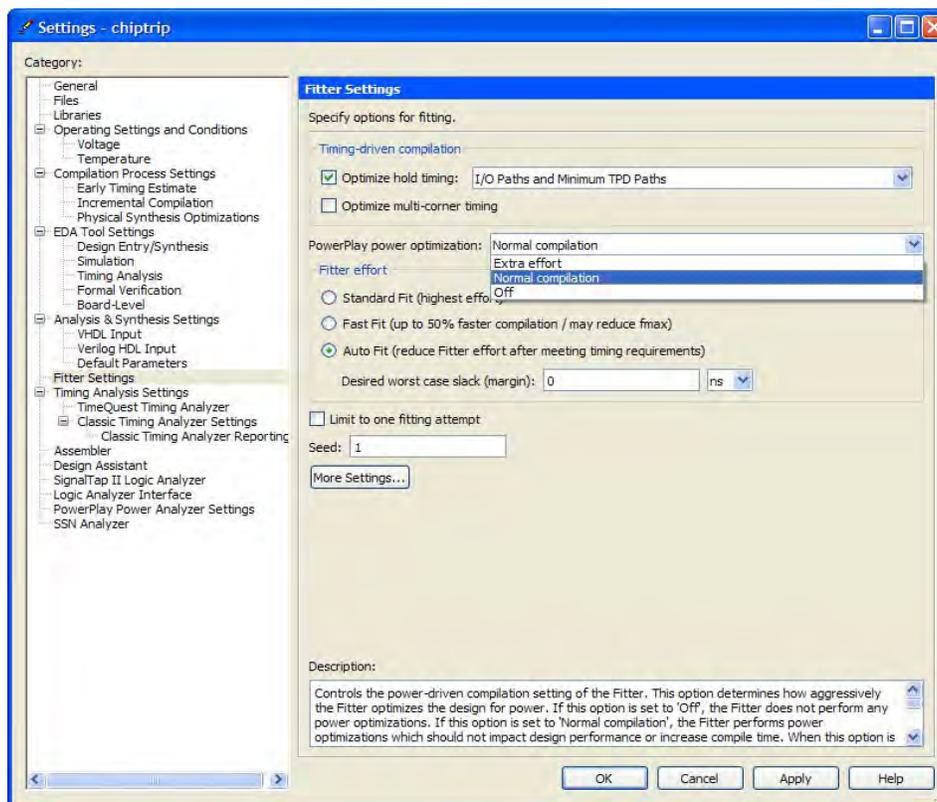


Table 14-2 lists the settings in the **PowerPlay power optimization** list. These settings can only be applied on a project-wide basis. The **Extra effort** setting for the Fitter requires extensive effort to optimize the design for power and can increase the compilation time.

Table 14-2. Power-Driven Fitter Option

Settings	Description
Off	No netlist, placement, or routing optimizations are performed to minimize power.
Normal compilation (Default)	Low compute effort algorithms are applied to minimize power through placement and routing optimizations as long as they are not expected to reduce design performance.
Extra effort	High compute effort algorithms are applied to minimize power through placement and routing optimizations. Max performance might be impacted.

The **Normal compilation** setting is selected by default and performs DSP optimization by creating power-efficient DSP block configurations for your DSP functions. For Stratix III, Stratix IV, and Stratix V devices, this setting, which is based on timing constraints entered for the design, enables the Programmable Power Technology to configure tiles as high-speed mode or low-power mode. Programmable Power Technology is always turned **ON** even when the **OFF** setting is selected for the **Fitter PowerPlay power optimization** option. Tiles are the combination of LAB and MLAB pairs (including the adjacent routing associated with LAB and MLAB), which can be configured to operate in high-speed or low-power mode. This level of power optimization does not have any affect on the fitting, timing results, or compile time. Also, for Stratix III devices, this setting enables the memory transformation as described in “Power-Driven Synthesis” on page 14-4.

 For more information about Stratix III power optimization, refer to [AN 437: Power Optimization in Stratix III FPGAs](#). For more information about Stratix IV power optimization, refer to [AN 514: Power Optimization in Stratix IV FPGAs](#).

The **Extra effort** setting performs the functions of the **Normal compilation** setting and other place-and-route optimizations during fitting to fully optimize the design for power. The Fitter applies an extra effort to minimize power even after timing requirements have been met by effectively moving the logic closer during placement to localize high-toggling nets, and using routes with low capacitance. However, this effort can increase the compilation time.

The **Extra effort** setting uses a Value Change Dump File (.vcd) that guides the Fitter to fully optimize the design for power, based on the signal activity of the design. The best power optimization during fitting results from using the most accurate signal activity information. Signal activities from full post-fit netlist (timing) simulation provide the highest accuracy because all node activities reflect the actual design behavior, provided that supplied input vectors are representative of typical design operation. If you do not have a .vcd file, the Quartus II software uses assignments, clock assignments, and vectorless estimation values (PowerPlay Power Analyzer Tool settings) to estimate the signal activities. This information is used to optimize your design for power during fitting. The benchmark data shows that the power-driven Fitter technique can reduce power consumption by as much as 19% in Stratix devices. On average, you can reduce core dynamic power by 16% with the Extra effort synthesis and Extra effort fitting settings, as compared to the **Off** settings in both synthesis and Fitter options for power-driven compilation.

 Only the **Extra effort** setting in the **PowerPlay power optimization** list for the Fitter option uses the signal activities (from .vcd files) during fitting. The settings made in the **PowerPlay Power Analyzer Settings** page in the **Settings** dialog box are used to calculate the signal activity of your design.

 For more information about .vcd files and how to create them, refer to the [PowerPlay Power Analysis](#) chapter in volume 3 of the *Quartus II Handbook*.

 For step-by-step instructions on how to perform power-driven fitting, refer to [Running a Power-Optimized Compilation](#) in Quartus II Help.

Area-Driven Synthesis

Using area optimization rather than timing or delay optimization during synthesis saves power because you use fewer logic blocks. Using less logic usually means less switching activity. The Quartus II integrated synthesis tool provides **Speed**, **Balanced**, or **Area** for the **Optimization Technique** option. You can also specify this logic option for specific modules in your design with the Assignment Editor in cases where you want to reduce area using the **Area** setting (potentially at the expense of register-to-register timing performance) while leaving the default **Optimization Technique** setting at **Balanced** (for the best trade-off between area and speed for certain device families). The **Speed Optimization Technique** can increase the resource usage of your design if the constraints are too aggressive, and can also result in increased power consumption.

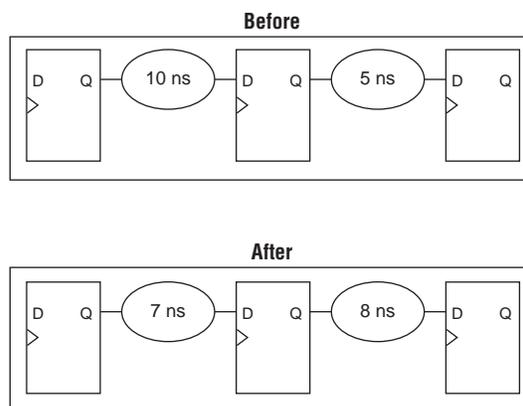
The benchmark data shows that the area-driven technique can reduce power consumption by as much as 31% in Stratix devices and as much as 15% in Cyclone devices.

Gate-Level Register Retiming

You can also use gate-level register retiming to reduce circuit switching activity. Retiming shuffles registers across combinational blocks without changing design functionality. The **Perform gate-level register retiming** option in the Quartus II software enables the movement of registers across combinational logic to balance timing, allowing the software to trade off the delay between timing critical and noncritical timing paths.

Retiming uses fewer registers than pipelining. [Figure 14-7](#) shows an example of gate-level register retiming, where the 10 ns critical delay is reduced by moving the register relative to the combinational logic, resulting in the reduction of data depth and switching activity.

Figure 14-7. Gate-Level Register Retiming



-  Gate-level register retiming makes changes at the gate level. If you are using an atom netlist from a third-party synthesis tool, you must also select the **Perform WYSIWYG primitive resynthesis** option to undo the atom primitives to gates mapping (so that register retiming can be performed), and then to remap gates to Altera primitives. When using Quartus II integrated synthesis, retiming occurs during synthesis before the design is mapped to Altera primitives. The benchmark data shows that the combination of WYSIWYG remapping and gate-level register retiming techniques can reduce power consumption by as much as 6% in Stratix devices and as much as 21% in Cyclone devices.
-  For more information about register retiming, refer to the *Netlist Optimizations and Physical Synthesis* chapter in volume 2 of the *Quartus II Handbook*.

Design Guidelines

Several low-power design techniques can reduce power consumption when applied during FPGA design implementation. This section provides detailed design techniques for Cyclone II, Cyclone III, Cyclone IV GX, Stratix II, and Stratix III devices that affect overall design power. The results of these techniques might be different from design to design.

Clock Power Management

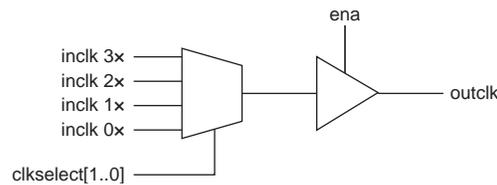
Clocks represent a significant portion of dynamic power consumption due to their high switching activity and long paths. [Figure 14-1 on page 14-2](#) shows a 14% average contribution to power consumption for global clock routing in Stratix III devices and 16% in Cyclone III devices. Actual clock-related power consumption is higher than this because the power consumed by local clock distribution within logic, memory, and DSP or multiplier blocks is included in the power consumption for the respective blocks.

Clock routing power is automatically optimized by the Quartus II software, which enables only those portions of the clock network that are required to feed downstream registers. Power can be further reduced by gating clocks when they are not required. It is possible to build clock-gating logic, but this approach is not recommended because it is difficult to generate a glitch free clock in FPGAs using ALMs or LEs.

Arria GX, Arria II GX, Cyclone III, Cyclone IV, Stratix II, Stratix III, Stratix IV, and Stratix V devices use clock control blocks that include an enable signal. A clock control block is a clock buffer that lets you dynamically enable or disable the clock network and dynamically switch between multiple sources to drive the clock network. You can use the Quartus II MegaWizard Plug-In Manager to create this clock control block with the ALTCLKCTRL megafunction. Arria GX, Arria II GX, Cyclone III, Cyclone IV, Stratix II, Stratix III, Stratix IV, and Stratix V devices provide clock control blocks for global clock networks. In addition, Stratix II, Stratix III,

Stratix IV, and Stratix V devices have clock control blocks for regional clock networks. The dynamic clock enable feature lets internal logic control the clock network. When a clock network is powered down, all the logic fed by that clock network does not toggle, thereby reducing the overall power consumption of the device. Figure 14-8 shows a 4-input clock control block diagram.

Figure 14-8. Clock Control Block Diagram

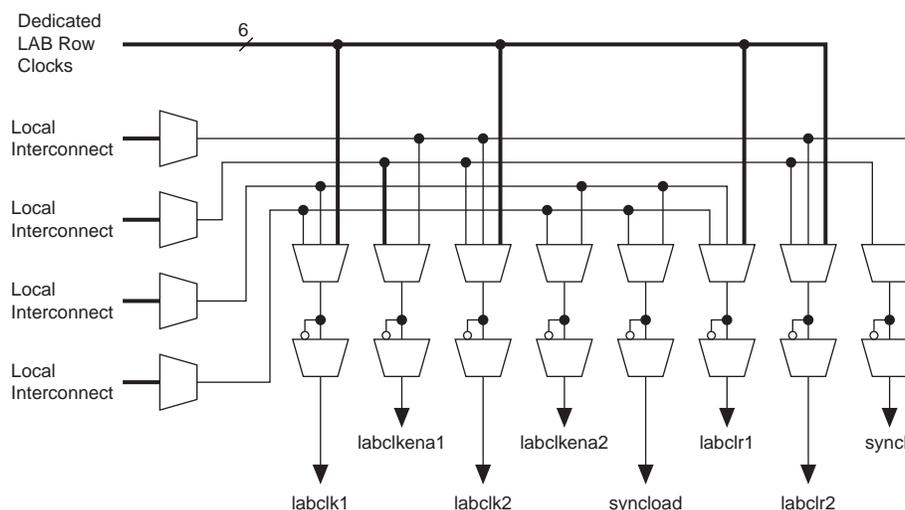


The enable signal is applied to the clock signal before being distributed to global routing. Therefore, the enable signal can either have a significant timing slack (at least as large as the global routing delay) or it can reduce the f_{MAX} of the clock signal.

For more information about using clock control blocks, refer to the *Clock Control Block Megafunction User Guide (ALTCLKCTRL)*.

Another contributor to clock power consumption is the LAB clock that distributes a clock to the registers within a LAB. LAB clock power can be the dominant contributor to overall clock power. For example, in Cyclone III devices, each LAB can use two clocks and two clock enable signals, as shown in Figure 14-9. Each LAB's clock signal and clock enable signal are linked. For example, an LE in a particular LAB using the labclk1 signal also uses the labclkena1 signal.

Figure 14-9. LAB-Wide Control Signals



To reduce LAB-wide clock power consumption without disabling the entire clock tree, use the LAB-wide clock enable to gate the LAB-wide clock. The Quartus II software automatically promotes register-level clock enable signals to the LAB-level. All registers within an LAB that share a common clock and clock enable are controlled by a shared gated clock. To take advantage of these clock enables, use a clock enable construct in the relevant HDL code for the registered logic.

LAB-Wide Clock Enable Example

The VHDL code in [Example 14-1](#) makes use of a LAB-wide clock enable. This clock-gating logic is automatically turned into an LAB-level clock enable signal.

Example 14-1.

```
IF clk'event AND clock = '1' THEN
  IF logic_is_enabled = '1' THEN
    reg <= value;
  ELSE
    reg <= reg;
  END IF;
END IF;
```



For more information about LAB-wide control signals, refer to the [Stratix II Architecture](#), [Cyclone III Device Family Overview](#), or [Cyclone II Architecture](#) chapters in the respective device handbook.

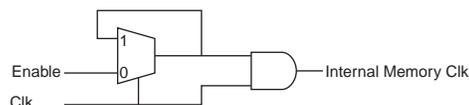
Reducing Memory Power Consumption

The memory blocks in FPGA devices can represent a large fraction of typical core dynamic power. Memory consumes approximately 20% of the core dynamic power in typical Cyclone III and Stratix III device designs. Memory blocks are unlike most other blocks in the device because most of their power is tied to the clock rate, and is insensitive to the toggle rate on the data and address lines.

When a memory block is clocked, there is a sequence of timed events that occur within the block to execute a read or write. The circuitry controlled by the clock consumes the same amount of power regardless of whether or not the address or data has changed from one cycle to the next. Thus, the toggle rate of input data and the address bus have no impact on memory power consumption.

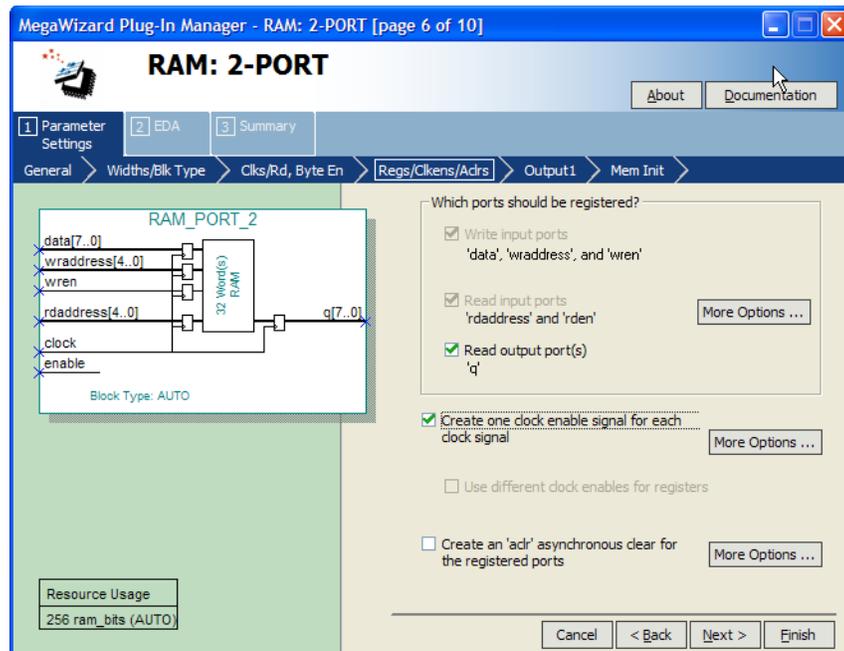
The key to reducing memory power consumption is to reduce the number of memory clocking events. You can achieve this through clock network-wide gating described in “[Clock Power Management](#)” on page 14-12, or on a per-memory basis through use of the clock enable signals on the memory ports. [Figure 14-10](#) shows the logical view of the internal clock of the memory block. Use the appropriate enable signals on the memory to make use of the clock enable signal instead of gating the clock.

Figure 14-10. Memory Clock Enable Signal



Using the clock enable signal enables the memory only when necessary and shuts it down for the rest of the time, reducing the overall memory power consumption. You can use the MegaWizard Plug-In Manager to create these enable signals by selecting the **Clock enable signal** option for the appropriate port when generating the memory block function (Figure 14-11).

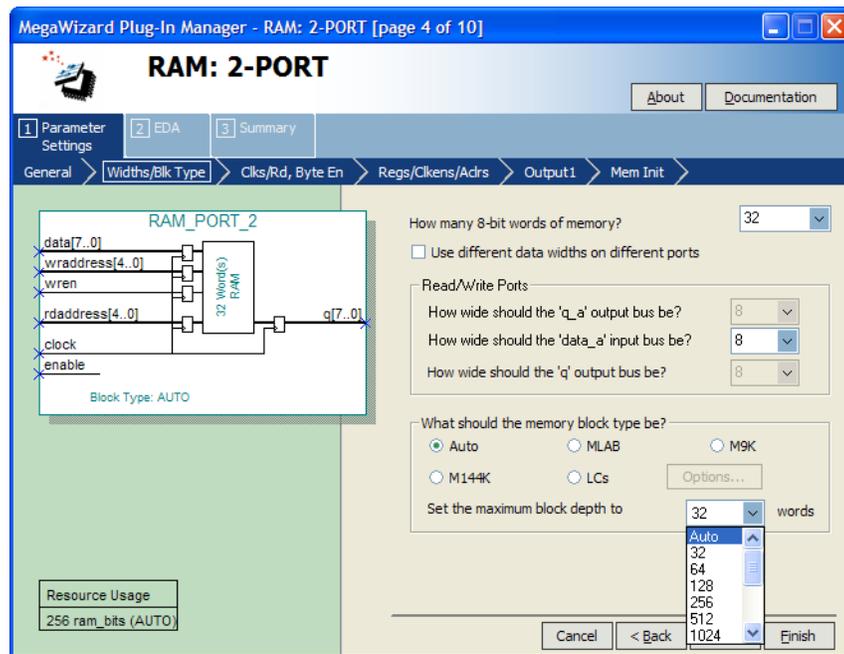
Figure 14-11. MegaWizard Plug-In Manager RAM 2-Port Clock Enable Signal Selectable Option



For example, consider a design that contains a 32-bit-wide M4K memory block in ROM mode that is running at 200 MHz. Assuming that the output of this block is only required approximately every four cycles, this memory block will consume 8.45 mW of dynamic power according to the demands of the downstream logic. By adding a small amount of control logic to generate a read clock enable signal for the memory block only on the relevant cycles, the power can be cut 75% to 2.15 mW.

You can also use the `MAXIMUM_DEPTH` parameter in your memory megafunction to save power in Cyclone II, Cyclone III, Cyclone IV GX, Stratix II, Stratix III, Stratix IV, and Stratix V devices; however, this approach might increase the number of LEs required to implement the memory and affect design performance.

You can set the `MAXIMUM_DEPTH` parameter for memory modules manually in the megafunction instantiation or in the MegaWizard Plug-In Manager (Figure 14-12). The Quartus II software automatically chooses the best design memory configuration for optimal power, as described in “Power-Driven Compilation” on page 14-4.

Figure 14-12. MegaWizard Plug-In Manager RAM 2-Port Maximum Depth Selectable Option

Memory Power Reduction Example

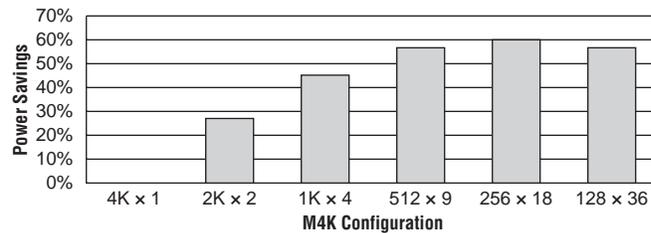
Table 14-3 shows power usage measurements for a $4\text{K} \times 36$ simple dual-port memory implemented using multiple M4K blocks in a Stratix II EP2S15 device. For each implementation, the M4K blocks are configured with a different memory depth.

Table 14-3. $4\text{K} \times 36$ Simple Dual-Port Memory Implemented Using Multiple M4K Blocks

M4K Configuration	Number of M4K Blocks	ALUTs
$4\text{K} \times 1$ (Default setting)	36	0
$2\text{K} \times 2$	36	40
$1\text{K} \times 4$	36	62
512×9	32	143
256×18	32	302
128×36	32	633

Figure 14-13 shows the amount of power saved using the `MAXIMUM_DEPTH` parameter. For all implementations, a user-provided read enable signal is present to indicate when read data is required. Using this power-saving technique can reduce power consumption by as much as 60%.

Figure 14-13. Power Savings Using the `MAXIMUM_DEPTH` Parameter



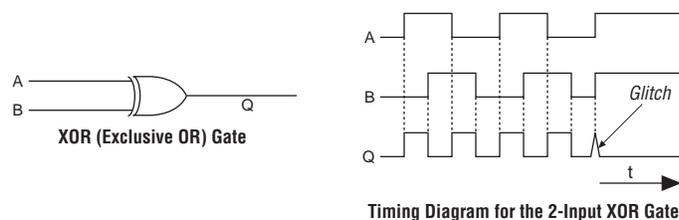
As the memory depth becomes more shallow, memory dynamic power decreases because unaddressed M4K blocks can be shut off using a decoded combination of address bits and the read enable signal. For a 128-deep memory block, power used by the extra LEs starts to outweigh the power gain achieved by using a more shallow memory block depth. The power consumption of the memory blocks and associated LEs depends on the memory configuration.

Pipelining and Retiming

Designs with many glitches consume more power because of faster switching activity. Glitches cause unnecessary and unpredictable temporary logic switches at the output of combinational logic. A glitch usually occurs when there is a mismatch in input signal timing leading to unequal propagation delay.

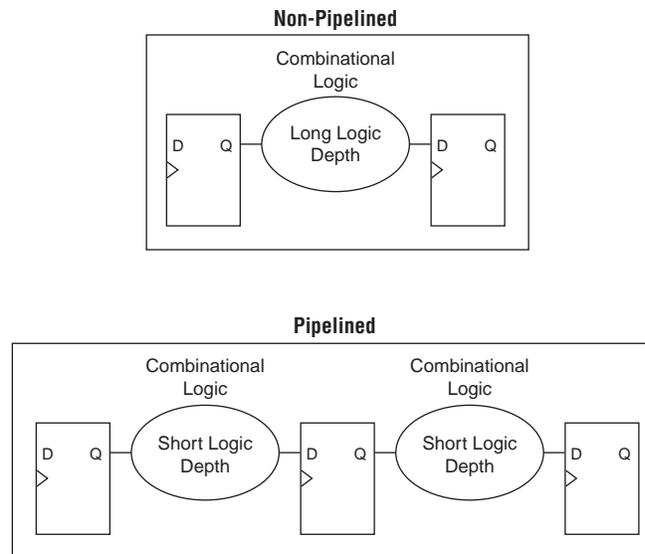
For example, consider an input change on one input of a 2-input XOR gate from 1 to 0, followed a few moments later by an input change from 0 to 1 on the other input. For a moment, both inputs become 1 (high) during the state transition, resulting in 0 (low) at the output of the XOR gate. Subsequently, when the second input transition takes place, the XOR gate output becomes 1 (high). During signal transition, a glitch is produced before the output becomes stable, as shown in Figure 14-14. This glitch can propagate to subsequent logic and create unnecessary switching activity, increasing power consumption. Circuits with many XOR functions, such as arithmetic circuits or cyclic redundancy check (CRC) circuits, tend to have many glitches if there are several levels of combinational logic between registers.

Figure 14-14. XOR Gate Showing Glitch at the Output



Pipelining can reduce design glitches by inserting flipflops into long combinational paths. Flipflops do not allow glitches to propagate through combinational paths. Therefore, a pipelined circuit tends to have less glitching. Pipelining has the additional benefit of generally allowing higher clock speed operations, although it does increase the latency of a circuit (in terms of the number of clock cycles to a first result). Figure 14-15 shows an example where pipelining is applied to break up a long combinational path.

Figure 14-15. Pipelining Example



Pipelining is very effective for glitch-prone arithmetic systems because it reduces switching activity, resulting in reduced power dissipation in combinational logic. Additionally, pipelining allows higher-speed operation by reducing logic-level numbers between registers. The disadvantage of this technique is that if there are not many glitches in your design, pipelining can increase power consumption by adding unnecessary registers. Pipelining can also increase resource utilization. The benchmark data shows that pipelining can reduce dynamic power consumption by as much as 30% in Cyclone and Stratix devices.

Architectural Optimization

You can use design-level architectural optimization by taking advantage of specific device architecture features. These features include dedicated memory and DSP or multiplier blocks available in FPGA devices to perform memory or arithmetic-related functions. You can use these blocks in place of LUTs to reduce power consumption. For example, you can build large shift registers from RAM-based FIFO buffers instead of building the shift registers from the LE registers.

The Stratix device family allows you to efficiently target small, medium, and large memories with the TriMatrix memory architecture. Each TriMatrix memory block is optimized for a specific function. The M512 memory blocks available in Stratix II devices are useful for implementing small FIFO buffers, DSP, and clock domain transfer applications. M512 memory blocks are more power-efficient than the distributed memory structures in some competing FPGAs. The M4K memory blocks

are used to implement buffers for a wide variety of applications, including processor code storage, large look-up table implementation, and large memory applications. The M-RAM blocks are useful in applications where a large volume of data must be stored on-chip. Effective utilization of these memory blocks can have a significant impact on power reduction in your design.

The latest Stratix and Cyclone device families have configurable M9K memory blocks that provide various memory functions such as RAM, FIFO buffers, and ROM.



For more information about using DSP and memory blocks efficiently, refer to the *Area and Timing Optimization* chapter in volume 2 of the *Quartus II Handbook*.

I/O Power Guidelines

Nonterminated I/O standards such as LVTTTL and LVCMOS have a rail-to-rail output swing. The voltage difference between logic-high and logic-low signals at the output pin is equal to the V_{CCIO} supply voltage. If the capacitive loading at the output pin is known, the dynamic power consumed in the I/O buffer can be calculated as shown in [Equation 14-1](#):

Equation 14-1. Capacitive loading at the output pin

$$P = 0.5 \times F \times C \times V^2$$

In this equation, F is the output transition frequency and C is the total load capacitance being switched. V is equal to V_{CCIO} supply voltage. Because of the quadratic dependence on V_{CCIO} , lower voltage standards consume significantly less dynamic power.

Transistor-to-transistor logic (TTL) I/O buffers consume very little static power. As a result, the total power consumed by a LVTTTL or LVCMOS output is highly dependent on load and switching frequency.

When using resistively terminated I/O standards like SSTL and HSTL, the output load voltage swings by a small amount around some bias point. The same dynamic power equation is used, where V is the actual load voltage swing. Because this is much smaller than V_{CCIO} , dynamic power is lower than for nonterminated I/O under similar conditions. These resistively terminated I/O standards dissipate significant static (frequency-independent) power, because the I/O buffer is constantly driving current into the resistive termination network. However, the lower dynamic power of these I/O standards means they often have lower total power than LVCMOS or LVTTTL for high-frequency applications. Use the lowest drive strength I/O setting that meets your speed and waveform requirements to minimize I/O power when using resistively terminated standards.

You can save a small amount of static power by connecting unused I/O banks to the lowest possible V_{CCIO} voltage of 1.2 V.

[Table 14-4](#) shows the total supply and thermal power consumed by outputs using different I/O standards for Stratix II devices. The numbers are for an I/O pin transmitting random data clocked at 200 MHz with a 10 pF capacitive load.

For this configuration, nonterminated standards generally use less power, but this is not always the case. If the frequency or the capacitive load is increased, the power consumed by nonterminated outputs increases faster than the power of terminated outputs.

Table 14-4. I/O Power for Different I/O Standards in Stratix II Devices

Standard	Total Supply Current Drawn from V_{CCIO} Supply (mA)	Total On-Chip Thermal Power Dissipation (mW)
3.3-V LVTTTL	2.42	9.87
2.5-V LVCMOS	1.9	6.69
1.8-V LVCMOS	1.34	4.18
1.5-V LVCMOS	1.18	3.58
3.3-V PCI	2.47	10.23
SSTL-2 class I	6.07	4.42
SSTL-2 class II	10.72	5.1
SSTL-18 class I	5.33	3.28
SSTL-18 class II	8.56	4.06
HSTL-15 class I	6.06	3.49
HSTL-15 class II	11.08	4.87
HSTL-18 class I	6.87	4.09
HSTL-18 class II	12.33	5.82

 For more information about I/O standards, refer to the *Selectable I/O Standards in Stratix II Devices and Stratix II GX Devices* chapter in volume 2 of the *Stratix II Device Handbook*, the *Stratix III Device I/O Features* chapter in volume 1 of the *Stratix III Device Handbook*, the *I/O Features in Stratix IV Devices* in volume 1 of the *Stratix IV Device Handbook*, or the *Selectable I/O Standards in Cyclone II Devices* chapter in the *Cyclone II Device Handbook*, the *Cyclone III Device Handbook*, or the *Cyclone IV GX Handbook*.

When calculating I/O power, the PowerPlay Power Analyzer uses the default capacitive load set for the I/O standard in the **Capacitive Loading** page of the **Device and Pin Options** dialog box. For Stratix II devices, if **Enable Advanced I/O Timing** is turned on, I/O power is measured using an equivalent load calculated as the sum of the near capacitance, the transmission line distributed capacitance, and the far-end capacitance as defined in the **Board Trace Model** page of the **Device and Pin Options** dialog box or the Board Trace Model view in the Pin Planner. Any other components defined in the board trace model are not taken into account for the power measurement.

For Cyclone III, Cyclone IV GX, Stratix III, Stratix IV, and Stratix V, devices, Advanced I/O Timing, which uses the full board trace model, is always used.

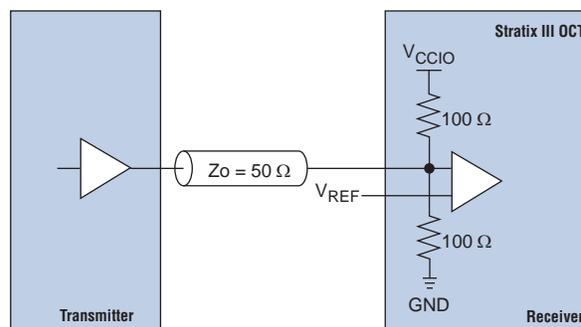
 For information about using Advanced I/O Timing and configuring a board trace model, refer to the *I/O Management* chapter in volume 2 of the *Quartus II Handbook*.

Dynamically Controlled On-Chip Terminations

Stratix V, Stratix IV and Stratix III FPGAs offer dynamic on-chip termination (OCT). Dynamic OCT enables series termination (RS) and parallel termination (RT) to dynamically turn on/off during the data transfer. This feature is especially useful when Stratix V, Stratix IV and Stratix III FPGAs are used with external memory interfaces, such as interfacing with DDR memories.

Compared to conventional termination, dynamic OCT reduces power consumption significantly as it eliminates the constant DC power consumed by parallel termination when transmitting data. Parallel termination is extremely useful for applications that interface with external memories where I/O standards, such as HSTL and SSTL, are used. Parallel termination supports dynamic OCT, which is useful for bidirectional interfaces (see Figure 14-16).

Figure 14-16. Stratix III On-Chip Parallel Termination



The following is an example of power saving for a DDR3 interface using on-chip parallel termination.

The static current consumed by parallel OCT is equal to the V_{CCIO} voltage divided by 100Ω . For DDR3 interfaces that use SSTL-15, the static current is $1.5 \text{ V} / 100 \Omega = 15 \text{ mA}$ per pin. Therefore, the static power is $1.5 \text{ V} \times 15 \text{ mA} = 22.5 \text{ mW}$. For an interface with 72 DQ and 18 DQS pins, the static power is $90 \text{ pins} \times 22.5 \text{ mW} = 2.025 \text{ W}$. Dynamic parallel OCT disables parallel termination during write operations, so if writing occurs 50% of the time, the power saved by dynamic parallel OCT is $50\% \times 2.025 \text{ W} = 1.0125 \text{ W}$.

 For more information about dynamic OCT in Stratix IV and Stratix III devices, refer to the *Stratix III Device I/O Features* chapter in the *Stratix III Device Handbook* and the *Stratix IV Device I/O Features* chapter in the *Stratix IV Device Handbook*, respectively.

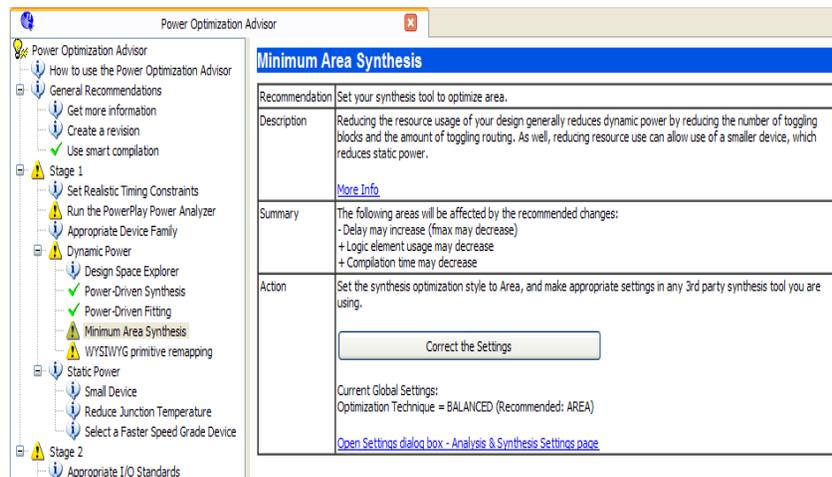
Power Optimization Advisor

The Quartus II software includes the Power Optimization Advisor, which provides specific power optimization advice and recommendations based on the current design project settings and assignments. The advisor covers many of the suggestions listed in this chapter. The following example shows how to reduce your design power with the Power Optimization Advisor.

Power Optimization Advisor Example

After compiling your design, run the PowerPlay Power Analyzer to determine your design power and to see where power is dissipated in your design. Based on this information, you can run the Power Optimization Advisor to implement recommendations that can reduce design power. Figure 14-17 shows the Power Optimization Advisor after compiling a design that is not fully optimized for power.

Figure 14-17. Power Optimization Advisor



The Power Optimization Advisor shows the recommendations that can reduce power in your design. The recommendations are split into stages to show the order in which you should apply the recommended settings. The first stage shows mostly CAD setting options that are easy to implement and highly effective in reducing design power. An icon indicates whether each recommended setting is made in the current project. In Figure 14-17, the checkmark icons for Stage 1 shows the recommendations that are already implemented. The warning icons indicate recommendations that are not followed for this compilation. The information icon shows the general suggestions. Each recommendation includes the description, summary of the effect of the recommendation, and the action required to make the appropriate setting.

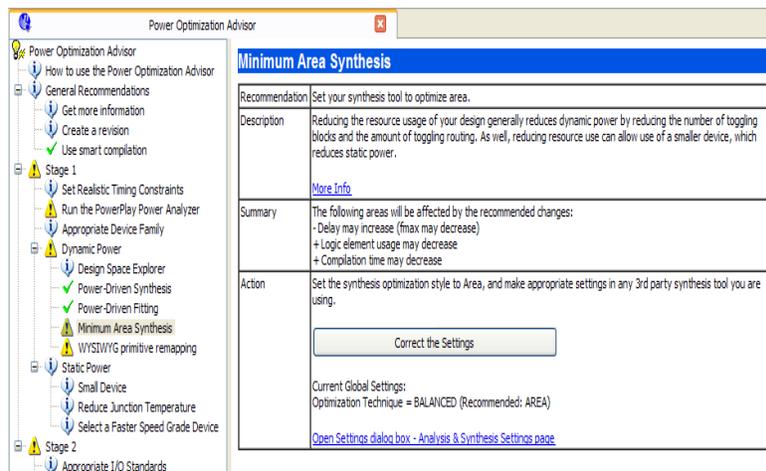
There is a link from each recommendation to the appropriate location in the Quartus II user interface where you can change the setting. You can change the **Power-Driven Synthesis** setting by clicking **Open Settings dialog box - Analysis & Synthesis Settings page** (Figure 14-18). The **Settings** dialog box is shown with the **Analysis & Synthesis Settings** page selected, where you can change the **PowerPlay** power optimization settings.

Figure 14-18. Analysis & Synthesis Settings Page

	<p>Logic element usage may increase + Completion time may decrease</p> <p>Action</p> <p>Set the synthesis optimization style to Area, and make appropriate settings in any 3rd party synthesis tool you are using.</p> <p>Correct the Settings</p> <p>Current Global Settings: Optimization Technique = BALANCED (Recommended: AREA)</p> <p>Open Settings dialog box - Analysis & Synthesis Settings page</p>
--	---

After making the recommended changes, recompile your design. The Power Optimization Advisor indicates with green check marks that the recommendations were implemented successfully (Figure 14–19). You can use the PowerPlay Power Analyzer to verify your design power results.

Figure 14–19. Implementation of Power Optimization Advisor Recommendations



The recommendations listed in Stage 2 generally involve design changes, rather than CAD settings changes as in Stage 1. You can use these recommendations to further reduce your design power consumption. Altera recommends that you implement Stage 1 recommendations first, then the Stage 2 recommendations.

Conclusion

The combination of a smaller process technology, the use of low-k dielectric material, and reduced supply voltage significantly reduces dynamic power consumption in the latest FPGAs. To further reduce your dynamic power, use the design recommendations presented in this chapter to optimize resource utilization and minimize power consumption.

Document Revision History

Table 14–5 shows the revision history for this chapter.

Table 14–5. Document Revision History (Part 1 of 2)

Date	Version	Changes
December 2010	10.0.1	Template update.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Was chapter 11 in the 9.1.0 release ■ Updated Figures 14-2, 14-3, 14-6, 14-18, 14-19, and 14-20 ■ Updated device support ■ Minor editorial updates

Table 14-5. Document Revision History (Part 2 of 2)

Date	Version	Changes
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Updated Figure 11-1 and associated references ■ Updated device support ■ Minor editorial update
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Was chapter 9 in the 8.1.0 release ■ Updated for the Quartus II software release ■ Added benchmark results ■ Removed several sections ■ Updated Figure 14-1, Figure 14-17, Figure 14-18, and Figure 14-19
November 2008	8.1.0	<ul style="list-style-type: none"> ■ Changed to 8½" × 11" page size ■ Changed references to altsyncram to RAM ■ Minor editorial updates
May 2008	8.0.0	<ul style="list-style-type: none"> ■ Added support for Stratix IV devices ■ Updated Table 9-1 and 9-9 ■ Updated "Architectural Optimization" on page 9-22 ■ Added "Dynamically-Controlled On-Chip Terminations" on page 9-26 ■ Updated "Referenced Documents" on page 9-29 ■ Updated references

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

As FPGA designs grow larger in density, the ability to analyze the design for performance, routing congestion, and logic placement to meet the design requirements becomes critical. This chapter discusses how to analyze the design floorplan with the Chip Planner.

You can perform design analysis and create and optimize the design floorplan with the Chip Planner. To make I/O assignments, use the Pin Planner.

-  For information about the Pin Planner, refer to the *I/O Management* chapter in volume 2 of the *Quartus II Handbook*.
-  You can use the Design Partition Planner with the Chip Planner to customize the floorplan of your design. For more information, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* and the *Best Practices for Incremental Compilation Partitions and Floorplan Assignments* chapters in volume 1 of the *Quartus II Handbook*.

This chapter includes the following topics:

- “Chip Planner Overview”
 - “LogicLock Regions” on page 15–4
 - “Using LogicLock Regions in the Chip Planner” on page 15–12
 - “Design Floorplan Analysis Using the Chip Planner” on page 15–13
 - “Scripting Support” on page 15–26
- ② For a list of devices supported by the Chip Planner, refer to *About the Chip Planner* in *Quartus II Help*.

Chip Planner Overview

The Chip Planner provides a visual display of chip resources. The Chip Planner can show logic placement, LogicLock regions, relative resource usage, detailed routing information, fan-in and fan-out connections between nodes, timing paths between registers, delay estimates for paths, and routing congestion information.

You can also make assignment changes with the Chip Planner, such as creating and deleting resource assignments, and you can perform post-compilation changes such as creating, moving, and deleting logic cells and I/O atoms. With the Chip Planner, you can view and create assignments for a design floorplan, perform power and design analyses, and implement ECOs. With the Chip Planner and Resource Property Editor, you can change connections between resources and make post-compilation changes to the properties of logic cells, I/O elements, PLLs, and RAM and digital signal processing (DSP) blocks.

- 
- For details about how to implement ECOs in your design using the Chip Planner in the Quartus II software, refer to the *Engineering Change Management with the Chip Planner* chapter in volume 2 of the *Quartus II Handbook*.

Starting the Chip Planner

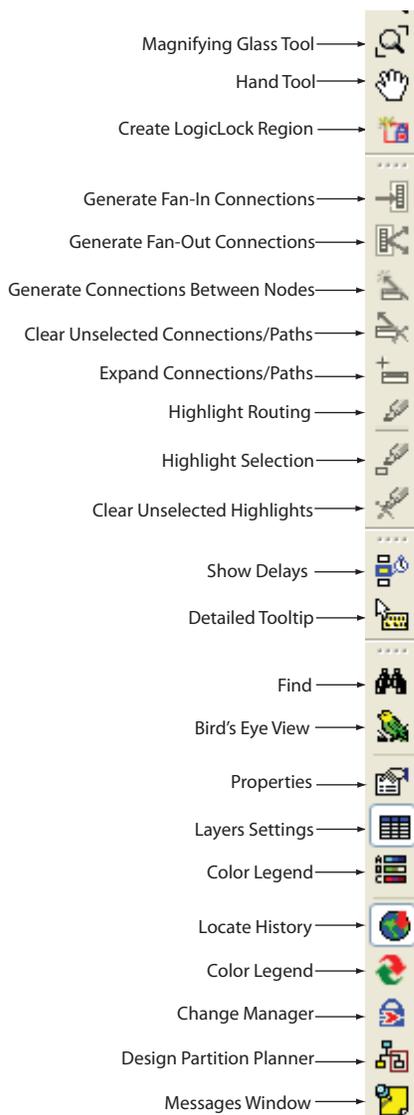
To start the Chip Planner, on the Tools menu, click **Chip Planner (Floorplan & Chip Editor)**. You can also start the Chip Planner by the following methods:

- Click the Chip Planner icon on the Quartus II software toolbar
- On the Shortcut menu in the following tools, click **Locate** and then click **Locate in Chip Planner (Floorplan and Chip Editor)**:
 - Design Partition Planner
 - Compilation Report
 - LogicLock Regions window
 - Technology Map Viewer
 - Project Navigator window
 - RTL source code
 - Node Finder
 - Simulation Report
 - RTL Viewer
 - Report Timing panel of the TimeQuest Timing Analyzer

Chip Planner Toolbar

The Chip Planner provides powerful tools for design analysis with a GUI. You can access Chip Planner commands from the View menu and the Shortcut menu, or by clicking the icons on the toolbar. Figure 15-1 shows an example of the Chip Planner toolbar and provides descriptions for commonly used icons.

Figure 15-1. Chip Planner Toolbar



Chip Planner Tasks, Layers, and Editing Modes

The Chip Planner models types of resource objects as unique display layers, and uses tasks—which are predefined sets of layer settings—to control the display of resources. The Chip Planner provides a set of default tasks, and you can create custom tasks to customize the display for your particular needs. The Basic, Detailed, and Floorplan Editing tasks provided with the Chip Planner are useful for general ECO and assignment-related activities, while the Partition Planner, Power, and Routing Congestion tasks are optimized for specific activities.

The Chip Planner has two editing modes, which determine the types of operations that you can perform. The Assignment editing mode allows you to make assignment changes that are applied by the Fitter during the next place and route operation. The ECO editing mode allows you to make post-compilation changes, commonly referred to as engineering change orders (ECOs).

You should choose the editing mode appropriate for the work that you want to perform, and a task that displays the resources that you want to view, in a level of detail appropriate for your design.

Locate History Window

As you optimize your design floorplan, you might have to locate a path or node in the Chip Planner many times. The Locate History window lists all the nodes and paths you have displayed using a **Locate in Chip Planner (Floorplan and Chip Editor)** command, providing easy access to the nodes and paths of interest to you. Double-clicking a node or path in the Locate History window displays the selected node or path in the Chip Planner.

 For more information about the Chip Planner, refer to *About the Chip Planner* and *Layers Settings Dialog Box* in Quartus II Help. For more information about the ECO editing mode, refer to the *Engineering Change Management with the Chip Planner* chapter in volume 2 of the *Quartus II Handbook*.

LogicLock Regions

LogicLock regions allow you to constrain resource placement in specified areas on the device. A LogicLock region is defined by its height, width, and location; you can specify the size or location of a region, or both, or the Quartus II software can generate these properties automatically. The Quartus II software bases the size and location of a region on the contents of the region and the timing requirements of the module.

You can use LogicLock regions to create a floorplan for your design. Your floorplan can contain several LogicLock regions. [Table 15-1](#) describes the options for creating LogicLock regions.

Table 15-1. Types of LogicLock Regions

Properties	Values	Behavior
State	Floating (default), Locked	Allows the Quartus II software to determine the location of the region on the device. Locked regions are areas that you define and are shown with a solid boundary in the floorplan. A locked region must have a fixed size.
Size	Auto (default), Fixed	Allows the Quartus II software to determine the appropriate size of a region given its contents. Fixed regions have a shape and size that you define.
Reserved	Off (default), On	Allows you to define whether the Fitter can use the resources within a region for entities that are not assigned to the region. If the reserved property is turned on, only items assigned to the region can be placed within its boundaries.
Origin	Any Floorplan Location	The origin of the LogicLock region's placement on the floorplan. For Arria series, Stratix, Cyclone series, MAX II, and MAX V devices, the origin is located in the lower left corner. For other Altera® device families, the origin is located in the upper left corner.

-  The Quartus II software cannot automatically define the size of a region if the location is locked. Therefore, if you want to specify the exact location of the region, you must also specify the size.

-  You can use the Design Partition Planner in conjunction with LogicLock regions to create a floorplan for your design. For more information about using the Design Partition Planner, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Designs* and the *Best Practices for Incremental Compilation Partition and Floorplan Assignments* chapters in volume 1 of the *Quartus II Handbook*.

Creating LogicLock Regions

You can create LogicLock Regions from the Project Navigator, the LogicLock Regions window, or the Chip Planner.

Creating LogicLock Regions with the Project Navigator

After you perform either a full compilation or analysis and elaboration on the design, the Quartus II software displays the hierarchy of the design. On the View menu, click **Project Navigator**. With the hierarchy of the design fully expanded, right-click on any design entity in the design, and click **Create New LogicLock Region** to create a LogicLock region.

Creating LogicLock Regions with the Chip Planner

To create a LogicLock region in the Chip Planner, click the **Create LogicLock Region** command on the View menu, then click and drag on the Chip Planner floorplan to create a region of your preferred location and size.

Placing LogicLock Regions

A fixed region must contain all resources required for the design block for which you define the region. Although the Quartus II software can automatically place and size LogicLock regions to meet resource and timing requirements, you can manually place and size regions to meet your design requirements, as follows:

- Place LogicLock regions with pin assignments on the periphery of the device, adjacent to the pins. For the Arria series, Cyclone series, MAX II, MAX V, and Stratix series devices, you must also include the I/O block within the LogicLock Region.
- Floating LogicLock regions can overlap with their ancestors or descendants, but not with other floating LogicLock regions.
- Avoid creating fixed and locked regions that overlap.

-  If you want to import multiple instances of a module into a top-level design, you must ensure that the device has two or more locations with exactly the same device resources. (You can determine this from the applicable device handbook.) If the device does not have another area with exactly the same resources, the Quartus II software generates a fitting error during compilation of the top-level design.

Placing Device Resources into LogicLock Regions

A LogicLock region includes all device resources within its boundaries, including memory and pins. You can assign pins to LogicLock regions; however, this placement puts location constraints on the region. When the Quartus II software places a floating auto-sized region, it places the region in an area that meets the requirements of the contents of the LogicLock region.

Only one LogicLock region can claim a device resource. If a LogicLock region boundary includes part of a device resource, the Quartus II software allocates the entire resource to that LogicLock region.



Pin assignments to LogicLock regions are effective only in fixed and locked regions. Pin assignments to floating regions do not influence the placement of the region.

LogicLock Regions Window

The LogicLock window consists of the LogicLock Regions window (Figure 15-2) and the **LogicLock Region Properties** dialog box. Use the LogicLock Regions window to create LogicLock regions and assign nodes and entities to them. The dialog box provides a summary of all LogicLock regions in your design. In the LogicLock Regions window, you can modify the properties of a LogicLock region such as size, state, width, height, origin, and whether the region is a reserved region. The LogicLock Regions window also has a recommendations toolbar at the bottom. Select a LogicLock region from the drop-down list in the recommendations toolbar to display the relevant suggestions to optimize that LogicLock region. You can customize the LogicLock Regions window by dragging and dropping the columns to change their order; you can also hide columns.

Figure 15-2. LogicLock Regions Window

Region Name	Size	Width	Height	State	Origin	Reserved	
LogicLock Regions							
Root Region	Fixed			Locked		Off	
<new>>							
Clk_ctrl:U_Clk_ctrl	Auto	1	1	Floating	X17_Y10	On	
Phy_int:U_Phys_int	Auto	3	3	Floating	X18_Y1	Off	
Reg_int:U_Reg_int	Auto	7	5	Floating	X1_Y13	Off	

LogicLock Regions window toolbar: Recommendation: 1 of 3 for Clk_ctrl:U_Clk_ctrl. Ensure that regions having timing-critical connections between them are placed near each other. Details...

Use the **LogicLock Region Properties** dialog box to obtain detailed information about your LogicLock region, such as which entities and nodes are assigned to your region and which resources are required. The **LogicLock Region Properties** dialog box shows the properties of the current selected regions. You can also modify the settings for LogicLock regions in the **LogicLock Region Properties** dialog box.

To open the **LogicLock Region Properties** dialog box, double-click any region in the LogicLock Regions window, or right-click the region and click **Properties**.

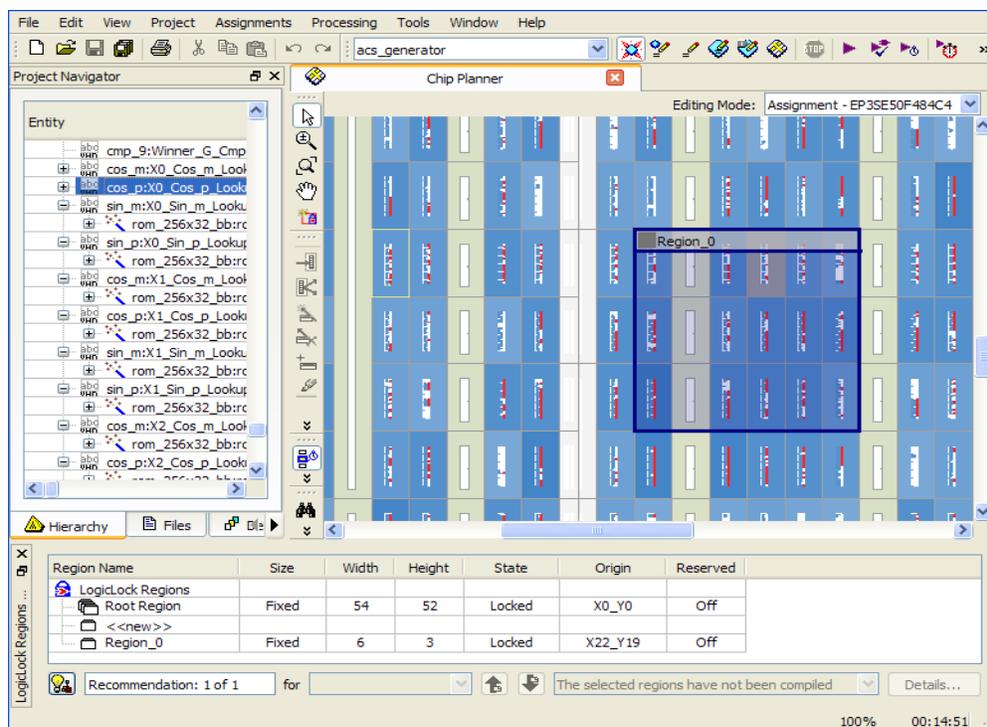
-  For designs that target Arria series, Cyclone series, Stratix series, MAX II, and MAX V devices, the Quartus II software automatically creates a LogicLock region that encompasses the entire device. This default region is labelled `Root_region`, and is locked and fixed.
-  For Arria series, Cyclone series, Stratix series, MAX II, and MAX V devices, the origin of the LogicLock region is located at the lower-left corner of the region. For all other supported devices, the origin is located at the upper-left corner of the region.

Assigning LogicLock Region Content

After you have created a LogicLock region, you must assign resources to the region using the Chip Planner, the **LogicLock Regions** dialog box, or a Tcl script.

You can drag selected logic displayed in the **Hierarchy** tab of the Project Navigator, in the Node Finder, or in a schematic design file, and drop it into the Chip Planner or the **LogicLock Regions** dialog box. [Figure 15-3](#) shows logic that has been dragged from the **Hierarchy** tab of the Project Navigator and dropped into a LogicLock region in the Chip Planner.

Figure 15-3. Drag and Drop Logic in the Chip Planner



You can also drag logic from the **Hierarchy** tab of the Project Navigator and drop the logic in the LogicLock Regions **Properties** dialog box. You can also drop logic into the **Design Element Assigned** column of the **Contents** tab of the **LogicLock Region Properties** box.

You must manually assign pins to a LogicLock region. The Quartus II software does not include pins automatically when you assign an entity to a region. The software only obeys pin assignments to locked regions that border the periphery of the device. For the Cyclone, Stratix, and MAX II series of devices, the locked regions must include the I/O pins as resources.

Hierarchical (Parent and Child) LogicLock Regions

You can define a hierarchy for a group of regions by declaring parent and child regions. The Quartus II software places a child region completely within the boundaries of its parent region, allowing you to further constrain module locations. Additionally, parent and child regions allow you to further improve the performance of a module by constraining the nodes in the critical path of the module.

To make one LogicLock region a child of another LogicLock region, in the LogicLock Regions window, select the new child region and drag and drop the new child region into its new parent region.



The LogicLock region hierarchy does not have to be the same as the design hierarchy.

You can create both fixed and floating LogicLock regions within a fixed parent LogicLock region. The location of a floating child region can float within its parent. If a child region is fixed, its location remains locked relative to its parent's origin. A locked parent region's location is locked relative to the device. If the child's location is locked and you change the parent's location, the child's origin changes, but maintains the same placement relative to the origin of its parent. Either you or the Quartus II software can determine a child region's size; however, the child region must fit entirely within the parent region. The levels of hierarchy in LogicLock regions are unlimited, but complicated hierarchical regions might result in some LABs not being utilized; thus, effectively increasing the resource utilization in the device.

Reserved LogicLock Region

The Quartus II software honors all entity and node assignments to LogicLock regions. Occasionally, entities and nodes do not occupy an entire region, which leaves some of the region's resources unoccupied. To increase the region's resource utilization and performance, the Quartus II software's default behavior fills the unoccupied resources with other nodes and entities that have not been assigned to another region. You can prevent this behavior by turning on **Reserved** on the **General** tab of the **LogicLock Region Properties** dialog box. When you turn on this option, your LogicLock region contains only the entities and nodes that you specifically assigned to your LogicLock region. When you set the reserved property for a LogicLock region, the Fitter does not place logic from the immediate parent LogicLock region in the assigned LogicLock area, but it might place logic from other parts of your design in that area.

Creating Non-Rectangular LogicLock Regions

When you create a floorplan for your design, you may want to create non-rectangular LogicLock regions to make some device resources accessible to design blocks outside a LogicLock region. You might also create a non-rectangular LogicLock region to place certain parts of your design around specific device resources to improve performance.

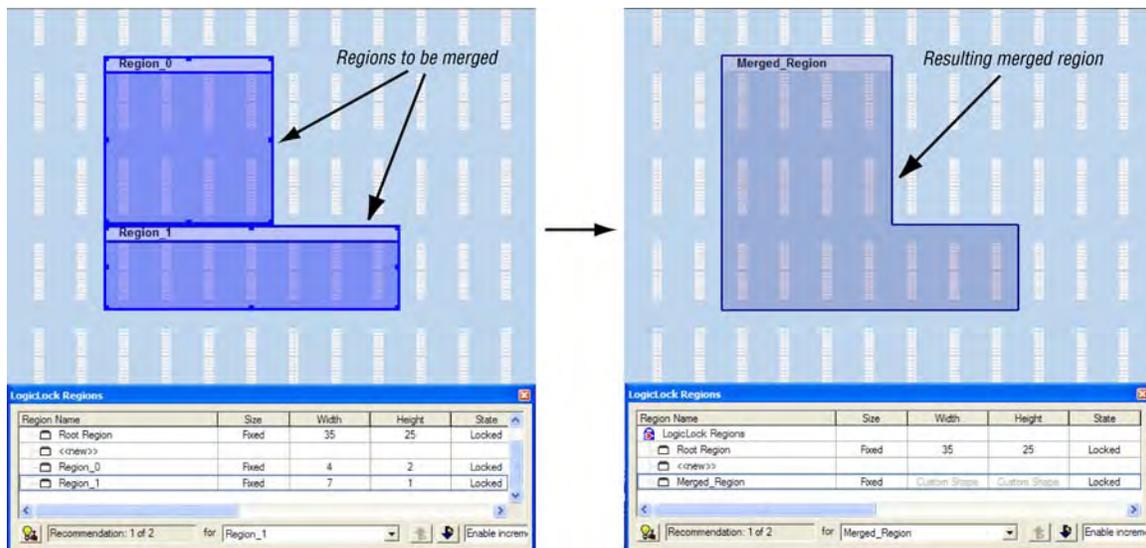
To create a non-rectangular region with the **Merge LogicLock Region** command, follow these steps:

1. In the Chip Planner, create two or more contiguous or non-contiguous rectangular regions as described in “[Creating LogicLock Regions](#)” on page 15-5.
2. Arrange the regions that you have created into the locations where you want the non-rectangular region to be.
3. Select all the individual regions that you want to merge by clicking each of them while holding the Shift key.
4. Right-click the title bar of any of the LogicLock regions that you want to merge, point to **LogicLock Regions**, and then click **Merge LogicLock Region**. The individual regions that you select merge to create a single new region.

By default, the new LogicLock region has the same name as the component region containing the greatest number of resources; however, you can rename the new region. In the **LogicLock Regions Window**, the new region is shown as having a custom shape.

[Figure 15-4](#) illustrates two autonomous LogicLock regions combined using the **Merge LogicLock Region** command to form a new non-rectangular region.

Figure 15-4. Using the Merge command to create a non-rectangular region



Excluded Resources

The Excluded Resources feature allows you to easily exclude specific device resources such as DSP blocks or M4K memory blocks from a LogicLock region. For example, you can specify resources that belong to a specific entity that are assigned to a LogicLock region, and specify that these resources be included with the exception of the DSP blocks. Use the Excluded Resources feature on a per-LogicLock region member basis.

To exclude certain device resources from an entity, in the **LogicLock Region Properties** dialog box, highlight the entity in the **Design Element** column, and click **Edit**. In the **Edit Node** dialog box, under **Excluded Element Types**, click the **Browse** button. In the **Excluded Resources Element Types** dialog box, you can select the device resources you want to exclude from the entity. When you have selected the resources to exclude, the **Excluded Resources** column is updated in the **LogicLock Region Properties** dialog box to reflect the excluded resources.

 The Excluded Resources feature prevents certain resource types from being included in a region, but it does not prevent the resources from being placed inside the region unless you set the region's **Reserved** property to **On**. To indicate to the Fitter that certain resources are not required inside a LogicLock region, define a resource filter.

Additional Quartus II LogicLock Design Features

To complement the **LogicLock Regions** window, the Quartus II software has additional features to help you design with LogicLock regions.

Analysis and Synthesis Resource Utilization by Entity

The Compilation Report contains an **Analysis and Synthesis Resource Utilization by Entity** section, which reports accurate resource usage statistics, including entity-level information. You can use this feature when you manually create LogicLock regions.

Quartus II Revisions Feature

When you evaluate different LogicLock regions in your design, you might want to experiment with different configurations to achieve your desired results. The Quartus II Revisions feature allows you to organize the same project with different settings until you find an optimum configuration.

To use the Revisions feature, on the Project menu, click **Revisions**. In the **Revisions** dialog box, you can create and specify revisions. You can base revision on the current design or any previously created revisions. Each revision can have an associated description. You can use revisions to organize the placement constraints created for your LogicLock regions.

LogicLock Assignment Precedence

You can encounter conflicts during the assignment of entities and nodes to LogicLock regions. For example, an entire top-level entity might be assigned to one region and a node within this top-level entity assigned to another region. To resolve conflicting assignments, the Quartus II software maintains an order of precedence for LogicLock assignments. The following order of precedence, from highest to lowest, applies:

- Exact node-level assignments
- Path-based and wildcard assignments
- Hierarchical assignments

 For more information about LogicLock assignment precedence, refer to *Understanding Assignment Priority* in Quartus II Help.

 Open the **Priority** dialog box by selecting **Priority** on the **General** tab of the **LogicLock Regions Properties** dialog box. You can change the priority of path-based and wildcard assignments with the **Up** and **Down** buttons in the **Priority** dialog box. To prioritize assignments between regions, you must select multiple LogicLock regions and then open the **Priority** dialog box from the **LogicLock Regions Properties** dialog box.

Virtual Pins

Usually, when you compile a design in the Quartus II software, all I/O ports are mapped directly to pins on the targeted device. However, you may not want to map all I/O ports to the device pins; use the Virtual Pin assignment then.

A virtual pin is an I/O element that you do not intend to connect to the chip pins. You create a virtual pin by assigning the Virtual Pin logic option to an I/O element. When you compile a design with some I/O elements assigned as virtual pins, those I/O elements become mapped to a logic element and not to a pin during compilation, and are then implemented as a LUT. You might use virtual pin assignments when you compile a partial design, because not all the I/Os from a partial design may drive chip pins at the top level.

The Virtual Pin assignment communicates to the Quartus II software which I/O ports of the design module are internal nodes in the top-level design. These assignments prevent the number of I/O ports in the lower-level modules from exceeding the total number of available device pins. Every I/O port that you designate as a virtual pin becomes mapped to either an LCELL or an adaptive logic module (ALM), depending on the target device.

 You cannot use virtual pins for bidirectional, registered I/O pins, or I/O pins with output enable signals.

In the top-level design, you connect these virtual pins to an internal node of another module. By making assignments to virtual pins, you can place those pins in the same location or region on the device as that of the corresponding internal nodes in the top-level module. You can use the Virtual Pin option when compiling a LogicLock module with more pins than the target device allows. The Virtual Pin option can enable timing analyses that more closely match the performance of the LogicLock module when you integrate it into the top-level design.

Apply the following guidelines when creating virtual pins in the Quartus II software:

- Do not declare clock pins as virtual pins
- You should not declare nodes or signals that drive physical device pins in the top-level design as virtual pins

 In the Node Finder, you can set **Filter Type** to **Pins: Virtual** to display all assigned virtual pins in the design. From the Assignment Editor, to access the Node Finder, double-click the **To** field; when the arrow appears on the right side of the field, click the arrow and select **Node Finder**.

Using LogicLock Regions in the Chip Planner

You can easily edit properties of LogicLock regions, and assign resources to them, in the Chip Planner. You can also create new LogicLock regions using the Chip Planner.

Viewing Connections Between LogicLock Regions in the Chip Planner

You can view and edit LogicLock regions using the Chip Planner. Select the **Floorplan Editing** task, or any task that has the **User-assigned LogicLock regions** setting enabled to manipulate LogicLock regions.

The Chip Planner shows the connections between LogicLock regions. By default, you can view each connection as an individual line drawn between LogicLock regions. You can choose to display connections between LogicLock regions as a single bundled connection rather than as individual connection lines. To use this option, open the Chip Planner and on the View menu, click **Inter-region Bundles**.

In the **Inter-region Bundles** dialog box, specify the **Maximum source node to region fanout** and the **Minimum Bundle Width** values.

- For more information about the **Inter-region Bundles** dialog box, refer to *Inter-region Bundles Dialog Box* in Quartus II Help.

Using LogicLock Regions with the Design Partition Planner

You can help optimize timing in a design by placing entities that share significant logical connectivity close to each other on the device. By default, the Fitter usually places closely connected entities in the same area of the device; however, you can use LogicLock regions, together with the Design Partition Planner and the Chip Planner, to help ensure that logically connected entities retain optimal placement from one compilation to the next.

You can view the logical connectivity between entities with the Design Partition Planner, and the physical placement of those entities with the Chip Planner. In the Design Partition Planner, you can identify entities that are highly interconnected, and place those entities in a partition. In the Chip Planner, you can create LogicLock regions and assign each partition to a LogicLock region, thereby preserving the placement of the entities.

- For more information about using LogicLock regions with design partitions, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* and the *Best Practices for Incremental Compilation Partition and Floorplan Assignments* chapters in volume 1 of the *Quartus II Handbook*. For more information about using the Design Partition Planner with the Chip Planner, refer to *About the Design Partition Planner* and *Using the Design Partition Planner* in Quartus II Help.

Design Floorplan Analysis Using the Chip Planner

The Chip Planner helps you visually analyze the floorplan of your design at any stage of your design cycle. With the Chip Planner, you can view post-compilation placement, connections, and routing paths. You can also create LogicLock regions and location assignments. The Chip Planner allows you to create new logic cells and I/O atoms and to move existing logic cells and I/O atoms using the architectural floorplan of your design. You can also see global and regional clock regions within the device, and the connections between both I/O atoms and PLLs and the different clock regions.

From the Chip Planner, you can launch the Resource Property Editor, which you can use to change the properties and parameters of device resources, and modify connectivity between certain types of device resources. The Change Manager records any changes that you make to your design floorplan, so that you can selectively undo changes if necessary.

 For more information about the Resource Property Editor and the Change Manager, refer to the *Engineering Change Management with the Chip Planner* chapter in volume 2 of the *Quartus II Handbook*, and to *About the Resource Property Editor* and *About the Change Manager* in Quartus II Help.

The following sections present Chip Planner floorplan views and design analysis procedures which you can use with any predefined task—unless a procedure requires a specific task or editing mode.

Chip Planner Floorplan Views

The Chip Planner uses a hierarchical zoom viewer that shows various abstraction levels of the targeted Altera device. As you zoom in, the level of abstraction decreases, revealing more detail about your design.

 For more information about Chip Planner floorplan views, refer to the *Engineering Change Management with the Chip Planner* chapter in volume 2 of the *Quartus II Handbook*.

Bird's Eye View

The Bird's Eye View displays a high-level picture of resource usage for the entire chip and provides a fast and efficient way to navigate between areas of interest in the Chip Planner.

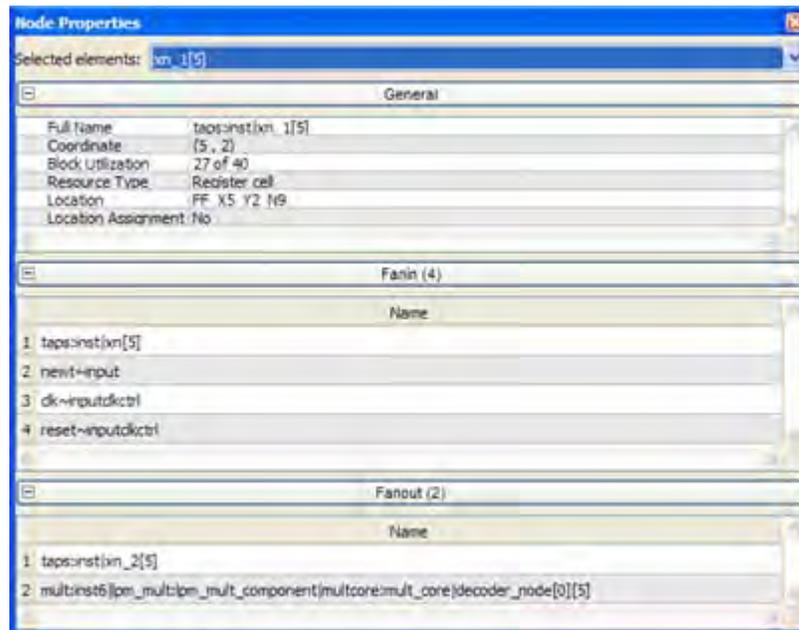
The Bird's Eye View is particularly useful when the parts of your design that you want to view are at opposite ends of the chip and you want to quickly navigate between resource elements without losing your frame of reference.

 For more information about the Bird's Eye View, refer to *Bird's Eye View* and *Displaying Resources and Information* in Quartus II Help.

Properties Window

The Properties Window displays detailed properties of the objects (such as atoms, paths, LogicLock regions, or routing elements) currently selected in the Chip Planner. To display the Properties Window, click **Properties** on the **View** menu in the Chip Planner.

Figure 15-5. Properties Window



Viewing Architecture-Specific Design Information

With the Chip Planner, you can view the following architecture-specific information related to your design:

- **Device routing resources used by your design**—View how blocks are connected, as well as the signal routing that connects the blocks.
- **LE configuration**—View logic element (LE) configuration in your design. For example, you can view which LE inputs are used; if the LE utilizes the register, the look-up table (LUT), or both; as well as the signal flow through the LE.
- **ALM configuration**—View ALM configuration in your design. For example, you can view which ALM inputs are used, if the ALM utilizes the registers, the upper LUT, the lower LUT, or all of them. You can also view the signal flow through the ALM.
- **I/O configuration**—View device I/O resource usage. For example, you can view which components of the I/O resources are used, if the delay chain settings are enabled, which I/O standards are set, and the signal flow through the I/O.
- **PLL configuration**—View phase-locked loop (PLL) configuration in your design. For example, you can view which control signals of the PLL are used with the settings for your PLL.
- **Timing**—View the delay between the inputs and outputs of FPGA elements. For example, you can analyze the timing of the DATAB input to the COMBOUT output.

In addition, you can modify the following device properties with the Chip Planner:

- LEs and ALMs
- I/O cells
- PLLs
- Registers in RAM and DSP blocks
- Connections between elements
- Placement of elements



For more information about LEs, ALMs, and other resources of an FPGA device, refer to the relevant device handbook.

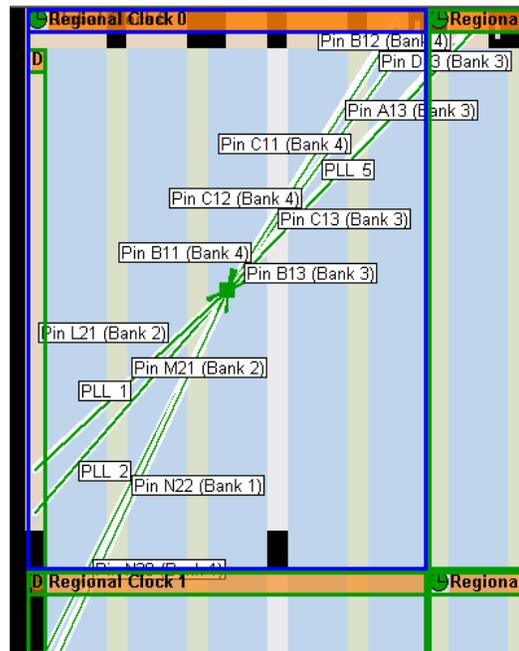
Viewing Available Clock Networks in the Device

When you select a task with clock region layers enabled, you can display the areas of the chip that are driven by global and regional clock networks. This global clock display feature is available for Arria GX, Arria II, Cyclone II, Cyclone III, HardCopy II, HardCopy III, Stratix II, Stratix II GX, Stratix III, Stratix IV, and Stratix V device families.

Depending on the clock layers activated in the selected task, the Chip Planner displays regional and global clocks and the regions they cover in the device, and the connectivity between clock regions, pins, and PLLs. Clock regions appear as rectangular overlay boxes with labels indicating the clock type and index. You can select each clock network region by clicking on the clock region. The clock-shaped icon at the top-left corner indicates that the region represents a clock network region. You can change the color in which the Chip Planner displays clock regions on the **Options** dialog box of the Tools menu.

The **Layer Settings** dialog box lists layers for different clock region types; when the selected device does not contain a given clock region, the option for that category is unavailable in the dialog box. You can customize the Chip Planner's display of clock regions by creating a custom task with selected clock layers enabled in the Layers Settings dialog box. **Figure 15-6** shows the potential fan-in in the Chip Planner.

Figure 15-6. Potential Fan-In



- ④ For more information about displaying potential fan-ins or fan-outs relative to a clock region, refer to *Displaying Resources and Information* in Quartus II Help.

Viewing Critical Paths

Critical paths are timing paths in your design that have a negative slack. These timing paths can span from device I/Os to internal registers, registers-to-registers, or from registers to device I/Os. The slack of a path determines its criticality; slack appears in the timing analysis report. Design analysis for timing closure is a fundamental requirement for optimal performance in highly complex designs. The analytical capability of the Chip Planner helps you close timing on complex designs.

Viewing critical paths in the Chip Planner helps you understand why a specific path is failing. You can see if any modification in the placement can reduce the negative slack. You can display details of a path (to expand/collapse the path to/from the connections in the path) by clicking **Expand Connections/Paths** in the toolbar, or by clicking on the "+/-" on the label.

You can locate failing paths from the timing report in the TimeQuest Timing Analyzer. To locate the critical paths, run the Report Timing task from the Custom Reports group in the Tasks pane of the TimeQuest Timing Analyzer. From the View pane, which lists the failing paths, right-click on any failing path or node, and select **Locate Path**. From the Locate dialog box, select **Chip Planner** to see the failing path in the Chip Planner.

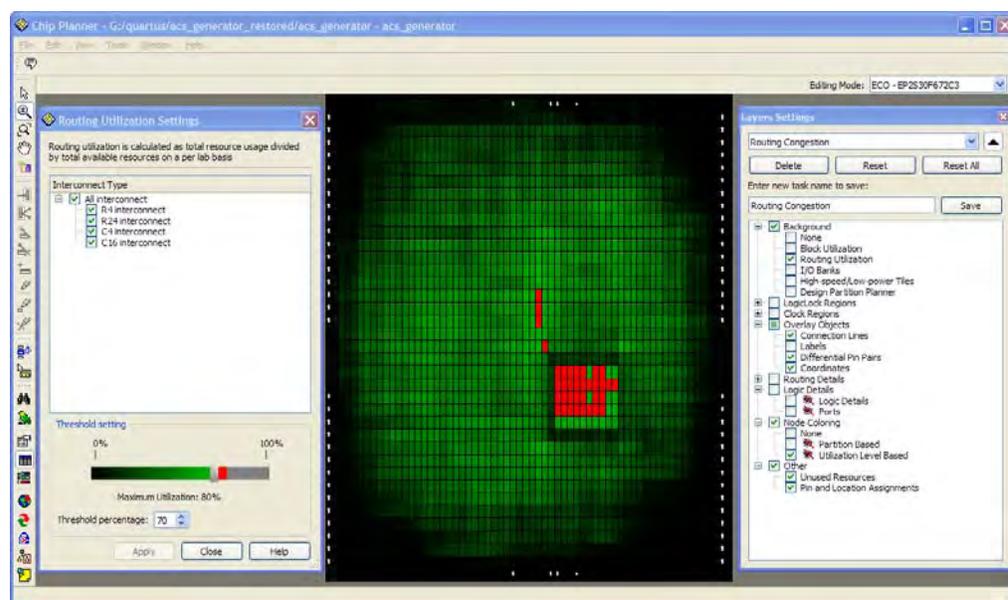
-  To display paths in the floorplan, you must first make timing settings and perform a timing analysis.
-  For more information about performing static timing analysis with the Quartus II TimeQuest Timing Analyzer, refer to *The Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

Viewing Routing Congestion

The **Routing Congestion** task allows you to determine the percentage of routing resources in use following a compilation. This feature identifies where there is a lack of routing resources, and helps you make design changes that might ease routing congestion and help meet design requirements.

To view routing congestion in the Chip Planner, select the **Routing Congestion** task. The **Routing Utilization Settings** dialog box appears whenever you select the **Routing Congestion** task; this dialog box allows you to set a congestion threshold value, and to specify the types of routing interconnects of interest (Figure 15-7).

Figure 15-7. Routing Utilization Settings dialog box



-  For more information about displaying routing congestion, refer to *Displaying Resources and Information* in Quartus II Help.

The routing congestion map uses the color and shading of logic resources to indicate relative resource utilization; darker shading represents a greater utilization of routing resources. Areas where routing utilization exceeds the threshold value specified in the Routing Utilization Settings dialog box appear in red. The congestion map can help you determine whether you can modify the floorplan, or make changes to the RTL to reduce routing congestion.

Viewing I/O Banks

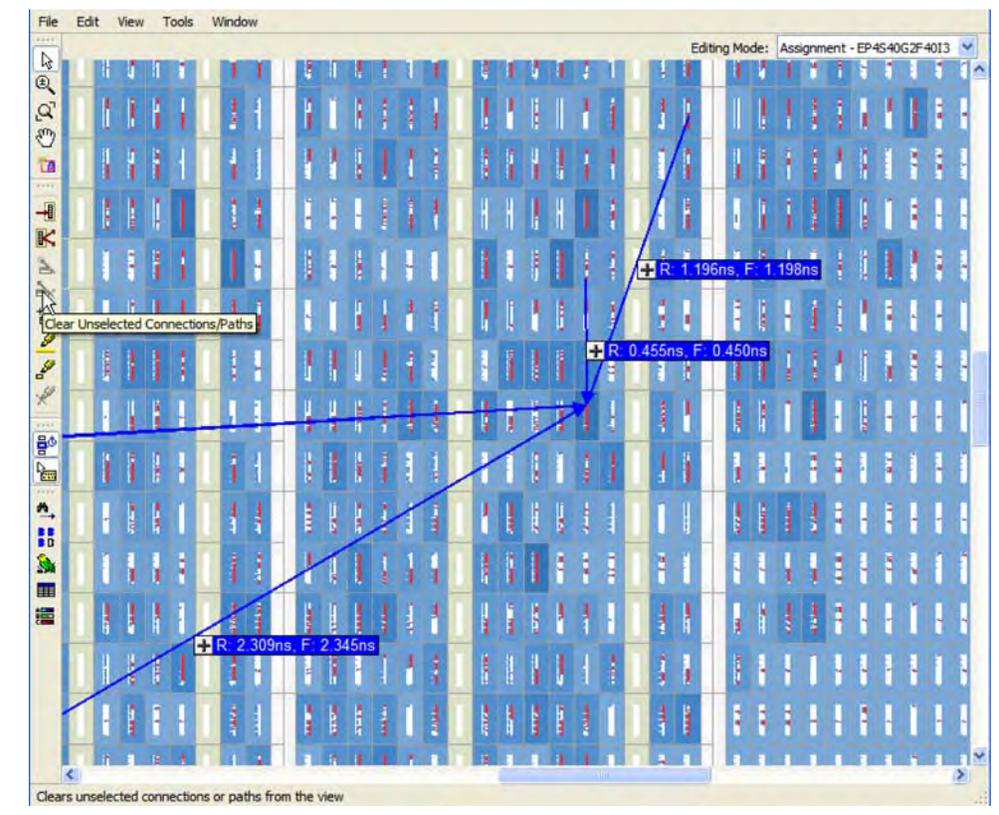
The Chip Planner can show all of the I/O banks of the device. To see the I/O bank map of the device, turn on the I/O Banks layer in the Layers Settings dialog box.

Generating Fan-In and Fan-Out Connections

The ability to display fan-in and fan-out connections enables you to view the atoms that fan-in to or fan-out from the selected atom. To remove the connections displayed, use the Clear Unselected Connections/Paths icon in the Chip Planner toolbar.

Figure 15-8 shows the fan-in connections for the selected resource.

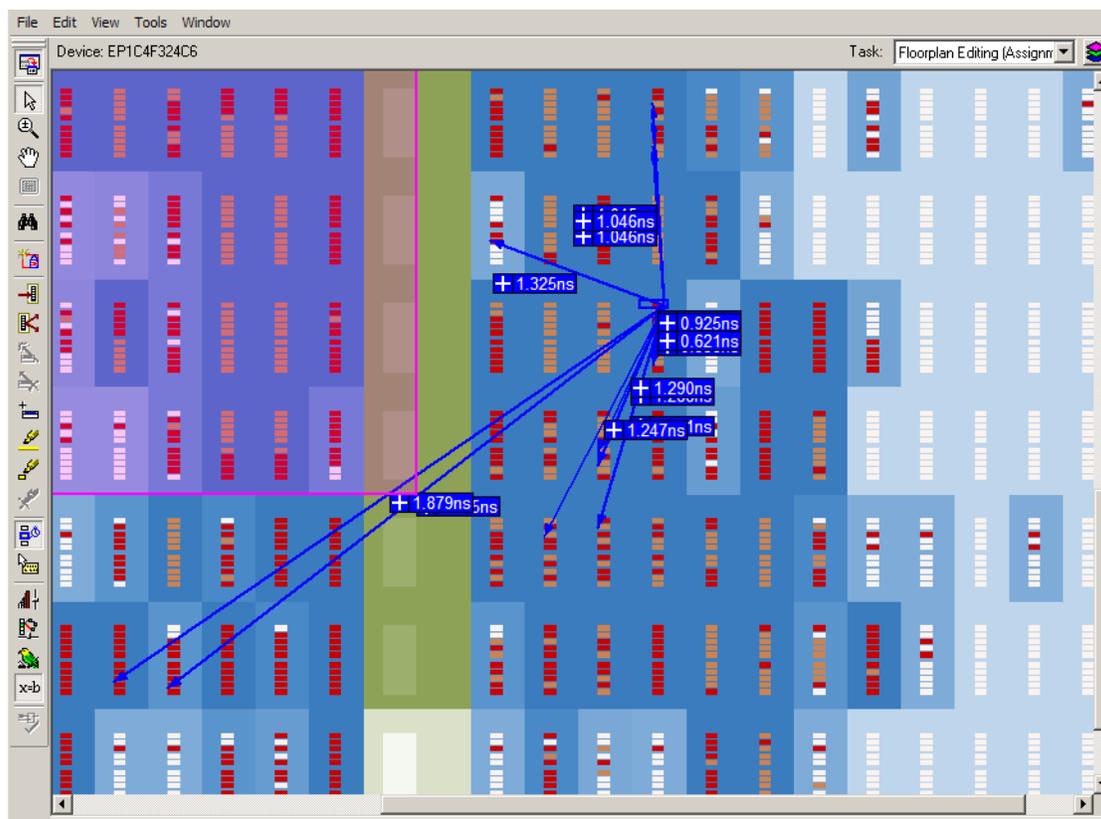
Figure 15-8. Generated Fan-In



Generating Immediate Fan-In and Fan-Out Connections

The ability to display immediate fan-in and fan-out connections enables you to view the resource that is the immediate fan-in or fan-out connection for the selected atom. For example, if you select a logic resource and choose to view the immediate fan-in for that resource, you can see the routing resource that drives the logic resource. You can generate immediate fan-in and fan-outs for all logic resources and routing resources. To remove the displayed connections from the screen, click the Clear Connections icon in the toolbar. Figure 15-9 shows the immediate fan-out connections for the selected resource.

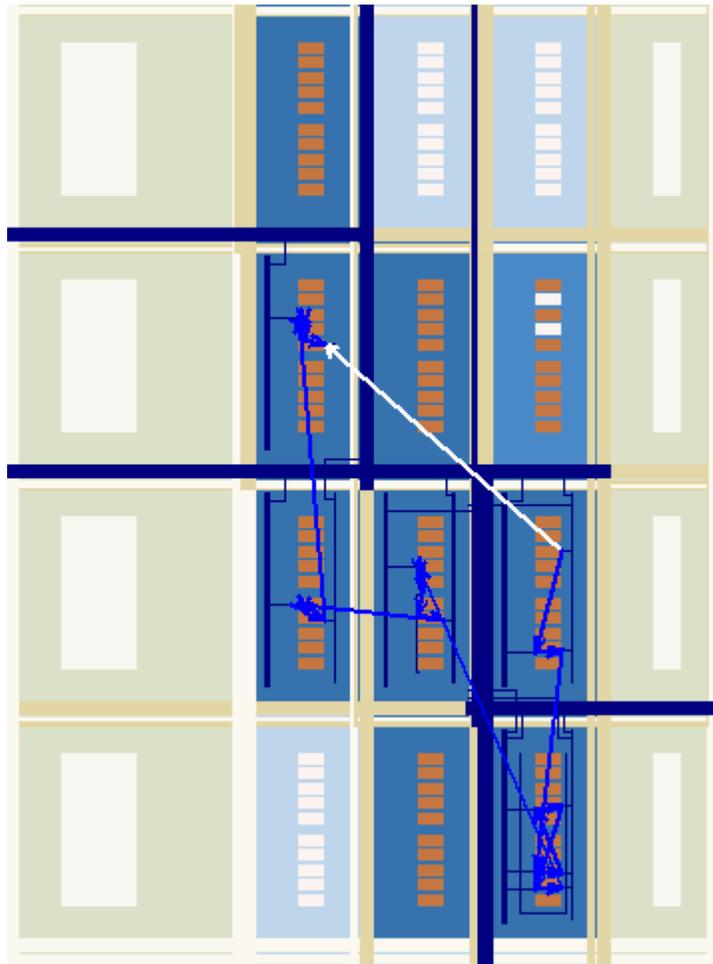
Figure 15-9. Immediate Fan-Out Connection



Highlight Routing

The **Highlight Routing** command enables you to highlight the routing resources used by a selected path or connection. [Figure 15-10](#) shows the routing resources in use between two logic elements.

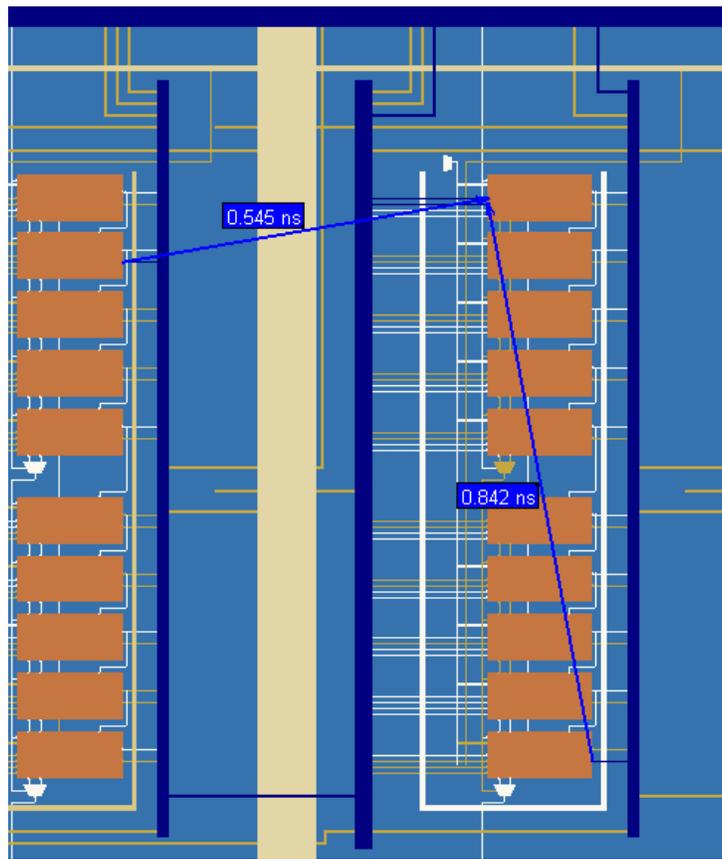
Figure 15-10. Highlight Routing



Show Delays

With the Show Delays command, you can view timing delays for connections displayed in the Chip Planner. For example, you can view the delay between two logic resources or between a logic resource and a routing resource. Figure 15-11 shows the delays between several logic elements.

Figure 15-11. Show Delays



Exploring Paths in the Chip Planner

You can use the Chip Planner to explore paths between logic elements. The following example uses the Chip Planner to traverse paths from the Timing Analysis report.

Locate Path from the Timing Analysis Report to the Chip Planner

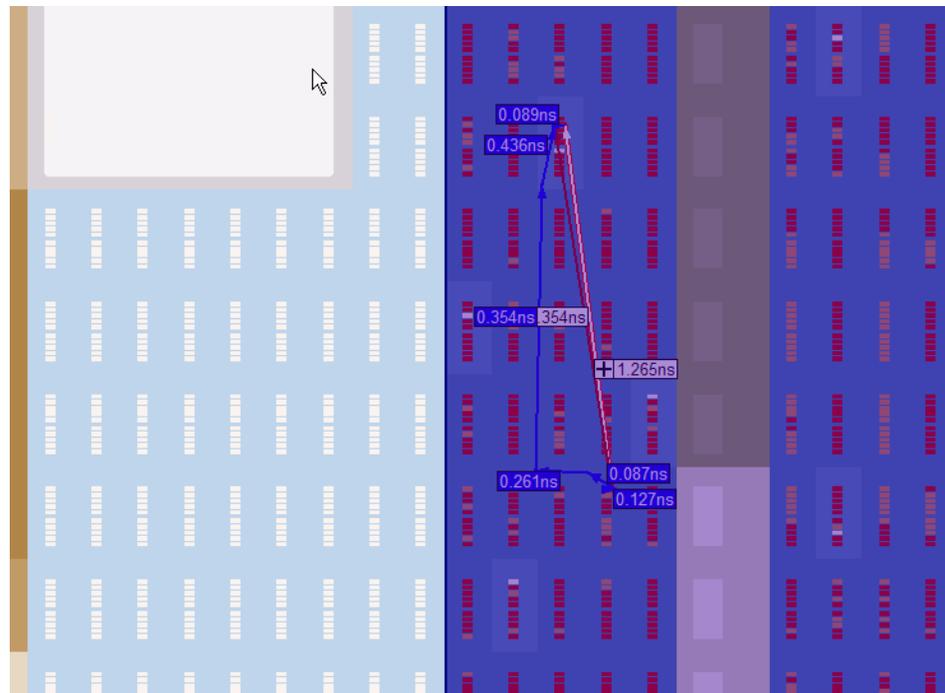
To locate a path from the Timing Analysis report to the Chip Planner, perform the following steps:

1. Select the path you want to locate.
2. Right-click the path in the Timing Analysis report, point to **Locate**, and click **Locate in Chip Planner (Floorplan & Chip Editor)**.
3. To view the routing resources taken for a path you have located in the Chip Planner, click the Highlight Routing icon in the Chip Planner toolbar, or from the View menu, click **Highlight Routing**.

Analyzing Connections for a Path

To determine the connections between items in the Chip Planner, click the Expand Connections/Paths icon on the toolbar. To add the timing delays between each connection, click the Show Delays icon on the toolbar. Figure 15-12 shows the connections for the selected path that are displayed in the Chip Planner. To see the constituent delays on the selected path, click on the “+” sign next to the path delay displayed in the Chip Planner.

Figure 15-12. Path Analysis

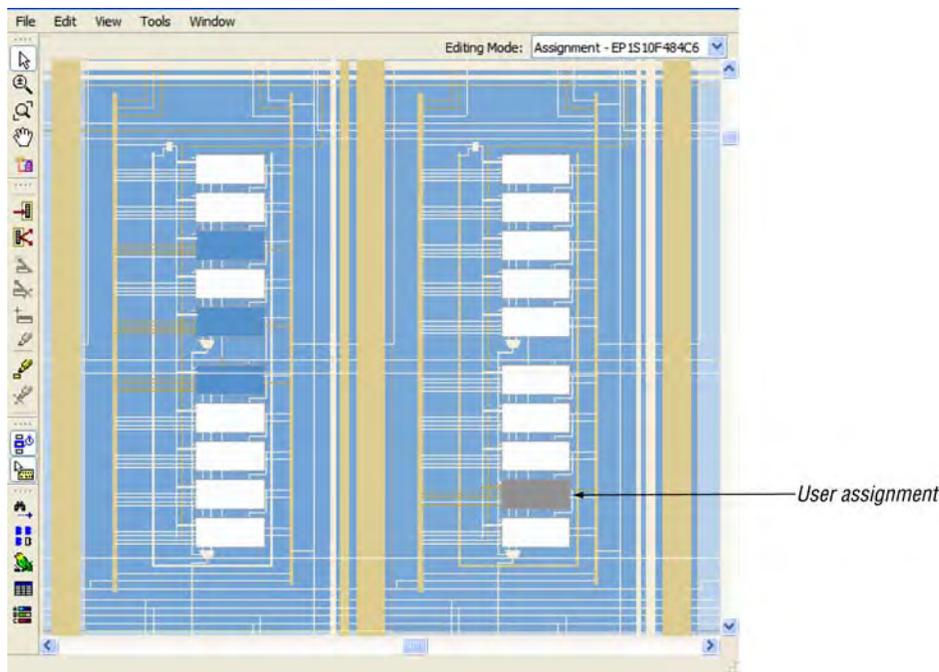


Viewing Assignments in the Chip Planner

You can view location assignments by selecting the appropriate layer set in the Chip Planner. To view location assignments, select the **Floorplan Editing** task or any custom task that displays block utilization, and the Assignment editing mode. See Figure 15-13.

The Chip Planner shows location assignments graphically, by displaying assigned resources in a particular color (gray, by default). You can create or move an assignment by dragging the selected resource to a new location.

Figure 15–13. Viewing Assignments in the Chip Planner



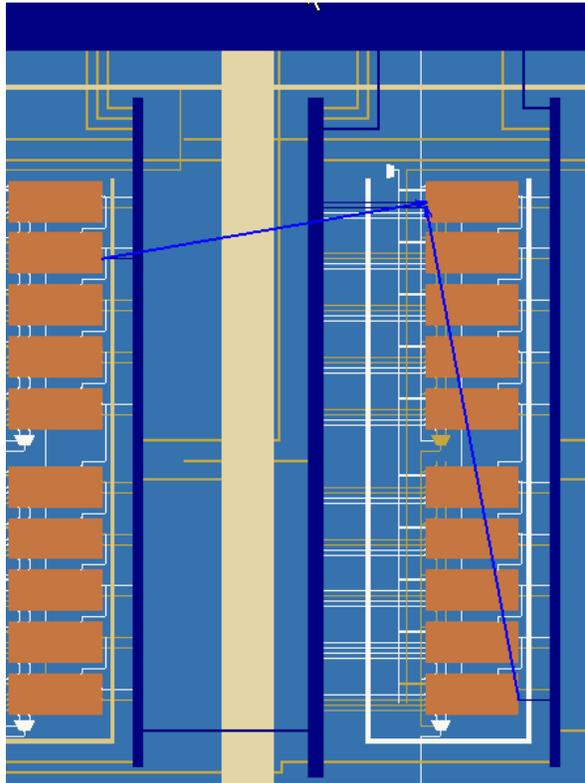
You can make node and pin location assignments and assignments to LogicLock regions and custom regions using the drag-and-drop method in the Chip Planner. The Fitter applies the assignments that you create during the next place-and-route operation.

- ④ For more information about managing assignments in the Chip Planner, refer to *Working With Assignments in the Chip Planner* in Quartus II Help.

Viewing Routing Channels for a Path in the Chip Planner

To determine the routing channels between connections, click the Highlight Routing icon on the toolbar. [Figure 15-14](#) shows the routing channels used for the selected path in the Chip Planner.

Figure 15-14. Highlight Routing

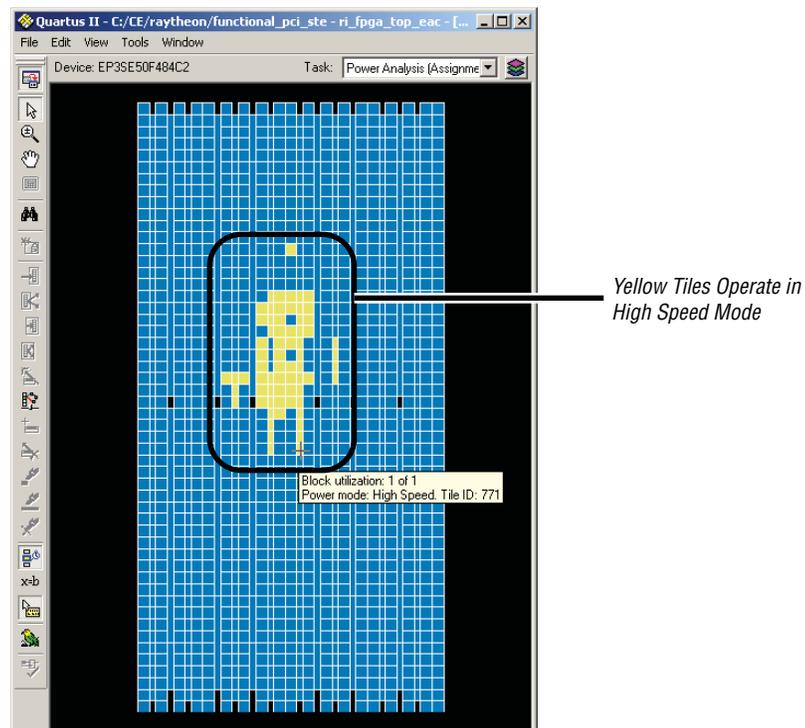


 You can view and edit resources in the FPGA using the Resource Property Editor. For more information, refer to the [Engineering Change Management with the Chip Planner](#) chapter in volume 2 of the *Quartus II Handbook*.

- To learn more about power analyses and optimizations in Stratix III devices, refer to *AN 437: Power Optimization in Stratix III FPGAs*. To learn more about power analyses and optimizations in Stratix IV devices, refer to *AN 514: Power Optimization in Stratix IV FPGAs*.

When you select the **Power** task in the Chip Planner for Stratix III, Stratix IV, or Stratix V devices, the Chip Planner displays low-power and high-speed tiles in contrasting colors; yellow tiles operate in a high-speed mode, while blue tiles operate in a low-power mode (see Figure 15-16). When you select the **Power** task, you can perform all floorplanner-related functions for this task; however, you cannot edit tiles to change the power mode.

Figure 15-16. Viewing High-Speed and Low Power Tiles in a Stratix III Device



Scripting Support

You can run procedures and specify the settings described in this chapter in a Tcl script. You can also run some procedures at a command prompt. For detailed information about scripting command options, refer to the Quartus II command-line and Tcl API Help browser. To run the Help browser, type the following command at the command prompt:

```
quartus_sh --qhelp ←
```

- Information about scripting command options is also available in *API Functions for Tcl* in Quartus II Help.

-  For more information about Tcl scripting, refer to the *Tcl Scripting* chapter in volume 2 of the *Quartus II Handbook*. For more information about command-line scripting, refer to the *Command-Line Scripting* chapter in volume 2 of the *Quartus II Handbook*.
-  For information about all settings and constraints in the Quartus II software, refer to the *Quartus II Settings File Manual*.

Initializing and Uninitializing a LogicLock Region

You must initialize the LogicLock data structures before creating or modifying any LogicLock regions and before executing any of the Tcl commands listed below.

Use the following Tcl command to initialize the LogicLock data structures:

```
initialize_logiclock
```

Use the following Tcl command to uninitialize the LogicLock data structures before closing your project:

```
uninitialize_logiclock
```

Creating or Modifying LogicLock Regions

Use the following Tcl command to create or modify a LogicLock region:

```
set_logiclock -auto_size true -floating true -region \ <my_region-name>
```

-  The command in the above example sets the size of the region to auto and the state to floating.

If you specify a region name that does not exist in the design, the command creates the region with the specified properties. If you specify the name of an existing region, the command changes all properties you specify and leaves unspecified properties unchanged.

For more information about creating LogicLock regions, refer to “[Creating LogicLock Regions](#)” on page 15-5.

Obtaining LogicLock Region Properties

Use the following Tcl command to obtain LogicLock region properties. This example returns the height of the region named `my_region`:

```
get_logiclock -region my_region -height
```

Assigning LogicLock Region Content

Use the following Tcl commands to assign or change nodes and entities in a LogicLock region. This example assigns all nodes with names matching `fifo*` to the region named `my_region`.

```
set_logiclock_contents -region my_region -to fifo*
```

You can also make path-based assignments with the following Tcl command:

```
set_logiclock_contents -region my_region -from fifo -to ram*
```

For more information about assigning LogicLock Region Content, refer to [“Assigning LogicLock Region Content” on page 15-7](#).

Save a Node-Level Netlist for the Entire Design into a Persistent Source File

Make the following assignments to cause the Quartus II Fitter to save a node-level netlist for the entire design into a **.vqm** file:

```
set_global_assignment -name LOGICLOCK_INCREMENTAL_COMPILE_ASSIGNMENT ON
set_global_assignment -name LOGICLOCK_INCREMENTAL_COMPILE_FILE <file name>
```

Any path specified in the file name is relative to the project directory. For example, specifying **atom_netlists/top.vqm** places **top.vqm** in the **atom_netlists** subdirectory of your project directory.

A **.vqm** file is saved in the directory specified at the completion of a full compilation.



The saving of a node-level netlist to a persistent source file is not supported for designs targeting newer devices such as Arria GX, Arria II, Cyclone III, MAX V, Stratix III, Stratix IV, or Stratix V.

Setting LogicLock Assignment Priority

Use the following Tcl code to set the priority for a LogicLock region's members. This example reverses the priorities of the LogicLock region in your design.

```
set reverse [list]
for each member [get_logiclock_member_priority] {
    set reverse [insert $reverse 0 $member]
}
set_logiclock_member_priority $reverse
```

Assigning Virtual Pins

Use the following Tcl command to turn on the virtual pin setting for a pin called **my_pin**:

```
set_instance_assignment -name VIRTUAL_PIN ON -to my_pin
```

For more information about assigning virtual pins, refer to [“Virtual Pins” on page 15-11](#).



For more information about Tcl scripting, refer to the [Tcl Scripting](#) chapter in volume 2 of the *Quartus II Handbook*.

Conclusion

Design floorplan analysis is a valuable method for achieving timing closure and optimal performance in highly complex designs. With analysis capability, the Quartus II Chip Planner helps you close timing quickly on your designs. Using the Chip Planner together with LogicLock and Incremental Compilation enables you to compile your designs hierarchically, preserving the timing results from individual compilation runs. You can use LogicLock regions as part of an incremental

compilation methodology to improve your productivity. You can also include a module in one or more projects while maintaining performance and reducing development costs and time-to-market. LogicLock region assignments give you complete control over logic and memory placement to improve the performance of non-hierarchical designs as well.

Document Revision History

Table 15-2 shows the revision history for this chapter.

Table 15-2. Document Revision History

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Updated for the 10.1 release.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Updated device support information ■ Removed references to Timing Closure Floorplan; removed “Design Analysis Using the Timing Closure Floorplan” section ■ Added links to online Help topics ■ Added “Using LogicLock Regions with the Design Partition Planner” section ■ Updated “Viewing Critical Paths” section ■ Updated several graphics ■ Updated format of Document revision History table
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Updated supported device information throughout ■ Removed deprecated sections related to the Timing Closure Floorplan for older device families. (For information on using the Timing Closure Floorplan with older device families, refer to previous versions of the <i>Quartus II Handbook</i>, available in the Quartus II Handbook Archive.) ■ Updated “Creating Non-Rectangular LogicLock Regions” section ■ Added “Selected Elements Window” section ■ Updated table 12-1
May 2008	8.0.0	<ul style="list-style-type: none"> ■ Updated the following sections: <ul style="list-style-type: none"> ■ “Chip Planner Tasks and Layers” ■ “LogicLock Regions” ■ “Back-Annotating LogicLock Regions” ■ “LogicLock Regions in the Timing Closure Floorplan” ■ Added the following sections: <ul style="list-style-type: none"> ■ “Reserve LogicLock Region” ■ “Creating Non-Retangular LogicLock Regions” ■ “Viewing Available Clock Networks in the Device” ■ Updated Table 10-1 ■ Removed the following sections: <ul style="list-style-type: none"> ■ Reserve LogicLock Region Design Analysis Using the Timing Closure Floorplan

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

The Quartus® II software offers physical synthesis optimizations to improve your design beyond the optimization performed in the normal course of the Quartus II compilation flow.

Physical synthesis optimizations can help improve the performance of your design regardless of the synthesis tool used, although the effect of physical synthesis optimizations depends on the structure of your design.

Netlist optimization options work with the atom netlist of your design, which describes a design in terms of Altera®-specific primitives. An atom netlist file can be an Electronic Design Interchange Format (.edf) file or a Verilog Quartus Mapping (.vqm) file generated by a third-party synthesis tool, or a netlist used internally by the Quartus II software. Physical synthesis optimizations are applied at different stages of the Quartus II compilation flow, either during synthesis, fitting, or both.

This chapter explains how the physical synthesis optimizations in the Quartus II software can modify your design's netlist to improve the quality of results. This chapter also provides information about preserving compilation results through back-annotation and writing out a new netlist, and provides guidelines for applying the various options.



Because the node names for primitives in the design can change when you use physical synthesis optimizations, you should evaluate whether your design flow requires fixed node names. If you use a verification flow that might require fixed node names, such as the SignalTap® II Logic Analyzer, formal verification, or the LogicLock based optimization flow (for legacy devices), you must turn off physical synthesis options.

WYSIWYG Primitive Resynthesis

If you use a third-party tool to synthesize your design, use the **Perform WYSIWYG primitive resynthesis** option to apply optimizations to the synthesized netlist.

The **Perform WYSIWYG primitive resynthesis** option directs the Quartus II software to un-map the logic elements (LEs) in an atom netlist to logic gates, and then re-map the gates back to Altera-specific primitives. Third-party synthesis tools generate either an .edf or .vqm atom netlist file using Altera-specific primitives. When you turn on the **Perform WYSIWYG primitive resynthesis** option, the Quartus II software can work on different techniques specific to the device architecture during the re-mapping process. This feature re-maps the design using the **Optimization Technique** specified for your project (**Speed, Area, or Balanced**).



The **Perform WYSIWYG primitive resynthesis** option has no effect if you are using Quartus II integrated synthesis to synthesize your design.

To turn on the **Perform WYSIWYG primitive resynthesis** option, perform the following steps:

1. On the Assignments menu, click **Settings**. The **Settings** dialog box appears.
2. In the **Category** list, select **Analysis and Synthesis Settings**. The **Analysis & Synthesis Settings** page appears.
3. Turn on **Perform WYSIWYG Primitive Resynthesis**, and click **OK**.

If you want to perform WYSIWYG resynthesis on only a portion of your design, you can use the Assignment Editor to assign the **Perform WYSIWYG primitive resynthesis** logic option to a lower-level entity in your design. This logic option is available for all Altera devices supported by the Quartus II software except MAX 3000 and MAX 7000 devices.

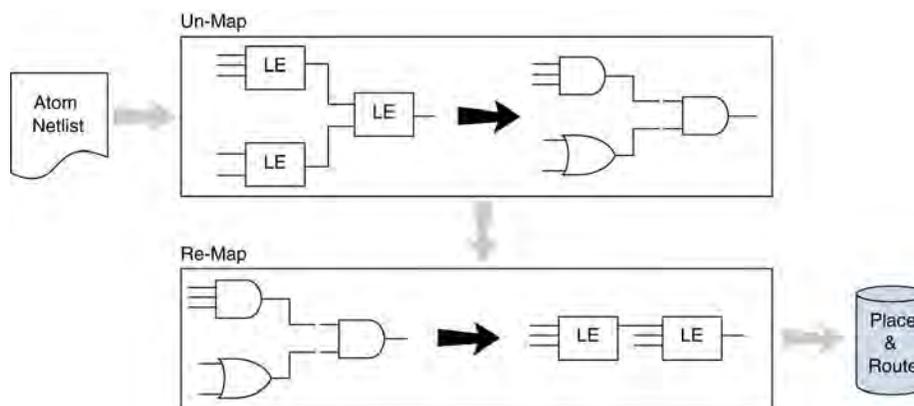
The results of the remapping depend on the **Optimization Technique** you choose. To select an **Optimization Technique**, perform the following steps:

1. In the **Category** list, select **Analysis & Synthesis Settings**. The **Analysis & Synthesis Settings** page appears.
2. Under **Optimization Technique**, select **Speed**, **Area**, or **Balanced** to specify how the Quartus II technology mapper optimizes the design. The **Balanced** setting is the default for many Altera device families; this setting optimizes the timing critical parts of the design for speed and the rest of the design for area.
3. Click **OK**.

Refer to the *Quartus II Integrated Synthesis* chapter in volume 1 of the *Quartus II Handbook* for details on the Optimization Technique option.

Figure 16-1 shows the Quartus II software flow for the WYSIWYG primitive resynthesis feature.

Figure 16-1. WYSIWYG Primitive Resynthesis



The **Perform WYSIWYG primitive resynthesis** option unmaps and remaps only logic cells, also referred to as LCELL or LE primitives, and regular I/O primitives (which may contain registers). Double data rate (DDR) I/O primitives, memory primitives, digital signal processing (DSP) primitives, and logic cells in carry/cascade chains are not remapped. Logic specified in an encrypted **.vqm** file or an **.edf** file, such as third-party intellectual property (IP), is not touched.

The **Perform WYSIWYG primitive resynthesis** option can change node names in the **.vqm** file or **.edf** file from your third-party synthesis tool, because the primitives in the atom netlist are broken apart and then remapped by the Quartus II software. The remapping process removes duplicate registers, but registers that are not removed retain the same name after remapping.

Any nodes or entities that have the **Netlist Optimizations** logic option set to **Never Allow** are not affected during WYSIWYG primitive resynthesis. You can use the Assignment Editor to apply the **Netlist Optimizations** logic option. This option disables WYSIWYG resynthesis for parts of your design.



Primitive node names are specified during synthesis. When netlist optimizations are applied, node names might change because primitives are created and removed. HDL attributes applied to preserve logic in third-party synthesis tools cannot be maintained because those attributes are not written into the atom netlist read by the Quartus II software.

If you use the Quartus II software to synthesize, you can use the **Preserve Register (preserve)** and **Keep Combinational Logic (keep)** attributes to maintain certain nodes in the design.



For more information about using these attributes during synthesis in the Quartus II software, refer to the *Quartus II Integrated Synthesis* chapter in volume 1 of the *Quartus II Handbook*.

Performing Physical Synthesis Optimizations

The Quartus II design flow involves separate steps of synthesis and fitting. The synthesis step optimizes the logical structure of a circuit for area, speed, or both. The Fitter then places and routes the logic cells to ensure critical portions of logic are close together and use the fastest possible routing resources. While you are using this push-button flow, the synthesis stage is unable to anticipate the routing delays seen in the Fitter. Because routing delays are a significant part of the typical critical path delay, the physical synthesis optimizations available in the Quartus II software take those routing delays into consideration and focus timing-driven optimizations at those parts of the design. This tight integration of the fitting and synthesis processes is known as physical synthesis.

The following sections describe the physical synthesis optimizations available in the Quartus II software, and how they can help improve your performance results. Physical synthesis optimization options can be used with Arria series, Cyclone, HardCopy, and Stratix series device families.

If you are migrating your design to a HardCopy II device, you can target physical synthesis optimizations to the FPGA architecture in the FPGA-first flow or to the HardCopy II architecture in the HardCopy-first flow. The optimizations are mapped to the other device architecture during the migration process.



You cannot target optimizations to both device architectures individually because doing so results in a different post-fitting netlist for each device.



For more information about physical synthesis optimizations, refer to *Physical Synthesis Optimizations Page (Settings Dialog Box)* in Quartus II Help. For more information about using physical synthesis with HardCopy devices, refer to the *Quartus II Support for HardCopy Series Devices* chapter in volume 1 of the *Quartus II Handbook*.

You can choose the physical synthesis optimization options you want for your design during synthesis and fitting in the **Physical Synthesis Optimizations** page under the **Compilation Process Settings** page in the **Settings** dialog box. The settings include optimizations for improving performance and fitting in the selected device.

You can also set the effort level for physical synthesis optimizations. Normally, physical synthesis optimizations increase the compilation time; however, you can select the **Fast** effort level if you want to limit the increase in compilation time. When you select the **Fast** effort level, the Quartus II software performs limited register retiming operations during fitting. The **Extra** effort level runs additional algorithms to get the best circuit performance, but results in increased compilation time.

To optimize performance, the following options are available:

- **Perform physical synthesis for combinational logic**
- **Perform register retiming**
- **Perform automatic asynchronous signal pipelining**
- **Perform register duplication**

To optimize for better fitting, you can choose from the following options:

- **Perform physical synthesis for combinational logic**
- **Perform logic to memory mapping**

To view and modify the physical synthesis optimization options, perform the following steps:

1. On the Assignments menu, click **Settings**. The **Settings** dialog box appears.
2. In the **Category** list, select **Physical Synthesis Optimizations** under **Compilation Process Settings**. The **Physical Synthesis Optimizations** page appears.
3. Specify the options for performing physical synthesis optimizations.

Some physical synthesis options affect only registered logic and some options affect only combinational logic. Select options based on whether you want to keep the registers intact or not. For example, if your verification flow involves formal verification, you might have to keep the registers intact.

All Physical Synthesis optimizations write results to the **Netlist Optimizations** report, which provides a list of atom netlist files that were modified, created, and deleted during physical synthesis. To access the **Netlist Optimizations** report, perform the following steps:

1. On the Processing menu, click **Compilation Report**.
2. In the **Compilation Report** list, select **Netlist Optimizations** under **Fitter**.

Similarly, physical synthesis optimizations performed during synthesis write results to the synthesis report. To access this report, perform the following steps:

1. On the Processing menu, click **Compilation Report**.
2. In the **Compilation Report** list, select **Analysis & Synthesis**.
3. In the **Optimization Results** folder, select **Netlist Optimizations**. The **Physical Synthesis Netlist Optimizations** table appears, listing the physical synthesis netlist optimizations performed during synthesis.

Nodes or entities that have the **Netlist Optimizations** logic option set to **Never Allow** are not affected by the physical synthesis algorithms. You can use the Assignment Editor to apply the **Netlist Optimizations** logic option. Use this option to disable physical synthesis optimizations for parts of your design.

Automatic Asynchronous Signal Pipelining

The **Perform automatic asynchronous signal pipelining** option on the **Physical Synthesis Optimizations** page in the **Compilation Process Settings** section of the **Settings** dialog box allows the Quartus II Fitter to perform automatic insertion of pipeline stages for asynchronous clear and asynchronous load signals during fitting when these signals negatively affect performance. You can use this option if asynchronous control signal recovery and removal times are not achieving their requirements.

The **Perform automatic asynchronous signal pipelining** option improves performance for designs in which asynchronous signals in very fast clock domains cannot be distributed across the chip fast enough due to long global network delays. This optimization performs automatic pipelining of these signals, while attempting to minimize the total number of registers inserted.



The **Perform automatic asynchronous signal pipelining** option adds registers to nets driving the asynchronous clear or asynchronous load ports of registers. These additional registers add register delays (adds latency) to the reset, adding the same number of register delays for each destination using the reset. The additional register delays can change the behavior of the signal in the design; therefore, you should use this option only if additional latency on the reset signals does not violate any design requirements. This option also prevents the promotion of signals to global routing resources.

The Quartus II software performs automatic asynchronous signal pipelining only if **Enable Recovery/Removal analysis** is turned on. If you use the TimeQuest Timing Analyzer, **Enable Recovery/Removal analysis** is turned on by default. Pipelining is allowed only on asynchronous signals that have the following properties:

- The asynchronous signal is synchronized to a clock (a synchronization register drives the signal)

- The asynchronous signal fans-out only to asynchronous control ports of registers
- The Quartus II software does not perform automatic asynchronous signal pipelining on asynchronous signals that have the **Netlist Optimization** logic option set to **Never Allow**.

Physical Synthesis for Combinational Logic

To optimize the design and reduce delay along critical paths, you can turn on the **Perform physical synthesis for combinational logic** option, which swaps the look-up table (LUT) ports within LEs so that the critical path has fewer layers through which to travel. The **Perform physical synthesis for combinational logic** option also allows the duplication of LUTs to enable further optimizations on the critical path.

- ② For more information about using the **Perform physical synthesis for combinational logic** option, refer to *Physical Synthesis Optimizations Page (Settings Dialog Box)* and to *Setting Up and Running the Fitter* in Quartus II Help.

The **Perform physical synthesis for combinational logic** option affects only combinational logic in the form of LUTs. These transformations might occur during the synthesis stage or the Fitter stage during compilation. The registers contained in the affected logic cells are not modified. Inputs into memory blocks, DSP blocks, and I/O elements (IOEs) are not swapped.

The Quartus II software does not perform combinational optimization on logic cells that have the following properties:

- Are part of a chain
- Drive global signals
- Are constrained to a single logic array block (LAB) location
- Have the **Netlist Optimizations** option set to **Never Allow**

If you want to consider logic cells with any of these conditions for physical synthesis, you can override these rules by setting the **Netlist Optimizations** logic option to **Always Allow** on a given set of nodes.

Physical Synthesis for Registers—Register Duplication

The **Perform register duplication** option on the **Physical Synthesis Optimizations** page in the **Compilation Process Settings** section of the **Settings** dialog box allows the Quartus II Fitter to duplicate registers based on Fitter placement information. You can also duplicate combinational logic when this option is enabled. A logic cell that fans out to multiple locations can be duplicated to reduce the delay of one path without degrading the delay of another. The new logic cell can be placed closer to critical logic without affecting the other fan-out paths of the original logic cell.

- ② For more information about the **Perform register duplication** option, refer to *Physical Synthesis Optimizations Page (Settings Dialog Box)* and to *Setting Up and Running the Fitter* in Quartus II Help.

The Quartus II software does not perform register duplication on logic cells that have the following properties:

- Are part of a chain
- Contain registers that drive asynchronous control signals on another register
- Contain registers that drive the clock of another register
- Contain registers that drive global signals
- Contain registers that are constrained to a single LAB location
- Contain registers that are driven by input pins without a t_{SU} constraint
- Contain registers that are driven by a register in another clock domain
- Are considered virtual I/O pins
- Have the **Netlist Optimizations** option set to **Never Allow**

 For more information about virtual I/O pins, refer to the *Analyzing and Optimizing the Design Floorplan* chapter in volume 2 of the *Quartus II Handbook*.

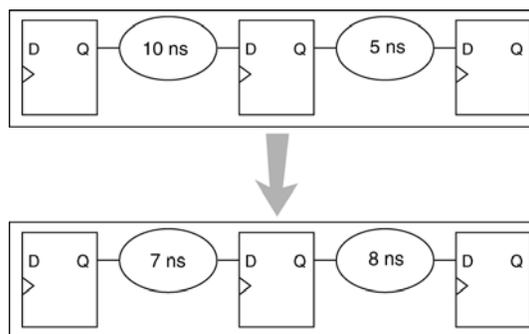
If you want to consider logic cells that meet any of these conditions for physical synthesis, you can override these rules by setting the **Netlist Optimizations** logic option to **Always Allow** on a given set of nodes.

Physical Synthesis for Registers—Register Retiming

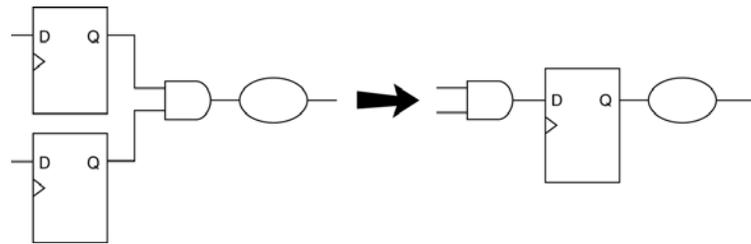
The **Perform Register Retiming** option enables the movement of registers across combinational logic, allowing the Quartus II software to trade off the delay between timing-critical paths and non-critical paths. Register retiming can be done during Quartus II integrated synthesis or during the Fitter stages of design compilation.

Figure 16-2 shows an example of register retiming in which the 10-ns critical delay is reduced by moving the register relative to the combinational logic.

Figure 16-2. Register Retiming Diagram



Retiming can create multiple registers at the input of a combinational block from a register at the output of a combinational block. In this case, the new registers have the same clock and clock enable. The asynchronous control signals and power-up level are derived from previous registers to provide equivalent functionality. Retiming can also combine multiple registers at the input of a combinational block to a single register (Figure 16-3).

Figure 16-3. Combining Registers with Register Retiming

To move registers across combinational logic to balance timing, perform the following steps:

1. On the Assignments menu, click **Settings**. The **Settings** dialog box appears.
 2. In the **Category** list, select **Physical Synthesis Optimizations under Compilation Process Settings**. The **Physical Synthesis Optimizations** page appears.
 3. Specify your preferred option under **Optimize for performance (physical synthesis)** and **Effort level**.
 4. Click **OK**.
- ② For more information about the **Optimize for performance (physical synthesis)** options and effort levels, refer to [Physical Synthesis Optimizations Page \(Settings Dialog Box\)](#) in Quartus II Help.

If you want to prevent register movement during register retiming, you can set the **Netlist Optimizations** logic option to **Never Allow**. You can apply this option to either individual registers or entities in the design using the Assignment Editor.

In digital circuits, synchronization registers are instantiated on cross clock domain paths to reduce the possibility of metastability. The Quartus II software detects such synchronization registers and does not move them, even if register retiming is turned on.

The following sets of registers are not moved during register retiming:

- Both registers in a direct connection from input pin-to-register-to-register if both registers have the same clock and the first register does not fan-out to anywhere else. These registers are considered synchronization registers.
- Both registers in a direct connection from register-to-register if both registers have the same clock, the first register does not fan out to anywhere else, and the first register is fed by another register in a different clock domain (directly or through combinational logic). These registers are considered synchronization registers.

The Quartus II software assumes that a synchronization register chain consists of two registers. If your design has synchronization register chains with more than two registers, you must indicate the number of registers in your synchronization chains so that they are not affected by register retiming. To do this, perform the following steps:

1. On the Assignments menu, click **Settings**. The **Settings** dialog box appears.
2. In the **Category** list, select **Analysis & Synthesis Settings**. The **Analysis & Synthesis Setting** page appears.

3. Click **More Settings**. The **More Analysis & Synthesis Settings** dialog box appears.
4. In the **Name** list, select **Synchronization Register Chain Length** and modify the setting to match the synchronization register length used in your design. If you set a value of 1 for the **Synchronization Register Chain Length**, it means that any registers connected to the first register in a register-to-register connection can be moved during retiming. A value of $n > 1$ means that any registers in a sequence of length 1, 2, ... n are not moved during register retiming.

The Quartus II software does not perform register retiming on logic cells that have the following properties:

- Are part of a cascade chain
- Contain registers that drive asynchronous control signals on another register
- Contain registers that drive the clock of another register
- Contain registers that drive a register in another clock domain
- Contain registers that are driven by a register in another clock domain

 The Quartus II software does not usually retime registers across different clock domains; however, if you use the Classic Timing Analyzer and specify a global f_{MAX} requirement, the Quartus II software interprets all clocks as related. Consequently, the Quartus II software might try to retime register-to-register paths associated with different clocks.

To avoid this circumstance, provide individual f_{MAX} requirements to each clock when using Classic Timing Analysis. When you constrain each clock individually, the Quartus II software assumes no relationship between different clock domains and considers each clock domain to be asynchronous to other clock domains; hence no register-to-register paths crossing clock domains are retimed.

When you use the TimeQuest Timing Analyzer, register-to-register paths across clock domains are never retimed, because the TimeQuest Timing Analyzer treats all clock domains as asynchronous to each other unless they are intentionally grouped.

- Contain registers that are constrained to a single LAB location
- Contain registers that are connected to SERDES
- Are considered virtual I/O pins
- Registers that have the **Netlist Optimizations** logic option set to **Never Allow**

 For more information about virtual I/O pins, refer to the *Analyzing and Optimizing the Design Floorplan* chapter in volume 2 of the *Quartus II Handbook*.

If you want to consider logic cells that meet any of these conditions for physical synthesis, you can override these rules by setting the **Netlist Optimizations** logic option to **Always Allow** on a given set of registers.

Preserving Your Physical Synthesis Results

The Quartus II software generates the same results on every compilation for the same source code and settings on a given system, hence you do not need to preserve your results from compilation to compilation. When you make changes to the source code or to the settings, you usually get the best results by allowing the software to compile without using previous compilation results or location assignments. In some cases, if you avoid performing analysis and synthesis or `quartus_map`, and run the Fitter or another desired Quartus II executable instead, you can skip the synthesis stage of the compilation.

When you use the Quartus II incremental compilation flow, you can preserve synthesis results for a particular partition of your design by choosing a netlist type of post-synthesis. If you want to preserve fitting results between compilation runs, choose a netlist type of post-fit during incremental compilation.

The rest of this section is relevant only for those designs using older devices that do not support incremental compilation.

 For information about the incremental compilation design methodology, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*, and to *About Incremental Compilation* in Quartus II Help.

You can preserve the resulting nodes from physical synthesis in older devices that do not support incremental compilation. You might need to preserve nodes if you use the LogicLock flow to back-annotate placement, import one design into another, or both. For all device families that support incremental compilation, use that feature to preserve results.

To preserve the nodes from Quartus II physical synthesis optimization options for older devices that do not support incremental compilation (such as Max II devices), perform the following steps:

1. On the Assignments menu, click **Settings**. The **Settings** dialog box appears.
2. In the **Category** list, select **Compilation Process Settings**. The **Compilation Process Settings** page appears.
3. Turn on **Save a node-level netlist of the entire design into a persistent source file**. This setting is not available for Cyclone III, Stratix III, and newer devices.
4. Click **OK**.

The **Save a node-level netlist of the entire design into a persistent source file** option saves your final results as an atom-based netlist in `.vqm` file format. By default, the Quartus II software places the `.vqm` file in the `atom_netlists` directory under the current project directory. To create a different `.vqm` file using different Quartus II settings, in the **Compilation Process Settings** page, change the **File name** setting.

If you use the physical synthesis optimizations and want to lock down the location of all LEs and other device resources in the design with the **Back-Annotate Assignments** command, a **.vqm** file netlist is required. The **.vqm** file preserves the changes that you made to your original netlist. Because the physical synthesis optimizations depend on the placement of the nodes in the design, back-annotating the placement changes the results from physical synthesis. Changing the results means that node names are different, and your back-annotated locations are no longer valid.

You should not use a Quartus II-generated **.vqm** file or back-annotated location assignments with physical synthesis optimizations unless you have finalized the design. Making any changes to the design invalidates your physical synthesis results and back-annotated location assignments. If you require changes later, use the new source HDL code as your input files, and remove the back-annotated assignments corresponding to the Quartus II-generated **.vqm** file.

To back-annotate logic locations for a design that was compiled with physical synthesis optimizations, first create a **.vqm** file. When recompiling the design with the hard logic location assignments, use the new **.vqm** file as the input source file and turn off the physical synthesis optimizations for the new compilation.

If you are importing a **.vqm** file and back-annotated locations into another project that has any **Netlist Optimizations** turned on, you must apply the **Never Allow** constraint to make sure node names don't change; otherwise, the back-annotated location or LogicLock assignments are invalid.

 For newer devices, such as the Arria, Cyclone, or Stratix series, use incremental compilation to preserve compilation results instead of using logic back-annotation.

Physical Synthesis Options for Fitting

The Quartus II software provides physical synthesis optimization options for improving fitting results. To access these options, perform the following steps:

1. On the Assignments menu, click **Settings**. The **Settings** dialog box appears.
2. In the **Category** list, select **Physical Synthesis Optimizations under Compilation Process Settings**. The **Physical Synthesis Optimizations** page appears.
3. Under **Optimize for fitting** (physical synthesis for density), there are two physical synthesis options available to improve fitting your design in the target device: **Physical synthesis for combinational logic** and **Perform logic to memory mapping** (Table 16-1).

Table 16-1. Physical Synthesis Optimizations Options

Option	Function
Physical Synthesis for Combinational Logic	When you select this option, the Fitter detects duplicate combinational logic and optimizes combinational logic to improve the fit.
Perform Logic to Memory Mapping	When you select this option, the Fitter can remap registers and combinational logic in your design into unused memory blocks and achieves a fit.

 For more information about physical synthesis optimization options, refer to *Physical Synthesis Optimizations Page (Settings Dialog Box)* in Quartus II Help.

Applying Netlist Optimization Options

The improvement in performance when using netlist optimizations is design dependent. If you have restructured your design to balance critical path delays, netlist optimizations might yield minimal improvement in performance. You may have to experiment with available options to see which combination of settings works best for a particular design. Refer to the messages in the compilation report to see the magnitude of improvement with each option, and to help you decide whether you should turn on a given option or specific effort level.

Turning on more netlist optimization options can result in more changes to the node names in the design; bear this in mind if you are using a verification flow, such as the SignalTap II Logic Analyzer or formal verification that requires fixed or known node names.

Applying all of the physical synthesis options at the **Extra** effort level generally produces the best results for those options, but adds significantly to the compilation time. You can also use the **Physical synthesis effort level** options to decrease the compilation time. The WYSIWYG primitive resynthesis option does not add much compilation time relative to the overall design compilation time.

To find the best results, you can use the Quartus II Design Space Explorer (DSE) to apply various sets of netlist optimization options.

- ❓ For more information about DSE, refer to *About Design Space Explorer* in Quartus II Help.

Scripting Support

You can run procedures and make settings described in this chapter in a Tcl script. You can also run some procedures at a command prompt. For detailed information about scripting command options, refer to the Quartus II Command-Line and Tcl API Help browser. To run the Help browser, type the following command at the command prompt:

```
quartus_sh --qhelp ←
```

-  For more information about Tcl scripting, refer to the *Tcl Scripting* chapter in volume 2 of the *Quartus II Handbook* and *API Functions for Tcl* in Quartus II Help. Refer to the *Quartus II Settings File Manual* for information about all settings and constraints in the Quartus II software. For more information about command-line scripting, refer to the *Command-Line Scripting* chapter in volume 2 of the *Quartus II Handbook*.

You can specify many of the options described in this section on either an instance or global level, or both.

Use the following Tcl command to make a global assignment:

```
set_global_assignment -name <QSF variable name> <value> ←
```

Use the following Tcl command to make an instance assignment:

```
set_instance_assignment -name <QSF variable name> <value> \  
-to <instance name> ←
```

Synthesis Netlist Optimizations

Table 16-2 lists the Quartus II Settings File (.qsf) variable names and applicable values for the settings discussed in “WYSIWYG Primitive Resynthesis” on page 16-1. The .qsf file variable name is used in the Tcl assignment to make the setting along with the appropriate value. The **Type** column indicates whether the setting is supported as a global setting, an instance setting, or both.

Table 16-2. Synthesis Netlist Optimizations and Associated Settings

Setting Name	Quartus II Settings File Variable Name	Values	Type
Perform WYSIWYG Primitive Resynthesis	ADV_NETLIST_OPT_SYNTH_WYSIWYG_REMAP	ON, OFF	Global, Instance
Optimization Technique	<Device Family Name>_OPTIMIZATION_TECHNIQUE	AREA, SPEED, BALANCED	Global, Instance
Power-Up Don't Care	ALLOW_POWER_UP_DONT_CARE	ON, OFF	Global
Save a node-level netlist into a persistent source file	LOGICLOCK_INCREMENTAL_COMPILE_ASSIGNMENT	ON, OFF	Global
	LOGICLOCK_INCREMENTAL_COMPILE_FILE	<file name>	
Allow Netlist Optimizations	ADV_NETLIST_OPT_ALLOWED	"ALWAYS ALLOW", DEFAULT, "NEVER ALLOW"	Instance

Physical Synthesis Optimizations

Table 16-3 lists the .qsf file variable name and applicable values for the settings discussed in “Performing Physical Synthesis Optimizations” on page 16-3. The .qsf file variable name is used in the Tcl assignment to make the setting, along with the appropriate value. The **Type** column indicates whether the setting is supported as a global setting, an instance setting, or both.

Table 16-3. Physical Synthesis Optimizations and Associated Settings (Part 1 of 2)

Setting Name	Quartus II Settings File Variable Name	Values	Type
Physical Synthesis for Combinational Logic	PHYSICAL_SYNTHESIS_COMBO_LOGIC	ON, OFF	Global
Automatic Asynchronous Signal Pipelining	PHYSICAL_SYNTHESIS_ASYNCHRONOUS_SIGNAL_PIPELINING	ON, OFF	Global
Perform Register Duplication	PHYSICAL_SYNTHESIS_REGISTER_DUPLICATION	ON, OFF	Global
Perform Register Retiming	PHYSICAL_SYNTHESIS_REGISTER_RETIMING	ON, OFF	Global
Power-Up Don't Care	ALLOW_POWER_UP_DONT_CARE	ON, OFF	Global, Instance
Power-Up Level	POWER_UP_LEVEL	HIGH, LOW	Instance

Table 16-3. Physical Synthesis Optimizations and Associated Settings (Part 2 of 2)

Setting Name	Quartus II Settings File Variable Name	Values	Type
Allow Netlist Optimizations	ADV_NETLIST_OPT_ALLOWED	"ALWAYS ALLOW", DEFAULT, "NEVER ALLOW"	Instance
Save a node-level netlist into a persistent source file	LOGICLOCK_INCREMENTAL_COMPILE_ASSIGNMENT	ON, OFF	Global
	LOGICLOCK_INCREMENTAL_COMPILE_FILE	<file name>	

Incremental Compilation

For information about scripting and command line usage for incremental compilation as mentioned in [“Preserving Your Physical Synthesis Results” on page 16-10](#), refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.

Back-Annotating Assignments

You can use the `logiclock_back_annotate` Tcl command to back-annotate resources in your design. This command can back-annotate resources in LogicLock regions, and resources in designs without LogicLock regions.

For more information about back-annotating assignments, refer to [“Preserving Your Physical Synthesis Results” on page 16-10](#).

The following Tcl command back-annotates all registers in your design:

```
logiclock_back_annotate -resource_filter "REGISTER"
```

The `logiclock_back_annotate` command is in the `backannotate` package.

Conclusion

Physical synthesis optimizations restructure and optimize your design netlist. You can take advantage of these Quartus II netlist optimizations to help improve your quality of results.

Document Revision History

[Table 16-4](#) shows the revision history for this chapter.

Table 16-4. Document Revision History (Part 1 of 2)

Date	Version	Changes
December 2010	10.0.1	Template update.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Added links to Quartus II Help in several sections. ■ Removed Referenced Documents section. ■ Reformatted Document Revision History

Table 16-4. Document Revision History (Part 2 of 2)

Date	Version	Changes
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Added information to “Physical Synthesis for Registers—Register Retiming” ■ Added information to “Applying Netlist Optimization Options” ■ Made minor editorial updates
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Was chapter 11 in the 8.1.0 release. ■ Updated the “Physical Synthesis for Registers—Register Retiming” and “Physical Synthesis Options for Fitting” ■ Updated “Performing Physical Synthesis Optimizations” ■ Deleted Gate-Level Register Retiming section. ■ Updated the referenced documents
November 2008	8.1.0	Changed to 8½” × 11” page size. No change to content.
May 2008	8.0.0	<ul style="list-style-type: none"> ■ Updated “Physical Synthesis Optimizations for Performance on page 11-9 ■ Added Physical Synthesis Options for Fitting on page 11-16

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

Programmable logic can accommodate changes to a system specification late in the design cycle. Last-minute design changes, commonly referred to as engineering change orders (ECOs), are small changes to the functionality of a design after the design has been fully compiled. This section describes how the Chip Planner feature in the Quartus® II software supports ECOs by allowing quick and efficient changes to your logic late in the design cycle.

This section includes the following chapter:

- [Chapter 17, Engineering Change Management with the Chip Planner](#)

This chapter addresses the impact that ECOs have on the design cycle, discusses the design flow for performing ECOs, and describes how you can use the Chip Planner to perform ECOs.

Programmable logic can accommodate changes to a system specification late in the design cycle. In a typical engineering project development cycle, the specification of the programmable logic portion is likely to change after engineering begins or while integrating all system elements. Last-minute design changes, commonly referred to as engineering change orders (ECOs), are small targeted changes to the functionality of a design after the design has been fully compiled. This chapter discusses the design flow for making ECOs, addresses the impact that ECOs have on the design cycle, and describes how you can use the Chip Planner to make ECOs.

The Chip Planner supports ECOs by allowing quick and efficient changes to your logic late in the design cycle. The Chip Planner provides a visual display of your post-place-and-route design mapped to the device architecture of your chosen FPGA and allows you to create, move, and delete logic cells and I/O atoms.

 For a list of supported devices, refer to *About the Chip Planner* in Quartus® II Help.

 In addition to making ECOs, the Chip Planner allows you to perform detailed analysis on routing congestion, relative resource usage, logic placement, LogicLock™ regions, fan-ins and fan-outs, paths between registers, and delay estimates for paths.

 For more information about using the Chip Planner for design analysis, refer to the *Analyzing and Optimizing the Design Floorplan* chapter in volume 2 of the *Quartus II Handbook*.

 ECOs directly apply to atoms in the target device. As such, performing an ECO relies on your understanding of the device architecture of the target device. For more information about the architecture of your device, refer to the appropriate device handbook on the [Literature](#) page of the Altera website.

This chapter includes the following topics:

- “Engineering Change Orders” on page 17–2
- “ECO Design Flow” on page 17–4
- “The Chip Planner Overview” on page 17–5
- “Performing ECOs with the Chip Planner (Floorplan View)” on page 17–6
- “Performing ECOs in the Resource Property Editor” on page 17–7
- “Change Manager” on page 17–21
- “Scripting Support” on page 17–22
- “Common ECO Applications” on page 17–22
- “Post ECO Steps” on page 17–27

Engineering Change Orders

In the context of an FPGA design, you can apply an ECO directly to a physical resource on the device to modify its behavior. ECOs are typically made during the verification stage of a design cycle. When a small change is required on a design (such as modifying a PLL for a different clock frequency or routing a signal out to a pin for analysis) recompilation of the entire design can be time consuming, especially for larger designs. Because several iterations of small design changes can occur during the verification cycle, recompilation times can quickly add up. Furthermore, a full recompilation due to a small design change can result in the loss of previous design optimizations. Making ECOs, instead of performing a full recompilation on your design, limits the change only to the affected portions of logic.

This section discusses the areas in which ECOs have an impact on a system design and how the Quartus II software can help you optimize the design in these areas. The following topics are discussed in this section:

- “Performance Preservation”
- “Compilation Time”
- “Verification”
- “Change Modification Record”

Performance Preservation

You can preserve the results of previous design optimizations when you make changes to an existing design with one of the following methods:

- Incremental compilation
- Rapid recompile
- ECOs

Choose the method to modify your design based on the scope of the change. The methods above are arranged from the larger scale change to the smallest targeted change to a compiled design.

The incremental compilation feature allows you to preserve compilation results at an RTL component or module level. After the initial compilation of your design, you can assign modules in your design hierarchy to partitions. Upon subsequent compilations, incremental compilation recompiles changed partitions based on the chosen preservation levels.

The rapid recompilation feature leverages results from the latest post-fit netlist to determine the changes required to honor modifications you have made to the source code. If you turn on the rapid recompilation feature, the Compiler attempts to refit only the portion of the netlist that is related to the code modification.

ECOs provide a finer granularity of control compared to the incremental compilation and the rapid recompilation feature. All modifications are performed directly on the architectural elements of the device. You should use ECOs for targeted changes to the post-fit netlist.

 In the Quartus II software versions 10.0 and later, the software does not preserve ECO modifications to the netlist when you recompile a design with the incremental compilation feature turned on. You can reapply ECO changes made during a previous compilation with the Change Manager.

 For more information about the incremental compilation feature, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.

Compilation Time

In the traditional programmable logic design flow, a small change in the design requires a complete recompilation of the design. A complete recompilation of the design consists of synthesis and place-and-route. Making small changes to the design to reach the final implementation on a board can be a long process. Because the Chip Planner works only on the post-place-and-route database, you can implement your design changes in minutes without performing a full compilation.

Verification

After you make a design change, you can verify the impact on your design. To verify that your changes do not violate timing requirements, perform static timing analysis with the Quartus II TimeQuest Timing Analyzer after you check and save your netlist changes in the Chip Planner.

 For more information about the TimeQuest analyzer, refer to the *Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

Additionally, you can perform a gate-level or timing simulation of the ECO-modified design with the post-place-and-route netlist generated by the Quartus II software.

Change Modification Record

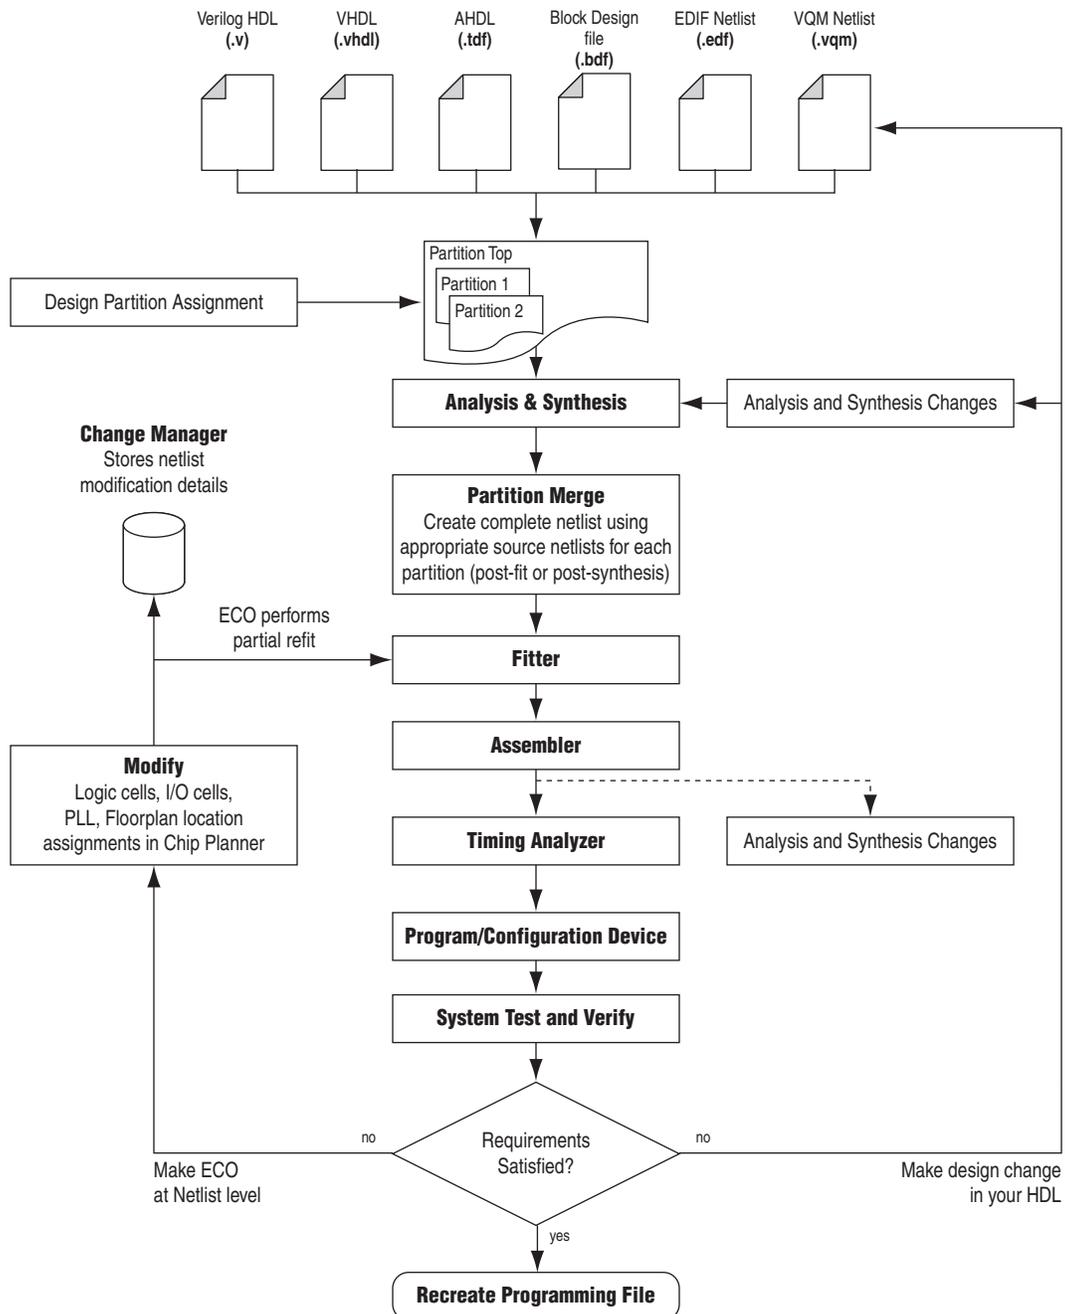
All ECOs made with the Chip Planner are logged in the Change Manager to track all changes. With the Change Manager, you can easily revert to the original post-fit netlist or you can pick and choose which ECOs to apply.

Additionally, the Quartus II software provides support for multiple compilation revisions of the same project. You can use ECOs made with the Chip Planner in conjunction with revision support to compare several different ECO changes and revert back to previous project revisions when required.

ECO Design Flow

Figure 17-1 shows the design flow for making ECOs.

Figure 17-1. Design Flow to Support ECOs



For iterative verification cycles, implementing small design changes at the netlist level can be faster than making an RTL code change. As such, making ECO changes are especially helpful when you debug the design on silicon and require a fast turnaround time to generate a programming file for debugging the design.

A typical ECO application occurs when you uncover a problem on the board and isolate the problem to the appropriate nodes or I/O cells on the device. You must be able to correct the functionality quickly and generate a new programming file. By making small changes with the Chip Planner, you can modify the post-place-and-route netlist directly without having to perform synthesis and logic mapping, thus decreasing the turnaround time for programming file generation during the verification cycle. If the change corrects the problem, no modification of the HDL source code is necessary. You can use the Chip Planner to perform the following ECO-related changes to your design:

- Document the changes made with the Change Manager
- Easily recreate the steps taken to produce design changes
- Generate EDA simulation netlists for design verification



For more complex changes that require HDL source code modifications, the incremental compilation feature can help reduce recompilation time.

The Chip Planner Overview

The Chip Planner provides a visual display of device resources. It shows the arrangement and usage of the resource atoms in the device architecture that you are targeting. Resource atoms are the building blocks for your device, such as ALMs, LEs, PLLs, DSP blocks, memory blocks, or I/O elements.

The Chip Planner also provides an integrated platform for design analysis and for making ECOs to your design after place-and-route. The toolset consists of the Chip Planner (providing a device floorplan view of your mapped design) and two integrated subtools—the Resource Property Editor and the Change Manager.

For analysis, the Chip Planner can show logic placement, LogicLock regions, relative resource usage, detailed routing information, routing congestion, fan-ins and fan-outs, paths between registers, and delay estimates for paths. Additionally, the Chip Planner allows you to create location constraints or resource assignment changes, such as moving or deleting logic cells or I/O atoms with the device floorplan. For ECO changes, the Chip Planner enables you to create, move, or delete logic cells in the post-place-and-route netlist for fast programming file generation. Additionally, you can open the Resource Property Editor from the Chip Planner to edit the properties of resource atoms or to edit the connections between resource atoms. All changes to resource atoms and connections are logged automatically with the Change Manager.

Opening the Chip Planner

To open the Chip Planner, on the Tools menu, click **Chip Planner**. Alternatively, click the **Chip Planner** icon on the Quartus II software toolbar.

Optionally, you can open the Chip Planner by cross-probing from the shortcut menu in the following tools:

- Design Partition Planner
- Compilation Report
- LogicLock Regions window

- Technology Map Viewer
- Project Navigator window
- RTL source code
- Node Finder
- Simulation Report
- RTL Viewer
- Report Timing panel of the TimeQuest Timing Analyzer

The Chip Planner Tasks and Layers

The Chip Planner allows you to set up tasks to quickly implement ECO changes or manipulate assignments for the floorplan of the device. Each task consists of an editing mode and a set of customized layer settings.

 For more information about tasks and layers in the Chip Planner, refer to *About the Chip Planner* in Quartus II Help.

 For more information about creating assignments and performing analysis with the Chip Planner, as well as the Chip Planner floorplan views, refer to the *Analyzing and Optimizing the Design Floorplan* chapter in volume 2 of the *Quartus II Handbook*.

For more information about making ECOs with the ECO mode, refer to “*Performing ECOs with the Chip Planner (Floorplan View)*” on page 17-6.

Performing ECOs with the Chip Planner (Floorplan View)

You can manipulate resource atoms in the Chip Planner when you select the ECO editing mode. The following ECO changes can be made with the Chip Planner Floorplan view:

- Create atoms
- Delete atoms
- Move existing atoms

 To configure the properties of atoms, such as managing the connections between different LEs/ALMs, use the Resource Property Editor.

For more information about editing atom resource properties, refer to “*Performing ECOs in the Resource Property Editor*” on page 17-7.

To select the ECO editing mode in the Chip Planner, in the **Editing Mode** list at the top of the Chip Planner, select the ECO editing mode.

Creating, Deleting, and Moving Atoms

You can use the Chip Planner to create, delete, and move atoms in the post-compilation design.

- ❓ For more information about creating, deleting, and moving atoms, refer to *Creating, Deleting, and Moving Atoms* in Quartus II Help.

Check and Save Netlist Changes

After making all the ECOs, you can run the Fitter to incorporate the changes by clicking the **Check and Save Netlist Changes** icon in the Chip Planner toolbar. The Fitter compiles the ECO changes, performs design rule checks on the design, and generates a programming file.

Performing ECOs in the Resource Property Editor

You can view and edit the following resources with the Resource Property Editor:

- “Logic Elements” on page 17-7
- “Adaptive Logic Modules” on page 17-10
- “FPGA I/O Elements” on page 17-12
- “PLL Properties” on page 17-24
- “FPGA RAM Blocks” on page 17-19
- “FPGA DSP Blocks” on page 17-20

Logic Elements

An Altera® LE contains a four-input LUT, which is a function generator that can implement any function of four variables. In addition, each LE contains a register fed by the output of the LUT or by an independent function generated in another LE.

You can use the Resource Property Editor to view and edit any LE in the FPGA. To open the Resource Property Editor for an LE, on the Project menu, point to **Locate**, and then click **Locate in Resource Property Editor** in one of the following views:

- RTL Viewer
- Technology Map Viewer
- Node Finder
- Chip Planner

-  For more information about LE architecture for a particular device family, refer to the device family handbook or data sheet.

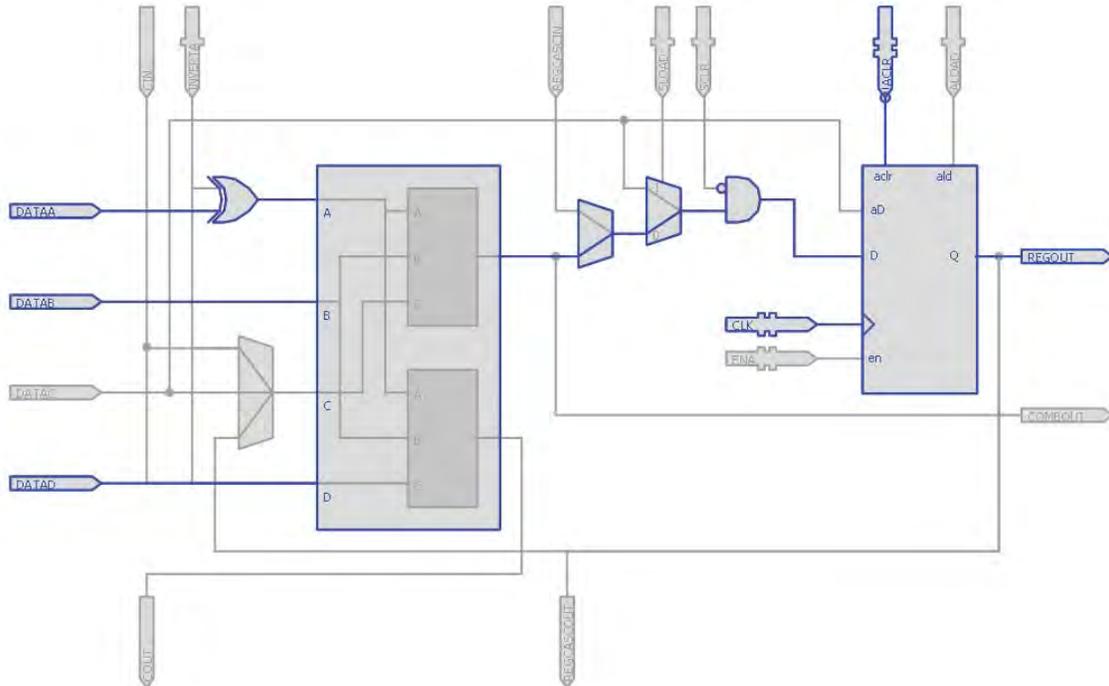
You can use the Resource Property Editor to change the following LE properties:

- Data input to the LUT
- LUT mask or LUT equation

Logic Element Schematic View

Figure 17-2 shows how the LE appears in the Resource Property Editor. By default, the Quartus II software displays the used resources in blue and the unused resources in gray.

Figure 17-2. Stratix LE Architecture (Note 1)



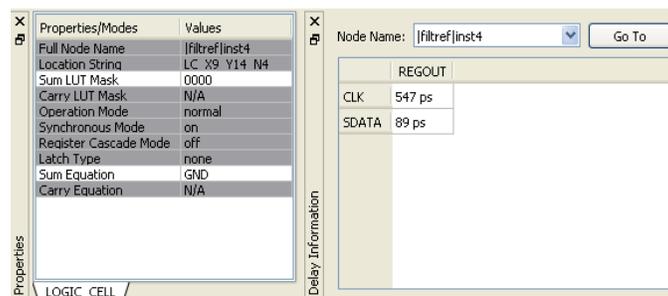
Notes to Figure 17-2:

- (1) For more information about the Stratix device's LE architecture, refer to the *Stratix Device Handbook*.

Logic Element Properties

Figure 17-3 shows an example of the properties that can be viewed for a selected LE in the Resource Property Editor. To view LE properties, on the View menu, click **View Properties**.

Figure 17-3. LE Properties



Modes of Operation

LUTs in an LE can operate in either normal or arithmetic mode.

When an LE is configured in normal mode, the LUT in the LE can implement a function of four inputs.

When the LE is configured in arithmetic mode, the LUT in the LE is divided into two 3-input LUTs. The first LUT generates the signal that drives the output of the LUT, while the second LUT generates the carry-out signal. The carry-out signal can drive only a carry-in signal of another LE.



For more information about LE modes of operation, refer to volume 1 of the appropriate device handbook.

Sum and Carry Equations

You can change the logic function implemented by the LUT by changing the sum and carry equations. When the LE is configured in normal mode, you can change only the sum equation. When the LE is configured in arithmetic mode, you can change both the sum and the carry equations.

The LUT mask is the hexadecimal representation of the LUT equation output. When you change the LUT equation, the Quartus II software automatically changes the LUT mask. Conversely, when you change the LUT mask, the Quartus II software automatically computes the LUT equation.

sload and sclr Signals

Each LE register contains a synchronous load (`sload`) signal and a synchronous clear (`sclr`) signal. You can invert either the `sload` or `sclr` signal feeding into the LE. If the design uses the `sload` signal in an LE, the signal and its inversion state must be the same for all other LEs in the same LAB. For example, if two LEs in a LAB have the `sload` signal connected, both LEs must have the `sload` signal set to the same value. This is also true for the `sclr` signal.

Register Cascade Mode

When register cascade mode is enabled, the cascade-in port feeds the input to the register. The register cascade mode is used most often when the design implements shift registers. You can change the register cascade mode by connecting (or disconnecting) the cascade in the port. However, if you create this port, you must ensure that the source port LE is directly above the destination LE.

Cell Delay Table

The cell delay table describes the propagation delay from all inputs to all outputs for the selected LE.

Logic Element Connections

To view the connections that feed in and out of an LE, on the View menu, click **View Port Connections**. Figure 17-4 shows the LE connections in the Connectivity window.

Figure 17-4. View LE Connections

Input Port Name	Signal Name	Inverted	Output Port Name	Signal Name
DATAA	<Disconnected>	False	COU	<Disconnected>
DATAB	<Disconnected>	False	COMBOUT	<Disconnected>
SDATA	filter state_m:inst1 filter_tap4	False	REGOUT	filter inst4
DATAD	<Disconnected>	False		
CIN	<Disconnected>	False		
INVERTA	<Disconnected>	False		
REGCASIN	<Disconnected>	False		
SLOAD	VCC	False		
SCLR	<Disconnected>	False		
IACL	VCC	False		
ALOAD	<Disconnected>	False		
CLK	filter clk	False		
ENA	<Disconnected>	False		

Delete a Logic Element

To delete an LE, follow these steps:

1. Right-click the desired LE in the Chip Planner, point to **Locate**, and click **Locate in Resource Property Editor**.
2. You must remove all fan-out connections from an LE prior to deletion. To delete fan-out connections, right-click each connected output signal, point to **Remove**, and click **Fanouts**. Select all of the fan-out signals in the **Remove Fan-outs** dialog box and click **OK**.
3. To delete an atom after all fan-out connections are removed, right-click the atom in the Chip Planner and click **Delete Atom**.

Adaptive Logic Modules

Each ALM contains LUT-based resources that can be divided between two adaptive LUTs (ALUTs). With up to eight inputs to the two ALUTs, each ALM can implement various combinations of two functions. This adaptability allows the ALM to be completely backward-compatible with four-input LUT architectures. One ALM can implement any function with up to six inputs and certain seven-input functions. In addition to the ALUT-based resources, each ALM contains two programmable registers, two dedicated full adders, a carry chain, a shared arithmetic chain, and a register chain. The ALM can efficiently implement various arithmetic functions and shift registers with these dedicated resources.

You can implement the following types of functions in a single ALM:

- Two independent 4-input functions
- An independent 5-input function and an independent 3-input function
- A 5-input function and a 4-input function, if they share one input
- Two 5-input functions, if they share two inputs
- An independent 6-input function
- Two 6-input functions, if they share four inputs and share the same functions

- Certain 7-input functions

You can use the Resource Property Editor to change the following ALM properties:

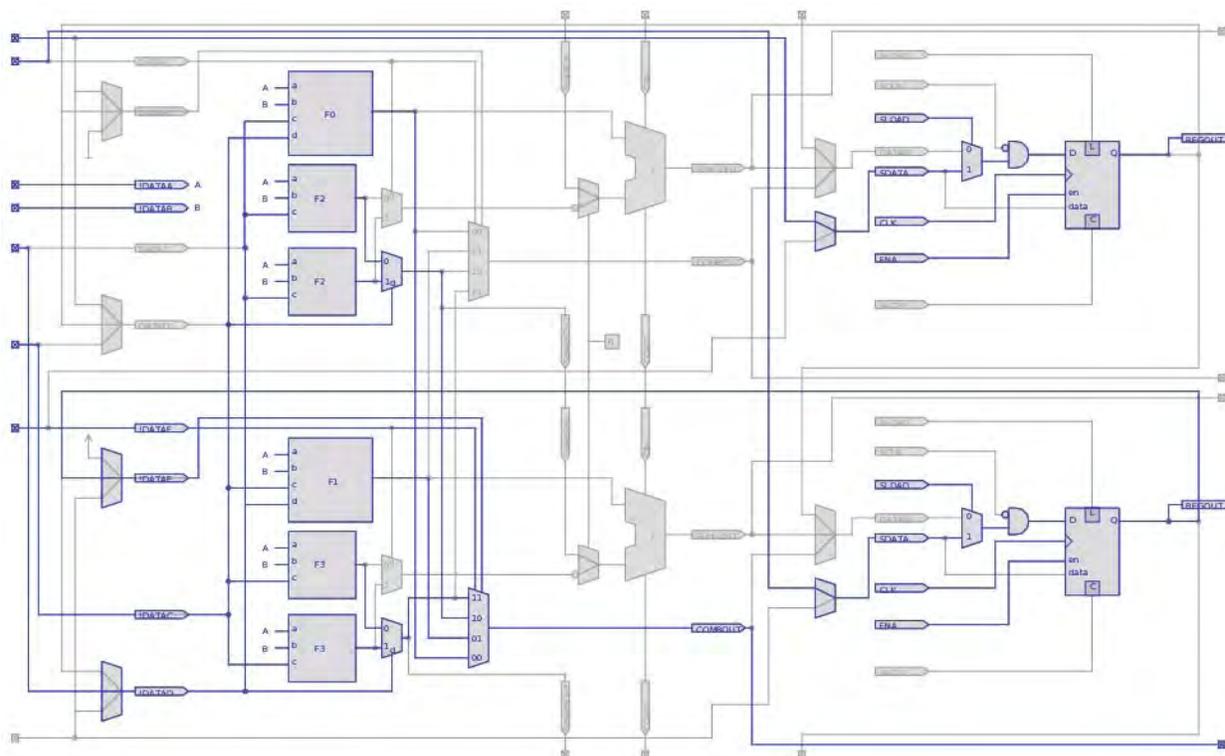
- Data input to the LUT
- LUT mask or LUT equation

Adaptive Logic Module Schematic

You can view and edit any ALM atom with the Resource Property Editor by right-clicking the ALM in the RTL Viewer, the Node Finder, or the Chip Planner, and clicking **Locate in Resource Property Editor** (Figure 17-5).

- For a detailed description of the ALM, refer to the device handbooks of devices based on an ALM architecture.

Figure 17-5. ALM Schematic (Note 1)



Note to Figure 17-5:

- (1) By default, the Quartus II software displays the used resources in blue and the unused in gray. For Figure 17-5, the used resources are in blue and the unused resources are in gray.

Adaptive Logic Module Properties

The properties that you can display for the ALM include an equations table that shows the name and location of each of the two combinational nodes and two register nodes in the ALM, the individual LUT equations for each of the combinational nodes, and the combout, sumout, carryout, and shareout equations for each combinational node.

Adaptive Logic Module Connections

On the View menu, click **View Connectivity** to view the input and output connections for the ALM.

FPGA I/O Elements

Altera FPGAs that have high-performance I/O elements, including up to six registers, are equipped with support for a number of I/O standards that allow you to run your design at peak speeds. Use the Resource Property Editor to view, change connectivity, and edit the properties of the I/O elements. Use the Chip Planner (Floorplan view) to change placement, delete, and create new I/O elements.



For a detailed description of the device I/O elements, refer to the applicable device handbook.

You can change the following I/O properties:

- Delay chain
- Bus hold
- Weak pull up
- Slow slew rate
- I/O standard
- Current strength
- Extend OE disable
- PCI I/O
- Register reset mode
- Register synchronous reset mode
- Register power up
- Register mode

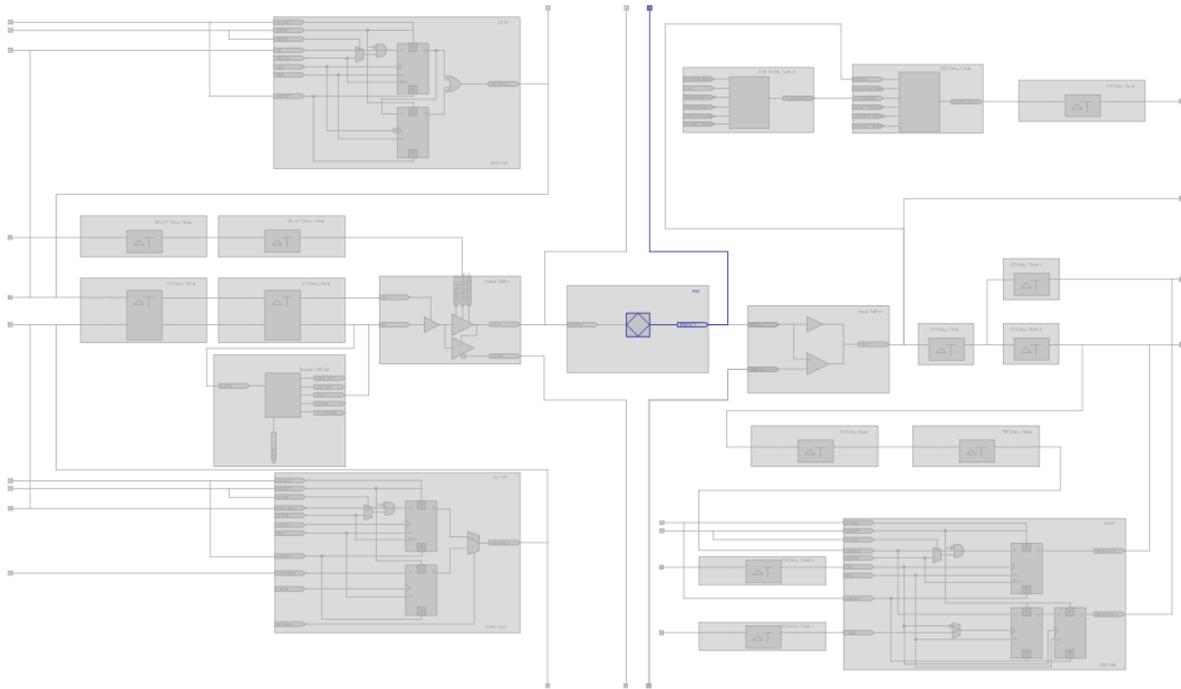
Stratix V I/O Elements

The I/O elements in Stratix® V devices contain a bidirectional I/O buffer and I/O registers to support a complete embedded bidirectional single data rate (SDR) or double data rate (DDR) transfer (shown in [Figure 17-6](#)).

I/O registers are composed of the input path for handling data from the pin to the core, the output path for handling data from the core to the pin, and the output enable path for handling the output enable signal to the output buffer. These registers allow faster source-synchronous register-to-register transfers and resynchronization. The input path consists of the DDR input registers, alignment and synchronization registers, and half data rate blocks; you can bypass each block in the input path. The input path uses the deskew delay to adjust the input register clock delay across process, voltage, and temperature (PVT) variations.

By default, the Quartus II software displays the used resources in blue and the unused resources in gray.

Figure 17-6. Stratix V Device I/O Element Structure



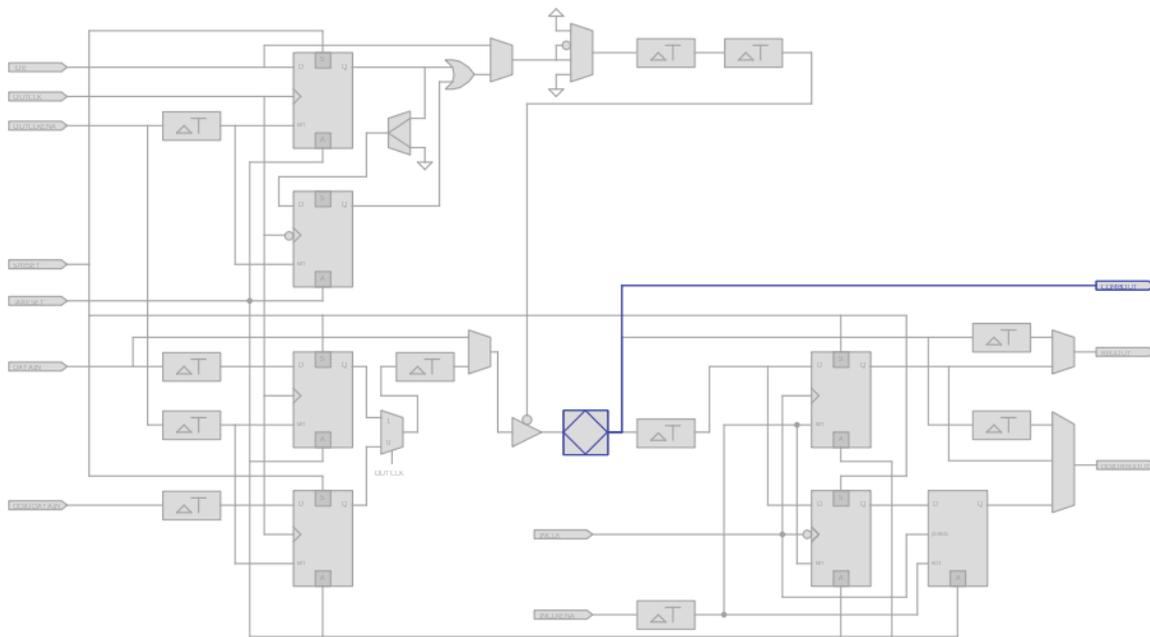
 For more information about I/O elements in Stratix V devices, refer to the *Stratix V Device Handbook*.

Arria GX, Stratix, Stratix II, and Stratix GX I/O Elements

The I/O elements in Arria® GX, Stratix, Stratix II, and Stratix GX devices contain a bidirectional I/O buffer, six registers, and a latch for a complete bidirectional SDR or DDR transfer.

Figure 17-7 shows the Stratix and Stratix GX I/O element structure. The I/O element structure contains two input registers (plus a latch), two output registers, and two output enable registers. By default, the Quartus II software displays the used resources in blue and the unused resources in gray.

Figure 17-7. Stratix and Stratix GX Device I/O Element and Structure

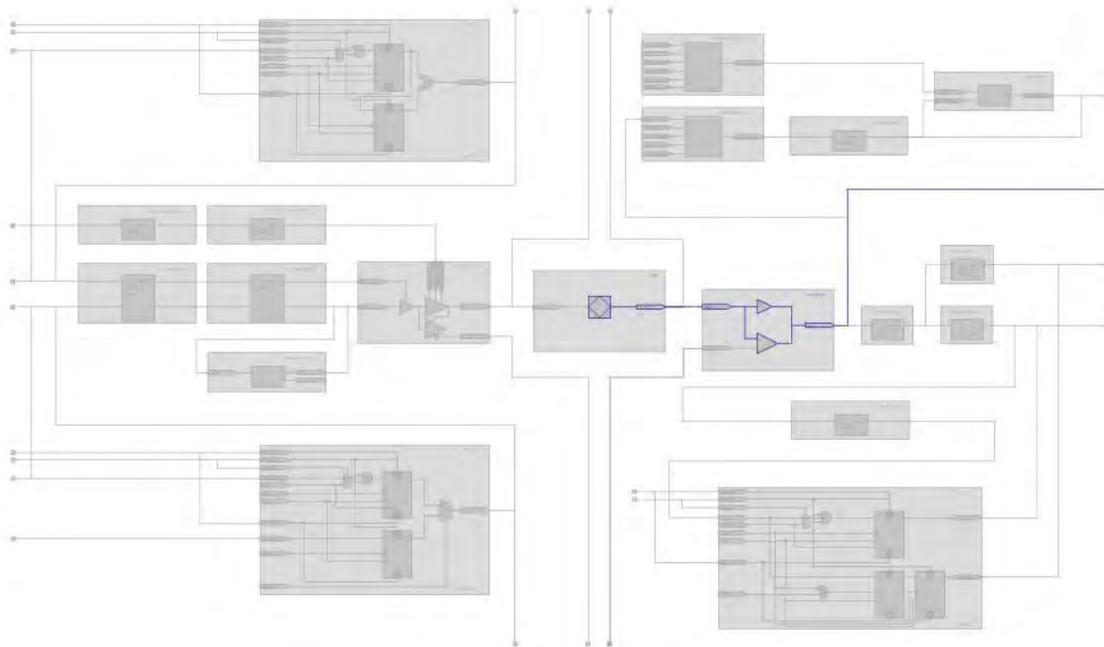


 For more information about I/O elements in Stratix and Stratix GX devices, refer to the *Stratix Device Handbook* and the *Stratix GX Device Handbook*.

Arria II GX, Stratix III, and Stratix IV I/O Elements

The I/O elements in Arria II GX, Stratix III, and Stratix IV devices contain a bidirectional I/O buffer and I/O registers to support a complete embedded bidirectional SDR or DDR transfer (shown in Figure 17-8). The I/O registers are composed of the input path for handling data from the pin to the core, the output path for handling data from the core to the pin, and the output enable path for handling the output enable signal for the output buffer. Each path consists of a set of delay elements that allow you to fine-tune the timing characteristics of each path for skew management. By default, the Quartus II software displays the used resources in blue and the unused resources in gray.

Figure 17-8. Stratix III Device I/O Element and Structure

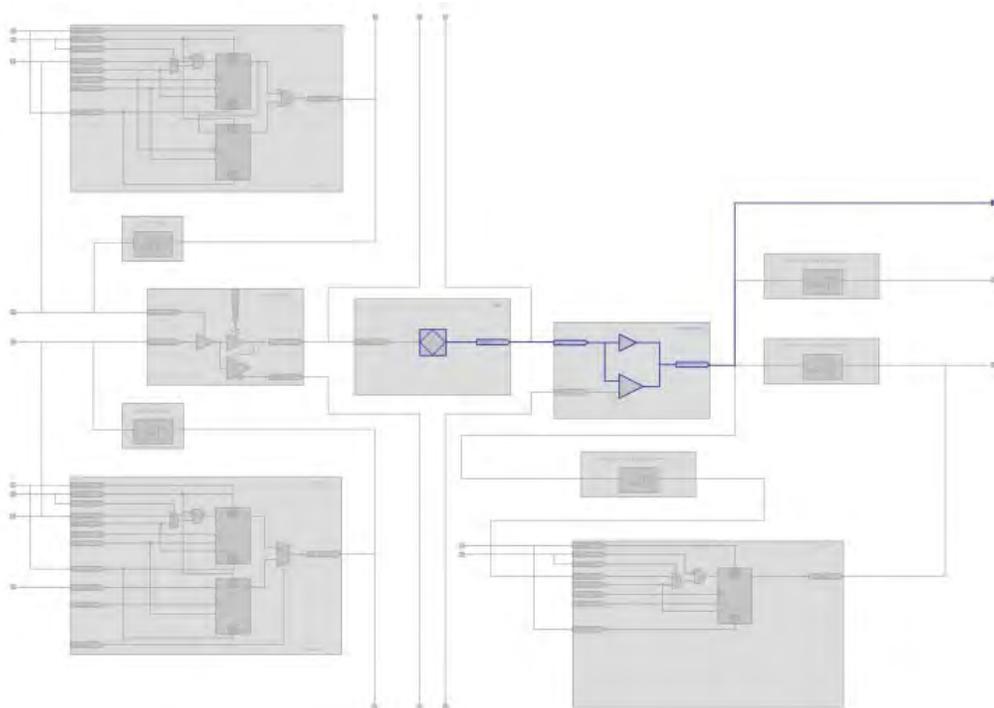


- For more information about I/O elements in Stratix III and Stratix IV devices, refer to the [Literature](#) page of the Altera website.
- For more information about programmable I/O elements in Stratix III devices, refer to [AN 474: Implementing Stratix III Programmable I/O Delay Settings in the Quartus II Software](#).

Cyclone III I/O Elements

The I/O elements in Cyclone III devices contain a bidirectional I/O buffer and five registers for complete embedded bidirectional single data rate transfer. Figure 17-10 shows the Cyclone III I/O element structure. The I/O element contains one input register, two output registers, and two output-enable registers. The two output registers and two output-enable registers are utilized for double-data rate (DDR) applications. You can use the input registers for fast setup times and the output registers for fast clock-to-output times. Additionally, you can use the output-enable (OE) registers for fast clock-to-output enable timing. You can use I/O elements for input, output, or bidirectional data paths. By default, the Quartus II software displays the used resources in blue and the unused resources in gray.

Figure 17-10. Cyclone III Device I/O Elements and Structure

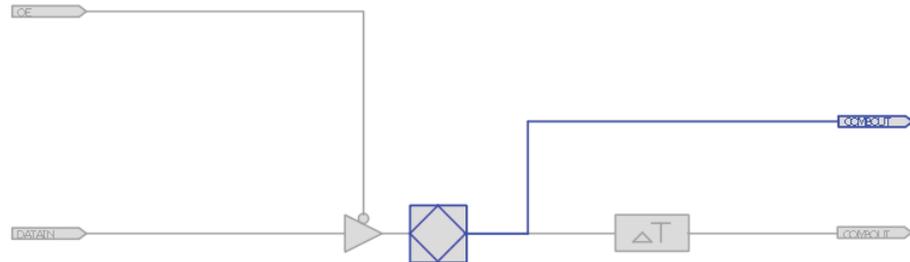


 For more information about I/O elements in Cyclone III devices, refer to the [Cyclone III Device Handbook](#).

MAX II I/O Elements

The I/O elements in MAX[®] II devices contain a bidirectional I/O buffer. [Figure 17-11](#) shows the MAX II I/O element structure. You can drive registers from adjacent LABs to or from the bidirectional I/O buffer of the I/O element. By default, the Quartus II software displays the used resources in blue and the unused resources in gray.

Figure 17-11. MAX II Device I/O Elements and Structure

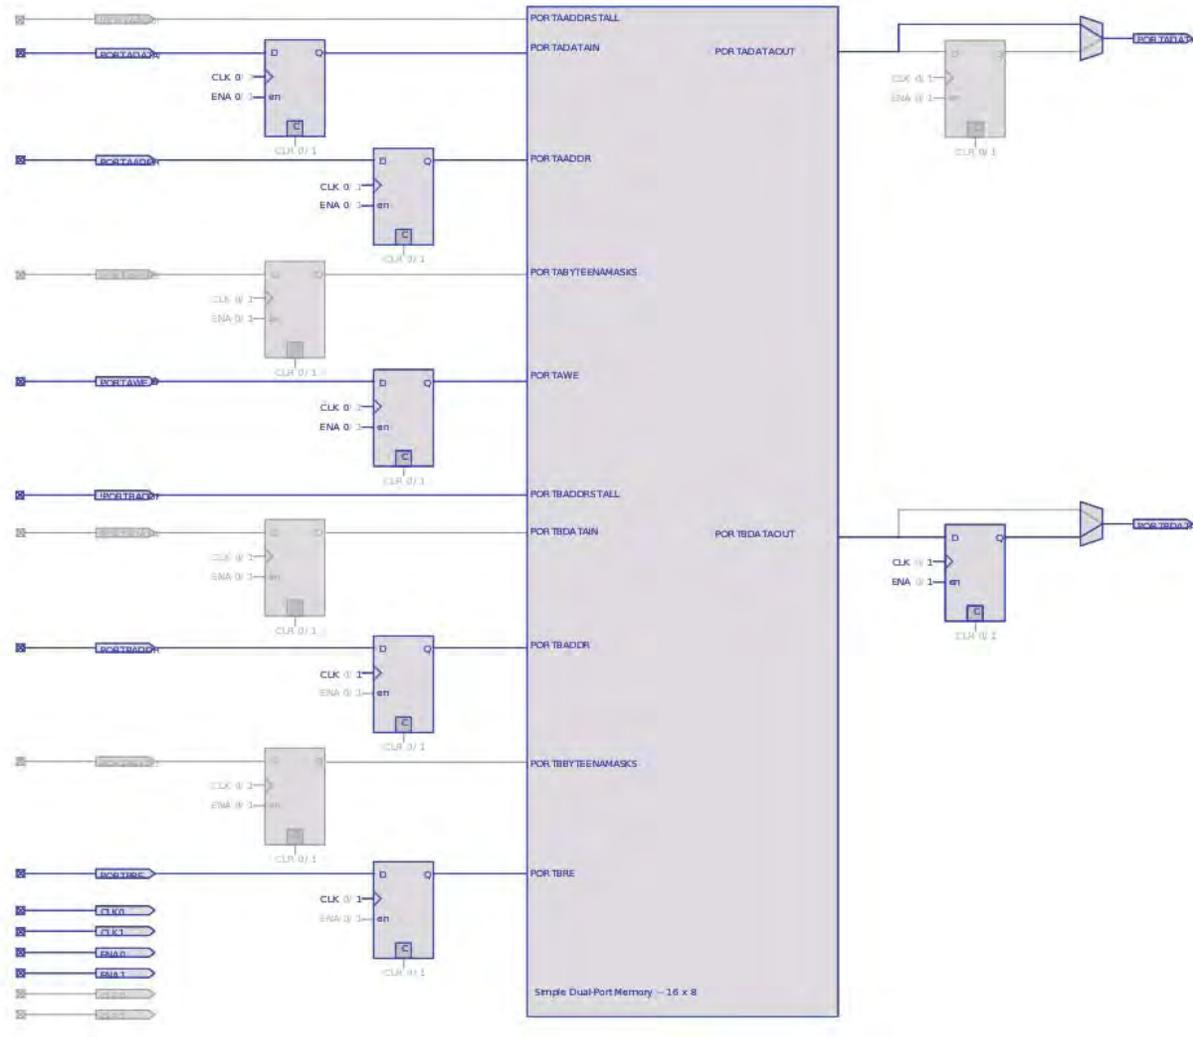


For more information about I/O elements in MAX II devices, refer to the [MAX II Device Handbook](#).

FPGA RAM Blocks

With the Resource Property Editor, you can view the architecture of different RAM blocks in the device, modify the input and output registers to and from the RAM blocks, and modify the connectivity of the input and output ports. Figure 17-12 shows an M9K RAM view in a Stratix III device.

Figure 17-12. M9K RAM View in a Stratix III Device (Note 1)



Note to Figure 17-12:

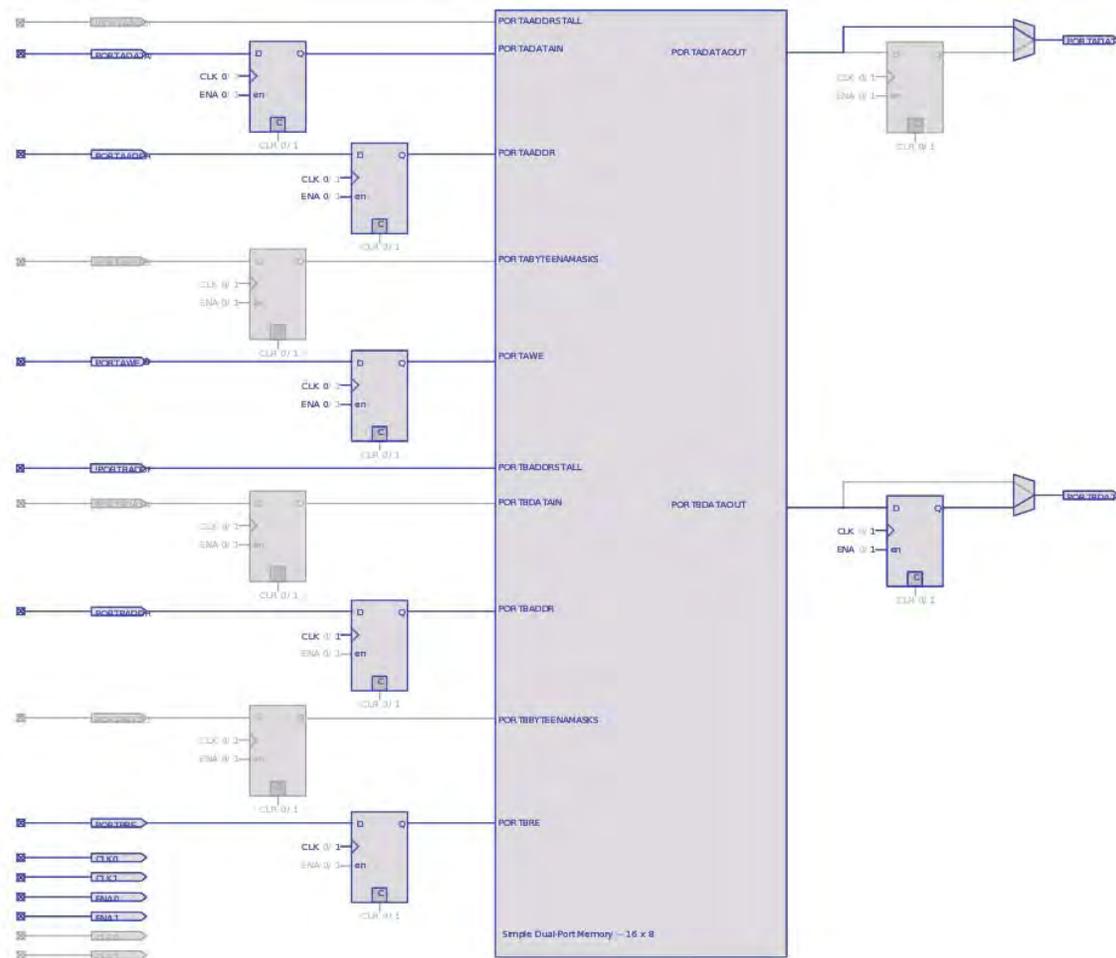
- (1) By default, the Quartus II software displays the used resources in blue and the unused resources in gray. In Figure 17-12, the used resources are in blue and the unused resources are in gray.

FPGA DSP Blocks

Dedicated hardware DSP circuit blocks in Altera devices provide performance benefits for the critical DSP functions in your design. The Resource Property Editor allows you to view the architecture of DSP blocks in the Resource Property Editor for the Cyclone and Stratix series of devices. The Resource Property Editor also allows you to modify the signal connections to and from the DSP blocks and modify the input and output registers to and from the DSP blocks.

Figure 17-13 shows the DSP architecture in a Stratix III device.

Figure 17-13. DSP Block View in a Stratix III Device (Note 1)



Note to Figure 17-13:

- (1) By default, the Quartus II software displays the used resources in blue and the unused resources in gray. In Figure 17-13, the used resources are in blue and the unused resources are in red.

Change Manager

The Change Manager maintains a record of every change you perform with the Chip Planner, the Resource Property Editor, the SignalProbe feature, or a Tcl script. Each row of data in the Change Manager represents one ECO.

The Change Manager allows you to apply changes, roll back changes, delete changes, and export change records to a Text File (.txt), a Comma-Separated Value File (.csv), or a Tcl Script File (.tcl). The Change Manager tracks dependencies between changes, so that when you apply, roll back, or delete a change, any prerequisite or dependent changes are also applied, rolled back, or deleted.

- For more information about the Change Manager, refer to *About the Change Manager* in Quartus II Help.

Complex Changes in the Change Manager

Certain changes (including creating or deleting atoms and changing connectivity) can appear to be self-contained, but are actually composed of multiple actions. The Change Manager marks such complex changes with a plus icon in the **Index** column.

You can click the plus icon to expand the change record and show all the component actions performed as part of that complex change.

- For more information about complex change records and about managing changes with the Change Manager, refer to *Examples of Managing Changes With the Change Manager* in Quartus II Help.

Managing SignalProbe Signals

The SignalProbe pins that you create from the **SignalProbe Pins** dialog box are recorded in the Change Manager. After you have made a SignalProbe assignment, you can use the Change Manager to quickly disable SignalProbe assignments by selecting **Revert to Last Saved Netlist** on the shortcut menu in the Change Manager.

- For more information about SignalProbe pins, refer to the *Quick Design Debugging Using SignalProbe* chapter in volume 3 of the *Quartus II Handbook*.

Exporting Changes

You can export changes to a .txt, a .csv, or a .tcl. Tcl scripts allow you to reapply changes that were deleted during compilation.

- For more information about exporting changes, refer to *Managing Changes With the Change Manager* in Quartus II Help.
- For more information about netlist types and the Quartus II incremental compilation feature, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.

Scripting Support

You can run procedures and make settings described in this chapter in a Tcl script. You can also run some procedures at a command prompt. The Tcl commands for controlling the Chip Planner are located in the `chip_planner` package of the `quartus_cdb` executable.

 A comprehensive list of Tcl commands for the Chip Planner can be found in *About Quartus II Scripting* in Quartus II Help.

 For more information about Tcl scripting, refer to the *Tcl Scripting* chapter in volume 2 of the *Quartus II Handbook*. For more information about all settings and constraints in the Quartus II software, refer to the *Quartus II Settings File Manual*. For more information about command-line scripting, refer to the *Command-Line Scripting* chapter in volume 2 of the *Quartus II Handbook*.

Common ECO Applications

This section provides examples of how you might use an ECO to make a post-compilation change to your design. To help build your system quickly, you can use Chip Planner functions to perform the following activities:

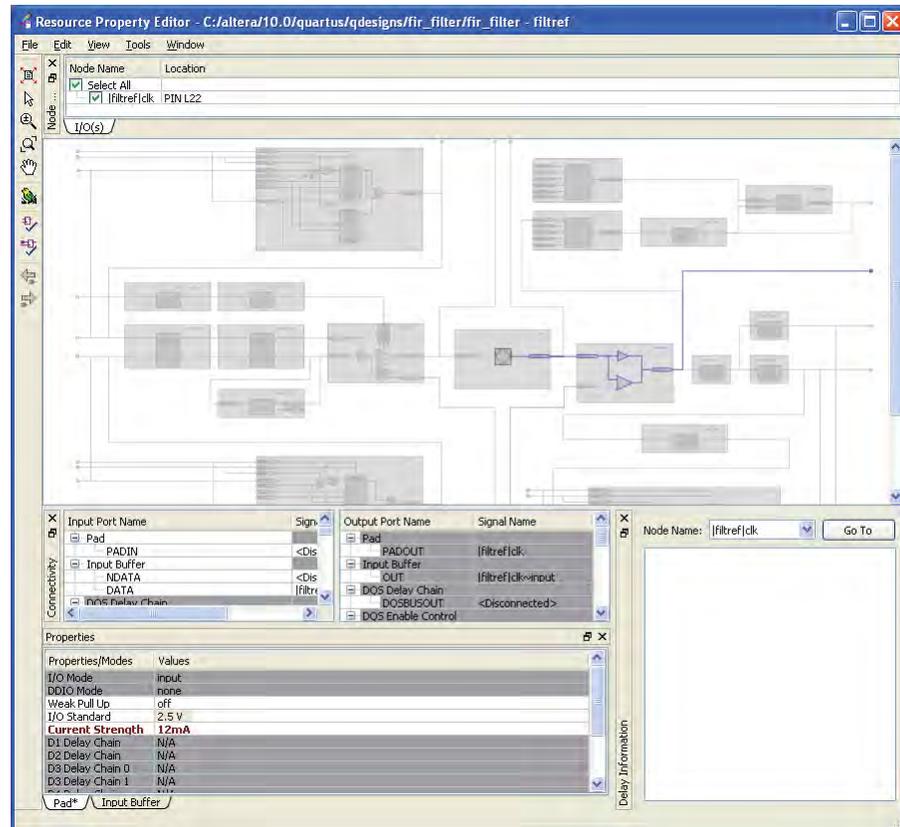
- Adjust the drive strength of an I/O with the Chip Planner
- Modify the PLL properties with the Resource Property Editor (see “[Modify the PLL Properties With the Chip Planner](#)” on page 17-23)
- Modify the connectivity between new resource atoms with the Chip Planner and Resource Property Editor

Adjust the Drive Strength of an I/O with the Chip Planner

To adjust the drive strength of an I/O, follow the steps in this section to incorporate the ECO changes into the netlist of the design.

1. In the **Editing Mode** list at the top of the Chip Planner, select the ECO editing mode.
2. Locate the I/O in the **Resource Property Editor**, as shown in Figure 17-14.

Figure 17-14. I/O in the Resource Property Editor



3. In the **Resource Property Editor**, point to the **Current Strength** option in the **Properties** pane and double-click the value to enable the drop-down list.
4. Change the value for the **Current Strength** option.
5. Right-click the ECO change in the Change Manager and click **Check & Save All Netlist Changes** to apply the ECO change.

 You can change the pin locations of input or output ports with the ECO flow. You can drag and move the signal from an existing pin location to a new location while in the Post Compilation Editing (ECO) task in the Chip Planner. You can then click **Check & Save All Netlist Changes** to compile the ECO.

Modify the PLL Properties With the Chip Planner

You use PLLs to modify and generate clock signals to meet design requirements. Additionally, you can use PLLs to distribute clock signals to different devices in a design, reducing clock skew between devices, improving I/O timing, and generating internal clock signals.

- M2 value
- N2 value
- SS counter
- Charge pump current
- Loop filter resistance
- Loop filter capacitance
- Counter delay
- Counter high
- Counter low
- Counter mode
- Counter initial
- V_{CO} tap

You can also view post-compilation PLL properties in the Compilation Report. To do so, in the Compilation Report, select **Fitter** and then select **Resource Section**.

Adjusting the Duty Cycle

Use [Equation 17-1](#) to adjust the duty cycle of individual output clocks.

Equation 17-1.

$$\text{High \%} = \frac{\text{Counter High}}{(\text{Counter High} + \text{Counter Low})}$$

Adjusting the Phase Shift

Use [Equation 17-2](#) to adjust the phase shift of an output clock of a PLL.

Equation 17-2.

$$\text{Phase Shift} = (\text{Period } V_{CO} \times 0.125 \times \text{Tap } V_{CO}) + (\text{Initial } V_{CO} \times \text{Period } V_{CO})$$

For normal mode, Tap V_{CO}, Initial V_{CO}, and Period V_{CO} are governed by the following settings:

$$\text{Tap } V_{CO} = \text{Counter Delay} - M \text{ Tap } V_{CO}$$

$$\text{Initial } V_{CO} = \text{Counter Initial} - M \text{ Initial}$$

$$\text{Period } V_{CO} = \text{In Clock Period} \times N \div M$$

For external feedback mode, Tap V_{CO}, Initial V_{CO}, and Period V_{CO} are governed by the following settings:

$$\text{Tap } V_{CO} = \text{Counter Delay} - M \text{ Tap } V_{CO}$$

$$\text{Initial } V_{CO} = \text{Counter Initial} - M \text{ Initial}$$

$$\text{Period } V_{CO} = \frac{\text{In Clock Period} \times N}{(M + \text{Counter High} + \text{Counter Low})}$$

 For a detailed description of the settings, refer to the Quartus II Help. For more information about Stratix device PLLs, refer to the *Stratix Architecture* chapter in volume 1 of the *Stratix Device Handbook*. For more information about PLLs in Arria GX, Cyclone, Cyclone II, and Stratix II devices, refer to the appropriate device handbook.

Adjusting the Output Clock Frequency

Use Equation 17-3 to adjust the PLL output clock in normal mode.

Equation 17-3.

$$\text{Output Clock Frequency} = \text{Input Frequency} \cdot \frac{M \text{ value}}{N \text{ Value} + \text{Counter High} + \text{Counter Low}}$$

Use Equation 17-4 to adjust the PLL output clock in external feedback mode.

Equation 17-4.

$$\text{OUTCLK} = \frac{M \text{ value} + \text{External Feedback Counter High} + \text{External Feedback Counter Low}}{N \text{ value} + \text{Counter High} + \text{Counter Low}}$$

Adjusting the Spread Spectrum

Use Equation 17-5 to adjust the spread spectrum for your PLL.

Equation 17-5.

$$\% \text{ spread} = \frac{M_2 N_1}{M_1 N_2}$$

Modify the Connectivity between Resource Atoms

The Chip Planner and Resource Property Editor allow you to create new resource atoms and manipulate the existing connection between resource atoms in the post-fit netlist. These features are useful for small changes when you are debugging a design, such as manually inserting pipeline registers into a combinational path that fails timing, or routing a signal to a spare I/O pin for analysis. Use the following procedure to create a new register in a Cyclone III device and route register output to a spare I/O pin. This example illustrates how to create a new resource atom and modify the connections between resource atoms.

To create new resource atoms and manipulate the existing connection between resource atoms in the post-fit netlist, follow these steps:

1. Create a new register in the Chip Planner.
2. Locate the atom in the Resource Property Editor.
3. To assign a clock signal to the register, right-click the clock input port for the register, point to **Edit connection**, and click **Other**. Use the Node Finder to assign a clock signal from your design.
4. To tie the SLOAD input port to V_{CC} , right-click the clock input port for the register, point to **Edit connection**, and click **VCC**.

5. Assign a data signal from your design to the SDATA port.
6. In the Connectivity window, under the output port names, copy the port name of the register.
7. In the Chip Planner, locate a free I/O resource and create an output buffer.
8. Locate the new I/O atom in the Resource Property Editor.
9. On the input port to the output buffer, right-click, point to **Edit connection**, and click **Other**.
10. In the **Edit Connection** dialog box, type the output port name of the register you have created.
11. Run the ECO Fitter to apply the changes by clicking **Check and Save Netlist Changes**.



A successful ECO connection is subject to the available routing resources. You can view the relative routing utilization by selecting **Routing Utilization** as the Background Color Map in the **Layers Settings** dialog box of the Chip Planner. Also, you can view individual routing channel utilization from local, row, and column interconnects with the tooltips created when you position your mouse pointer over the appropriate resource. Refer to the device data sheet for more information about the architecture of the routing interconnects of your device.

Post ECO Steps

After you make an ECO change with the Chip Planner, you must perform static timing analysis of your design with the TimeQuest analyzer to ensure that your changes did not adversely affect the timing performance of your design.

For example, when you turn on one of the delay chain settings for a specific pin, you change the I/O timing. Therefore, to ensure that the design still meets all timing requirements, you should perform static timing analysis.



For more information about performing a static timing analysis of your design, refer to *The Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the Quartus II Handbook.

Conclusion

The Chip Planner allows you to analyze and modify your design floorplan. Also, ECO changes made with the Chip Planner do not require a full recompilation, eliminating the lengthy process of RTL modification, resynthesis, and another place-and-route cycle. In summary, the Chip Planner speeds design verification and timing closure.

Document Revision History

Table 17-1 shows the revision history for this chapter.

Table 17-1. Document Revision History

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Updated chapter to new template ■ Removed “The Chip Planner FloorPlan Views” section ■ Combined “Creating Atoms”, “Deleting Atoms”, and “Moving Atoms” sections, and linked to Help. ■ Added Stratix V I/O elements in “FPGA I/O Elements” on page 17-12.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Added information to page 17-1. ■ Added information to “Engineering Change Orders” on page 17-2. ■ Changed heading from “Performance” to “Performance Preservation” on page 17-2. ■ Updated information in “Performance Preservation” on page 17-2. ■ Changed heading from “Documentation” to “Change Modification Record” on page 17-3. ■ Changed heading from “Resource Property Editor” to “Performing ECOs in the Resource Property Editor” on page 17-15. ■ Removed “Using Incremental Compilation in the ECO Flow” section. Preservation support for ECOs with the incremental compilation flow has been removed in the Quartus II software version 10.0. ■ Removed “Referenced Documents” section.
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Updated device support list ■ Made minor editorial updates
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Updated Figure 17-17. ■ Made minor editorial updates. ■ Chapter 15 was previously Chapter 13 in the 8.1.0 release.
November 2008	8.1.0	<ul style="list-style-type: none"> ■ Corrected preservation attributes for ECOs in the section “Using Incremental Compilation in the ECO Flow” on page 15-32. ■ Minor editorial updates. ■ Changed to 8½” x 11” page size.
May 2008	8.0.0	<ul style="list-style-type: none"> ■ Updated device support list ■ Modified description for ECO support for block RAMs and DSP blocks ■ Corrected Stratix PLL ECO example ■ Added an application example to show modifying the connectivity between resource atoms

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter provides additional information about the document and Altera.

About this Handbook

This handbook provides comprehensive information about the Altera® Quartus® II design software, version 10.1.

How to Contact Altera

To locate the most up-to-date information about Altera products, refer to the following table.

Contact (1)	Contact Method	Address
Technical support	Website	www.altera.com/support
Technical training	Website	www.altera.com/training
	Email	custrain@altera.com
Product literature	Website	www.altera.com/literature
Non-technical support (General) (Software Licensing)	Email	nacomp@altera.com
	Email	authorization@altera.com

Note to Table:

(1) You can also contact your local Altera sales office or sales representative.

Third-Party Software Product Information

Third-party software products described in this handbook are not Altera products, are licensed by Altera from third parties, and are subject to change without notice. Updates to these third-party software products may not be concurrent with Quartus II software releases. Altera has assumed responsibility for the selection of such third-party software products and its use in the Quartus II 10.1 software release. To the extent that the software products described in this handbook are derived from third-party software, no third party warrants the software, assumes any liability regarding use of the software, or undertakes to furnish you any support or information relating to the software. EXCEPT AS EXPRESSLY SET FORTH IN THE APPLICABLE ALTERA PROGRAM LICENSE SUBSCRIPTION AGREEMENT UNDER WHICH THIS SOFTWARE WAS PROVIDED TO YOU, ALTERA AND THIRD-PARTY LICENSORS DISCLAIM ALL WARRANTIES WITH RESPECT TO THE USE OF SUCH THIRD-PARTY SOFTWARE CODE OR DOCUMENTATION IN THE SOFTWARE, INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, AND NONINFRINGEMENT. For more information, including the latest available version of specific third-party software products, refer to the documentation for the software in question.

Typographic Conventions

The following table shows the typographic conventions this document uses.

Visual Cue	Meaning
Bold Type with Initial Capital Letters	Indicate command names, dialog box titles, dialog box options, and other GUI labels. For example, Save As dialog box. For GUI elements, capitalization matches the GUI.
bold type	Indicates directory names, project names, disk drive names, file names, file name extensions, software utility names, and GUI labels. For example, \qdesigns directory, D: drive, and chiptrip.gdf file.
<i>Italic Type with Initial Capital Letters</i>	Indicate document titles. For example, <i>Stratix IV Design Guidelines</i> .
<i>italic type</i>	Indicates variables. For example, $n + 1$. Variable names are enclosed in angle brackets (< >). For example, <file name> and <project name>.pof file.
Initial Capital Letters	Indicate keyboard keys and menu names. For example, the Delete key and the Options menu.
“Subheading Title”	Quotation marks indicate references to sections within a document and titles of Quartus II Help topics. For example, “Typographic Conventions.”
Courier type	Indicates signal, port, register, bit, block, and primitive names. For example, data1, tdi, and input. The suffix n denotes an active-low signal. For example, resetn. Indicates command line commands and anything that must be typed exactly as it appears. For example, c:\qdesigns\tutorial\chiptrip.gdf. Also indicates sections of an actual file, such as a Report File, references to parts of files (for example, the AHDL keyword SUBDESIGN), and logic function names (for example, TRI).
	An angled arrow instructs you to press the Enter key.
1., 2., 3., and a., b., c., and so on	Numbered steps indicate a list of items when the sequence of the items is important, such as the steps listed in a procedure.
	Bullets indicate a list of items when the sequence of the items is not important.
	The hand points to information that requires special attention.
	A question mark directs you to a software help system with related information.
	The feet direct you to another document or website with related information.
	A caution calls attention to a condition or possible situation that can damage or destroy the product or your work.
	A warning calls attention to a condition or possible situation that can cause you injury.
	The envelope links to the Email Subscription Management Center page of the Altera website, where you can sign up to receive update notifications for Altera documents.



Quartus II Handbook Version 10.1 Volume 3: Verification



101 Innovation Drive
San Jose, CA 95134
www.altera.com

QII5V3-10.1.0

© 2010 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, HARDCOPY, MAX, MEGACORE, NIOS, QUARTUS and STRATIX are Reg. U.S. Pat. & Tm. Off. and/or trademarks of Altera Corporation in the U.S. and other countries. All other trademarks and service marks are the property of their respective holders as described at www.altera.com/common/legal.html. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.



Chapter Revision Dates	xxi
-------------------------------------	-----

Section I. Simulation

Chapter 1. Simulating Altera Designs

Design Flow	1-2
RTL Functional Simulation Flow	1-3
Gate-Level Timing Simulation Flow	1-4
Simulation Netlist Files	1-5
EDA Simulation Library Compiler	1-8
Running the EDA Simulation Library Compiler Through the GUI	1-8
Running the EDA Simulation Library Compiler from the Command Line	1-9
Using the NativeLink Feature	1-10
Setting Up the EDA Simulator Execution Path	1-10
Configuring NativeLink Settings	1-11
Setting Up Testbench Files Using the NativeLink Feature	1-13
Simulating Altera IP Cores	1-14
IP Simulation Flows	1-15
IP Variant Directory Structure	1-15
Synthesis Files	1-15
Simulation Model Files	1-16
Instantiate the IP Core	1-17
Perform Functional Simulation	1-18
Verilog HDL and VHDL IP Functional Simulation Models	1-18
Simulating Altera IP Cores Using the Quartus II NativeLink Feature	1-19
Using the EDA Simulation Library Compiler	1-20
Running RTL Functional Simulation Using the NativeLink Feature	1-20
Running Gate-Level Timing Simulation Using the NativeLink Feature	1-21
Simulating Qsys and SOPC Builder System Designs	1-21
Document Revision History	1-22

Chapter 2. Mentor Graphics ModelSim and QuestaSim Support

Software Requirements	2-2
Design Flow with ModelSim-Altera, ModelSim, or QuestaSim Software	2-2
Simulation Libraries	2-2
Precompiled Simulation Libraries in the ModelSim-Altera Software	2-3
Simulation Library Files in the Quartus II Software	2-3
Disabling Timing Violation on Registers	2-3
Performing Simulation Using the ModelSim-Altera Software	2-3
Performing Functional Simulation	2-4
Setting Up a Quartus II Project for the ModelSim-Altera Software	2-4
Compiling and Loading Designs with the ModelSim-Altera Software	2-4
Performing the Simulation	2-4
Performing Post-Synthesis Simulation	2-4
Performing Gate-Level Timing Simulation	2-5
Performing Simulation Using the ModelSim and QuestaSim Software	2-5
Simulating VHDL Designs Using the GUI	2-5
Performing Functional Simulation	2-5

Performing Post-Synthesis Simulation	2-6
Performing Gate-Level Simulation	2-6
Simulating Verilog HDL Designs Using the GUI	2-7
Performing Functional Simulation	2-7
Performing Post-Synthesis Simulation	2-8
Performing Gate-Level Simulation	2-8
Simulating VHDL Designs From the Command Line	2-8
Performing Functional Simulation	2-9
Performing Post-Synthesis Simulation	2-10
Performing Gate-Level Simulation	2-11
Simulating Verilog HDL Designs from the Command Line	2-12
Performing Functional Simulation	2-12
Performing Post-Synthesis Simulation	2-13
Performing Gate-Level Simulation	2-14
Passing Parameter Information from Verilog to VHDL	2-15
Speeding Up Simulation	2-15
Simulating Designs that Include Transceivers	2-15
Functional Simulation for Stratix GX Devices	2-16
Performing Functional Simulation in VHDL (ModelSim-Altera)	2-16
Performing Functional Simulation in VHDL (ModelSim and QuestaSim)	2-16
Performing Functional Simulation in Verilog HDL (ModelSim-Altera)	2-17
Performing Functional Simulation in Verilog HDL (ModelSim and QuestaSim)	2-17
Gate-Level Timing Simulation for Stratix GX Devices	2-17
Performing Gate-Level Timing Simulation in VHDL (ModelSim-Altera)	2-17
Performing Gate-Level Timing Simulation in VHDL (ModelSim and QuestaSim)	2-17
Performing Gate-Level Timing Simulation in Verilog HDL (ModelSim-Altera)	2-18
Performing Gate-Level Timing Simulation in Verilog HDL (ModelSim and QuestaSim)	2-18
Functional Simulation for Stratix II GX Devices	2-18
Performing Functional Simulation in VHDL (ModelSim-Altera)	2-18
Performing Functional Simulation in VHDL (ModelSim and QuestaSim)	2-19
Performing Functional Simulation in Verilog HDL (ModelSim-Altera)	2-19
Performing Functional Simulation in Verilog HDL (ModelSim and QuestaSim)	2-19
Gate-Level Timing Simulation for Stratix II GX Devices	2-19
Performing Gate-Level Timing Simulation in VHDL (ModelSim-Altera)	2-19
Performing Gate-Level Timing Simulation in VHDL (ModelSim and QuestaSim)	2-20
Performing Gate-Level Timing Simulation in Verilog HDL ModelSim-Altera)	2-20
Performing Gate-Level Timing Simulation in Verilog HDL ModelSim and QuestaSim)	2-20
Functional Simulation for Stratix IV GX Devices	2-20
Performing Functional Simulation in VHDL (ModelSim-Altera)	2-20
Performing Functional Simulation in VHDL (ModelSim and QuestaSim)	2-21
Performing Functional Simulation in Verilog HDL (ModelSim-Altera)	2-21
Performing Functional Simulation in Verilog HDL (ModelSim and QuestaSim)	2-21
Gate-Level Timing Simulation for Stratix IV GX Devices	2-21
Performing Gate-Level Timing Simulation in VHDL (ModelSim-Altera)	2-21
Performing Gate-Level Timing Simulation in VHDL (ModelSim and QuestaSim)	2-22
Performing Gate-Level Timing Simulation in Verilog HDL (ModelSim-Altera)	2-22
Performing Gate-Level Timing Simulation in Verilog HDL (ModelSim and QuestaSim)	2-22
Functional Simulation for Stratix V GX Devices	2-22
Performing Functional Simulation in VHDL (ModelSim and QuestaSim)	2-22
Performing Functional Simulation in Verilog HDL (ModelSim-Altera)	2-23
Performing Functional Simulation in Verilog HDL (ModelSim and QuestaSim)	2-23
Transport Delays	2-23
+transport_path_delays	2-23
+transport_int_delays	2-23

Using the NativeLink Feature with ModelSim-Altera, ModelSim, or QuestaSim Software	2-23
ModelSim and QuestaSim Error Message Verification	2-24
Generating a Timing Value Change Dump (.vcd) File for the PowerPlay Power Analyzer	2-24
Viewing a Waveform from a .wlf File	2-25
Simulating with ModelSim-Altera Waveform	2-25
Scripting Support	2-26
Generating a Post-Synthesis Simulation Netlist for ModelSim and QuestaSim	2-26
Tcl Commands	2-26
Command Prompt	2-26
Generating a Gate-Level Timing Simulation Netlist for ModelSim and QuestaSim	2-26
Tcl Commands	2-27
Command Line	2-27
Software Licensing and Licensing Setup in ModelSim-Altera Subscription Edition	2-27
LM_LICENSE_FILE Variable	2-28
Conclusion	2-29
Document Revision History	2-29

Chapter 3. Synopsys VCS and VCS MX Support

Software Requirements	3-1
Using the VCS or VCS MX Software in the Quartus II Design Flow	3-1
Compiling Libraries Using the EDA Simulation Library Compiler	3-2
Functional Simulations	3-2
Functional Simulation for Verilog HDL Designs	3-2
Functional Simulation for VHDL Designs	3-3
Post-Synthesis Simulation	3-4
Post-Synthesis Simulation for Verilog HDL Designs	3-5
Post-Synthesis Simulation for VHDL Designs	3-5
Gate-Level Timing Simulation	3-6
Gate-Level Timing Simulation for Verilog HDL Designs	3-6
Gate-Level Timing Simulation for VHDL Designs	3-6
Disabling Timing Violation on Registers	3-7
Performing Timing Simulation Using the Post-Synthesis Netlist	3-7
Common VCS and VCS MX Software Compiler Options	3-8
Using DVE	3-8
Debugging Support Command-Line Interface	3-8
Simulating Designs that Include Transceivers	3-9
Functional Simulation for Stratix GX Devices	3-9
Compiling Library Files for Functional Simulation in Verilog HDL	3-9
Gate-Level Timing Simulation for Stratix GX Devices	3-9
Compiling Library Files for Gate-Level Timing Simulation in Verilog HDL	3-10
Functional Simulation for Stratix II GX Devices	3-10
Compiling Library Files for Functional Simulation in Verilog HDL	3-10
Gate-Level Timing Simulation for Stratix II GX Devices	3-10
Compiling Library Files for Gate-Level Timing Simulation in Verilog HDL	3-11
Functional Simulation for Stratix IV GX Devices	3-11
Compiling Library Files for Functional Simulation in Verilog HDL	3-11
Gate-Level Timing Simulation for Stratix IV GX Devices	3-11
Compiling Library Files for Gate-Level Timing Simulation in Verilog HDL	3-12
Functional Simulation for Stratix V GX Devices	3-12
Compiling Library Files for Functional Simulation	3-12
Transport Delays	3-13
+transport_path_delays	3-13
+transport_int_delays	3-13
Using NativeLink with the VCS or VCS MX Software	3-13

Generating a Timing .vcd File for the PowerPlay Power Analyzer	3-13
Viewing a Waveform from a .vpd or .vcd File	3-14
Scripting Support	3-15
Generating a Post-Synthesis Simulation Netlist for VCS	3-15
Tcl Commands	3-15
Command Prompt	3-15
Generating a Gate-Level Timing Simulation Netlist for VCS	3-16
Tcl Commands	3-16
Command Prompt	3-16
Conclusion	3-16
Document Revision History	3-16

Chapter 4. Cadence NC-Sim Support

Software Requirements	4-1
Simulation Flow Overview	4-1
Operation Modes	4-2
Quartus II Software and NC Simulation Flow Overview	4-2
Compiling Libraries Using the EDA Simulation Library Compiler	4-3
Functional Simulation	4-3
Creating Libraries	4-4
For VHDL Designs	4-4
For Verilog HDL Designs	4-4
Compiling Source Code	4-5
Elaborating Your Design	4-6
Simulating Your Design	4-6
Post-Synthesis Simulation	4-7
Quartus II Simulation Output Files	4-7
Creating Libraries	4-7
Compiling Project Files and Libraries	4-7
Elaborating Your Design	4-8
Simulating Your Design	4-8
Gate-Level Timing Simulation	4-8
Generating a Gate-Level Timing Simulation Netlist	4-8
Disabling Timing Violation on Registers	4-8
Creating Libraries	4-9
Compiling Project Files and Libraries	4-9
Elaborating Your Design	4-9
Compiling the .sdo File (VHDL Only) in Command-Line Mode	4-10
Compiling the .sdo File (VHDL Only) in GUI Mode	4-10
Simulating Your Design	4-10
Simulating Designs that Include Transceivers	4-10
Functional Simulation for Stratix GX Devices	4-11
Compiling Library Files for Functional Simulation in VHDL	4-11
Compiling Library Files for Functional Simulation in Verilog HDL	4-11
Gate-Level Timing Simulation for Stratix GX Devices	4-11
Compiling Library Files for Gate-Level Timing Simulation in VHDL	4-12
Compiling Library Files for Gate-Level Timing Simulation in Verilog HDL	4-12
Functional Simulation for Stratix II GX Devices	4-12
Compiling Library Files for Functional Simulation in VHDL	4-13
Compiling Library Files for Functional Simulation in Verilog HDL	4-13
Gate-Level Timing Simulation for Stratix II GX Devices	4-13
Compiling Library Files for Gate-Level Timing Simulation in VHDL	4-14
Compiling Library Files for Gate-Level Timing Simulation in Verilog HDL	4-14
Functional Simulation for Stratix IV GX Devices	4-14

Compiling Library Files for Functional Simulation in VHDL	4-15
Compiling Library Files for Functional Simulation in Verilog HDL	4-15
Gate-Level Timing Simulation for Stratix IV GX Devices	4-15
Compiling Library Files for Gate-Level Timing Simulation in VHDL	4-16
Compiling Library Files for Gate-Level Timing Simulation in Verilog HDL	4-16
Functional Simulation for Stratix V GX Devices	4-16
Compiling Library Files for Functional Simulation	4-17
Compiling Library Files for Functional Simulation in VHDL	4-17
Compiling Library Files for Functional Simulation in Verilog HDL	4-17
Pulse Reject Delays	4-18
-PULSE_R	4-18
-PULSE_INT_R	4-18
Using the NativeLink Feature with NC-Sim	4-18
Generating a Timing VCD File for the PowerPlay Power Analyzer	4-18
Viewing a Waveform from a .trn File	4-19
Scripting Support	4-20
Generating NC-Sim Simulation Output Files	4-20
Tcl Commands	4-20
Command Prompt	4-21
Conclusion	4-21
Document Revision History	4-21

Chapter 5. Aldec Active-HDL and Riviera-PRO Support

Software Requirements	5-1
Using Active-HDL or Riviera-PRO Software in Quartus II Design Flows	5-1
Simulation Libraries	5-2
Simulation Library Files in the Quartus II Software	5-2
Disabling Timing Violation on Registers	5-2
Compiling Libraries Using the EDA Simulation Library Compiler	5-3
Performing Simulation Using the Active-HDL Software (GUI Mode)	5-3
Simulating VHDL Designs	5-3
Performing Functional Simulation	5-3
Simulating Verilog HDL Designs	5-4
Performing Functional Simulation	5-4
Performing Simulation Using the Riviera-PRO Software (GUI Mode)	5-5
Performing Simulation Using the Active-HDL or Riviera-PRO Software (Command-Line Mode)	5-6
Simulating Verilog HDL Designs Using the Active-HDL or Riviera-PRO Software	5-7
Performing Functional Simulation	5-7
Performing Post-Synthesis Simulation	5-7
Performing Gate-Level Timing Simulation	5-8
Simulating VHDL Designs Using the Active-HDL or Riviera-PRO Software	5-9
Performing Functional Simulation	5-9
Performing Post-Synthesis Simulation	5-10
Performing Gate-Level Timing Simulation	5-10
Compiling System Verilog Files	5-11
Simulating Designs that Include Transceivers	5-11
Functional Simulation for Stratix II GX Devices	5-11
Performing Functional Simulation in VHDL	5-12
Performing Functional Simulation in Verilog HDL	5-12
Gate-Level Timing Simulation for Stratix II GX Devices	5-12
Performing Gate-Level Timing Simulation in VHDL	5-13
Performing Gate-Level Timing Simulation in Verilog HDL	5-13
Functional Simulation for Stratix GX Devices	5-13
Performing Functional Simulation in VHDL	5-13

Performing Functional Simulation in Verilog HDL	5-14
Gate-Level Timing Simulation for Stratix GX Devices	5-14
Performing Gate-Level Timing Simulation in VHDL	5-14
Performing Gate-Level Timing Simulation in Verilog HDL	5-14
Functional Simulation for Stratix IV GX Devices	5-14
Performing Functional Simulation in VHDL	5-15
Performing Functional Simulation in Verilog HDL	5-15
Gate-Level Timing Simulation for Stratix IV GX Devices	5-15
Performing Gate-Level Timing Simulation in VHDL	5-16
Performing Gate-Level Timing Simulation in Verilog HDL	5-16
Functional Simulation for Stratix V GX Devices	5-16
Performing Functional Simulation in VHDL	5-17
Performing Functional Simulation in Verilog HDL	5-17
Transport Delays	5-17
Using the NativeLink Feature in Active-HDL or Riviera-PRO Software	5-18
Generating .vcd Files for the PowerPlay Power Analyzer	5-18
Scripting Support	5-19
Generating a Post-Synthesis Simulation Netlist for Active-HDL or Riviera-PRO	5-19
Tcl Commands	5-19
Command Line	5-19
Generating a Gate-Level Timing Simulation Netlist for Active-HDL or Riviera-PRO	5-20
Tcl Commands	5-20
Command Line	5-20
Conclusion	5-20
Document Revision History	5-21

Section II. Timing Analysis

Chapter 6. The Quartus II TimeQuest Timing Analyzer

Understanding Timing Analysis with the TimeQuest Analyzer	6-1
Timing Netlists and Timing Paths	6-2
The Timing Netlist	6-2
Timing Paths	6-3
Launch and Latch Edges	6-4
Data and Clock Arrival Times	6-4
Clock Setup Check	6-5
Clock Hold Check	6-6
Recovery and Removal Time	6-7
Multicycle Paths	6-9
Metastability	6-11
Common Clock Path Pessimism Removal	6-11
Clock-As-Data Analysis	6-13
Multicorner Analysis	6-15
Getting Started with the TimeQuest Analyzer	6-16
Running the TimeQuest Analyzer	6-17
Locating Timing Paths in Other Tools	6-18
Recommended Flows	6-19
Creating and Setting Up your Design	6-20
Performing an Initial Compilation	6-20
Specifying Timing Requirements	6-21
Fitting and Timing Analysis with .sdc Files	6-22
Performing a Full Compilation	6-22
Verifying Timing	6-22

SDC File Precedence	6-22
Using Tcl Commands	6-23
Wildcard Characters	6-24
Collection Commands	6-24
Adding and Removing Collection Items	6-24
Refining Collections with Wildcards	6-25
Removing Constraints and Exceptions	6-27
Creating Clocks and Clock Constraints	6-27
Creating Clocks	6-27
Creating Virtual Clocks	6-28
Creating Multifrequency Clocks	6-28
Creating Generated Clocks	6-29
Automatically Detecting Clocks and Creating Default Clock Constraints	6-30
Deriving PLL Clocks	6-31
Creating Clock Groups	6-33
Exclusive Clock Groups	6-33
Asynchronous Clock Groups	6-33
Accounting for Clock Effect Characteristics	6-34
Clock Latency	6-34
Clock Uncertainty	6-34
I/O Interface Uncertainty	6-35
Creating I/O Constraints	6-37
Creating Delay and Skew Constraints	6-38
Net Delay	6-38
Advanced I/O Timing and Board Trace Model Delay	6-38
Maximum Skew	6-38
Creating Timing Exceptions	6-39
Precedence	6-39
False Paths	6-39
Minimum and Maximum Delays	6-39
Delay Annotation	6-40
Timing Reports	6-40
Document Revision History	6-42

Chapter 7. Best Practices for the Quartus II TimeQuest Timing Analyzer

Creating Clock Requirements	7-2
Base Clocks	7-2
Derived Clocks	7-3
Virtual Clocks	7-3
Creating I/O Requirements	7-5
Input Constraints	7-6
Output Constraints	7-6
Creating Timing Exceptions	7-7
False Paths	7-7
Minimum and Maximum Delays	7-8
Creating Multicycle Exceptions	7-8
Multicycle Clock Setup Check and Hold Check Analysis	7-9
Multicycle Clock Setup	7-10
Multicycle Clock Hold	7-12
Examples of Basic Multicycle Exceptions	7-14
Default Settings	7-15
End Multicycle Setup = 2 and End Multicycle Hold = 0	7-17
End Multicycle Setup = 1 and End Multicycle Hold = 1	7-20
End Multicycle Setup = 2 and End Multicycle Hold = 1	7-23

Start Multicycle Setup = 2 and Start Multicycle Hold = 0	7-26
Start Multicycle Setup = 1 and Start Multicycle Hold = 1	7-29
Start Multicycle Setup = 2 and Start Multicycle Hold = 1	7-32
Application of Multicycle Exceptions	7-35
Same Frequency Clocks with Destination Clock Offset	7-35
The Destination Clock Frequency is a Multiple of the Source Clock Frequency	7-37
The Destination Clock Frequency is a Multiple of the Source Clock Frequency with an Offset	7-40
The Source Clock Frequency is a Multiple of the Destination Clock Frequency	7-42
The Source Clock Frequency is a Multiple of the Destination Clock Frequency with an Offset	7-45
Document Revision History	7-47

Chapter 8. Switching to the Quartus II TimeQuest Timing Analyzer

Benefits of Switching to the TimeQuest Analyzer	8-1
Switching Your Design	8-2
Open Your Compiled Design	8-2
Create an SDC Constraints	8-2
Create SDC Constraints Manually	8-2
Create SDC Constraints from Existing Timing Assignments	8-2
Start the TimeQuest Analyzer	8-3
Differences Between the Quartus II Timing Analyzers	8-3
Terminology	8-3
Netlists	8-3
Collections	8-4
Constraints	8-5
Constraint Files	8-5
Constraint Entry	8-5
Time Units	8-6
MegaCore Functions	8-6
Bus Name Format	8-6
Constraint File Priority	8-7
Constraint Priority	8-8
Ambiguous Constraints	8-9
Clocks	8-10
Related and Unrelated Clocks	8-10
Clock Offset	8-10
Clock Latency	8-11
Offset and Latency Example	8-11
Clock Uncertainty	8-13
Derived and Generated Clocks	8-14
Automatic Clock Detection	8-14
Hold Relationship	8-17
Clock Objects	8-18
Hold Multicycle	8-18
Fitter Performance and Behavior	8-19
Reporting	8-20
Paths and Pairs	8-20
Default Reports	8-21
Netlist Names	8-21
Non-Integer Clock Periods	8-21
Other Features	8-23
Timing Assignment Conversion	8-24
Setup Relationship	8-24

Hold Relationship	8–25
Clock Latency	8–25
Clock Uncertainty	8–25
Inverted Clock	8–25
Not a Clock	8–26
Default Required f_{MAX} Assignment	8–26
Virtual Clock Reference	8–26
Clock Settings	8–27
Multicycle	8–27
Clock Enable Multicycle	8–28
I/O Constraints	8–28
Input and Output Delay	8–29
t_{SU} Requirement	8–30
t_H Requirement	8–32
t_{CO} Requirement	8–34
Minimum t_{CO} Requirement	8–36
t_{PD} Requirement	8–38
Minimum t_{PD} Requirement	8–40
Cut Timing Path	8–40
Maximum Delay	8–40
Minimum Delay	8–41
Maximum Clock Arrival Skew	8–41
Maximum Data Arrival Skew	8–41
Constraining Skew on an Output Bus	8–41
Conversion Utility	8–43
Unsupported Global Assignments	8–43
Recommended Global Assignments	8–43
Clock Conversion	8–45
Instance Assignment Conversion	8–46
PLL Phase Shift Conversion	8–48
t_{CO} Requirement Conversion	8–49
Entity-Specific Assignments	8–49
Paths Between Unrelated Clock Domains	8–49
Unsupported Instance Assignments	8–50
Reviewing Conversion Results	8–50
Warning Messages	8–50
Clocks	8–52
Clock Transfers	8–52
Path Details	8–52
Unconstrained Paths	8–52
Bus Names	8–53
Other	8–53
Rerunning the Conversion Utility	8–53
Notes	8–53
Output Pin Load Assignments	8–53
Constraint Target Types	8–53
DDR Constraints with the DDR Timing Wizard	8–54
Unsupported SDC Features	8–54
Constraint Passing and Optimization	8–54
Clock Network Delay Reporting	8–54
Project Management	8–55
Conversion Utility	8–55
t_{PD} and Minimum t_{PD} Requirement Conversion	8–55
Document Revision History	8–55

Chapter 9. Synopsys PrimeTime Support

Quartus II Settings for Generating the PrimeTime Software Files	9-1
Files Generated for the PrimeTime Software Environment	9-2
The Netlist	9-2
The .sdo File	9-2
Generating Multiple Operating Conditions with the TimeQuest Analyzer	9-3
The Tcl Script	9-4
Generated File Summary	9-6
Running the PrimeTime Software	9-6
Analyzing Quartus II Projects	9-7
Other pt_shell Commands	9-7
PrimeTime Timing Reports	9-7
Sample PrimeTime Software Timing Report	9-7
Comparing Timing Reports from the Classic Timing Analyzer and the PrimeTime Software	9-8
Clock Setup Relationship and Slack	9-9
Clock Hold Relationship and Slack	9-12
Input Delay and Output Delay Relationships and Slack	9-16
Static Timing Analyzer Differences	9-18
Classic Timing Analyzer and PrimeTime Software	9-18
Rise/Fall Support	9-18
Minimum and Maximum Delays	9-18
Recovery/Removal Analysis	9-18
Encrypted Intellectual Property Blocks	9-18
Registered Clock Signals	9-19
Multiple Source and Destination Register Pairs	9-19
Latches	9-19
LVDS I/O	9-20
Clock Latency	9-20
Input and Output Delay Assignments	9-20
Generated Clocks Derived from Generated Clocks	9-20
TimeQuest Timing Analyzer and PrimeTime Software	9-20
Encrypted Intellectual Property Blocks	9-20
Latches	9-21
LVDS I/O	9-21
The TimeQuest Timing Analyzer .sdc File and PrimeTime Compatibility	9-21
Clock and Data Paths	9-21
Inverting and Non-Inverting Propagation	9-21
Multiple Rise/Fall Numbers For a Timing Arc	9-21
Virtual Generated Clocks	9-22
Generated Clocks Derived from Generated Clocks	9-22
Conclusion	9-22
Document Revision History	9-22

Section III. Power Estimation and Analysis

Chapter 10. PowerPlay Power Analysis

Types of Power Analyses	10-2
Factors Affecting Power Consumption	10-2
Device Selection	10-2
Environmental Conditions	10-3
Airflow	10-3
Heat Sink and Thermal Compound	10-3
Junction Temperature	10-3

Board Thermal Model	10-3
Device Resource Usage	10-3
Number, Type, and Loading of I/O Pins	10-4
Number and Type of Logic Elements, Multiplier Elements, and RAM Blocks	10-4
Number and Type of Global Signals	10-4
Signal Activities	10-4
Creating PowerPlay EPE Spreadsheets	10-5
PowerPlay EPE File Generator Compilation Report	10-6
PowerPlay Power Analyzer Flow	10-7
Operating Settings and Conditions	10-7
Signal Activities Data Sources	10-8
Simulation Results	10-9
Using Simulation Files in Modular Design Flows	10-10
Complete Design Simulation	10-11
Modular Design Simulation	10-11
Multiple Simulations on the Same Entity	10-12
Overlapping Simulations	10-12
Partial Simulations	10-13
Node Name Matching Considerations	10-13
Glitch Filtering	10-13
Node and Entity Assignments	10-15
Timing Assignments to Clock Nodes	10-16
Default Toggle Rate Assignment	10-16
Vectorless Estimation	10-16
Using the PowerPlay Power Analyzer	10-16
Common Analysis Flows	10-17
Signal Activities from Full Post-Fit Netlist (Timing) Simulation	10-17
Signal Activities from Full Post-Fit Netlist (Zero Delay) Simulation	10-17
Signal Activities from RTL (Functional) Simulation, Supplemented by Vectorless Estimation	10-17
Signal Activities from Vectorless Estimation and User-Supplied Input Pin Activities	10-17
Signal Activities from User Defaults Only	10-17
Generating a .vcd	10-18
Generating a .vcd from ModelSim Software	10-19
Generating a .vcd from Full Post-Fit Netlist (Zero Delay) Simulation	10-19
Running the PowerPlay Power Analyzer Using the Quartus II GUI	10-20
PowerPlay Power Analyzer Compilation Report	10-20
Summary	10-20
Settings	10-20
Simulation Files Read	10-20
Operating Conditions Used	10-21
Thermal Power Dissipated by Block	10-21
Thermal Power Dissipation by Block Type (Device Resource Type)	10-21
Thermal Power Dissipation by Hierarchy	10-21
Core Dynamic Thermal Power Dissipation by Clock Domain	10-21
Current Drawn from Voltage Supplies	10-21
Confidence Metric Details	10-22
Signal Activities	10-22
Messages	10-22
Specific Rules for Reporting	10-22
Scripting Support	10-22
Running the PowerPlay Power Analyzer from the Command-Line	10-23
Conclusion	10-24
Document Revision History	10-24

Section IV. System Debugging Tools

Chapter 11. System Debugging Tools Overview

System Debugging Tools	11-1
Analysis Tools for RTL Nodes	11-4
Resource Usage	11-4
Pin Usage	11-5
Usability Enhancements	11-6
Stimulus-Capable Tools	11-8
In-System Sources and Probes	11-8
In-System Memory Content Editor	11-8
Virtual JTAG Interface Megafunction	11-9
System Console	11-9
Conclusion	11-9
Document Revision History	11-10

Chapter 12. Analyzing and Debugging Designs with the System Console

Introduction	12-1
System Console Overview	12-1
Finding and Referring To Services	12-1
Accessing the Service Life Cycle	12-2
Applying Services	12-2
Setting Up the System Console	12-3
Interactive Help	12-3
Using the System Console	12-3
Qsys and SOPC Builder Communications	12-4
Console Commands	12-6
Plugins	12-8
Design Service Commands	12-8
Data Pattern Generator Commands	12-9
Data Pattern Checker Commands	12-10
Programmable Logic Device (PLD) Commands	12-10
Board Bring-Up Commands	12-11
JTAG Debug Commands	12-11
Clock and Reset Signal Commands	12-12
Avalon-MM Commands	12-12
Processor Commands	12-14
Bytestream Commands	12-14
Transceiver Toolkit Commands	12-15
In-System Sources and Probes Commands	12-20
Monitor Commands	12-21
Dashboard Commands	12-22
Specifying Widgets	12-23
Customizing Widgets	12-24
Assigning Dashboard Widget Properties	12-24
System Console Examples	12-27
LED Light Show Example	12-28
Loading and Linking a Design	12-30
JTAG Examples	12-30
Verify JTAG Chain	12-30
Verify Clock	12-32
Checksum Example	12-33
Nios II Processor Example	12-36

Device Support	12-37
Conclusion	12-37
Document Revision History	12-37

Chapter 13. Transceiver Link Debugging Using the System Console

Transceiver Toolkit Overview	13-1
Auto Sweep	13-1
EyeQ	13-2
Control Links	13-2
Transceiver Link Debugging Design Examples	13-2
Setting Up Tests for Link Debugging	13-3
Custom PHY IP Core	13-4
Low Latency PHY IP Core	13-5
Avalon-ST Data Pattern Generator	13-5
Data Checker	13-6
Compiling Design Examples	13-7
Changing Pin Assignments	13-7
Transceiver Toolkit Link Test Setup	13-8
Loading the Project in System Console	13-8
Linking the Hardware Resource	13-8
Creating the Channels	13-9
Running the Link Tests	13-9
Viewing Results in the EyeQ Feature	13-10
Tcl Script in System Console	13-11
Running Tcl Scripts	13-12
Usage Scenarios	13-12
Linking One Design to One Device Connected By One USB Blaster Cable	13-13
Linking Two Designs to Two Separate Devices on Same Board (JTAG Chained), Connected By One USB Blaster Cable	13-13
Linking Two Designs to Two Separate Devices on Separate Boards, Connected to Separate USB Blaster Cables	13-13
Linking Same Design on Two Separate Devices	13-14
Linking Unrelated Designs	13-14
Saving Your Setup As Tcl Scripts	13-14
Verifying Channels Are Correct When Creating Link	13-14
Using the Recommended DFE Flow	13-15
Running Simultaneous Tests	13-15
Conclusion	13-16
Document Revision History	13-16

Chapter 14. Quick Design Debugging Using SignalProbe

Debugging Using the SignalProbe Feature	14-1
Reserve the SignalProbe Pins	14-2
Perform a Full Compilation	14-2
Assign a SignalProbe Source	14-2
Add Registers to the Pipeline Path to SignalProbe Pin	14-3
Perform a SignalProbe Compilation	14-3
Analyze the Results of the SignalProbe Compilation	14-4
Performing a SignalProbe Compilation	14-4
Understanding the Results of a SignalProbe Compilation	14-5
Analyzing SignalProbe Routing Failures	14-6
Scripting Support	14-6
Make a SignalProbe Pin	14-6

Delete a SignalProbe Pin	14-7
Enable a SignalProbe Pin	14-7
Disable a SignalProbe Pin	14-7
Perform a SignalProbe Compilation	14-7
Script Example	14-7
Reserving SignalProbe Pins	14-7
Common Problems When Reserving a SignalProbe Pin	14-7
Adding SignalProbe Sources	14-8
Assigning I/O Standards	14-8
Adding Registers for Pipelining	14-8
Run SignalProbe Automatically	14-9
Run SignalProbe Manually	14-9
Enable or Disable All SignalProbe Routing	14-9
Allow SignalProbe to Modify Fitting Results	14-9
Conclusion	14-10
Document Revision History	14-10

Chapter 15. Design Debugging Using the SignalTap II Logic Analyzer

Hardware and Software Requirements	15-2
Design Flow Using the SignalTap II Logic Analyzer	15-4
SignalTap II Logic Analyzer Task Flow	15-5
Add the SignalTap II Logic Analyzer to Your Design	15-5
Configure the SignalTap II Logic Analyzer	15-6
Define Trigger Conditions	15-6
Compile the Design	15-6
Program the Target Device or Devices	15-6
Run the SignalTap II Logic Analyzer	15-7
View, Analyze, and Use Captured Data	15-7
Embedding Multiple Analyzers in One FPGA	15-7
Monitoring FPGA Resources Used by the SignalTap II Logic Analyzer	15-7
Using the MegaWizard Plug-In Manager to Create Your Logic Analyzer	15-8
Configure the SignalTap II Logic Analyzer	15-8
Assigning an Acquisition Clock	15-9
Adding Signals to the SignalTap II File	15-9
Signal Preservation	15-10
Assigning Data Signals Using the Technology Map Viewer	15-11
Node List Signal Use Options	15-11
Untappable Signals	15-12
Adding Signals with a Plug-In	15-12
Adding Finite State Machine State Encoding Registers	15-13
Modifying and Restoring Mnemonic Tables for State Machines	15-14
Additional Considerations	15-14
Specifying the Sample Depth	15-14
Capturing Data to a Specific RAM Type	15-15
Choosing the Buffer Acquisition Mode	15-15
Non-Segmented Buffer	15-16
Segmented Buffer	15-16
Using the Storage Qualifier Feature	15-17
Input Port Mode	15-19
Transitional Mode	15-20
Conditional Mode	15-20
Start/Stop Mode	15-22
State-Based	15-23
Showing Data Discontinuities	15-23

Disable Storage Qualifier	15–23
Managing Multiple SignalTap II Files and Configurations	15–23
Define Triggers	15–25
Creating Basic Trigger Conditions	15–25
Creating Advanced Trigger Conditions	15–25
Examples of Advanced Triggering Expressions	15–27
Trigger Condition Flow Control	15–28
Sequential Triggering	15–28
State-Based Triggering	15–29
SignalTap II Trigger Flow Description Language	15–32
State Labels	15–33
Boolean_expression	15–33
Action_list	15–34
Resource Manipulation Action	15–34
Buffer Control Action	15–35
State Transition Action	15–35
Using the State-Based Storage Qualifier Feature	15–35
Specifying the Trigger Position	15–39
Creating a Power-Up Trigger	15–40
Enabling a Power-Up Trigger	15–41
Managing and Configuring Power-Up and Runtime Trigger Conditions	15–41
Using External Triggers	15–42
Using the Trigger Out of One Analyzer as the Trigger In of Another Analyzer	15–42
Compile the Design	15–44
Faster Compilations with Quartus II Incremental Compilation	15–45
Enabling Incremental Compilation for Your Design	15–45
Using Incremental Compilation with the SignalTap II Logic Analyzer	15–45
Preventing Changes Requiring Recompile	15–48
Timing Preservation with the SignalTap II Logic Analyzer	15–48
Performance and Resource Considerations	15–48
Program the Target Device or Devices	15–49
Run the SignalTap II Logic Analyzer	15–50
Runtime Reconfigurable Options	15–52
SignalTap II Status Messages	15–55
View, Analyze, and Use Captured Data	15–55
Capturing Data Using Segmented Buffers	15–56
Differences in Pre-fill Write Behavior Between Different Acquisition Modes	15–57
Creating Mnemonics for Bit Patterns	15–59
Automatic Mnemonics with a Plug-In	15–59
Locating a Node in the Design	15–60
Saving Captured Data	15–61
Exporting Captured Data to Other File Formats	15–61
Creating a SignalTap II List File	15–61
Other Features	15–61
Using the SignalTap II MATLAB MEX Function to Capture Data	15–62
Using SignalTap II in a Lab Environment	15–63
Remote Debugging Using the SignalTap II Logic Analyzer	15–63
Equipment Setup	15–64
Using the SignalTap II Logic Analyzer in Devices with Configuration Bitstream Security	15–64
Backward Compatibility with Previous Versions of Quartus II Software	15–64
SignalTap II Command-Line Options	15–65
SignalTap II Tcl Commands	15–66
Design Example: Using SignalTap II Logic Analyzers in SOPC Builder Systems	15–66
Custom Triggering Flow Application Examples	15–67

Design Example 1: Specifying a Custom Trigger Position	15–67
Design Example 2: Trigger When triggercond1 Occurs Ten Times between triggercond2 and triggercond3	15–68
SignalTap II Scripting Support	15–69
Conclusion	15–69
Document Revision History	15–70

Chapter 16. In-System Debugging Using External Logic Analyzers

Choosing a Logic Analyzer	16–1
Required Components	16–2
Debugging Your Design Using the LAI	16–4
Working with LAI Files	16–4
Configuring the File Core Parameters	16–5
Mapping the LAI File Pins to Available I/O Pins	16–5
Mapping Internal Signals to the LAI Banks	16–5
Using the Node Finder	16–6
Compiling Your Quartus II Project	16–6
Programming Your Altera-Supported Device Using the LAI	16–6
Controlling the Active Bank During Runtime	16–7
Acquiring Data on Your Logic Analyzer	16–7
Using the LAI with Incremental Compilation	16–7
Conclusion	16–8
Document Revision History	16–8

Chapter 17. In-System Modification of Memory and Constants

Overview	17–1
Updating Memory and Constants in Your Design	17–2
Creating In-System Modifiable Memories and Constants	17–2
Running the In-System Memory Content Editor	17–2
Instance Manager	17–3
Editing Data Displayed in the Hex Editor Pane	17–3
Importing and Exporting Memory Files	17–3
Scripting Support	17–4
Programming the Device with the In-System Memory Content Editor	17–4
Example: Using the In-System Memory Content Editor with the SignalTap II Logic Analyzer ..	17–4
Conclusion	17–5
Document Revision History	17–5

Chapter 18. Design Debugging Using In-System Sources and Probes

Overview	18–1
Hardware and Software Requirements	18–3
Design Flow Using the In-System Sources and Probes Editor	18–4
Configuring the ALTSOURCE_PROBE Megafunction	18–4
Instantiating the ALTSOURCE_PROBE Megafunction	18–6
Compiling the Design	18–6
Running the In-System Sources and Probes Editor	18–7
Programming Your Device With JTAG Chain Configuration	18–7
Instance Manager	18–8
In-System Sources and Probes Editor Pane	18–8
Reading Probe Data	18–8
Writing Data	18–9
Organizing Data	18–9
Tcl interface for the In-System Sources and Probes Editor	18–9

Design Example: Dynamic PLL Reconfiguration	18–13
Conclusion	18–16
Document Revision History	18–16

Section V. Formal Verification

Chapter 19. Cadence Encounter Conformal Support

Formal Verification Versus Simulation	19–1
Formal Verification: What You Need to Know	19–2
Formal Verification Design Flow	19–2
Quartus II Integrated Synthesis	19–3
EDA Tool Support for Quartus II Integrated Synthesis	19–3
Synplify Pro	19–3
RTL Coding Guidelines for Quartus II Integrated Synthesis	19–4
Synthesis Directives and Attributes	19–4
Stuck-at Registers	19–5
ROM, LPM_DIVIDE, and Shift Register Inference	19–6
RAM Inference	19–7
Latch Inference	19–7
Combinational Loops	19–7
Finite State Machine Coding Styles	19–8
Black Boxes in the Conformal LEC Flow	19–8
Tcl Command	19–9
GUI	19–9
Generating the Post-Fit Netlist Output File and the Conformal LEC Setup Files	19–10
Quartus II Software Generated Files, Formal Verification Scripts, and Directories	19–11
Understanding the Formal Verification Scripts for Conformal LEC	19–12
Conformal LEC Commands within the Quartus II Software-Generated Scripts	19–12
Comparing Designs Using Conformal LEC	19–15
Running the Conformal LEC Software from the GUI	19–15
Running the Conformal LEC Software From a System Command Prompt	19–16
Known Issues and Limitations	19–16
Black Box Models	19–18
Conformal Dofile/Script Example	19–19
Conclusion	19–21
Document Revision History	19–21

Section VI. Device Programming

Chapter 20. Quartus II Programmer

Programming Flow	20–1
Quartus II Programmer GUI	20–3
Hardware Setup	20–4
JTAG Settings	20–4
JTAG Chain Debugger Tool	20–4
JTAG Chain Debugger Example	20–4
Other Programming Tools	20–5
Stand-Alone Quartus II Programmer	20–5
Programming and Configuration Modes	20–5
Configuration Modes	20–5
Design Security Key Programming	20–6
Generating Secondary Programming Files	20–7
Convert Programming Files Dialog Box	20–7

Flash Loaders	20-10
Scripting Support	20-10
The jtagconfig Debugging Tool	20-11
Conclusion	20-11
Document Revision History	20-11

Additional Information

How to Contact Altera	Info-1
Typographic Conventions	Info-2

The chapters in this document, Quartus II Handbook Version 10.1 Volume 3: Verification, were revised on the following dates. Where chapters or groups of chapters are available separately, part numbers are listed.

- Chapter 1. Simulating Altera Designs
Revised: *December 2010*
Part Number: *QII53025-10.1.0*

- Chapter 2. Mentor Graphics ModelSim and QuestaSim Support
Revised: *December 2010*
Part Number: *QII53001-10.1.0*

- Chapter 3. Synopsys VCS and VCS MX Support
Revised: *December 2010*
Part Number: *QII53002-10.0.1*

- Chapter 4. Cadence NC-Sim Support
Revised: *December 2010*
Part Number: *QII53003-10.0.1*

- Chapter 5. Aldec Active-HDL and Riviera-PRO Support
Revised: *December 2010*
Part Number: *QII53023-10.0.1*

- Chapter 6. The Quartus II TimeQuest Timing Analyzer
Revised: *December 2010*
Part Number: *QII53018-10.1.0*

- Chapter 7. Best Practices for the Quartus II TimeQuest Timing Analyzer
Revised: *December 2010*
Part Number: *QII53024-10.1.0*

- Chapter 8. Switching to the Quartus II TimeQuest Timing Analyzer
Revised: *December 2010*
Part Number: *QII53019-10.1.0*

- Chapter 9. Synopsys PrimeTime Support
Revised: *December 2010*
Part Number: *QII53005-10.0.1*

- Chapter 10. PowerPlay Power Analysis
Revised: *December 2010*
Part Number: *QII53013-10.1.0*

- Chapter 11. System Debugging Tools Overview
Revised: *December 2010*
Part Number: *QII53027-10.1.0*

- Chapter 12. Analyzing and Debugging Designs with the System Console
Revised: *December 2010*
Part Number: *QII53028-10.1.0*
- Chapter 13. Transceiver Link Debugging Using the System Console
Revised: *December 2010*
Part Number: *QII53029-10.1.0*
- Chapter 14. Quick Design Debugging Using SignalProbe
Revised: *December 2010*
Part Number: *QII53008-10.0.1*
- Chapter 15. Design Debugging Using the SignalTap II Logic Analyzer
Revised: *December 2010*
Part Number: *QII53009-10.0.1*
- Chapter 16. In-System Debugging Using External Logic Analyzers
Revised: *December 2010*
Part Number: *QII53016-10.1.0*
- Chapter 17. In-System Modification of Memory and Constants
Revised: *December 2010*
Part Number: *QII53012-10.0.2*
- Chapter 18. Design Debugging Using In-System Sources and Probes
Revised: *December 2010*
Part Number: *QII53021-10.1.0*
- Chapter 19. Cadence Encounter Conformal Support
Revised: *December 2010*
Part Number: *QII53011-10.1.0*
- Chapter 20. Quartus II Programmer
Revised: *December 2010*
Part Number: *QII53022-10.1.0*

As the design complexity of FPGAs continues to rise, verification engineers are finding it increasingly difficult to simulate their system-on-a-programmable-chip (SOPC) designs in a timely manner. The verification process is now the bottleneck in the FPGA design flow. The Quartus II software provides a wide range of features for performing functional and timing simulation of designs in EDA simulation tools.

This section includes the following chapters:

■ **Chapter 1, Simulating Altera Designs**

This chapter provides guidelines to help you perform simulation for your Altera® designs using EDA simulators and the Quartus II NativeLink feature. This chapter also describes the process for instantiating the IP megafunctions in your design and simulating their functional simulation models.

■ **Chapter 2, Mentor Graphics ModelSim and QuestaSim Support**

This chapter provides detailed instructions about how to simulate your design in the ModelSim-Altera® software or the Mentor Graphics® ModelSim software.

■ **Chapter 3, Synopsys VCS and VCS MX Support**

This chapter describes how to use the Synopsys VCS and VCS MX software to simulate designs that target Altera FPGAs.

■ **Chapter 4, Cadence NC-Sim Support**

This chapter describes the basic NC-Sim, NC-Verilog, and NC-VHDL functional, post-synthesis, and gate-level timing simulations.

■ **Chapter 5, Aldec Active-HDL and Riviera-PRO Support**

This chapter describes how to use the Active-HDL software to simulate designs that target Altera FPGAs.

This chapter provides guidelines to help simulate your Altera® designs using third-party EDA simulators. You can simulate complex designs that include Altera or third-party intellectual property (IP) cores. Simulation is the process of verifying the design behavior before configuring the device.

You can use either of the following simulation tool flows:

- Automatically create scripts to set up and launch an EDA simulator with the Quartus® II NativeLink feature
- Manually set up a simulation in your EDA simulator

The simulation tools Altera supports include ModelSim, ModelSim-Altera, QuestaSim, VCS, VCS MX, NCSim, Active-HDL, and Riviera-PRO.

This chapter includes the following topics:

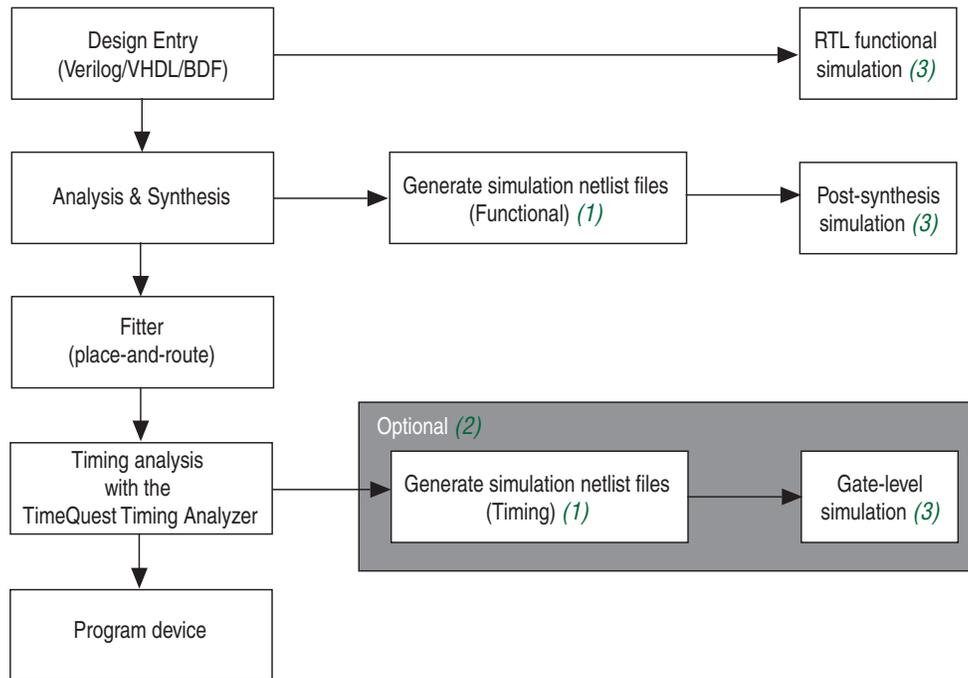
- “Design Flow” on page 1–2
- “EDA Simulation Library Compiler” on page 1–8
- “Using the NativeLink Feature” on page 1–10
- “Simulating Altera IP Cores” on page 1–14
- “Simulating Qsys and SOPC Builder System Designs” on page 1–21

Design Flow

This section describes the simulation flows supported by the Quartus II software, including RTL functional, post-synthesis, and gate-level simulation.

Figure 1-1 shows how these simulation flows fit within a typical design flow.

Figure 1-1. Altera Design Flow Incorporating Simulation



Notes to Figure 1-1:

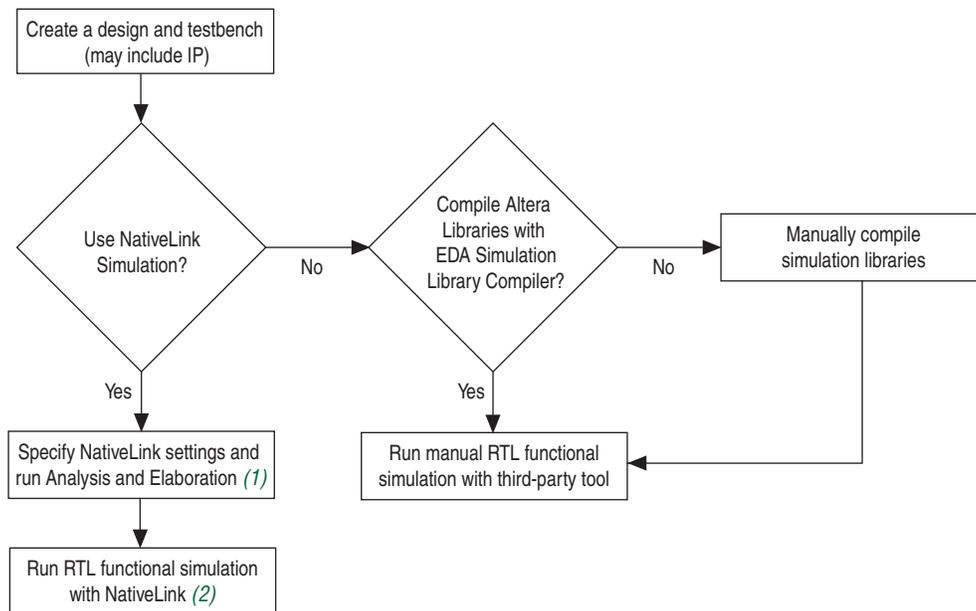
- (1) Generate Verilog Output Files (.vo), VHDL Output Files (.vho), and Standard Delay Format Output Files (.sdo) with the Quartus II Netlist Writer.
- (2) Altera recommends that you use the Quartus II TimeQuest Timing Analyzer to help you verify and close timing. You also have the option to run post-fit timing simulation, which can be slow for a large design. For more information about the TimeQuest Timing Analyzer, refer to *The TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.
- (3) You can use the NativeLink feature to automate the process of running EDA simulations from the Quartus II software.

RTL Functional Simulation Flow

RTL functional simulation allows you to simulate the behavior of your design without timing information. Figure 1-2 shows the overall RTL functional simulation flow supported by the Quartus II software. For information about running EDA simulation automatically through the Quartus II software, refer to “Using the NativeLink Feature” on page 1-10.

For information about manual simulation, refer to “Simulating Altera IP Cores Manually” on page 1-20.

Figure 1-2. : RTL Functional Simulation Flow



Notes to Figure 1-2:

- (1) For more information, refer to “Using the NativeLink Feature” on page 1-10.
 (2) For more information, refer to “Running RTL Functional Simulation Using the NativeLink Feature” on page 1-20.

❓ For more information about performing functional simulation with ModelSim, Questa, VCS, VCS-MX, NCSim, Active-HDL, and Riviera-PRO, refer to the following, respectively, in Quartus II Help:

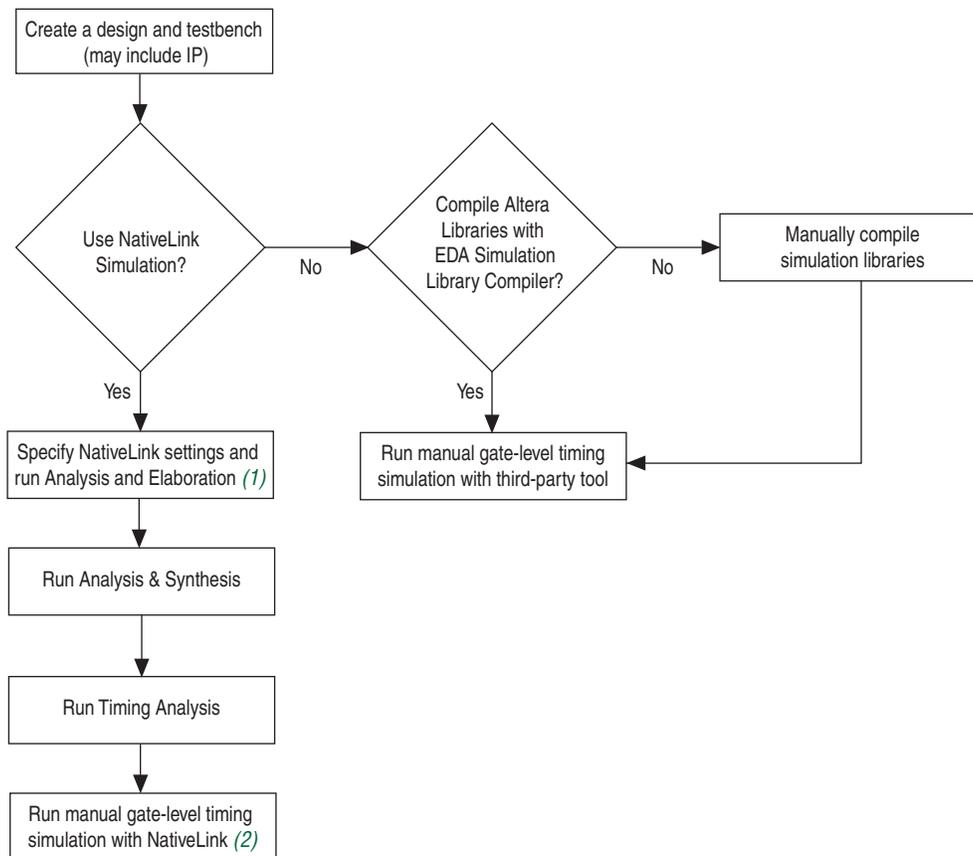
- *Performing a Functional Simulation with the ModelSim Software*
- *Performing a Functional Simulation with the Questa Software*
- *Performing a Functional Simulation with the VCS Software*
- *Performing a Functional Simulation with the VCS-MX Software*
- *Performing a Functional Simulation with the NCSim Software*
- *Performing a Simulation of a Verilog HDL Design with the Active-HDL Software*
- *Performing a Simulation of a VHDL Design with the Active-HDL Software*
- *Performing an RTL Functional Simulation with the Riviera-PRO Software*

Gate-Level Timing Simulation Flow

Gate-level timing simulation allows you to simulate your design with post-fit timing information. [Figure 1-3](#) shows the overall gate-level timing simulation flow using the Quartus II software. For information about running your simulation tool automatically within the Quartus II software, refer to [“Using the NativeLink Feature” on page 1-10](#).

For information about manual simulation, refer to [“Simulating Altera IP Cores Manually” on page 1-20](#).

Figure 1-3. Gate-Level Timing Simulation Flow



Notes to Figure 1-3:

- (1) For more information, refer to [“Using the NativeLink Feature” on page 1-10](#).
- (2) For more information, refer to [“Running RTL Functional Simulation Using the NativeLink Feature” on page 1-20](#).

❓ For more information about performing timing simulation with ModelSim, Questa, VCS, VCS-MX, NCSim, Active-HDL, and Riviera-PRO, refer to the following, respectively, in Quartus II Help:

- [Performing a Timing Simulation with the ModelSim Software](#)
- [Performing a Timing Simulation with the Questa Software](#)
- [Performing a Timing Simulation with the VCS Software](#)
- [Performing a Timing Simulation with the VCS-MX \(VHDL\) Software](#)
- [Performing a Timing Simulation with the NCSim Software](#)
- [Performing a Simulation of a Verilog HDL Design with the Active-HDL Software](#)
- [Performing a Simulation of a VHDL Design with the Active-HDL Software](#)
- [Performing a Gate-Level Simulation with the Riviera-PRO Software](#)



Altera recommends that you use the TimeQuest analyzer as an alternative to gate-level timing simulation to achieve timing closure.

Simulation Netlist Files

Simulation netlist files are required to perform post-synthesis simulation or gate-level timing simulation. These simulation netlist files are generated using the EDA Netlist Writer.

If you are performing post-synthesis simulation, the Verilog HDL Output File (**.vo**) or the VHDL Output File (**.vho**) is required. For the steps required to generate post-synthesis simulation netlist files for **.vo** or **.vho** files, refer to “[Generating Post-Synthesis Simulation Netlist Files](#)”.

If you are performing gate-level timing simulation, the **.vo** or **.vho** and the Standard Delay Format Output File (**.sdo**) are required. The **.sdo** is used to annotate the delay for the elements found in the **.vo** or **.vho**.

❓ To specify options for generating **.vo**, **.vho**, and **.sdo** files, refer to [Specifying HDL Output Settings](#) in Quartus II Help.

Generating Post-Synthesis Simulation Netlist Files

If you want to generate a post-synthesis simulation netlist with just the cell delays, you can generate an **.sdo** file without running the Fitter. In this case, the **.sdo** file includes all timing values for only the device cells. Interconnect delays are not included because fitting (placement and routing) was not performed. To generate the post-synthesis netlist and the **.sdo** file, type the following commands at a command prompt:

```
quartus_map <project name> -c <revision name> ␣  
quartus_sta <project name> -c <revision name> --post_map ␣  
quartus_eda <project name> -c <revision name> --simulation \  
--tool= <3rd party EDA tool> --format=<HDL language> ␣
```

For more information about the **-format** and **-tool** options, type the following command at a command prompt:

```
quartus_eda --help=<option> ␣
```

- ② For more information about how to generate post-synthesis simulation netlist files in the Quartus II software, refer to *Generating Simulation Netlist Files* in Quartus II Help.

Generating Gate-Level Timing Simulation Netlist Files

To perform gate-level timing simulation, the EDA simulators require information about how the design was placed into device-specific architectural blocks. The Quartus II software provides this information in the form of **.vo** for Verilog HDL designs and **.vho** for VHDL designs. The accompanying timing information is stored in the **.sdo**, which annotates the delay for the elements found in the **.vo** or **.vho**. To generate a gate-level timing simulation netlist in the Quartus II software, follow these steps:

1. Configure the EDA Netlist Writer to generate functional simulation netlist. Refer to “*Generating Post-Synthesis Simulation Netlist Files*”.
2. If you have not run a full compilation prior to the fitting process, perform a full compilation. On the Processing menu, click **Start Compilation**.
3. On the Processing menu, point to **Start** and click **Start EDA Netlist Writer**.

During the full compilation or EDA Netlist Writer stage, the Quartus II software produces a **.vo**, a **.vho**, and an **.sdo** used for gate-level timing simulations in the EDA simulators. This netlist file is mapped to architecture-specific primitives. The timing information for the netlist is included in the **.sdo**. The resulting netlist is located in the output directory you specified in the **Settings** dialog box, which by default is the `<project directory>/simulation/<EDA Simulator> directory` (`<EDA Simulator>` can be **modelsim**, **questasim**, **vcs**, **vcs mx**, **rivierapro**, **ncsim**, or **activehdl**).

- ② For more information about how to generate a gate-level timing simulation netlist in the Quartus II software, refer to *Generating Simulation Netlist Files* in Quartus II Help.

Generating Timing Simulation Netlist Files with Different Timing Models

In Stratix III and later devices, you can specify different temperature and voltage parameters to generate the timing simulation netlist files with the Quartus II TimeQuest Timing Analyzer. When you generate the timing simulation netlist files (**.vo**, **.vho**, and **.sdo** files), different timing models for different operating conditions are used by default. Multi-corner timing analysis is run by default during the full compilation.

- ② For more information about generating timing simulation netlist files with different timing models, refer to *EDA Gate Level Simulation (Tools Menu)* in Quartus II Help.

To manually generate the simulation netlist files (**.vo** or **.vho** and **.sdo**) for the three different operating conditions, follow these steps:

1. Generate all available corner models at all operating conditions by typing the following command at a command prompt:


```
quartus_sta <project name> --multicorner ↵
```
2. Generate the timing simulation netlist files for all three corners by performing steps 1 through 3 in “*Generating Gate-Level Timing Simulation Netlist Files*” on page 1–6. The output files are generated in the simulation output directory.

The following examples show the timing simulation netlist files are generated for the operating conditions of the preceding steps when Verilog HDL is selected as the output netlist format.

First Slow Corner (slow, 1100 mV, 85° C)

- .vo file—*<revision name>.vo*
- .sdo file—*<revision name>_v.sdo*

The *<revision name>.vo* and *<revision name>_v.sdo* are generated for backward compatibility in case there are existing scripts that still use them.

- .vo file—*<revision name>_<speedgrade>_1100mv_85c_slow.vo*
- .sdo file—*<revision name>_<speedgrade>_1100mv_85c_v_slow.sdo*

Second Slow Corner (slow, 1100 mV, 0° C)

- .vo file—*<revision name>_<speedgrade>_1100mv_0c_slow.vo*
- .sdo file—*<revision name>_<speedgrade>_1100mv_0c_v_slow.sdo*

Fast Corner (fast, 1100 mV, 0° C)

- .vo file—*<revision name>_min_1100mv_0c_fast.vo*
- .sdo file—*<revision name>_min_1100mv_0c_v_fast.sdo*

For older devices, a slow-corner (worst case) timing model is used by default. There are only two timing models available: slow-corner and fast-corner. To generate the timing simulation netlist files using a different timing model, you must run the TimeQuest analyzer with a different timing model before you start the EDA Netlist Writer.

To run the TimeQuest analyzer with a best-case model, use the **-fast_model** option after you create the timing netlist.

The following command enables the fast timing model:

```
create_timing_netlist --fast_model ↵
```

After running the TimeQuest analyzer, perform steps 1 through 3 in “[Generating Gate-Level Timing Simulation Netlist Files](#)” on page 1–6 to generate the timing simulation netlist file. For fast corner timing models, the **-fast** post fix is added to the .vo or .vho and .sdo file (for example, **my_project_fast.vo** or **my_project_fast.vho** and **my_project_fast.sdo**).



For more information about performing multi-corner timing analysis, refer to the [Quartus II TimeQuest Timing Analyzer](#) chapter in volume 3 of the *Quartus II Handbook*.

EDA Simulation Library Compiler

The EDA Simulation Library Compiler compiles Verilog HDL and VHDL simulation libraries for all Altera devices and supported third-party simulators. You can use this tool to automatically compile all libraries required for RTL functional and gate-level timing simulation.

The compiled libraries are stored in the directory you specify. When you perform the simulation, you can reuse the compiled libraries to avoid the overhead associated with redundant library compilations.

If the compilation targets the VCS simulator, then the VCS options file **simlib_comp.vcs** is generated after compilation. You can then include your design and testbench files in the option files and invoke them with the **vcs** command.

Before using the EDA Simulation Library Compiler, ensure that the appropriate simulation tools are already installed and their execution paths are specified. To specify the path, refer to [“Setting Up the EDA Simulator Execution Path” on page 1–10](#).

- For more information about compiling Verilog HDL and VHDL simulation libraries for all Altera devices and supported third-party simulators, refer to [Compiling Simulation Libraries in the Quartus II Software](#) in Quartus II Help.

Running the EDA Simulation Library Compiler Through the GUI

 The EDA Simulation Library Compiler does not support ModelSim-Altera because ModelSim-Altera includes precompiled libraries.

To compile libraries with the EDA Simulation Library Compiler GUI, follow these steps:

1. On the Tools menu, click **Launch EDA Simulation Library Compiler**. The **EDA Simulation Library Compiler** dialog box appears.
2. In the **Tool name** list under **EDA simulation tool**, select your preferred EDA simulator. The **Executable location** box displays the location of the simulation tool you specified. This location must be set before running the EDA Simulation Library Compiler.
3. Under **Library families**, select one or more device families for your design compilation and move them to the **Selected families** box.
4. Under **Library language**, select **VHDL**, **Verilog HDL**, or both.
5. Specify the output netlist destination by editing or browsing to a directory in the **Output directory** box.

You can then link to the directory so that the NativeLink feature reuses the compiled libraries rather recompile them. Refer to step 9 under [“Configuring NativeLink Settings” on page 1–11](#).

6. Click **Start Compilation**.

When the EDA Simulation Library Compiler finishes, for ModelSim, QuestaSim, Active-HDL, Riviera-PRO, VCS MX, and NCSim, all required libraries are compiled and stored in *<output location you specified>/<verilog_libs or vhdl_libs>*. The library files (for example, *modelsim.ini*, *cds.lib*, and *synopsys_sim.setup*) are also generated and stored in *<output location you specified>*.

When the EDA Simulation Library Compiler finishes, for VCS, the *simlib_comp.vcs* file is generated. [Example 1-1](#) shows the contents of a simple *simlib_comp.vcs* file.

Example 1-1.

```
+cli+1 -line -timescale=1ps/1ps \  
-v /apps/quartus/10.1/quartus/eda/sim_lib/altera_primitives.v \  
-v /apps/quartus/10.1/quartus/eda/sim_lib/220model.v \  
-v /apps/quartus/10.1/quartus/eda/sim_lib/sgate.v \  
-v /apps/quartus/10.1/quartus/eda/sim_lib/altera_mf.v \  
-v /apps/quartus/10.1/quartus/eda/sim_lib/stratixiii_atoms.v \  
+incdir+/apps/quartus/10.1/quartus/eda/sim_lib
```

If you manually run a simulation using the Synopsys VCS simulator, you must include your design files and testbench file in the *.synopsys_sim.setup* option file, as shown in [Example 1-2](#).

Example 1-2.

```
+cli+1 -line -timescale=1ps/1ps design_file.v test_bench_file.v\  
-v /apps/quartus/10.1/quartus/eda/sim_lib/altera_primitives.v \  
-v /apps/quartus/10.1/quartus/eda/sim_lib/220model.v \  
-v /apps/quartus/10.1/quartus/eda/sim_lib/sgate.v \  
-v /apps/quartus/10.1/quartus/eda/sim_lib/altera_mf.v \  
-v /apps/quartus/10.1/quartus/eda/sim_lib/stratixiii_atoms.v \  
+incdir+/apps/quartus/10.1/quartus/eda/sim_lib
```

To compile all of the libraries, design files, and testbench files, type the following command:

```
vcs -file simlib_comp.vcs ↵
```

Running the EDA Simulation Library Compiler from the Command Line

For Linux operating systems, you can force the EDA Simulation Library Compiler to use the EDA simulator executables from the search path by typing the following command at shell prompt:

```
export QUARTUS_INIT_PATH=$PATH
```

To run the EDA Simulation Library Compiler from the command line, type the following command:

```
quartus_sh --simlib_comp -family <device> -tool <simulation tool name>  
-tool_path <simulator executable location>  
-language <language> -directory <directory> ↵
```



For more information about the command line options and how to use them, type the following command:

```
quartus_sh --help=simlib_comp ↵
```

Using the NativeLink Feature

The NativeLink feature in the Quartus II software facilitates the seamless transfer of information between the Quartus II software and EDA tools by allowing you to launch an EDA simulator from the Quartus II software.

The Quartus II software contains all the libraries required for setting up and running a successful simulation of Altera IP cores. If the IP core you are using supports the Quartus II NativeLink feature, it is easy to use the NativeLink feature to set up your simulation. However, you can also simulate Altera IP cores directly with third-party simulators. To determine whether the NativeLink feature is supported, refer to the applicable IP core user guide.

Setting Up the EDA Simulator Execution Path

To run an EDA simulator automatically from the Quartus II software using the NativeLink feature, specify the path to your simulation tool by performing the following steps:

1. On the Tools menu, click **Options**. The **Options** dialog box appears.
2. In the **Category** list, select **EDA Tool Options**.
3. Double-click the entry under **Location of executable** beside the name of your EDA tool.
4. Type the path or browse to the directory containing the executables of your EDA tool.

Table 1-1 lists the execution paths for each EDA simulator.

Table 1-1. Execution Paths for EDA Simulators

Simulator	Path
ModelSim-Altera	c:\<ModelSim-Altera installation path>\win32aloem (Windows) /<ModelSim-Altera installation path>/linuxaloem (Linux)
ModelSim	c:\<ModelSim installation path>\win32 (Windows) /<ModelSim installation path>/bin (Linux)
QuestaSim	c:\<QuestaSim installation path>\win32 (Windows) /<QuestaSim installation path>/bin (Linux)
VCS/VCS MX	/<VCS MX installation path>/bin (Linux)
NCSim	/<NCSim installation path>/tools/bin (Linux)
Active-HDL	c:\<Active-HDL installation path>\bin (Windows)
Riviera-PRO	c:\<Riviera-PRO installation path>\bin (Windows)

5. Click **OK**.

You can also specify the path to the simulator's executables by typing the **set_user_option** Tcl command, as follows:

```
set_user_option -name EDA_TOOL_PATH_MODELSIM <path to executables> ←  
set_user_option -name EDA_TOOL_PATH_MODELSIM_ALTERA <path to \  
executables> ←  
set_user_option -name EDA_TOOL_PATH_QUESTASIM <path to executables> ←  
set_user_option -name EDA_TOOL_PATH_VCS <path to executables> ←  
set_user_option -name EDA_TOOL_PATH_VCS_MX <path to executables> ←  
set_user_option -name EDA_TOOL_PATH_NCSIM <path to executables> ←  
set_user_option -name EDA_TOOL_PATH_ACTIVEHDL <path to executables> ←  
set_user_option -name EDA_TOOL_PATH_RIVIERAPRO <path to executables> ←
```

Configuring NativeLink Settings

To configure NativeLink settings, follow these steps:

1. On the Assignments menu, click **Settings**. The **Settings** dialog box appears.
2. In the **Category** list, select **Simulation**. The **Simulation** page appears.
3. In the **Tool name** list, select your preferred EDA simulator.
4. For gate-level simulation, if you want to run simulation in your EDA simulator automatically after Quartus II full compilation, turn on **Run gate-level simulation automatically after compilation**.
5. If you have testbench files or macro scripts, enter the information under **NativeLink settings**.

For more information about setting up a testbench file with NativeLink, refer to [“Setting Up Testbench Files Using the NativeLink Feature” on page 1-13](#).

6. If you want to run the EDA simulator in command-line mode, follow these steps:
 - a. On the **Simulation** page, click **More NativeLink Settings**. The **More NativeLink Settings** dialog box appears.
 - b. Under **Existing option settings**, click **Launch third-party EDA tool in command-line mode**.
 - c. In the **Setting** field, select **On**.
 - d. Click **OK**.
7. If you want to generate only the simulation script without launching the EDA simulator during NativeLink simulation, follow these steps:
 - a. On the **Simulation** page, click **More NativeLink Settings**. The **More NativeLink Settings** dialog box appears.
 - b. Under **Existing option settings**, click **Generate third-party EDA tool command scripts without running the EDA tool**.
 - c. In the **Setting** field, select **On**.
 - d. Click **OK**.

If you turn this option on and run NativeLink, only the simulation command script is generated. The file names of simulation command scripts for various simulators are as follows:

- `<project_name>_run_msim_<rtl/gate>_level_<verilog/vhdl>.do` (ModelSim)
 - `<project_name>_run_questasim_<rtl/gate>_level_<verilog/vhdl>.do` (QuestaSim)
 - `<project_name>_sim_<rtl/gate>_<verilog/vhdl>.do` (Riviera-PRO and Active-HDL)
 - `script_file.sh` and `<project_name>_rtl.vcs` (VCS)
 - `<project_name>_vcsmx_<rtl/gate>_<vhdl/verilog>.tcl` (VCS MX)
 - `<project_name>_ncsim_<rtl/gate>_<verilog/vhdl>.tcl` (NCSim)
8. Depending on the simulator, perform the simulation by typing one of the following commands:
- `do <script>.do` ← (ModelSim Macro File)
 - `quartus_sh -t <script>.tcl` ← (Tcl Script File)
 - `sh <script>.sh` ← (Shell script)
9. If you have compiled libraries using the EDA Simulation Library Compiler, follow these steps:
- a. On the **Simulation** page, click **More EDA Netlist Writer Settings**. The **More EDA Netlist Writer Settings** dialog box appears.
 - b. Under **Existing option settings**, click **Location of user compiled simulation library**.
 - c. In the **Setting** field, type the path that contains the user-compiled libraries that are generated from the EDA Simulation Library Compiler. The path should be the same as the path you have set in the Output Directory in the EDA Simulation Library Compiler.

 Step 9 is not applicable for Active-HDL and Riviera-PRO.

For more information about the EDA Simulation Library Compiler, refer to “[EDA Simulation Library Compiler](#)” on page 1–8.

- ❓ For more information about using the Quartus II software with other EDA tools, refer to [About Using the Quartus II Software with Other EDA Tools](#) in Quartus II Help.

Setting Up Testbench Files Using the NativeLink Feature

You can use the NativeLink feature to compile your design files and testbench files, and run an EDA simulation tool to automatically perform a simulation.

To set up the NativeLink feature for simulation, follow these steps:

1. On the Assignments menu, click **Settings**. The **Settings** dialog box appears.
2. In the **Category** list, under **EDA Tool Settings**, click **Simulation**. The **Simulation** page appears.
3. In the **Tool name** list, select your preferred EDA simulator.
4. Under **NativeLink settings**, select **None**, **Compile test bench**, or **Script to compile test bench** (Table 1-2).

Table 1-2. NativeLink Testbench Settings

Setting	Description
None	NativeLink compiles simulation models and design files.
Compile test bench	NativeLink compiles simulation models, design files, testbench files, and starts simulation.
Script to compile test bench	NativeLink compiles the simulation models and design files. The script you provide is sourced after design files are compiled. Use this option when you want to create your own script to compile your testbench file and perform simulation.

If you select **Compile test bench**, select your testbench setup from the **Compile test bench** list. You can use different testbench setups to specify different test scenarios. If there are no testbench setups entered, create a testbench setup by following these steps:

1. Click **Test Benches**. The **Test Benches** dialog box appears.
2. Click **New**. The **New Test Bench Settings** dialog box appears.
3. In the **Top level module in test bench** box, type the top-level testbench entity or module name. For example, for a Quartus II-generated VHDL testbench, type `<Vector Waveform File name>_vhd_vec_tst`. The testbench name in the **Test bench name** box automatically follows the top-level testbench entity or module name.
4. For gate-level simulation, turn on **Use test bench to perform VHDL timing simulation**. In the **Design instance name in test bench** box, type the full instance path to the top level of your FPGA design. For example, for a Quartus II-generated VHDL testbench, type `i1`.
5. Under **Simulation period**, select **Run simulation until all vector stimuli are used** or specify the end time of the simulation.
6. Under **Test bench files**, browse and add all of your testbench files in the **File name** box. Use the **Up** and **Down** buttons to reorder your files. The script used by the NativeLink feature compiles the files in order from top to bottom.



You can also specify the library name and HDL version to compile the testbench file. The NativeLink feature compiles the testbench file to a library name using the specified HDL version.

7. Click **OK**.
8. In the **Test benches** dialog box, click **OK**.
9. Under **NativeLink settings**, turn on **Use script to set up simulation** and browse to your script. Your script is executed to set up and run simulation after loading the design using the **vsim** command.

If you select **Script to compile test bench**, browse to your script and click **OK**.



You can also use the NativeLink feature in your script to compile your design files and testbench files with customized settings.

Simulating Altera IP Cores

This section describes the process of simulating Altera IP cores in your design.

Even when the IP source code is encrypted or otherwise restricted, the Quartus II software allows you to easily simulate designs that contain Altera IP cores. You can customize Altera IP cores, then generate a VHDL or Verilog HDL functional simulation model.

Altera IP cores support both Verilog HDL and VHDL simulation. The manner in which dual-language simulation is supported for specific IP cores might differ.

For more information about IP simulation models, refer to [“Simulation Model Files” on page 1-16](#).

When IEEE 1364-2005 encrypted Verilog HDL simulation models are provided, they are encrypted separately for each Altera-supported simulation vendor. You need a simulator that is capable of VHDL/Verilog HDL co-simulation if you want to simulate the model in a VHDL design. The IEEE 1364-2005 encrypted Verilog HDL model is only used for Stratix V devices.

By special arrangement with Mentor Graphics®, you can simulate Altera IEEE encrypted Verilog HDL models for your VHDL designs using the VHDL-only version of ModelSim Altera Edition, as well as single-language VHDL versions of ModelSim PE. Additionally, Altera IEEE encrypted Verilog models for all Mentor Graphics simulators do not consume an additional runtime license. For example, if you are simulating a VHDL design that contains Altera IEEE encrypted Verilog models, no Verilog license for tools such as ModelSim SE/LNL is checked out. If your design code is written in both Verilog HDL and VHDL, standard ModelSim license consumption rules apply.

Some AMPPSM megafunctions might also use IP Functional Simulation (IPFS) models for functional simulation. An IPFS model is a cycle-accurate VHDL or Verilog HDL model produced by the Quartus II software. The model allows for fast functional simulation of IP using industry-standard VHDL and Verilog HDL simulators.



IEEE encrypted Verilog models are generally faster than IPFS models.



Use IPFS models for simulation only. Do not use them for synthesis or any other purpose. Using these models for synthesis results in a nonfunctional design.

IP Simulation Flows

The parameter editor for each IP core allows you to quickly and easily view documentation, specify parameters, and generate simulation models and other output files necessary to integrate the IP core into your design. When you use the MegaWizard™ Plug-In Manager to parameterize your IP core, the Quartus II software generates a Quartus II IP File (**.qip**) for inclusion in your Quartus II project. For IP cores that use IPFS models, the Quartus II software can also generate a **.vo** or **.vho** that contains an IPFS model. For IP cores that use IEEE encrypted Verilog HDL models or plain-text HDL, the generated directory structure is shown in [Figure 1-4 on page 1-17](#).

- For a list of the RTL functional simulation library files, refer to [Altera Functional Simulation Libraries](#) in Quartus II Help. For a list of the gate-level timing simulation library files, refer to [Altera Post-Fit Libraries](#) in Quartus II Help.

IP Variant Directory Structure

You can use the parameter editor to help you parameterize IP cores. The location of simulation output files varies by IP core. The following sections describe the related output files generated by the IP generation process.

- For information about how to parameterize IP cores, refer to the appropriate IP core user guide, available on the [User Guides](#) literature page of the Altera website.

Synthesis Files

For synthesis purposes, a `<variant_name>.qip` is generated in your project directory for each IP core variation in your design. If you specify Verilog HDL or VHDL as the output file type, an IP variant file (`<variant_name>.v` or `<variant_name>.vhd`) is generated, respectively. The **.qip** is a single file that contains all synthesis information required for processing by the Quartus II Compiler.

For some IP cores, a `<variant_name>` subdirectory is created in the project directory. This directory contains files needed for synthesis, and might also contain Synopsys Design Constraints Files (**.sdc**), Tcl Script Files (**.tcl**), and Pin Planner Files (**.ppf**).

To compile your design in the Quartus II software, you should add the `<variant_name>.qip` files to your project, along with your design files. Additional steps might be necessary, depending on the IP core.

- For more information, refer to the appropriate IP core user guide, available on the [User Guides](#) literature page of the Altera website.

Simulation Model Files

The form and structure of simulation models that support Altera IP cores vary. IP cores use the simulation models that are provided in the `quartus/eda/sim_lib` directory of your Quartus II installation. Altera provides some of the models in this directory as plain-text VHDL and Verilog HDL, while other models are IEEE encrypted Verilog HDL for each of the Altera-supported EDA simulation vendors. You can simulate the IEEE encrypted Verilog HDL models in VHDL designs if you have a VHDL/Verilog HDL co-simulator. You can also simulate the IEEE encrypted Verilog HDL models in any version of ModelSim, including the VHDL-only versions of ModelSim Altera and ModelSim PE, because of the special arrangement Altera has made with Mentor Graphics.

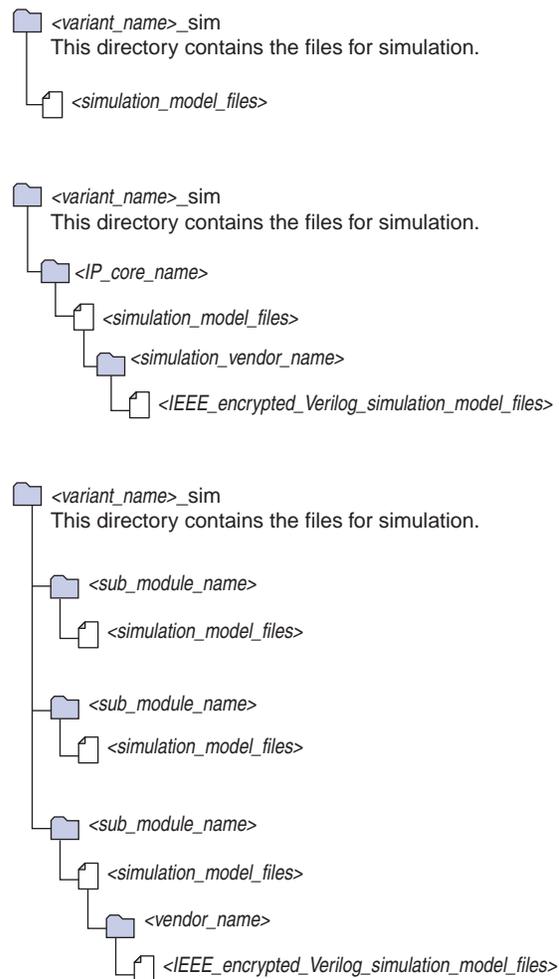
Depending on the specific IP core, Altera may provide simulation models in one or more of the following formats:

- Plain-text RTL in Verilog HDL, VHDL, or both
- Mixed structural/RTL IPFS model in `.vho` and `.vo`
- IEEE encrypted Verilog RTL models for each supported EDA simulation vendor

The simulation models can be generated to one or more of the following locations:

- Your project directory—the simulation model is named `<variant_name>.vo` or `<variant_name>.vho`
- A directory hierarchy under your project directory—[Figure 1-4](#) shows the examples of directory hierarchies for simulation models.

Figure 1-4. Examples of Directory Hierarchies



Some Altera IP cores include design examples, testbenches, simulator script examples, and/or Quartus II project examples within one of the directory structures in [Figure 1-4](#).

- For core-specific information, refer to the appropriate IP core user guide, available on the [User Guides](#) literature page of the Altera website.

Instantiate the IP Core

A fully parameterized IP core is called a variant. A `.qip` for each IP variant must be in your Quartus II project. The Quartus II software may add the `.qip` files automatically. For an overview about how to do this for synthesis and simulation, refer to “[IP Variant Directory Structure](#)” on page 1-15.

- For information about instantiating IP cores, refer to the [Megafunction Overview User Guide](#), or the [User Guide](#) literature page of the Altera website.

-  For information about synthesis and compilation with the Quartus II software, refer to the applicable chapters in *Volume 1: Design Synthesis* of the *Quartus II Handbook*.

Perform Functional Simulation

To perform functional simulation, in addition to adding your design files and testbench files, you must also add the IP core variation file and other corresponding IP simulation files and model libraries to your simulation project. For more information, refer to “[IP Variant Directory Structure](#)” on page 1–15.

-  IP cores that use device transceiver resources require the Altera transceiver libraries for simulation. For more information, refer to *Altera Functional Simulation Libraries* in Quartus II Help.
-  The hard IP implementation of the PCI Express® IP core requires the PCIe libraries for simulation. For more information, refer to *Altera Functional Simulation Libraries* in Quartus II Help.

The Quartus II software includes all the simulation model libraries required for running a successful simulation of Altera IP cores. If the IP core you are using supports the Quartus II NativeLink feature, you can easily use the NativeLink feature to set up your simulation. However, you can also simulate Altera IP cores directly with third-party simulators. To determine whether the NativeLink feature is supported, refer to the applicable IP core user guide.

Verilog HDL and VHDL IP Functional Simulation Models

Some IP cores use IP Functional Simulation (IPFS) models for functional simulation. The IPFS models in Verilog HDL or VHDL format differ from the low-level synthesized netlist in Verilog HDL or VHDL format generated by the Quartus II software for post-synthesis or post place-and-route simulations. The IPFS models generated by the Quartus II software are much faster than the low-level post-synthesis or post place-and-route netlists of your design because they are mapped to higher-level primitives such as adders, multipliers, and multiplexers. You can use these IPFS models together with the rest of your design in any Altera-supported simulator.

-  Simulator-independent IPFS primitives are located in the `quartus/eda/sim_lib` directory. You must compile the files that correspond to the device you are using and your simulation language.
-  Generating an IPFS model for Altera IP cores does not require a license. However, generating an IPFS model for AMPP megafunctions may require a license. For more information about AMPP licensing requirements, refer to “[Obtaining and Licensing an AMPP Megafunction](#)” in *AN 343: OpenCore Evaluation of AMPP Megafunctions*.

Stratix V Simulation Model Libraries

For Stratix V devices, Altera provides a set of IEEE encrypted Verilog RTL models for use in both VHDL and Verilog HDL designs. In general, the RTL models simulate faster than IPFS models.

- ② For more information about IEEE encrypted Verilog RTL models and for a list of IEEE encrypted libraries for Stratix V devices, refer to *Guidelines for Compiling Stratix V Libraries* in Quartus II Help.

In addition to Stratix V libraries, the altera_insim library which contains the altera_pll model, currently only supports Stratix V devices. New device-independent models will be added to the altera_insim library in later software releases. The altera_insim library is provided only in SystemVerilog, but the version in the **quartus/eda/sim_lib/mentor** directory can be simulated in all versions (6.6 c or later) of Mentor Graphics simulators.

-  For more information about the ALTERA_PLL megafunction, refer to the *Altera Phase-Locked Loop (ALTERA_PLL) Megafunction User Guide*.

To use the Stratix V (and the altera_insim) libraries in VHDL, compile the VHDL wrappers and Verilog HDL files into the same library.

If your design uses a mix of VHDL and Verilog HDL that references Altera simulation models, you should use the Altera Verilog HDL models with the Verilog HDL portion of your design and the Altera VHDL models with the VHDL portion of your design. You can do this by specifying, mapping, and searching the Verilog HDL and VHDL files of your simulator. For instructions about how to compile models into logical libraries and how to map logical libraries to physical libraries, refer to your EDA simulator's documentation.

Simulating Altera IP Cores Using the Quartus II NativeLink Feature

The Quartus II NativeLink feature eases the tasks of setting up and running a simulation. The NativeLink feature launches the supported simulator of your choice from within the Quartus II software. The NativeLink feature also automates the compilation and simulation of testbenches.

For more information about using the NativeLink feature to simulate Altera IP cores, refer to *“Using the NativeLink Feature”* on page 1-10. For more information about the RTL functional simulation flow, refer to *Figure 1-2* on page 1-3.

Before running a NativeLink simulation, you must specify the NativeLink settings and perform analysis and elaboration, as described in the following section.

Perform Analysis and Elaboration on Your Design

To perform Analysis and Elaboration on your design, on the Quartus II Processing menu, point to **Start**, then click **Start Analysis & Elaboration**.

If you are using the Quartus II NativeLink feature and your Quartus II project contains IP cores that require IPFS models for simulation, you do not have to manually add the IPFS models to the Quartus II project for these IP cores. When the Quartus II NativeLink feature launches the third-party simulator tool and starts the simulation, it automatically adds the IPFS model files required for simulation if they are present in the Quartus II project directory.

You can now perform functional simulation in your third-party simulator.

Run Simulation with the Quartus II NativeLink Feature

For more information about using the NativeLink feature to simulate Altera IP cores, refer to [“Using the NativeLink Feature” on page 1–10](#).

Using the EDA Simulation Library Compiler

If you do not use the NativeLink feature, you must compile the simulation libraries with the EDA Simulation Library Compiler or manually. The EDA Simulation Library Compiler compiles all required libraries automatically.

For more information about the EDA Simulation Library Compiler, refer to [“EDA Simulation Library Compiler” on page 1–8](#).

Simulating Altera IP Cores Manually

You can alternatively simulate Altera IP cores in your third-party simulator by adding its variation file to your simulation project. If the IP core requires IPFS model files, do not add the IP core variation file to your simulation project. Instead, add its IPFS model files (either Verilog HDL or VHDL) to your simulation project.

If your IP core generates any other type of simulation models as described in [“IP Variant Directory Structure” on page 1–15](#), you must compile all of the IP simulation files (including the appropriate Altera simulation model library files) along with your design and testbench.

To properly compile, load, and simulate the IP cores, you must first compile the libraries in your simulation tool.

- ❓ For a list of library files, refer to [Altera Functional Simulation Libraries](#) in Quartus II Help.

Running RTL Functional Simulation Using the NativeLink Feature

To run RTL functional simulation using the NativeLink feature, follow these steps:

1. Configure the NativeLink settings. Refer to [“Configuring NativeLink Settings” on page 1–11](#).
2. On the Processing menu, point to **Start** and click **Start Analysis & Elaboration** to perform an Analysis and Elaboration. This command collects all your file name information and builds your design hierarchy in preparation for simulation.

3. On the Tools menu, point to **Run EDA Simulation Tool** and click **EDA RTL Simulation** to automatically run the EDA simulator, compile all necessary design files, and complete a simulation.

Running Gate-Level Timing Simulation Using the NativeLink Feature

To run a gate-level timing simulation using the NativeLink feature, follow these steps:

1. Configure the EDA Netlist Writer settings. Refer to “[Generating Post-Synthesis Simulation Netlist Files](#)” on page 1–5.
2. Configure the NativeLink settings. Refer to “[Configuring NativeLink Settings](#)” on page 1–11.
3. On the Processing menu, click **Start Compilation** to perform Quartus II full compilation, including generation of an EDA netlist file.
4. On the Tools menu, point to **Run EDA Simulation Tool** and click **EDA Gate Level Simulation** to automatically run the EDA simulator, compile all necessary design files, and complete a simulation.



If you have turned on **Run gate-level simulation automatically after compilation** while configuring NativeLink settings, you can skip step 4.

Simulating Qsys and SOPC Builder System Designs



Altera’s Qsys system integration tool is now available as beta for evaluation in the Quartus II software subscription edition version 10.1. Altera does not recommend using the beta release of Qsys in the Quartus II software version 10.1 for designs that are close to completion and are meeting design requirements. Before using Qsys, review the [Quartus II Software Version 10.1 Release Notes](#) and [AN 632: SOPC Builder to Qsys Migration Guidelines](#) for known issues and limitations. To submit general feedback or technical support on the beta release of Qsys, submit a service request through mysupport.altera.com. Alternatively, to submit general feedback, click **Feedback** on the Quartus II software Help menu.

You can use the Qsys system integration tool or SOPC Builder in the Quartus II software to create your system and generate simulation models for functional simulation. You can also use SOPC Builder to generate system-level testbenches that help you debug your system design.

If your design uses the Nios II processor, you must first initialize any memories that contain software prior to simulation. You can create Memory Initialization Files (**.mif**) for this purpose with the Nios II Software Build Tools.



For more information about SOPC Builder, refer to the [SOPC Builder User Guide](#). For more information about using bus functional models (BFMs) to simulate Avalon® standard interfaces in your system, refer to [Avalon Verification IP Suite User Guide](#).



For more information about simulating Qsys designs, refer to the [Creating a System with Qsys](#) chapter of the [Quartus II Handbook](#).

 For more information about simulating designs that contain a Nios II processor, refer to [AN 351: Simulating Nios II Embedded Processors Designs](#).

Document Revision History

Table 1-3 shows the revision history for this chapter.

Table 1-3. Document Revision History

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Title changed from “Simulating Designs with EDA Tools” ■ Merged content from “Simulating Altera IP in Third-Party Simulation Tools” chapter to “Simulating Altera IP Cores” on page 1-14 ■ Added new section “IP Variant Directory Structure” on page 1-15 ■ Added new section “Simulating Qsys and SOPC Builder System Designs” on page 1-21 ■ Added information about simulating designs with Stratix V devices ■ Updated chapter to new template
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Linked to Quartus II Help where appropriate ■ Removed Referenced Documents section ■ Removed Creating Testbench Files ■ Added VCS and QuestaSim as third-party simulation tools ■ Updated “Running the EDA Simulation Library Compiler Through the GUI” on page 1-18 ■ Updated “Setting Up the EDA Simulator Execution Path” on page 1-19 ■ Updated “Configuring NativeLink Settings” on page 1-20 ■ Updated “Setting Up Testbench Files Using the NativeLink Feature” on page 1-22
November 2009	9.1.0	Initial release

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter provides detailed instructions about how to simulate your design in the ModelSim-Altera® software, Mentor Graphics® ModelSim software, and Mentor Graphics QuestaSim software.

An Altera Quartus® II software subscription includes the ModelSim-Altera Starter Edition, which is a no-cost entry-level version of the ModelSim-Altera Subscription Edition software. The ModelSim-Altera Subscription Edition software offers support for all Altera devices. Both versions are available on PC and Linux platforms. You can use the ModelSim-Altera software to perform functional, post-synthesis, and gate-level timing simulations for either Verilog HDL or VHDL designs that target an Altera FPGA.



In this chapter, ModelSim refers to ModelSim SE, PE, DE, and QuestaSim. ModelSim-Altera refers to ModelSim-Altera Starter Edition and ModelSim-Altera Subscription Edition software.

This chapter includes the following topics:

- “Software Requirements” on page 2–2
- “Design Flow with ModelSim-Altera, ModelSim, or QuestaSim Software” on page 2–2
- “Simulation Libraries” on page 2–2
- “Performing Simulation Using the ModelSim-Altera Software” on page 2–3
- “Performing Simulation Using the ModelSim and QuestaSim Software” on page 2–5
- “Simulating Designs that Include Transceivers” on page 2–15
- “Using the NativeLink Feature with ModelSim-Altera, ModelSim, or QuestaSim Software” on page 2–23
- “Generating a Timing Value Change Dump (.vcd) File for the PowerPlay Power Analyzer” on page 2–24
- “Viewing a Waveform from a .wlf File” on page 2–25
- “Simulating with ModelSim-Altera Waveform” on page 2–25
- “Scripting Support” on page 2–26
- “Software Licensing and Licensing Setup in ModelSim-Altera Subscription Edition” on page 2–27

Software Requirements

To simulate your design using the ModelSim-Altera, ModelSim, and QuestaSim software, you must first set up the Altera libraries. These libraries are installed with the Quartus II software.

-  For more information about installing the software and directories created during the Quartus II software installation, refer to the *Altera Software Installation and Licensing* manual.

Design Flow with ModelSim-Altera, ModelSim, or QuestaSim Software

You can perform the following types of simulations using the ModelSim-Altera, ModelSim, or QuestaSim software:

- Functional simulation
- Post-synthesis simulation
- Gate-level timing simulation

-  Some versions of ModelSim and QuestaSim support SystemVerilog, PSL assertions, SystemC, and more. For more information about the features supported in the different versions of ModelSim and QuestaSim, refer to Mentor Graphics literature or your salesperson.

-  The VHDL version of ModelSim-Altera and other single-language VHDL versions of ModelSim cannot simulate designs that target the Stratix V device family if you are using transceivers.

You need a version of ModelSim that supports VHDL/Verilog co-simulation to simulate designs that use Stratix V transceivers.

-  For more information about the Quartus II software design flow, refer to the “PLD Design Flow” section in the *Simulating Altera Designs* chapter in volume 3 of the *Quartus II Handbook*.
-  For additional documentation about ModelSim-Altera, refer to the ModelSim-Altera Help that ships with the product. Click the **Help** button on the ModelSim-Altera toolbar.

Simulation Libraries

Simulation model libraries are required to run a simulation whether you are running a functional simulation, post-synthesis simulation, or gate-level timing simulation. In general, running a functional simulation requires the functional simulation model libraries, while running a post-synthesis or gate-level timing simulation requires the gate-level timing simulation model libraries. Unless you are using ModelSim-Altera, you must compile the necessary library files before you can run the simulation. The ModelSim-Altera software has the Altera libraries pre-compiled and built in. Do not compile these libraries again.

There are a few exceptions in which you must compile gate-level timing simulation library files to perform functional simulation. For example, some Altera megafunctions require gate-level libraries to perform a functional simulation using third-party simulators.

Precompiled Simulation Libraries in the ModelSim-Altera Software

Precompiled libraries for both functional and gate-level simulations are provided for the ModelSim-Altera software. You should not compile these library files before running a simulation.

The precompiled libraries provided in `<ModelSim Altera path>/altera` must be compatible with the version of the Quartus II software that is used to create the simulation netlist. To check whether the precompiled libraries are compatible with your version of the Quartus II software, refer to the `<ModelSim Altera path>/altera/version.txt` file. This file shows which version and build of the Quartus II software was used to create the precompiled libraries.

- ② For a list of precompiled library names for all functional and gate-level simulation models, refer to *ModelSim-Altera Precompiled Libraries* in Quartus II Help.

Simulation Library Files in the Quartus II Software

In ModelSim and QuestaSim, no precompiled libraries are available. You must compile the necessary libraries to perform functional or gate-level simulation.

- ② For a list of all functional simulation library files in the Quartus II directory, refer to *Altera Functional Simulation Libraries* in Quartus II Help. For a list of all post-synthesis and post-fit (gate-level) library files in the Quartus II directory, refer to *Altera Post-Fit Libraries* in Quartus II Help. For a list of logical library names to compile for simulation models, refer to *Libraries For Altera Simulation Models* in Quartus II Help.

Disabling Timing Violation on Registers

In certain situations, a timing violation can be ignored and you can disable timing violations on registers (for example, timing violations that occur in internal synchronization registers used for asynchronous clock domain crossing).

By default, the `x_on_violation_option logic` option is **On**, which means simulation shows “x” whenever a timing violation occurs. To disable showing the timing violation on certain registers, set the `x_on_violation_option logic` option to **Off** on those registers. The following Quartus II Tcl command disables timing violation on registers. This Tcl command is also stored in the `.qsf` file.

```
set_instance_assignment -name X_ON_VIOLATION_OPTION OFF -to <register_name>
```

Performing Simulation Using the ModelSim-Altera Software

You can perform simulation of Verilog HDL or VHDL designs with the ModelSim-Altera software at three levels: functional, post-synthesis, and gate-level.

For high-speed simulation, you must select **ps** in the **Resolution** list for your simulator resolutions (**Design** tab of the **Start Simulation** dialog box). If you choose slower than **ps**, the high-speed simulation may fail.

Performing Functional Simulation

Functional simulation verifies code syntax and design functionality. The following sections describe how to perform functional simulation in the ModelSim-Altera software for a Verilog HDL or VHDL design.

 The ModelSim-Altera software includes precompiled simulation libraries for Altera-provided models. You should not create simulation libraries and compile simulation models for the pre-compiled Altera libraries.

Setting Up a Quartus II Project for the ModelSim-Altera Software

The first steps in performing a simulation are starting the ModelSim-Altera software, changing to your project/simulation directory, and creating libraries for your design.

 For more information, refer to *Setting Up a Project with the ModelSim-Altera Software* in Quartus II Help.

Compiling and Loading Designs with the ModelSim-Altera Software

 For information about compiling and loading your design files and testbench files, refer to *Mapping to Libraries and Compiling Design Files with the ModelSim-Altera Software* in Quartus II Help.

Performing the Simulation

 For information about performing a functional simulation with the ModelSim-Altera software, refer to *Performing a Functional Simulation with the ModelSim-Altera Software* in Quartus II Help.

Performing Post-Synthesis Simulation

You can perform post-synthesis simulation to verify that design functionality is preserved after synthesis. You can create the post-synthesis netlist in the Quartus II software and use the netlist to perform post-synthesis simulation with the ModelSim-Altera software.

 Before running post-synthesis simulation, generate post-synthesis simulation netlist files. For more information, refer to the “Generating Post-Synthesis Simulation Netlist Files” section in the *Simulating Altera Designs* chapter in volume 3 of the *Quartus II Handbook*.

 The ModelSim-Altera software includes precompiled simulation libraries. You should not create simulation libraries and compile simulation models.

- For information about performing a post-synthesis simulation with the ModelSim-Altera software, refer to *Performing a Timing Simulation with the ModelSim-Altera Software* in Quartus II Help.

Performing Gate-Level Timing Simulation

Gate-level timing simulation is an important step in ensuring that the device functionality is correct and meets all timing requirements following place and route. You can create the gate-level netlist in the Quartus II software and use the netlist to perform gate-level timing simulation with the ModelSim-Altera software.

- Before running gate-level timing simulation, generate gate-level timing simulation netlist files. For more information, refer to the “Generating Gate-Level Timing Simulation Netlist Files” section in the *Simulating Altera Designs* chapter in volume 3 of the *Quartus II Handbook*.
- The ModelSim-Altera software includes precompiled simulation libraries. It is not necessary to create simulation libraries and compile simulation models.
- For information about performing a gate-level simulation with the ModelSim-Altera software, refer to *Performing a Timing Simulation with the ModelSim-Altera Software* in Quartus II Help.
- For additional documentation about ModelSim-Altera, refer to the ModelSim-Altera Help that ships with the product. Click the **Help** button on the ModelSim-Altera toolbar.

Performing Simulation Using the ModelSim and QuestaSim Software

You can perform simulation of Verilog HDL or VHDL designs with the ModelSim and QuestaSim software at three levels: functional, post-synthesis, and gate-level.

You can perform the simulation through the GUI or from the command line. The following sections provide instructions to perform the simulation through the GUI and from the command line. You can proceed to the specific section that meets your needs.

For high-speed simulation, you must select **ps** in the **Resolution** list for your simulator resolutions (**Design** tab of the **Start Simulation** dialog box). If you choose slower than **ps**, the high-speed simulation may fail.

Simulating VHDL Designs Using the GUI

This section provides information about performing functional, post-synthesis, and gate-level simulations of VHDL designs using the GUI.

Performing Functional Simulation

This section provides information about compiling simulation models and performing a functional simulation.

Compiling Simulation Models into Simulation Libraries

- ② For more information about how to compile simulation models into simulation libraries if you are not using the EDA Simulation Library Compiler, refer to *Compiling Libraries and Design Files with the ModelSim Software* in Quartus II Help.
- ② For more information about targeting a Stratix V device, refer to *Guidelines for Compiling Stratix V Libraries* in Quartus II Help.
-  The `altera_mf_components.vhd` and `altera_mf.vhd` model files should be compiled into the `altera_mf` library. The `220pack.vhd` and `220model.vhd` model files should be compiled into the `lpm` library.

Performing the Simulation

- ② For information about simulating VHDL designs using the GUI, refer to *Performing a Functional Simulation with the ModelSim Software* and *Performing a Functional Simulation with the QuestaSim Software* in Quartus II Help.
- ② To see all of the functional simulation library files, refer to *Altera Functional Simulation Libraries* in Quartus II Help.

Performing Post-Synthesis Simulation

You can perform post-synthesis simulation to verify that design functionality is preserved after synthesis. You can create the post-synthesis netlist in the Quartus II software and use the netlist to perform post-synthesis simulation with the ModelSim and QuestaSim software.

-  Before running post-synthesis simulation, generate post-synthesis simulation netlist files. For more information, refer to the “Generating Post-Synthesis Simulation Netlist Files” section in the *Simulating Altera Designs* chapter in volume 3 of the *Quartus II Handbook*.
-  You cannot perform post-synthesis or post-fit (gate-level) simulation if you are targeting the Stratix V device family.
- ② For information about performing a post-synthesis simulation using the GUI, refer to *Performing a Timing Simulation with the ModelSim Software* and *Performing a Timing Simulation with the QuestaSim Software* in Quartus II Help.

Performing Gate-Level Simulation

Gate-level simulation is a very important step in ensuring that the FPGA device’s functionality is still correct and meets all required timing requirements after the design was placed and routed. You can create the gate-level netlist in the Quartus II software and use the netlist to perform gate-level simulation with the ModelSim and QuestaSim software.

-  You cannot perform post-synthesis or post-fit (gate-level) simulation if you are targeting the Stratix V device family.

-  Before running gate-level simulation, generate gate-level timing simulation netlist files. For more information, refer to the “Generating Gate-Level Timing Simulation Netlist Files” section in the *Simulating Altera Designs* chapter in volume 3 of the *Quartus II Handbook*.
-  For information about performing a gate-level simulation using the GUI, refer to *Performing a Timing Simulation with the ModelSim Software* and *Performing a Timing Simulation with the QuestaSim Software* in Quartus II Help.

Simulating Verilog HDL Designs Using the GUI

This section provides information about performing functional, post-synthesis, and gate-level simulations of Verilog HDL designs using the GUI.

Performing Functional Simulation

This section provides information about compiling simulation models and performing a functional simulation.

Compiling Simulation Models into Simulation Libraries

-  For more information about how to compile simulation models into simulation libraries if you are not using the EDA Simulation Library Compiler, refer to *Compiling Libraries and Design Files with the ModelSim Software* in Quartus II Help.
-  For more information about targeting a Stratix V device, refer to *Guidelines for Compiling Stratix V Libraries* in Quartus II Help.
-  The PCIe file is required only if you are using the PCIe hard IP.
-  The `altera_mf.v` model files should be compiled into the `altera_mf_ver` library. The `220model.v` model files should be compiled into the `lpm_ver` library.

Performing the Simulation

-  For information about performing a functional simulation using the GUI, refer to *Performing a Functional Simulation with the ModelSim Software* and *Performing a Functional Simulation with the QuestaSim Software* in Quartus II Help.
-  To see all of the functional simulation library files, refer to *Altera Functional Simulation Libraries* in Quartus II Help.

Performing Post-Synthesis Simulation

Perform post-synthesis simulation to verify that design functionality is preserved after synthesis. Create the post-synthesis netlist in the Quartus II software and use the netlist to perform post-synthesis simulation with the ModelSim and QuestaSim software.

-  Before running post-synthesis simulation, generate post-synthesis simulation netlist files. For more information, refer to the “Generating Post-Synthesis Simulation Netlist Files” section in the *Simulating Altera Designs* chapter in volume 3 of the *Quartus II Handbook*.
-  You cannot perform post-synthesis or post-fit (gate-level) simulation if you are targeting the Stratix V device family.
-  For information about performing a post-synthesis simulation using the GUI, refer to *Performing a Timing Simulation with the ModelSim Software* and *Performing a Timing Simulation with the QuestaSim Software* in Quartus II Help.

Performing Gate-Level Simulation

Gate-level simulation is a very important step in ensuring that the FPGA device’s functionality is still correct and meets all required timing requirements after the design was placed and routed. You can create the gate-level netlist in the Quartus II software and use the netlist to perform gate-level simulation with the ModelSim and QuestaSim software.

-  You cannot perform post-synthesis or post-fit (gate-level) simulation if you are targeting the Stratix V device family.
-  Before running gate-level simulation, generate gate-level timing simulation netlist files. For more information, refer to the “Generating Gate-Level Timing Simulation Netlist Files” section in the *Simulating Altera Designs* chapter in volume 3 of the *Quartus II Handbook*.
-  For information about performing a gate-level simulation using the GUI, refer to *Performing a Timing Simulation with the ModelSim Software* and *Performing a Timing Simulation with the QuestaSim Software* in Quartus II Help.

Simulating VHDL Designs From the Command Line

This section provides information about performing functional, post-synthesis, and gate-level simulations of VHDL designs from the command line.

Simulating VHDL designs from the ModelSim and QuestaSim command line gives you more flexibility and control in compiling the libraries and loading and simulating the VHDL design files. All simulation commands are Tcl commands that can be included in the *.do file. Using the *.do file allows you to run simulation in batch mode. You have to execute only the *.do file, and the ModelSim and QuestaSim tool automatically executes all commands in the *.do script macro file.

Performing Functional Simulation

Function simulation verifies code syntax and design functionality.

Type the following commands to perform a functional simulation for VHDL designs with one of the libraries (lib1) listed in the *Altera Functional Simulation Libraries* in Quartus II Help.

To create and compile Altera libraries, type the following commands:

```
vlib <lib1> ↵  
vmap <lib1> <lib1> ↵  
vcom -work <lib1> <lib1>.vhd ↵  
vlib <lib2> ↵  
vmap <lib2> <lib2> ↵  
vcom -work <lib2> <lib2>.vhd ↵
```

To create the work library and compile the design and testbench files, type the following commands:

```
vlib work ↵  
vmap work work ↵  
vcom -work work <design_file1>.vhd <design_file2>.vhd <testbench \  
file>.vhd ↵
```

To load the design, type the following command:

```
vsim -L work -L <lib1> -L <lib2> work.<testbench module name> ↵
```

To add signals to the waveform viewer and run the simulation, type the following commands:

```
add wave * ↵  
run ↵
```

Example

```
# Create and compile Altera libraries  
  
vlib altera_mf  
vmap altera_mf altera_mf  
vcom -work altera_mf altera_mf_components.vhd altera_mf.vhd  
vlib lpm  
vmap lpm lpm  
vcom -work lpm 220pack.vhd 220model.vhd  
  
# Create work library and compile design files and testbench file  
  
vlib work  
vmap work work  
vcom -work work top_level.vhd adder.vhd testbench.vhd  
  
# Load design  
  
vsim -L work -L altera_mf -L lpm work.testbench  
  
# add signals to the waveform viewer and run simulation  
  
add wave *  
run
```

- ② For more information about targeting a Stratix V device, refer to *Guidelines for Compiling Stratix V Libraries* in Quartus II Help.

Performing Post-Synthesis Simulation

Perform post-synthesis simulation to verify that design functionality is preserved after synthesis. Create the post-synthesis netlist in the Quartus II software and use the netlist to perform post-synthesis simulation with the ModelSim and QuestaSim software. Before running post-synthesis simulation, generate post-synthesis simulation netlist files.

-  For more information, refer to the “Generating Post-Synthesis Simulation Netlist Files” section in the *Simulating Altera Designs* chapter in volume 3 of the Quartus II Handbook.
-  You cannot perform post-synthesis or post-fit (gate-level) simulation if you are targeting the Stratix V device family.

Type the following commands to perform a post-synthesis simulation for VHDL designs with one of the libraries (lib1) listed in *Altera Post-Fit Libraries* in Quartus II Help.

To create and compile Altera libraries, type the following commands:

```
vlib <lib1> ←
vmap <lib1> <lib1> ←
vcom -work <lib1> <lib1>.vhd ←
vlib <lib2> ←
vmap <lib2> <lib2> ←
vcom -work <lib2> <lib2>.vhd ←
```

To create the work library and compile design and testbench files, type the following commands:

```
vlib work ←
vmap work work ←
vcom -work work <output_netlist>.vho <testbench file>.vhd ←
```

To load the design, type the following command:

```
vsim +transport_int_delays +transport_path_delays -L work -L \ <lib1>
-L <lib2> work.<testbench module name> ←
```

To add signals to the waveform viewer and run simulation, type the following commands:

```
add wave * ←
run ←
```

Example

```
# Create and compile Altera libraries

vlib altera
vmap altera altera
vcom -work altera altera_primitives_components.vhd \
altera_primitives.vhd
vlib stratixiii
vmap stratixiii stratixiii
vcom -work stratixiii stratixiii.atoms.vhd stratixiii_components.vhd

# Create work library and compile design files and testbench file

vlib work
vmap work work
vcom -work work top_level.vho testbench.vhd
```

```
# Load design
vsim +transport_int_delays +transport_path_delays -L work -L \ altera
-L stratixiii work.testbench
# add signals to the waveform viewer and run simulation
add wave *
run
```

Performing Gate-Level Simulation

The steps for gate-level timing simulation are almost same as the steps for post-synthesis simulation.

The only difference is that the **.sdo** file must be back-annotated for gate level-timing simulation.

For VHDL designs, add the **-sdftyp** option for back-annotating.

Example

```
vsim +transport_int_delays +transport_path_delays -sdftyp \ <instance
path to design>= <path to SDO file> -L work -L stratixiii -L \ altera
work.testbench
```

You do not have to set the value (minimum, average, maximum) for the ***.sdo** file, because the Quartus II EDA Netlist Writer generates the ***.sdo** file using the same value for the triplet (minimum, average, and maximum timing values).

If your design under test is instantiated in the testbench file under the **i1** label, the **<design instance>** should be "i1" (for example, /i1=<my design>.sdo).



You cannot perform post-synthesis or post-fit (gate-level) simulation if you are targeting the Stratix V device family.

Simulating Verilog HDL Designs from the Command Line

This section provides information about performing functional, post-synthesis, and gate-level simulations of Verilog HDL designs from the command line.

Simulating Verilog HDL designs from the ModelSim and QuestaSim command line gives you more flexibility and control in compiling the libraries and loading and simulating the Verilog HDL design files. All simulation commands are Tcl commands that can be included in the *.do file. Using the *.do file allows you to run simulation in batch mode. You have to execute only the *.do file, and the ModelSim and QuestaSim tool automatically executes all commands in the *.do script macro file.

Performing Functional Simulation

Functional simulation verifies code syntax and design functionality.

Type the following commands to perform a functional simulation for Verilog HDL designs with one of the libraries (lib1) listed in *Altera Functional Simulation Libraries* in Quartus II Help.

To create and compile Altera libraries, type the following commands:

```
vlib <lib1> ↵
vmap <lib1> <lib1> ↵
vlog -work <lib1> <lib1>.v ↵
vlib <lib2> ↵
vmap <lib2> <lib2> ↵
vlog -work <lib2> <lib2>.v ↵
```

To create the work library and compile design and testbench files, type the following commands:

```
vlib work ↵
vmap work work ↵
vlog -work work <design_file1>.v <design_file2>.v <testbench_file>.v ↵
```

To load the design, type the following command:

```
vsim -L work -L <lib1> -L <lib2> work.<testbench module name> ↵
```

To add signals to the waveform viewer and run simulation, type the following commands:

```
add wave * ↵
run ↵
```

Example

```
# Create and compile Altera libraries
vlib altera_mf_ver
vmap altera_mf_ver altera_mf_ver
vlog -work altera_mf_ver altera_mf.v
vlib lpm_ver
vmap lpm_ver lpm_ver
vlog -work lpm_ver 220model.v

# Create work library and compile design files and testbench file
vlib work
vmap work work
vlog -work work top_level.v adder.v testbench.v

# Load design
vsim -L work -L altera_mf_ver -L lpm_ver work.testbench
```

```
# add signals to the waveform viewer and run simulation
add wave *
run
```

-  For more information about targeting a Stratix V device, refer to [Guidelines for Compiling Stratix V Libraries](#) in Quartus II Help.

Performing Post-Synthesis Simulation

Perform post-synthesis simulation to verify that design functionality is preserved after synthesis. Create the post-synthesis netlist in the Quartus II software and use the netlist to perform post-synthesis simulation with the ModelSim and QuestaSim software. Before running post-synthesis simulation, generate post-synthesis simulation netlist files.

-  For more information, refer to the “Generating Post-Synthesis Simulation Netlist Files” section in the [Simulating Altera Designs](#) chapter in volume 3 of the Quartus II Handbook.

-  You cannot perform post-synthesis or post-fit (gate-level) simulation if you are targeting the Stratix V device family.

Type the following commands to perform a post-synthesis simulation for Verilog HDL designs with one of the libraries (lib1) listed in [Altera Post-Fit Libraries](#) in Quartus II Help.

To create and compile Altera libraries, type the following commands:

```
vlib <lib1> ↵
vmap <lib1> <lib1> ↵
vlog -work <lib1> <lib1>.v ↵
vlib <lib2> ↵
vmap <lib2> <lib2> ↵
vlog -work <lib2> <lib2>.v ↵
```

To create the work library and compile design and testbench files, type the following commands:

```
vlib work ↵
vmap work work ↵
vlog -work work <output_netlist>.vo <testbench file>.v ↵
```

To load the design, type the following command:

```
vsim -t ps +transport_int_delays +transport_path_delays -L work -L \
<lib1> -L <lib2> work.<testbench module name> ↵
```

To add signals to the waveform viewer and run the simulation, type the following commands:

```
add wave * ↵
run ↵
```

Example

```
# Create and compile Altera libraries
```

```
vlib altera_ver
vmap altera_ver altera_ver
vlog -work altera_ver altera_primitives.v
vlib stratixiii_ver

vmap stratixiii_ver stratixiii_ver
vlog -work stratixiii_ver stratixiii_atoms.v

# Create work library and compile design files and testbench file

vlib work
vmap work work
vlog -work work top_level.vo testbench.v

# Load design

vsim +transport_int_delays +transport_path_delays -L work -L
altera_ver -L stratixiii_ver work.testbench

#add signals to the waveform viwer and run simulation

add wave *
run
```

Performing Gate-Level Simulation

The steps for gate-level timing simulation are almost same as the steps for post-synthesis simulation.

The only difference is that the **.sdo** file must be back-annotated for gate level-timing simulation.

For Verilog HDL designs, the back-annotating process is done within the **output_netlist.vo** script. Therefore, you are not required to back-annotate the SDO file again.



You cannot perform post-synthesis or post-fit (gate-level) simulation if you are targeting the Stratix V device family.

Passing Parameter Information from Verilog to VHDL

You must use in-line parameters to pass values from Verilog HDL to VHDL. Using the `defparam` construct will cause an error in simulation. In the example below:

```
lpm_add_sub_component (
    .dataa (dataa),
    .datab (datab),
    .result (sub_wire0)
);
defparam
lpm_add_sub_component.lpm_direction = "ADD",
lpm_add_sub_component.lpm_hint =
"ONE_INPUT_IS_CONSTANT=NO,CIN_USED=NO",
lpm_add_sub_component.lpm_type = "LPM_ADD_SUB",
lpm_add_sub_component.lpm_width = 12;
```

You will see the following error message:

```
# ** Error: (vsim-3043)
/apps2/home/users/bhlee/SPR_ADOQS/ADOQS10000935_IN_LINE_PARAMETER/lpm_
add_sub1.v(67): Unresolved reference to 'lpm_add_sub_component' in
lpm_add_sub_component.lpm_direction.
# Region: /IN_LINE_PARAMETER_vlg_vec_tst/il/b2v_inst
```

This megafunction instantiation has been modified to use in-line parameters:

```
lpm_add_sub#(12,"SIGNED","ADD",0,"LPM_ADD_SUB","ONE_INPUT_IS_CONSTANT=
NO,CIN_USED=NO")
lpm_add_sub_component (
    .dataa (dataa),
    .datab (datab),
    .result (sub_wire0)
);
```



The sequence of the parameters depends on the sequence of the GENERIC in the VHDL component declaration.

Speeding Up Simulation

By default, the ModelSim and QuestaSim software runs in a debug-optimized mode. To run the ModelSim and QuestaSim software in speed-optimized mode, add the following two `vlog` command line switches:

```
vlog -fast -05
```

In this mode, module boundaries are flattened and loops are optimized. This eliminates levels of debugging hierarchy, which may result in faster simulation. This switch is not supported in the ModelSim-Altera simulator.

Simulating Designs that Include Transceivers

If your design includes an Arria GX, Arria II GX, Cyclone IV, HardCopy IV, Stratix GX, Stratix II GX, Stratix IV, or Stratix V transceiver, you must compile additional library files to perform functional or gate-level timing simulations.



You cannot perform post-synthesis or post-fit (gate-level) simulation if you are targeting the Stratix V device family.

- ❓ For Stratix V, you must compile the libraries listed in *Guidelines for Compiling Stratix V Libraries* in Quartus II Help.

Performing simulation with transceivers in Arria II, Cyclone IV, HardCopy IV, and Stratix IV device families are very similar. You have to replace only `stratixiigx_atoms` and `stratixiigx_hssi_atoms` model files with `arriaii_atoms` and `arriaii_hssi_atoms` model files for Arria II devices, `cycloneiv_atoms` and `cycloneiv_hssi_atoms` model files for Cyclone IV devices, and `stratixiv_atoms` and `stratixiv_hssi_atoms` model files for Stratix IV devices.

For high-speed simulation, you must select `ps` in the **Resolution** list for your simulator resolutions (**Design** tab of the **Start Simulation** dialog box). If you choose slower than `ps`, the high-speed simulation may fail.

- 👉 If your design contains PCI Express® hard IP, refer to the “Simulate the Design” section in the *PCI Express Compiler User Guide*.
- 👉 The VHDL version of ModelSim-Altera and other single language VHDL versions of ModelSim cannot simulate designs that target the Stratix V device family.

You need a version of ModelSim that supports VHDL/Verilog co-simulation to simulate designs that use Stratix V transceivers.

Functional Simulation for Stratix GX Devices

To perform a functional simulation of your design that instantiates the ALTGXB megafunction, which enables the gigabit transceiver block on Stratix GX devices, compile the `stratixgx_mf` model file into the `altgxb` library.

- 👉 The `stratixgx_mf` model file references the `lpm` and `sgate` libraries. If you are using ModelSim and QuestaSim, you must create these libraries to perform a simulation.

Performing Functional Simulation in VHDL (ModelSim-Altera)

To perform functional simulation for Stratix GX devices in VHDL, type the following commands:

```
vcom -work <my_design>.vhd <my_testbench>.vhd ←
vsim -L lpm -L altera_mf -L sgate -L altgxb work.<my_testbench> ←
```

Performing Functional Simulation in VHDL (ModelSim and QuestaSim)

To perform functional simulation for Stratix GX devices in VHDL, type the following commands:

```
vcom -work altera_mf altera_mf_components.vhd altera_mf.vhd ←
vcom -work lpm 220pack.vhd 220model.vhd ←
vcom -work sgate sgate_pack.vhd sgate.vhd ←
vcom -work altgxb stratixgx_mf.vhd stratixgx_mf_components.vhd ←
vcom -work <my_design>.vhd <my_testbench>.vhd ←
vsim -L lpm -L altera_mf -L sgate -L altgxb work.<my_testbench> ←
```

Performing Functional Simulation in Verilog HDL (ModelSim-Altera)

To perform functional simulation for Stratix GX devices in Verilog HDL, type the following commands:

```
vlog -work <my design>.v <my testbench>.v ←
vsim -L lpm_ver -L altera_mf_ver -L sgate_ver -L altgxb work.<my
testbench> ←
```

Performing Functional Simulation in Verilog HDL (ModelSim and QuestaSim)

To perform functional simulation for Stratix GX devices in Verilog HDL, type the following commands:

```
vlib work_ver ←
vlib lpm_ver ←
vlib altera_mf_ver ←
vlib sgate_ver ←
vlib altgxb_ver ←
vlog -work lpm_ver 220model.v ←
vlog -work altera_mf_ver altera_mf.v ←
vlog -work sgate_ver sgate.v ←
vlog -work altgxb_ver stratixgx_mf.v ←
vlog -work <my design>.v <my testbench>.v ←
vsim -L lpm_ver -L sgate_ver -L altgxb_ver work.<my testbench> ←
```

Gate-Level Timing Simulation for Stratix GX Devices

Perform a gate-level timing simulation of your design that includes a Stratix GX transceiver by compiling the **stratixgx_atoms** and **stratixgx_hssi_atoms** model files into the **stratixgx** and **stratixgx_gxb** libraries, respectively.



The **stratixgx_hssi_atoms** model file references the **lpm** and **sgate** libraries. If you are using ModelSim and QuestaSim, you must create these libraries to perform a simulation.

Performing Gate-Level Timing Simulation in VHDL (ModelSim-Altera)

To perform gate-level timing simulation for Stratix GX devices in VHDL, type the following commands:

```
vcom -work <my design>.vho <my testbench>.vhd ←
vsim -L lpm -L altera_mf -L sgate -L stratixgx -L stratixgx_gxb \
-sdftyp <design instance>=<path to .sdo file>.sdo work.<my testbench> \
-t ps - +transport_int_delays+transport_path_delays←
```

Performing Gate-Level Timing Simulation in VHDL (ModelSim and QuestaSim)

To perform gate-level timing simulation for Stratix GX devices in VHDL, type the following commands:

```
vcom -work lpm 220pack.vhd 220model.vhd ←
vcom -work altera_mf altera_mf_components.vhd altera_mf.vhd ←
vcom -work sgate sgate_pack.vhd sgate.vhd ←
vcom -work stratixgx stratixgx_atoms.vhd stratixgx_components.vhd ←
vcom -work stratixgx_gxb stratixgx_hssi_atoms.vhd \
stratixgx_hssi_components.vhd ←
vcom -work <my design>.vho <my testbench>.vhd ←
vsim -L lpm -L altera_mf -L sgate -L stratixgx -L stratixgx_gxb \
-sdftyp <design instance>=<path to .sdo file>.sdo work.<my testbench> \
-t ps +transport_int_delays +transport_path_delays ←
```

Performing Gate-Level Timing Simulation in Verilog HDL (ModelSim-Altera)

To perform gate-level timing simulation for Stratix GX devices in Verilog HDL, type the following commands:

```
vlog -work <my design>.vo <my testbench>.v ←
vsim -L lpm_ver -L altera_mf_ver -L sgate_ver -L stratixgx_ver -L \
stratixgx_gxb_ver work.<my testbench> -t ps +transport_int_delays \
+transport_path_delays ←
```

Performing Gate-Level Timing Simulation in Verilog HDL (ModelSim and QuestaSim)

To perform gate-level timing simulation for Stratix GX devices in Verilog HDL, type the following commands:

```
vlog -work lpm_ver 220model.v ←
vlog -work altera_mf_ver altera_mf.v ←
vlog -work sgate_ver sgate.v ←
vlog -work stratixgx_ver stratixgx_atoms.v ←
vlog -work stratixgx_gxb_ver stratixgx_hssi_atoms.v ←
vlog -work <my design>.vo <my testbench>.v ←
vsim -L lpm_ver -L altera_mf_ver -L sgate_ver -L stratixgx_ver \
-L stratixgx_gxb_ver work.<my testbench> -t ps +transport_int_delays \
+transport_path_delays ←
```

Functional Simulation for Stratix II GX Devices

To perform functional simulation of your design that instantiates the ALT2GXB megafunction, which enables the gigabit transceiver block on Stratix II GX devices, compile the **stratixiigx_hssi** model file into the **stratixiigx_hssi** library.



The **stratixiigx_hssi_atoms** model file references the **lpm** and **sgate** libraries. If you are using ModelSim and QuestaSim, you must create these libraries to perform a simulation.

Generate a functional simulation netlist by turning on **Generate Simulation Model** in the **Simulation Library** tab of the ALT2GXB MegaWizard Plug-In Manager. The **<alt2gxb entity name>.vho** or **<alt2gxb module name>.vo** is generated in the current project directory.



The ALT2GXB functional simulation library file generated by the Quartus II software references **stratixiigx_hssi** WYSIWYG atoms.

Performing Functional Simulation in VHDL (ModelSim-Altera)

To perform functional simulation for Stratix II GX devices in VHDL, type the following commands:

```
vcom -work work <alt2gxb entity name>.vho ←
vcom -work <my design>.vhd <my testbench>.vhd ←
vsim -L lpm -L altera_mf -L sgate -L stratixgx_hssi work.<my design> ←
```

Performing Functional Simulation in VHDL (ModelSim and QuestaSim)

To perform functional simulation for Stratix II GX devices in VHDL, type the following commands:

```
vcom -work lpm 220pack.vhd 220model.vhd ↵
vcom -work altera_mf altera_mf_components.vhd altera_mf.vhd ↵
vcom -work sgate sgate_pack.vhd sgate.vhd ↵
vcom -work stratixiigx_hssi stratixiigx_hssi_components.vhd \
stratixiigx_hssi_atoms.vhd ↵
vcom -work work <alt2gxb entity name>.vho ↵
vcom -work <my design>.vhd <my testbench>.vhd ↵
vsim -L lpm -L altera_mf -L sgate -L stratixgx_hssi work.<my testbench>
↵
```

Performing Functional Simulation in Verilog HDL (ModelSim-Altera)

To perform functional simulation for Stratix II GX devices in Verilog HDL, type the following commands:

```
vlog -work work <alt2gxb module name>.vo ↵
vlog -work <my design>.v <my testbench>.v ↵
vsim -L lpm_ver -L altera_mf_ver -L sgate_ver -L stratixgx_hssi_ver \
work.<my testbench> ↵
```

Performing Functional Simulation in Verilog HDL (ModelSim and QuestaSim)

To perform functional simulation for Stratix II GX devices in Verilog HDL, type the following commands:

```
vlog -work lpm_ver 220model.v ↵
vlog -work altera_mf_ver altera_mf.v ↵
vlog -work sgate_ver sgate.v ↵
vlog -work stratixiigx_hssi_ver stratixiigx_hssi_atoms.v ↵
vlog -work work <alt2gxb module name>.vo ↵
vlog -work <my design>.v <my testbench>.v ↵
vsim -L lpm_ver -L altera_mf_ver -L sgate_ver -L stratixgx_hssi_ver \
work.<my testbench> ↵
```

Gate-Level Timing Simulation for Stratix II GX Devices

To perform a gate-level timing simulation of your design that includes a Stratix II GX transceiver, compile **stratixiigx_atoms** and **stratixiigx_hssi_atoms** into the **stratixiigx** and **stratixiigx_hssi** libraries, respectively.



The **stratixiigx_hssi_atoms** model file references the **lpm** and **sgate** libraries. If you are using ModelSim and QuestaSim, you must create these libraries to perform a simulation.

Performing Gate-Level Timing Simulation in VHDL (ModelSim-Altera)

To perform gate-level timing simulation for Stratix II GX devices in VHDL, type the following commands:

```
vcom -work <my design>.vho <my testbench>.vhd ↵
vsim -L lpm -L altera_mf -L sgate -L stratixiigx -L stratixiigx_hssi \
-sdftyp <design instance>=<path to .sdo file>.sdo work.<my testbench> \
-t ps +transport_int_delays +transport_path_delays ↵
```

Performing Gate-Level Timing Simulation in VHDL (ModelSim and QuestaSim)

To perform gate-level timing simulation for Stratix II GX devices in VHDL, type the following commands:

```
vcom -work lpm 220pack.vhd 220model.vhd ←
vcom -work altera_mf altera_mf_components.vhd altera_mf.vhd ←
vcom -work sgate sgate_pack.vhd sgate.vhd ←
vcom -work stratixiigx stratixiigx_atoms.vhd \
stratixiigx_components.vhd ←
vcom -work stratixiigx_hssi stratixiigx_hssi_components.vhd \
stratixiigx_hssi_atoms.vhd ←
vcom -work <my design>.vho <my testbench>.vhd ←
vsim -L lpm -L altera_mf -L sgate -L stratixiigx -L stratixiigx_hssi \
-sdftyp <design instance>=<path to .sdo file>.sdo work.<my testbench> \
-t ps +transport_int_delays +transport_path_delays ←
```

Performing Gate-Level Timing Simulation in Verilog HDL ModelSim-Altera)

To perform gate-level timing simulation for Stratix II GX devices in Verilog HDL, type the following commands:

```
vlog -work <my design>.vo <my testbench>.v ←
vsim -L lpm_ver -L altera_mf_ver -L sgate_ver -L stratixiigx_ver \
-L stratixiigx_hssi_ver work.<my testbench> -t ps \
+transport_int_delays +transport_path_delays ←
```

Performing Gate-Level Timing Simulation in Verilog HDL ModelSim and QuestaSim)

To perform gate-level timing simulation for Stratix II GX devices in Verilog HDL, type the following commands:

```
vlog -work lpm_ver 220model.v ←
vlog -work altera_mf_ver altera_mf.v ←
vlog -work sgate_ver sgate.v ←
vlog -work stratixiigx_ver stratixiigx_atoms.v ←
vlog -work stratixiigx_hssi_ver stratixiigx_hssi_atoms.v ←
vlog -work <my design>.vo <my testbench>.v ←
vsim -L lpm_ver -L altera_mf_ver -L sgate_ver -L stratixiigx_ver \
-L stratixiigx_hssi_ver work.<my testbench> -t ps \
+transport_int_delays +transport_path_delays ←
```

Functional Simulation for Stratix IV GX Devices

To perform a functional simulation of your design that instantiates the ALTGX megafunction, which enables the gigabit transceiver block on Stratix IV devices, compile the `stratixiv_hssi` model file into the `altgx` library.



The `stratixiv_hssi` model file references the `lpm` and `sgate` libraries. If you are using ModelSim and QuestaSim, you must create these libraries to perform a simulation.

Performing Functional Simulation in VHDL (ModelSim-Altera)

To perform functional simulation for Stratix IV devices in VHDL, type the following commands:

```
vcom -work <my design>.vhd <my testbench>.vhd ←
vsim -L lpm -L altera_mf -L sgate -L stratixiv_hssi work.<my testbench>←
```

Performing Functional Simulation in VHDL (ModelSim and QuestaSim)

To perform functional simulation for Stratix IV devices in VHDL, type the following commands:

```
vcom -work altera_mf altera_mf_components.vhd altera_mf.vhd ↵
vcom -work lpm 220pack.vhd 220model.vhd ↵
vcom -work sgate sgate_pack.vhd sgate.vhd ↵
vcom -work stratixiv_hssi \
stratixiv_hssi_atoms.vhd stratixiv_hssi_components.vhd ↵
vcom -work <my design>.vhd <my testbench>.vhd ↵
vsim -L lpm -L altera_mf -L sgate -L stratixiv_hssi work.<my testbench>
↵
```

Performing Functional Simulation in Verilog HDL (ModelSim-Altera)

To perform functional simulation for Stratix IV devices in Verilog HDL, type the following commands:

```
vlog -work <my design>.v <my testbench>.v ↵
vsim -L lpm_ver -L altera_mf_ver -L sgate_ver -L \
stratixiv_hssi_ver work.<my testbench> ↵
```

Performing Functional Simulation in Verilog HDL (ModelSim and QuestaSim)

To perform functional simulation for Stratix IV devices in Verilog HDL, type the following commands:

```
vlog -work lpm_ver 220model.v ↵
vlog -work altera_mf_ver altera_mf.v ↵
vlog -work sgate_ver sgate.v ↵
vlog -work stratixiv_hssi_ver stratixiv_hssi_atoms.v ↵
vlog -work <my design>.v <my testbench>.v ↵
vsim -L lpm_ver -L sgate_ver -L stratixiv_hssi_ver work.<my testbench> ↵
```

Gate-Level Timing Simulation for Stratix IV GX Devices

Perform a gate-level timing simulation of your design that includes a Stratix IV transceiver by compiling the `stratixiv_atoms` and `stratixiv_hssi_atoms` model files into the `stratixiv` and `stratixiv_hssi` libraries, respectively.



The `stratixgx_hssi_atoms` model file references the `lpm` and `sgate` libraries. If you are using ModelSim and QuestaSim, you must create these libraries to perform a simulation.

Performing Gate-Level Timing Simulation in VHDL (ModelSim-Altera)

To perform gate-level timing simulation for Stratix IV devices in VHDL, type the following commands:

```
vcom -work <my design>.vho <my testbench>.vhd ↵
vsim -L lpm -L altera_mf -L sgate -L stratixiv -L stratixiv_hssi \
-sdftyp <design instance>=<path to .sdo file>.sdo work.<my testbench> \
-t ps - +transport_int_delays +transport_path_delays ↵
```

Performing Gate-Level Timing Simulation in VHDL (ModelSim and QuestaSim)

To perform gate-level timing simulation for Stratix IV devices in VHDL, type the following commands:

```
vcom -work lpm 220pack.vhd 220model.vhd ←
vcom -work altera_mf altera_mf_components.vhd altera_mf.vhd ←
vcom -work sgate sgate_pack.vhd sgate.vhd ←
vcom -work stratixiv stratixiv_atoms.vhd stratixiv_components.vhd ←
vcom -work stratixiv_hssi stratixiv_hssi_atoms.vhd \
stratixiv_hssi_components.vhd ←
vcom -work <my design>.vho <my testbench>.vhd ←
vsim -L lpm -L altera_mf -L sgate -L stratixiv -L stratixiv_hssi \
-sdftyp <design instance>=<path to .sdo file>.sdo work.<my testbench> \
-t ps +transport_int_delays +transport_path_delays ←
```

Performing Gate-Level Timing Simulation in Verilog HDL (ModelSim-Altera)

To perform gate-level timing simulation for Stratix IV devices in Verilog HDL, type the following commands:

```
vlog -work <my design>.vo <my testbench>.v ←
vsim -L lpm_ver -L altera_mf_ver -L sgate_ver -L stratixiv_ver -L \
stratixiv_hssi_ver work.<my testbench> -t ps +transport_int_delays \
+transport_path_delays ←
```

Performing Gate-Level Timing Simulation in Verilog HDL (ModelSim and QuestaSim)

To perform gate-level timing simulation for Stratix IV devices in Verilog HDL, type the following commands:

```
vlog -work lpm_ver 220model.v ←
vlog -work altera_mf_ver altera_mf.v ←
vlog -work sgate_ver sgate.v ←
vlog -work stratixiv_ver stratixiv_atoms.v ←
vlog -work stratixiv_hssi_ver stratixiv_hssi_atoms.v ←
vlog -work <my design>.vo <my testbench>.v ←
vsim -L lpm_ver -L altera_mf_ver -L sgate_ver -L stratixiv_ver \
-L stratixiv_hssi_ver work.<my testbench> -t ps +transport_int_delays \
+transport_path_delays ←
```

Functional Simulation for Stratix V GX Devices

To perform a functional simulation of your design that instantiates the Custom PHY megafunction, which enables the gigabit transceiver block on Stratix V devices, compile the `stratixv_hssi` model file.



The `stratixv_hssi` model file references the `lpm` and `sgate` libraries. You must create these libraries to perform a simulation using ModelSim and QuestaSim.



The transceiver module from the MegaWizard Plug-In Manager is created in `Interfaces/Transceiver PHY`. Select `Custom PHY`.

Performing Functional Simulation in VHDL (ModelSim and QuestaSim)

For information about how to perform functional simulation for Stratix V devices in VHDL, refer to [“Performing Functional Simulation” on page 2-9](#).

Performing Functional Simulation in Verilog HDL (ModelSim-Altera)

- For information about how to perform functional simulation for Stratix V devices in Verilog HDL, refer to *Performing a Functional Simulation with the ModelSim-Altera Software* and *Guidelines for Compiling Stratix V Libraries* in Quartus II Help.

Performing Functional Simulation in Verilog HDL (ModelSim and QuestaSim)

For information about how to perform functional simulation for Stratix V devices in Verilog HDL, refer to “Performing Functional Simulation” on page 2–12.

Transport Delays

By default, the ModelSim and QuestaSim software filters out all pulses that are shorter than the propagation delay between primitives. Turning on the transport delay options in the ModelSim and QuestaSim software prevents the simulation tool from filtering out these pulses. Use the following options to ensure that all signal pulses are seen in the simulation results.

+transport_path_delays

Use this option when the pulses in your simulation are shorter than the delay within a gate-level primitive.

+transport_int_delays

Use this option when the pulses in your simulation are shorter than the interconnect delay between gate-level primitives.

- The **+transport_path_delays** and **+transport_int_delays** options are also used by default in the NativeLink feature for gate-level timing simulation.

- For more information about either of these options, refer to the ModelSim-Altera Command Reference installed with the ModelSim and QuestaSim software.

The following ModelSim and QuestaSim software command shows the command line syntax to perform a gate-level timing simulation with the device family library:

```
vsim -t lps -L stratixii -sdftyp /il=filtref_vhd.sdo work.filtref_vhd_vec_tst \  
+transport_int_delays +transport_path_delays
```

Using the NativeLink Feature with ModelSim-Altera, ModelSim, or QuestaSim Software

The NativeLink feature in the Quartus II software facilitates the seamless transfer of information between the Quartus II software and EDA tools and allows you to run ModelSim and QuestaSim within the Quartus II software.

- For more information, refer to the “Using the NativeLink Feature” section in the *Simulating Altera Designs* chapter in volume 3 of the *Quartus II Handbook*.

ModelSim and QuestaSim Error Message Verification

ModelSim and QuestaSim error and warning messages are tagged with a `vsim` or `vcom` code. To determine the cause and resolution for a `vsim` or `vcom` error or warning, use the `verror` command.

For example, ModelSim and QuestaSim may display the following error message:

```
# ** Error:
C:/altera_trn/DUALPORT_TRY/simulation/modelsim/DUALPORT_TRY.vho (31):
(vcom-1136) Unknown identifier "stratixiii".
```

In this case, type the following command:

```
verror 1136 ←
```

At that point, the error message appears as follows:

```
# vcom Message # 1136:
# The specified name was referenced but was not found. This indicates
# that either the name specified does not exist or is not visible at
# this point in the code.
```

Generating a Timing Value Change Dump (.vcd) File for the PowerPlay Power Analyzer

To generate a timing Value Change Dump (*.vcd) file for the PowerPlay Power Analyzer, you must first generate a *.vcd script file in the Quartus II software and run the *.vcd script file from the ModelSim, QuestaSim, or ModelSim-Altera software to generate a timing *.vcd file. This timing *.vcd file can then be used by PowerPlay for power analysis. The following instructions show you step-by-step how to generate a timing *.vcd file.

To generate timing VCD Scripts in the Quartus II software, perform the following steps:

1. In the Quartus II software, on the Assignments menu, click **Settings**. The **Settings** dialog box appears.
2. In the **Category** list, click the "+" icon to expand **EDA Tool Settings** and select **Simulation**. The **Simulation** page appears.
3. Choose the appropriate third-party simulation tool (ModelSim, QuestaSim, or ModelSim-Altera) in the **Tool name** list. Turn on the **Generate Value Change Dump (VCD) file script** option.
4. To generate the *.vcd script file, perform a full compilation.

To generate a timing *.vcd file in the ModelSim-Altera, ModelSim, or QuestaSim software, perform the following steps:

1. In the ModelSim-Altera, ModelSim, or QuestaSim software, before simulating your design, source the `<revision_name>_dump_all_vcd_nodes.tcl` script. To source the Tcl script, run the following command before running the `vsim` command. For example:

```
source <revision_name>_dump_all_vcd_nodes.tcl ←
```

2. Continue to run the simulation as usual until the end of the simulation. Exit the ModelSim-Altera, ModelSim, or QuestaSim software. If you do not exit the software, the ModelSim and QuestaSim software may end the writing process of the timing *.vcd files improperly, resulting in a corrupted timing *.vcd file.

 For more information about using the timing *.vcd file for power estimation, refer to the *PowerPlay Power Analysis* chapter in volume 3 of the *Quartus II Handbook*.

Viewing a Waveform from a .wlf File

A .wlf file is automatically generated when your simulation is done. The .wlf file is used for generating the waveform view through ModelSim-Altera, ModelSim, or QuestaSim.

To view a waveform from a .wlf file through ModelSim-Altera, ModelSim, or QuestaSim, perform the following steps:

1. Type `vsim` on the command line. The **ModelSim/QuestaSim** or **ModelSim-Altera** dialog box appears.
2. On the File menu, click **Datasets**. The **Datasets Browser** dialog box appears.
3. Click **Open** and browse to the directory that contains your .wlf file.
4. Select the .wlf file and click **Open**, then click **OK**.
5. Click **Done**.
6. In the Object browser, select the signals that you want to observe.
7. On the Add menu, click **Wave** and then click **Selected Signals**.

You cannot view a waveform from a .vcd file in ModelSim-Altera, ModelSim, or QuestaSim directly. The .vcd file must first be converted to a .wlf file.

1. Use the `vcd2wlf` command to convert the file. For example, type the following on a command-line:

```
vcd2wlf <example>.vcd <example>.wlf ↵
```
2. After you convert the .vcd file to a .wlf file, follow the procedures for viewing a waveform from a .wlf file through ModelSim and QuestaSim.

You can also convert your .wlf file to a .vcd file by using the `wlf2vcd` command.

Simulating with ModelSim-Altera Waveform

You can use the ModelSim-Altera Waveform Editor as a simple method to create design stimulus for simulation. You can create this design stimulus via interactive manipulation of waveforms from the wave window in ModelSim-Altera. With the ModelSim-Altera waveform editor, you can create and edit waveforms, drive simulation directly from created waveforms, and save created waveforms into a stimulus file.

For more information, refer to the *Generating Stimulus with Waveform Editor* chapter in the *ModelSim SE User's Manual* available on the ModelSim website (model.com).

Scripting Support

You can run procedures and create settings described in this chapter in a Tcl script. You can also run some procedures at the command line prompt.

-  For more information about Tcl scripting, refer to the *Tcl Scripting* chapter in volume 2 of the *Quartus II Handbook*.
-  For more information about command line scripting, refer to the *Command Line Scripting* chapter in volume 2 of the *Quartus II Handbook*.

For detailed information about scripting command options, refer to the Quartus II Help command line and Tcl API help browser. To access this information, type the following command to start a help browser:

```
quartus_sh --qhelp ←
```

Generating a Post-Synthesis Simulation Netlist for ModelSim and QuestaSim

You can use the Quartus II software to generate a post-synthesis simulation netlist with Tcl commands or with a command at the command-line prompt. The following example assumes that you are selecting ModelSim and QuestaSim (Verilog HDL output from the Quartus II software).

Tcl Commands

Use the following Tcl commands to set the output format to Verilog HDL, the simulation tool to ModelSim and QuestaSim for Verilog HDL, and to generate a functional netlist:

```
set_global_assignment-name EDA_SIMULATION_TOOL "ModelSim (Verilog)" ←
set_global_assignment-name EDA_GENERATE_FUNCTIONAL_NETLIST ON ←
```

or

```
set_global_assignment-name EDA_SIMULATION_TOOL "QuestaSim (Verilog)" ←
set_global_assignment-name EDA_GENERATE_FUNCTIONAL_NETLIST ON ←
```

Command Prompt

Use the following command to generate a simulation output file for the ModelSim and QuestaSim simulator. Specify VHDL or Verilog HDL for the format:

```
quartus_eda <project name> --simulation=on --format=<format> \
--tool=ModelSim --functional ←
```

or

```
quartus_eda <project name> --simulation=on --format=<format> \
--tool=QuestaSim --functional ←
```

Generating a Gate-Level Timing Simulation Netlist for ModelSim and QuestaSim

Use the Quartus II software to generate a gate-level timing simulation netlist with Tcl commands or with a command at the command prompt.

Tcl Commands

Use one of the following Tcl commands:

- `set_global_assignment -name EDA_SIMULATION_TOOL \ "ModelSim-Altera (Verilog)" ←`

or
`set_global_assignment -name EDA_SIMULATION_TOOL \ "QuestaSim-Altera (Verilog)" ←`
- `set_global_assignment -name EDA_SIMULATION_TOOL \ "ModelSim-Altera (VHDL)" ←`

or
`set_global_assignment -name EDA_SIMULATION_TOOL \ "QuestaSim-Altera (VHDL)" ←`
- `set_global_assignment -name EDA_SIMULATION_TOOL \ "ModelSim (Verilog)" ←`

or
`set_global_assignment -name EDA_SIMULATION_TOOL \ "QuestaSim (Verilog)" ←`
- `set_global_assignment -name EDA_SIMULATION_TOOL \ "ModelSim (VHDL)" ←`

or
`set_global_assignment -name EDA_SIMULATION_TOOL \ "QuestaSim (VHDL)" ←`

Command Line

Generate a simulation output file for the ModelSim and QuestaSim simulator by specifying VHDL or Verilog HDL for the format by typing the following command at the command prompt:

```
quartus_eda <project name> --simulation=on --format=<format> \  
--tool=ModelSim ←
```

or

```
quartus_eda <project name> --simulation=on --format=<format> \  
--tool=QuestaSim ←
```

Software Licensing and Licensing Setup in ModelSim-Altera Subscription Edition

License the ModelSim-Altera Subscription Edition software subscription with a parallel port FIXEDPC license, or a network FLOATNET or FLOATPC license. Each Altera software subscription includes a license for both VHDL and Verilog HDL. The ModelSim-Altera Subscription Edition software supports both VHDL and Verilog HDL, but the software does not support mixed language simulation.

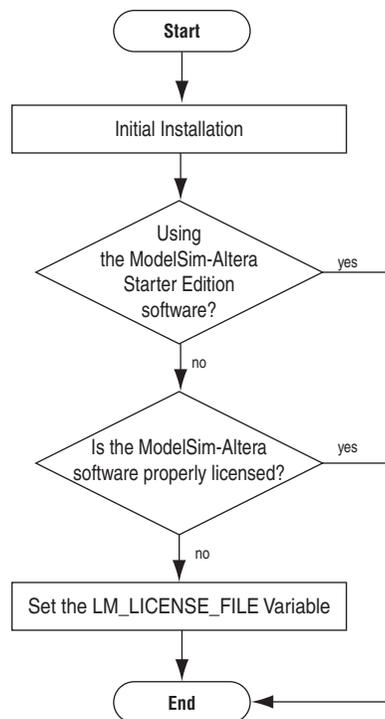
 The USB software guard is not supported by versions earlier than Mentor Graphics ModelSim software 5.8d.

You can obtain a license for the ModelSim-Altera Subscription Edition software from the Altera website at www.altera.com. Get licensing information for the Mentor Graphics ModelSim software directly from Mentor Graphics. Refer to [Figure 2-1](#) for the set-up process.

 For ModelSim-Altera software versions prior to 5.5b, use the PCLS utility included with the software to set up the license.

For the Quartus II software version 8.1 and later, the no-cost entry level of the ModelSim-Altera software does not require a license file. However, you must request a license file to use the ModelSim-Altera Subscription Edition software.

Figure 2-1. ModelSim-Altera Subscription Edition Software Licensing Set Up Process



LM_LICENSE_FILE Variable

Altera recommends setting the LM_LICENSE_FILE environment variable to the location of the license file. For example, the value for the LM_LICENSE_FILE environment variable should point to *<path to license file>\license.dat*.

 For more information about setting up the license for ModelSim-Altera Subscription Edition software, refer to the [Altera Software Installation and Licensing](#) manual.

Conclusion

Using the ModelSim, QuestaSim, and ModelSim-Altera simulation software within the Altera FPGA design flow enables Altera software users to easily and accurately perform functional simulations, post-synthesis simulations, and gate-level simulations on their designs. Proper verification of designs at the functional, post-synthesis, and post place-and-route stages using the ModelSim, QuestaSim, and ModelSim-Altera software helps ensure design functionality and success and, ultimately, a quick time-to-market.

Document Revision History

Table 2-1 shows the revision history for this chapter.

Table 2-1. Document Revision History (Part 1 of 2)

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Changed to new document template ■ Referenced Simulating Altera Designs chapter ■ Added new section, “Simulating with ModelSim-Altera Waveform” on page 2-25 ■ Removed Stratix V compilation information and linked to Quartus II Help
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Removed simulation library tables and linked to Quartus II Help ■ Added other links to Quartus II Help and ModelSim-Altera Help where appropriate and removed redundant information ■ Added QuestaSim support ■ Added Stratix V simulation information ■ Minor editorial changes throughout ■ Removed Referenced Documents section
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Removed NativeLink information and referenced new <i>Simulating Designs with EDA Tools</i> chapter ■ Added Stratix IV transceiver simulation section ■ Reformatted transceiver simulation sections ■ Text edits throughout chapter
March 2009	9.0.0	<p>Added the following sections:</p> <ul style="list-style-type: none"> ■ “Compile Libraries Using the EDA Simulation Library Compiler” on page 2-17 ■ “Generate Simulation Script from EDA Netlist Writer” on page 2-77 ■ “Viewing a Waveform from a .wlf File” on page 2-78 <p>Updated the following:</p> <ul style="list-style-type: none"> ■ Table 2-1, Table 2-2, Table 2-5, Table 2-6, Table 2-7, Table 2-8, Table 2-9, Table 2-10 ■ Figure 2-4 on page 2-81 ■ All sections titled “Loading the Design”

Table 2-1. Document Revision History (Part 2 of 2)

Date	Version	Changes
November 2008	8.1.0	Updated the following: <ul style="list-style-type: none"> ■ Table 2-2, Table 2-3, Table 2-4, Table 2-5, Table 2-6 ■ Removed <code>--zero_ic_delays</code> from <code>quartus_sta</code> option in “Generate Post-Synthesis Simulation Netlist Files” on page 2-11 ■ Removed steps to include the library when the simulation is run in VHDL mode from all procedures; this is no longer necessary ■ Added information about the Altera Simulation Library Compiler throughout the chapter ■ Added “Compile Libraries Using the Altera Simulation Library Compiler” on page 2-15 ■ Added “Disabling Simulation” on page 2-72 ■ Minor editorial updates ■ Updated entire chapter using 8½” × 11” chapter template
May 2008	8.0.0	Updated the following: <ul style="list-style-type: none"> ■ “Altera Design Flow with ModelSim-Altera or ModelSim Software” on page 2-3 ■ “Simulation Libraries” on page 2-4 ■ “Simulation Netlist Files” on page 2-11 ■ “Perform Simulation Using ModelSim-Altera Software” on page 2-15 ■ “Perform Simulation Using ModelSim Software” on page 2-33 ■ “Simulating Designs that Include Transceivers” on page 2-57 ■ “Using the NativeLink Feature with ModelSim-Altera or ModelSim Software” on page 2-63 ■ “Generating a Timing VCD File for PowerPlay” on page 2-68

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter describes how to use the Synopsys VCS and VCS MX software to simulate designs that target Altera® FPGAs. This chapter provides step-by-step instructions about how to perform functional simulations, post-synthesis simulations, and gate-level timing simulations.

 Verilog HDL design simulation is the same in both the VCS and VCS MX software. In this chapter, VCS MX is used in VHDL design simulation examples.

This chapter includes the following topics:

- “Software Requirements”
- “Using the VCS or VCS MX Software in the Quartus II Design Flow” on page 3–1
- “Common VCS and VCS MX Software Compiler Options” on page 3–8
- “Using DVE” on page 3–8
- “Debugging Support Command-Line Interface” on page 3–8
- “Simulating Designs that Include Transceivers” on page 3–9
- “Transport Delays” on page 3–13
- “Using NativeLink with the VCS or VCS MX Software” on page 3–13
- “Generating a Timing .vcd File for the PowerPlay Power Analyzer” on page 3–13
- “Viewing a Waveform from a .vpd or .vcd File” on page 3–14
- “Scripting Support” on page 3–15

Software Requirements

To simulate your design using the Synopsys VCS or VCS MX software, you must first set up the Altera libraries. These libraries are installed with the Quartus® II software.

 For more information about installing the software and the directories created during the Quartus II software installation, refer to the *Altera Software Installation and Licensing* manual.

Using the VCS or VCS MX Software in the Quartus II Design Flow

You can perform the following types of simulations using the VCS and VCS MX software:

- Functional simulation
- Post-synthesis simulation
- Gate-level timing simulation

- Refer to the “PLD Design Flow” section in the *Simulating Designs with EDA Tools* chapter in volume 3 of the *Quartus II Handbook* for the Quartus II software design flow.

Compiling Libraries Using the EDA Simulation Library Compiler

The EDA Simulation Library Compiler is used to compile Verilog HDL and VHDL simulation libraries for all Altera devices and supported third-party simulators. You can compile all libraries required by functional and gate-level timing simulations.

- For more information, refer to the “EDA Simulation Library Compiler” section in the *Simulating Designs with EDA Tools* chapter in volume 3 of the *Quartus II Handbook*.

Functional Simulations

Functional simulations verify the functionality of the design before synthesis, placement, and routing. These simulations are independent of any Altera FPGA architecture implementation. After the HDL designs are verified to be functionally correct, the next step is to synthesize the design and use the Quartus II software to place-and-route the design in an Altera device.

To functionally simulate an Altera FPGA design in the VCS or VCS MX software that uses Altera intellectual property (IP) megafunctions or a library of parameterized modules (LPM) functions, you must include certain libraries during the compilation.

- For a list of the functional simulation library files in the Quartus II directory, refer to *Altera Functional Simulation Libraries* in Quartus II Help.

Functional Simulation for Verilog HDL Designs

The following VCS command performs a functional simulation for Verilog HDL designs with one of the libraries listed in *Altera Functional Simulation Libraries* in Quartus II Help:

```
vcs -R <testbench>.v <design name>.v -v <altera_library1>.v -v <altera_library2>.v ←
```

If you have already generated the option file (**simlib_comp.vcs**) from “*Compiling Libraries Using the EDA Simulation Library Compiler*” on page 3-2, type the following command:

```
vcs -f simlib_comp.vcs ←
```

Be sure to include the design files and testbench files in **simlib_comp.vcs**.

Alternatively, you can use the following commands to perform functional simulation for Verilog HDL designs:

```
# Create Libraries Directories
mkdir <Directory_to_store_compiled_altera_library1>
mkdir <Directory_to_store_compiled_altera_library2>

# Create Work Directory
mkdir <Directory_to_store_compiled_design_and_testbench_files>

# Compilation
# (Before this step, make sure the mapped file "synopsys_sim.setup" was created.)
# Libraries Compilation
vlogan -work <altera_library1_name> <altera_library1>.v
```

```
vlogan -work <altera_library2_name> <altera_library2>.v

# Design and Test bench files Compilation
vlogan -work <work_library_name> <design>.v <testbench>.v

# Elaboration
vcs -debug_all <work_library_name>.<testbench_top_level_module>

# Run Simulation
simv -gui
```

The `synopsys_sim.setup` file contains the following mapping commands to map the libraries:

```
<altera_library1_name> : <Directory_to_store_compiled_altera_library1>
<altera_library2_name> : <Directory_to_store_compiled_altera_library2>
<work_library_name> : <Directory_to_store_compiled_design_and_testbench_files>
```

 The `altera_mf.v` model files should be compiled into the `altera_mf_ver` library. The `220model.v` model files should be compiled into the `lpm_ver` library.

If you are targeting a Stratix V device, compile the following files in the `quartus/eda/sim_lib/synopsys` directory:

```
stratixv_atoms_ncrypt.v
stratixv_hssi_atoms_ncrypt.v
stratixv_pcie_hip_atoms_ncrypt.v
```

These files contain IEEE encrypted Verilog models.

 Compile `stratixv_pcie_hip_atoms_ncrypt.v` with the SystemVerilog option.

Also, compile the following files in the `quartus/eda/sim_lib` directory:

```
stratixv_atoms.v
stratixv_hssi_atoms.v
stratixv_pcie_hip_atoms.v
```

 The PCIe® files are required only if you are using the PCIe HIP.

Functional Simulation for VHDL Designs

For VHDL designs, you need to use VCS MX software to perform all three types of simulations. The following commands are for performing a functional simulation for VHDL designs with one of the libraries listed in *Altera Functional Simulation Libraries* in Quartus II Help:

```
# Create Libraries Directories
mkdir <Directory_to_store_compiled_altera_library1>
mkdir <Directory_to_store_compiled_altera_library2>

# Create Work Directory
mkdir <Directory_to_store_compiled_design_and_testbench_files>

# Compilation
# (Before this step, make sure the mapped file "synopsys_sim.setup" is created.)
# Libraries Compilation
vhdlan -work <altera_library1_name> <altera_library1>.vhd
vhdlan -work <altera_library2_name> <altera_library2>.vhd
```

```
# Design and Test bench files Compilation
vhdlan -work <work_library_name> <design>.vhd <testbench>.vhd

# Elaboration
vcs -debug_all <work_library_name>.<testbench_top_level_module>
# Run Simulation

simv -gui
```

The **synopsys_sim.setup** file contains the following mapping commands to map the libraries:

```
<altera_library1_name> : <Directory_to_store_compiled_altera_library1>
<altera_library2_name> : <Directory_to_store_compiled_altera_library2>
<work_library_name> : <Directory_to_store_compiled_design_and_testbench_files>
```

 The **altera_mf.v** model files should be compiled into the **altera_mf_ver** library. The **220model.v** model files should be compiled into the **lpm_ver** library.

 If you have generated the Altera libraries with the EDA Simulation Library Compiler, ignore the steps **# Create Libraries Directories** and **# Libraries Compilation**.

If you are targeting a Stratix V device, compile the following files in the **quartus/eda/sim_lib/synopsys** directory:

```
stratixv_atoms_ncrypt.v
stratixv_hssi_atoms_ncrypt.v
stratixv_pcie_hip_atoms_ncrypt.v
```

These files contain IEEE encrypted Verilog models suitable for VHDL/Verilog co-simulation. You need a co-simulation license from Synopsys to use these models.

 Compile these encrypted Verilog files before you compile any VHDL files.

 Compile **stratixv_pcie_hip_atoms_ncrypt.v** with the SystemVerilog option.

Also, compile the following files in the **quartus/eda/sim_lib** directory:

```
stratixv_atoms.vhd
stratixv_components.vhd
stratixv_hssi_components.vhd
stratixv_hssi_atoms.vhd
stratixv_pcie_components.vhd
stratixv_pcie_hip_atoms.vhd
```

 The PCIe files are required only if you are using the PCIe HIP.

Post-Synthesis Simulation

A post-synthesis simulation verifies the functionality of a design after synthesis has been performed. You can create a post-synthesis netlist in the Quartus II software and use this netlist to perform a post-synthesis simulation in the VCS or VCS MX software. When the post-synthesis version of the design has been verified, the next step is to place-and-route the design in the target architecture using the Quartus II software.

 For more information, refer to the “Generating Post-Synthesis Simulation Netlist Files” section in the *Simulating Altera Designs* chapter in volume 3 of the *Quartus II Handbook*.

 You cannot perform post-synthesis or post-fit simulation if you are targeting the Stratix V device family.

Post-Synthesis Simulation for Verilog HDL Designs

The following VCS command shows the command-line syntax used to perform a post-synthesis simulation with the appropriate device family library listed in *Altera Functional Simulation Libraries* in Quartus II Help.

```
vcs -R <testbench> <post_synthesis_netlist> -v <altera_library1> ←
```

If you have already generated the option file (**simlib_comp.vcs**) as described in “Compiling Libraries Using the EDA Simulation Library Compiler” on page 3–2, modify the **simlib_comp.vcs** file to add the testbench and post synthesis netlist, and then type the following command:

```
vcs -f simlib_comp.vcs ←
```

Be sure to include the post synthesis netlist file and testbench files in **simlib_comp.vcs**.

Alternatively, you can use the following commands to perform Post-Synthesis simulation for Verilog HDL designs:

```
# Create Libraries Directories
mkdir <Directory_to_store_compiled_altera_library1>
mkdir <Directory_to_store_compiled_altera_library2>
# Create Work Directory
mkdir <Directory_to_store_compiled_post_synthesis_netlist_and_testbench_files>
# Compilation
# (Before this step, make sure that mapped file "synopsys_sim.setup" is created.)
# Libraries Compilation
vlogan -work <altera_library1_name> <altera_library1>.v
vlogan -work <altera_library2_name> <altera_library2>.v
# Design and Test bench files Compilation
vlogan -work <work_library_name> <post_synthesis_netlist>.vo <test_bench>.v
# Elaboration
vcs -debug_all <work_library_name>.<testbench_top_level_module>
# Run Simulation
simv -gui
```

The **synopsys_sim.setup** file contains the following commands to map the libraries:

```
<altera_library1_name> : <Directory_to_store_compiled_altera_library1>
<altera_library2_name> : <Directory_to_store_compiled_altera_library2>
<work_library_name> :
    <Directory_to_store_compiled_post_synthesis_netlist_and_testbench_files>
```

Post-Synthesis Simulation for VHDL Designs

The following VCS MX commands show the command-line syntax used to perform a post-synthesis simulation with the appropriate device family library listed in *Altera Post-Fit Libraries* in Quartus II Help.

```
# Create Libraries Directories
mkdir <Directory_to_store_compiled_altera_library1>
mkdir <Directory_to_store_compiled_altera_library2>
# Create Work Directory
```

```
mkdir <Directory_to_store_compiled_post_synthesis_netlist_and_testbench_files>
# Compilation
#(Before this step, make sure that mapped file "synopsys_sim.setup" is created.)
# Libraries Compilation
vhdlan -work <altera_library1_name> <altera_library1>.vhd
vhdlan -work <altera_library2_name> <altera_library2>.vhd
# Design and Test bench files Compilation
vhdlan -work <work_library_name> <post_synthesis_netlist>.vho <test_bench>.vhd
# Elaboration
vcs -debug_all <work_library_name>.<testbench_top_level_module>
# Run Simulation
simv -gui
```

The `synopsys_sim.setup` file contains the following commands to map the libraries:

```
<altera_library1_name> : <Directory_to_store_compiled_altera_library1>
<altera_library2_name> : <Directory_to_store_compiled_altera_library2>
<work_library_name> :
    <Directory_to_store_compiled_post_synthesis_netlist_and_testbench_files>
```

 If you have generated the Altera libraries as described in “[Compiling Libraries Using the EDA Simulation Library Compiler](#)” on page 3–2, ignore the steps # Create Libraries Directories and # Libraries Compilation.

Gate-Level Timing Simulation

A gate-level timing simulation verifies the functionality and timing of the design after place-and-route. You can create a post-fit netlist in the Quartus II software and use this netlist to perform a gate-level timing simulation in the VCS or VCS MX software.

 For information about how to generate a gate-level timing simulation netlist, refer to the “[Generating Gate-Level Timing Simulation Netlist Files](#)” section in the *Simulating Altera Designs* chapter in volume 3 of the *Quartus II Handbook*.

 For a list of the gate-level timing simulation library files in the Quartus II directory, refer to *Altera Post-Fit Libraries* in Quartus II Help.

 You cannot perform post-synthesis or post-fit simulation if you are targeting the Stratix V device family.

Gate-Level Timing Simulation for Verilog HDL Designs

For gate-level timing simulation, follow the steps in “[Post-Synthesis Simulation for Verilog HDL Designs](#)” on page 3–5.

You do not have to specify the Standard Delay Output File (*.sdo) file because it is already specified in the Verilog Output File (*.vo) file. However, the *.sdo file must be in the same directory as the simulator executable file `simv` generated by VCS.

Gate-Level Timing Simulation for VHDL Designs

For gate-level timing simulation, follow the steps in “[Post-Synthesis Simulation for VHDL Designs](#)” on page 3–5.

For VHDL, the *.sdo file must be specified in the `simv` command as follows:

```
simv -xlrn -gui -sdf typ:<testbench module name>/<design instance name>.sdo ←
```

 Adding the `-x1rm` switch avoids the errors that occur when the timing arcs in SDO do not match Altera VHDL simulation models as per the IEEE VITAL ASIC standard. However, adding this switch reduces timing accuracy, as it may cause some SDO delays to be ignored. Therefore, Altera recommends generating the Verilog HDL simulation output netlist (`.vo`) if you want to perform gate-level timing simulation.

Disabling Timing Violation on Registers

In certain situations, the timing violations can be ignored and you can disable the “X” propagation that happens when there are timing violations on registers (for example, timing violations that occur in internal synchronization registers used for asynchronous clock-domain crossing).

By default, the `x_on_violation_option logic` option that applies to all registers in the design is **On**, which means a register outputs “X” whenever a timing violation occurs. To disable “X” propagation due to a timing violation on certain registers, you can set the `x_on_violation_option logic` option to **Off** for those registers. The following command is an example of the Quartus II Settings File (`.qsf`):

```
set_instance_assignment -name X_ON_VIOLATION_OPTION OFF -to <register_name> ←
```

Performing Timing Simulation Using the Post-Synthesis Netlist

You can perform a timing simulation using the post-synthesis netlist instead of using a gate-level netlist, and you can generate an `.sdo` file without running the Fitter. In this case, the `.sdo` file includes all timing values for the device cells only. Interconnect delays are not included because fitting (placement and routing) has not been performed.

To generate the post-synthesis netlist and the `.sdo` file, type the following commands at a command prompt:

```
quartus_map <project name> -c <revision name> ←  
quartus_tan <project name> -c <revision name> --post_map \  
--zero_ic_delays ←  
quartus_eda <project name> -c <revision name> --simulation \  
--tool= <third-party EDA tool> --format=<HDL language> ←
```

For more information about the `-format` and `-tool` options, type the following command:

```
quartus_eda -help=<options> command ←
```

Common VCS and VCS MX Software Compiler Options

The VCS software has options that help you simulate your design. [Table 3-1](#) lists some of the available options.

Table 3-1. VCS Software Compiler Options

Library	Description
-R	Runs the executable file immediately.
-v <library filename>	Specifies a Verilog HDL library file (for example, 220model.v or altera_mf.v). The VCS software looks in this file for module definitions that are found in the source code.
-y <library directory>	Specifies a Verilog HDL library directory. The VCS software looks for library files in this folder that contain module definitions that are instantiated in the source code.
+compsdf	Indicates that the VCS software compiler includes the back-annotated Standard Delay File (.sdf) file in the compilation.
+cli	The VCS software enters Command-Line Interface (CLI) mode upon successful compilation completion.
+race	Specifies that the VCS software generate a report that indicates all of the race conditions in the design. The default report name is race.out.
-P	Compiles user-defined Programming Language Interface (PLI) table files.
-q	Indicates the VCS software runs in quiet mode. All messages are suppressed.

 For more information about VCS software options, refer to the [VCS User Guide](#).

Using DVE

Design Viewpoint Editor (DVE) is the graphical debugging system for the VCS and VCS MX software. This tool is included with the VCS MX software. It can be run by using the `-gui` (simulating option) when running a simulation.

The following VCS or VCS MX software command shows the command-line syntax for simulating in DVE:

```
simv -gui (for Verilog HDL simulations)
```

However, to open the GUI with these commands, you must enable the use of Unified Command Line Interface (UCLI) and DVE when performing elaboration. To enable UCLI and DVE, enter the following command:

```
vcs -debug_all (for Verilog HDL simulations)
```

 For more information about using DVE, refer to the [DVE User Guide](#) included in the VCS MX software installation.

Debugging Support Command-Line Interface

The Synopsys VCS software has an interactive non-graphical debugging capability that is very similar to other UNIX debuggers such as the GNU debugger (GDB). The VCS software UCLI can be used to halt simulations at user-defined break points, force registers with values, and display register values.

Enable the non-graphical capability by using the `+ucli` run-time option. Use the VCS software UCLI to debug your Altera FPGA design by typing the following command:

```
vcs -R <testbench>.v <design name>.vo -v <path to Quartus II \
installation> \eda\sim_lib\<device family>_atoms.v +compsdf +ucli ←
```

The `+ucli` command takes an optional number argument that specifies the level of debugging capability. As the optional debugging capability is increased, there is an increase in simulation time.

-  For more information about the `+ucli` options, refer to the *VCS User Guide* included in the VCS software installation.

For the design examples to run gate-level timing simulations in VHDL or Verilog HDL language, refer to *Synopsys VCS Simulation Design Example*.

Simulating Designs that Include Transceivers

If your design includes Arria®, Arria II, Cyclone® IV, HardCopy® IV, Stratix®, Stratix II, Stratix IV, or Stratix V transceivers, you must compile additional library files to perform functional or gate-level timing simulations.

For high-speed simulation, you must select **ps** in the **Resolution** list for your simulator resolutions (**Design** tab of the **Start Simulation** dialog box). If you choose slower than **ps**, the high-speed simulation might fail.

-  If your design contains PCI Express hard IP, refer to the “Simulate the Design” section in the *PCI Express Compiler User Guide*.

Functional Simulation for Stratix GX Devices

To perform a functional simulation of your design that instantiates the ALTGXB megafunction, enabling the gigabit transceiver block gigabit transceiver block on Stratix GX devices, compile the `stratixgx_mf` model file into the `altgxb` library.

-  The `stratixgx_mf` model file references the `lpm` and `sgate` libraries. You must create these libraries to perform a simulation.

Compiling Library Files for Functional Simulation in Verilog HDL

To compile the libraries necessary for functional simulation of a Verilog HDL design targeting a Stratix GX device, at the VCS command prompt, type the following command:

```
vcs -R <testbench>.v <design files>.v -v stratixgx_mf.v -v sgate.v \
-v 220model.v -v altera_mf.v ←
```

Gate-Level Timing Simulation for Stratix GX Devices

Perform a gate-level timing simulation of your design that includes a Stratix GX transceiver by compiling the `stratixgx_atoms` and `stratixgx_hssi_atoms` model files into the `stratixgx` and `stratixgx_gxb` libraries, respectively.

-  The `stratixgx_hssi_atoms` model file references the `lpm` and `sgate` libraries. You must create these libraries to perform a simulation.

Compiling Library Files for Gate-Level Timing Simulation in Verilog HDL

To compile the libraries necessary for timing simulation of a Verilog HDL design targeting a Stratix GX device, at the VCS command prompt, type the following command:

```
vcs -R <testbench>.v <gate-level netlist>.vo -v stratixgx_atoms.v -v \
stratixgx_hssi_atoms.v -v sgate.v -v 220model.v -v altera_mf.v +transport_int_delays \
+pulse_int_e/0 +pulse_int_r/0 +transport_path_delays +pulse_e/0 +pulse_r/0 ←
```

Functional Simulation for Stratix II GX Devices

Functional simulation for Stratix II GX devices is similar to functional simulation for Arria GX devices. To simulate the transceiver in Arria GX devices, you only have to replace the `stratixiigx_hssi` model file with the `arriagx_hssi` model file.

To perform a functional simulation of your design that instantiates the ALT2GXB megafunction, enabling the gigabit transceiver block on Stratix II GX devices, compile the `stratixiigx_hssi` model file into the `stratixiigx_hssi` library.

-  The `stratixiigx_hssi_atoms` model file references the `lpm` and `sgate` libraries. You must create these libraries to perform a simulation.

Generate a functional simulation netlist by turning on **Generate Simulation Model in the Simulation Library** in the ALT2GXB MegaWizard™ Plug-In Manager. The `<alt2gxb entity name>.vho` file or `<alt2gxb module name>.vo` file is generated in the current project directory.

-  The ALT2GXB functional simulation library file generated by the Quartus II software references `stratixiigx_hssi` `wysiwyg` `atoms`.

Compiling Library Files for Functional Simulation in Verilog HDL

To compile the libraries necessary for functional simulation of a Verilog HDL design targeting a Stratix II GX device, type the following command at the VCS command prompt:

```
vcs -R <testbench>.v <alt2gxb simulation netlist>.vo -v \
stratixgx_hssi_atoms.v -v sgate.v -v 220model.v -v altera_mf.v ←
```

Gate-Level Timing Simulation for Stratix II GX Devices

Gate-level timing simulation for Stratix II GX devices is similar to gate-level timing simulation for Arria GX devices. To simulate the transceiver in Arria GX devices, you only have to replace the `stratixiigx_hssi` model file with the `arriagx_hssi` model file.

To perform a gate-level timing simulation of your design that includes a Stratix II GX transceiver, compile `stratixiigx_atoms` and `stratixiigx_hssi_atoms` into the `stratixiigx` and `stratixiigx_hssi` libraries, respectively.

-  The `stratixiigx_hssi_atoms` model file references the `lpm` and `sgate` libraries. You must create these libraries to perform a simulation.

Compiling Library Files for Gate-Level Timing Simulation in Verilog HDL

To compile the libraries necessary for timing simulation of a Verilog HDL design targeting a Stratix II GX device, type the following command at the VCS command prompt:

```
vcs -R <testbench>.v <gate-level netlist>.vo -v stratixiigx_atoms.v -v \  
stratixiigx_hssi_atoms.v -v sgate.v -v 220model.v -v altera_mf.v \  
+transport_int_delays +pulse_int_e/0 +pulse_int_r/0 \  
+transport_path_delays +pulse_e/0 +pulse_r/0 ←
```

Functional Simulation for Stratix IV GX Devices

Functional simulation for Stratix IV devices is similar to functional simulation for Arria II, Cyclone IV, and HardCopy IV devices. To simulate transceivers in Arria II, Cyclone IV, and HardCopy IV devices, you only have to replace the **stratixiv_hssi** model file with the **arriaii_hssi**, **cycloneiv_hssi**, and **hardcopyiv_hssi** model files, respectively.

To perform a functional simulation of your design that instantiates the ALTGX megafunction, enabling the gigabit transceiver block on Stratix IV devices, compile the **stratixiv_hssi** model file into the **stratixiv_hssi** library,

The **stratixiv_hssi_atoms** model file references the **lpm** and **sgate** libraries. You must create these libraries to perform a simulation.

Compiling Library Files for Functional Simulation in Verilog HDL

To compile the libraries necessary for functional simulation of a Verilog HDL design targeting a Stratix IV device, type the following command at the VCS command prompt:

```
vcs -R <testbench>.v <altgx>.v -v \ stratixiv_hssi_atoms.v -v sgate.v \  
-v 220model.v -v altera_mf.v ←
```

Gate-Level Timing Simulation for Stratix IV GX Devices

Gate-level timing simulation for Stratix IV devices is similar to gate-level timing simulation for Arria II, Cyclone IV, and HardCopy IV devices. To simulate transceivers in Arria II, Cyclone IV, and HardCopy IV devices, you only have to replace the **stratixiv_hssi** model file with the **arriaii_hssi**, **cycloneiv_hssi**, and **hardcopyiv_hssi** model files, respectively.

To perform a gate-level timing simulation of your design that includes a Stratix IV transceiver, compile **stratixiv_atoms** and **stratixiv_hssi_atoms** into the **stratixiv** and **stratixiv_hssi** libraries, respectively.

To perform a gate-level timing simulation of your design that includes a Stratix IV transceiver, compile **stratixiv_atoms** and **stratixiv_hssi_atoms** into the **stratixiv** and **stratixiv_hssi** libraries, respectively.

The **stratixiv_hssi_atoms** model file references the **lpm** and **sgate** libraries. You must create these libraries to perform a simulation.

Compiling Library Files for Gate-Level Timing Simulation in Verilog HDL

To compile the libraries necessary for timing simulation of a Verilog HDL design targeting a Stratix IV device, type the following command at the VCS command prompt:

```
vcs -R <testbench>.v <gate-level netlist>.vo -v stratixiv_atoms.v \
-v stratixiv_hssi_atoms.v -v sgate.v -v 220model.v -v altera_mf.v \
+transport_int_delays +pulse_int_e/0 +pulse_int_r/0 \
+transport_path_delays +pulse_e/0 +pulse_r/0
```

Functional Simulation for Stratix V GX Devices

Functional simulation for Stratix V devices is similar to functional simulation for Arria II, Cyclone IV, HardCopy IV, and Stratix IV devices. To simulate transceivers in Arria II, Cyclone IV, HardCopy IV, and Stratix IV devices, you only have to replace the **stratixv_hssi** model file with the **arriaii_hssi**, **cycloneiv_hssi**, **hardcopyiv_hssi**, and **stratixiv_hssi** model files, respectively.

The **stratixv_hssi_atoms** model file references the **lpm** and **sgate** libraries. You must compile these libraries to perform a simulation.



The transceiver module from the MegaWizard Plug-In Manager is created in **Interfaces/Transceiver PHY**. Select **Custom PHY**.

Compiling Library Files for Functional Simulation

To compile the libraries necessary for functional simulation of a Verilog HDL or VHDL design targeting a Stratix V device, type the following commands at the VCS command prompt:

```
mkdir -p ./stratixv ←
mkdir -p ./stratixv_pcie_hip ←
mkdir -p ./stratixv_hssi ←
mkdir -p ./work ←

vlogan +v2k -work stratixv \
$QUARTUS_ROOTDIR/eda/sim_lib/synopsys/stratixv_atoms_ncrypt.v ←

vlogan +v2k -work stratixv_hssi \
$QUARTUS_ROOTDIR/eda/sim_lib/synopsys/stratixv_hssi_atoms_ncrypt.v ←

vlogan -sverilog -work stratixv_pcie_hip \
$QUARTUS_ROOTDIR/eda/sim_lib/synopsys/stratixv_pcie_hip_atoms_ncrypt.v ←

vhdlan -work stratixv_hssi \
$QUARTUS_ROOTDIR/eda/sim_lib/stratixv_hssi_components.vhd ←

vhdlan -work stratixv_hssi \
$QUARTUS_ROOTDIR/eda/sim_lib/stratixv_hssi_atoms.vhd ←

vcs test ←
./simv ←
```



The PCIe files are required only if you are using the PCIe HIP.



For VHDL, you must compile the Verilog HDL files first.

 In addition to the top-level variant wrapper, *<variant>.v*, you also get a simulation files subdirectory, *<variant>_sim/*. All Verilog (.v) and SystemVerilog (.sv) files in the simulation directory must also be compiled into the simulation project.

Transport Delays

By default, the VCS software filters out all pulses that are shorter than the propagation delay between primitives. Turning on the transport delay options in the VCS software prevents the simulation tool from filtering out these pulses. Use the following options to ensure that all signal pulses are seen in the simulation results.

+transport_path_delays

Use this option when the pulses in your simulation are shorter than the delay within a gate-level primitive. You must include the `+pulse_e/number` and `+pulse_r/number` options.

+transport_int_delays

Use this option when the pulses in your simulation are shorter than the interconnect delay between gate-level primitives. You must include the `+pulse_int_e/number` and `+pulse_int_r/number` options.

 The `+transport_path_delays` and `+transport_int_delays` options are also used by default in the NativeLink feature for gate-level timing simulation.

 For more information about either of these options, refer to the *VCS User Guide* installed with the VCS software.

The following VCS software command shows the command-line syntax to perform a post-synthesis simulation with the device family library:

```
vcs -R <testbench>.v <gate-level netlist>.v -v <Altera device family
library>.v \ +transport_int_delays +pulse_int_e/0 +pulse_int_r/0
+transport_path_delays \
+pulse_e/0 +pulse_r/0 ←
```

Using NativeLink with the VCS or VCS MX Software

The NativeLink feature in the Quartus II software facilitates the seamless transfer of information between the Quartus II software and EDA tools and allows you to run VCS or VCS MX within the Quartus II software.

 For more information, refer to the “Using the NativeLink Feature” section in the *Simulating Altera Designs* chapter in volume 3 of the *Quartus II Handbook*.

Generating a Timing .vcd File for the PowerPlay Power Analyzer

To generate a timing Verilog Value Change Dump File (*.vcd) for PowerPlay, you must first generate a VCD script in the Quartus II software, and then run the VCD script from the VCS software. This timing .vcd file can then be used by PowerPlay for power analysis.

To generate timing VCD scripts in the Quartus II software, perform the following steps:

1. In the Quartus II software, on the Assignments menu, click **Settings**. The **Settings** dialog box appears.
2. In the **Category** list, under **EDA Tool Settings**, click **Simulation**. On the **Simulation** page, in the **Tool name** list, select **VCS** and turn on the **Generate Value Change Dump (VCD) file script** option.
3. To generate the VCD script file, perform a full compilation.

Perform the following steps to generate a timing .vcd file in the VCS software:

1. Before compiling and simulating your design, include the script in your testbench file where the design under test (DUT) is instantiated:

```
include <revision_name>_dump_all_vcd_nodes.v ←
```



Include the script within the testbench module block. If you include the script outside of the testbench module block, syntax errors occur during compilation.

2. Run the simulation using the VCS command as usual. Exit the VCS software when the simulation is finished and the <revision_name>.vcd file is generated in the simulation directory.



The .vcd file is not supported in the VCS MX software.



For more detailed information about using the timing .vcd file for power analysis, refer to the *PowerPlay Power Analysis* chapter in volume 3 of the *Quartus II Handbook*.

Viewing a Waveform from a .vpd or .vcd File

A Virtual Panoramic Display (.vpd) file is automatically generated when your simulation is finished. The .vpd file is not readable. It is used for generating the waveform view through DVE. You can view your waveform result in DVE if you have created a .vpd or .vcd file.

To view a waveform from a .vpd file through DVE, perform the following steps:

1. Type `dve` on a command line. The **DVE** dialog box appears.
2. On the File menu, click **Open Database**. The **Open Database** dialog box appears.
3. Browse to the directory that contains your .vpd file (for example, **inter.vpd**).
4. Double-click the .vpd file.
5. In the **DVE** dialog box, select the signals that you want to observe from the Hierarchy.
6. On the Signal menu, click **Add To Waves**.
7. Click **New Wave View**. The waveform appears.

You cannot view a waveform from a `.vcd` file in DVE directly. The `.vcd` file must first be converted to a `.vpd` file. To convert the file, perform the following steps:

1. Use the `vcd2vpd` command to convert the file. For example, type the following on a command line:

```
vcd2vpd <example>.vcd <example>.vpd ↵
```

2. After you convert the `.vcd` file to a `.vpd` file, follow the procedures for viewing a waveform from a `.vpd` file through DVE.

You can also convert your `.vpd` file to a `.vcd` file by using the `vpd2vcd` command.

Scripting Support

You can run procedures and create settings described in this chapter in a Tcl script. You can also run some procedures at a command prompt.



For more information about Tcl scripting, refer to the *Tcl Scripting* chapter in volume 2 of the *Quartus II Handbook*. For more information about command-line scripting, refer to the *Command-Line Scripting* chapter in volume 2 of the *Quartus II Handbook*.

For detailed information about scripting command options, refer to the Qhelp utility.

To start the **Qhelp** utility, type the following command:

```
quartus_sh --qhelp ↵
```

Generating a Post-Synthesis Simulation Netlist for VCS

You can use the Quartus II software to generate a post-synthesis simulation netlist with Tcl commands or with a command at a command prompt.

Tcl Commands

Type the following Tcl commands to generate a post-synthesis simulation netlist when you compile your design or as part of a Tcl script that compiles your design:

```
set_global_assignment -name EDA_SIMULATION_TOOL "VCS" ↵  
set_global_assignment -name EDA_GENERATE_FUNCTIONAL_NETLIST ON ↵
```

Command Prompt

Type the following command to generate a simulation output file for the VCS software simulator; specify VHDL or Verilog HDL for the format:

```
quartus_eda <project name> --simulation=on --format=<format> \  
--tool=vcs --functional ↵
```

Generating a Gate-Level Timing Simulation Netlist for VCS

You can use the Quartus II software to generate a gate-level timing simulation netlist with Tcl commands or with a command at a command prompt.

Tcl Commands

Type the following Tcl command to generate a gate-level timing simulation netlist:

```
set_global_assignment -name EDA_SIMULATION_TOOL "VCS" ←
```

Command Prompt

Type the following command to generate a simulation output file for the VCS software simulator. Specify VHDL or Verilog HDL for the format.

```
quartus_eda <project name> --simulation=on --format=<format> --tool=vcs ←
```

Conclusion

You can use the Synopsys VCS or VCS MX software in your Altera FPGA design flow to easily and accurately perform functional simulations, post-synthesis simulations, and gate-level functional timing simulations. The seamless integration of the Quartus II software and VCS or VCS MX software make this simulation flow an ideal method for fully verifying an FPGA design.

Document Revision History

Table 3-2 shows the revision history for this chapter.

Table 3-2. Document Revision History (Part 1 of 2)

Date	Version	Changes
December 2010	10.0.1	Changed to new document template. No change to content.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Linked to Quartus II Help where appropriate ■ Added Stratix V simulation information ■ Minor text edits ■ Removed VirSim references ■ Removed Referenced Documents section
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Removed NativeLink information and referenced new <i>Simulating Designs with EDA Tools</i> chapter in volume 3 of the <i>Quartus II Handbook</i> ■ Added “RTL Functional Simulation for Stratix IV Devices” and “Gate-Level Timing Simulation for Stratix IV Devices” sections ■ Minor text edits

Table 3-2. Document Revision History (Part 2 of 2)

Date	Version	Changes
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Added support for Synopsys VCS MX software ■ Changed chapter title to “Synopsys VCS and VCS MX Support” ■ Major revision to “Compiling Libraries Using the EDA Simulation Library Compiler” on page 4-2 ■ Major revision to “RTL Functional Simulations” on page 4-2 ■ Added Table 3-4 on page 3-10 and Table 3-5 on page 3-11 ■ Added new section “Using DVE” on page 4-7 ■ Added new section “Generating a Simulation Script from the EDA Netlist Writer” on page 3-16 ■ Added new section “Viewing a Waveform from a .vpd or .vcd File” on page 4-13
November 2008	8.1.0	<ul style="list-style-type: none"> ■ Added “Compile Libraries Using the EDA Simulation Library Compiler” on page 3-3 ■ Added information about the <code>--simlib_comp</code> utility ■ Updated entire chapter using 8½” × 11” chapter template ■ Minor editorial updates

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter describes the basic NC-Sim, NC-Verilog, and NC-VHDL functional, post-synthesis, and gate-level timing simulations.

The Cadence Incisive verification platform includes NC-Sim, NC-Verilog, NC-VHDL, Verilog HDL, and VHDL desktop simulators.

This chapter is a “getting started” guide to using the Cadence Incisive verification platform in Altera® FPGA design flows. This chapter also describes the location of the simulation libraries and how to automate simulations.

This chapter includes the following topics:

- “Software Requirements”
- “Simulation Flow Overview”
- “Functional Simulation” on page 4–3
- “Post-Synthesis Simulation” on page 4–7
- “Gate-Level Timing Simulation” on page 4–8
- “Simulating Designs that Include Transceivers” on page 4–10
- “Using the NativeLink Feature with NC-Sim” on page 4–18
- “Generating a Timing VCD File for the PowerPlay Power Analyzer” on page 4–18
- “Viewing a Waveform from a .trn File” on page 4–19
- “Scripting Support” on page 4–20

Software Requirements

To simulate your design using the NC-Sim software, you must first set up the Altera libraries. These libraries are installed with the Quartus II software.

- 
 For more information about installing the software and directories created during the Quartus II software installation, refer to the *Altera Software Installation and Licensing* manual.

Simulation Flow Overview

The Incisive verification platform supports the following simulation flows:

- Functional Simulation
- Post-Synthesis Simulation
- Gate-Level Timing Simulation
- Using the NativeLink Feature with NC-Sim

Functional simulation verifies the functionality of your design. When you perform a functional simulation with Cadence Incisive simulators, you use your design files (Verilog HDL or VHDL) and the models provided with the Quartus II software. These Quartus II models are required if your design uses the library of parameterized modules (LPM) functions or Altera-specific megafunctions. Refer to “[Functional Simulation](#)” on page 4-3 for more information about how to perform this simulation.

A post-synthesis simulation verifies the functionality of a design after synthesis has been performed. You can create a post-synthesis netlist (.vo or .vho) in the Quartus II software and use this netlist to perform a post-synthesis simulation with the Incisive simulator. Refer to “[Post-Synthesis Simulation](#)” on page 4-7 for more information about how to perform this simulation.

After performing place-and-route, the Quartus II software generates a Verilog HDL Output File (.vo) or VHDL Output File (.vho) and a Standard Delay Output file (.sdo) for gate-level timing simulation. The netlist files map your design to architecture-specific primitives. The .sdo file contains the delay information of each architecture primitive and routing element specific to your design. Together, these files provide an accurate simulation of your design with the selected Altera FPGA architecture. Refer to “[Gate-Level Timing Simulation](#)” on page 4-8 for more information about how to perform this simulation.

Operation Modes

You can use either the GUI mode or the command-line mode to simulate your design in the NC simulators.

To start the Incisive simulators in GUI mode, type the following command at a command prompt:

```
nclaunch ←
```

To simulate in command-line mode, use the programs shown in [Table 4-1](#).

Table 4-1. Command-Line Programs

Program	Function
ncvlog or ncvhdl	NC-Verilog (ncvlog) compiles your Verilog HDL code into a Verilog Syntax Tree (.vst) file. ncvlog also performs syntax and static semantics checks. NC-VHDL (ncvhdl) compiles your VHDL code into a VHDL Syntax Tree (.ast) file. ncvhdl also performs syntax and static semantics checks.
ncelab	NC-Elab (ncelab) elaborates the design. ncelab constructs the design hierarchy and establishes signal connectivity. This program also generates a Signature File (.sig) and a Simulation SnapShot File (.sss).
ncsim	NC-Sim (ncsim) performs mixed-language simulation. This program is the simulation kernel that performs event scheduling and executes the simulation code.

Quartus II Software and NC Simulation Flow Overview

This section provides an overview of the Quartus II software and Cadence NC simulation flow. More detailed information is provided in “[Functional Simulation](#)” on page 4-3, “[Post-Synthesis Simulation](#)” on page 4-7, and “[Gate-Level Timing Simulation](#)” on page 4-8.

For high-speed simulation, you must select **ps** in the **Resolution** list for your simulator resolutions (**Design** tab of the **Start Simulation** dialog box). If you choose slower than **ps**, the high-speed simulation might fail.

Complete the following tasks:

1. Create user libraries.

Create a file that maps logical library names to their physical locations. These library mappings include your working directory and any design-specific libraries; for example, Altera LPM functions or megafunctions.

2. Compile source code and testbenches.

Compile your design files at the command-line using **ncvlog** (Verilog HDL files) or **ncvhdl** (VHDL files) or, on the Tools menu, click **Verilog Compiler** or **VHDL Compiler** in NCLaunch. During compilation, the NC software performs syntax and static semantic checks. If no errors are found, compilation produces an internal representation for each HDL design unit in the source files. By default, these intermediate objects are stored in a single, packed, library database file in your working directory.

3. Elaborate your design.

Before you can simulate your model, you must define the design hierarchy in a process called “elaboration”. Use **ncelab** in command-line mode or, on the Tools menu in NCLaunch, click **Elaborator**.

4. Add signals to your waveform.

Specify which signals to view in your waveform using a simulation history manager (SHM) database.

5. Simulate your design.

Run the simulator with the **ncsim** program (command-line mode) or by clicking **Run** in the SimVision Console window.

Compiling Libraries Using the EDA Simulation Library Compiler

The EDA Simulation Library Compiler compiles Verilog HDL and VHDL simulation libraries for all Altera devices and supported third-party simulators. You can compile all libraries required by functional and gate-level simulation with this tool.



For more information about this tool, refer to the “EDA Simulation Library Compiler” section in the *Simulating Altera Designs* chapter in volume 3 of the *Quartus II Handbook*.

Functional Simulation

The following sections provide detailed instructions for performing functional simulation using the Quartus II software and the Cadence Incisive verification platform tools.



For the Altera Behavioral Simulation Models, refer to *Altera Functional Simulation Libraries* in Quartus II Help.

Creating Libraries

To create libraries, perform the following steps:

1. Create a directory for the work library and any other libraries you require by typing the following command at a command prompt:

```
mkdir <library name> ↵
```

Examples:

```
mkdir worklib ↵
mkdir altera_mf ↵
```

2. Using a text editor, create a **cds.lib** file and add the following line to it:

```
DEFINE <library name> <physical directory path>
```

Examples:

```
DEFINE worklib ./worklib
DEFINE altera_mf ./altera_mf
```

-  For information about creating a **cds.lib** file in GUI mode, refer to *Performing a Functional Simulation with the NCSim Software* in Quartus II Help.

For VHDL Designs

If you are targeting a Stratix V device, compile the following files in the **quartus/eda/sim_lib/cadence** directory:

```
stratixv_atoms_ncrypt.v
stratixv_hssi_atoms_ncrypt.v
stratixv_pcie_hip_atoms_ncrypt.v
```

These files contain IEEE encrypted Verilog models suitable for VHDL/Verilog co-simulation. You need a co-simulation license from Cadence to use these models in VHDL.

-  Compile these encrypted Verilog files before you compile any VHDL files. Compile **stratixv_pcie_hip_atoms_ncrypt.v** with the SystemVerilog option.

Also, compile the following files in the **quartus/eda/sim_lib** directory:

```
stratixv_atoms.vhd
stratixv_components.vhd
stratixv_hssi_components.vhd
stratixv_pcie_components.vhd
stratixv_hssi_atoms.vhd
stratixv_pcie_hip_atoms.vhd
```

-  The PCIe® files are required only if you are using the PCIe HIP.

For Verilog HDL Designs

If you are targeting a Stratix V device, compile the following files in the **<quartus/eda/sim_lib/cadence** directory:

```
stratixv_atoms_ncrypt.v
stratixv_hssi_atoms_ncrypt.v
stratixv_pcie_hip_atoms_ncrypt.v
```

These files contain IEEE encrypted Verilog models.



Compile `stratixv_pcie_hip_atoms.v` with the SystemVerilog option.

Also, compile the following files in the `quartus/eda/sim_lib` directory:

```
stratixv_atoms.v
stratixv_hssi_atoms.v
stratixv_pcie_hip_atoms.v
```



The PCIe files are required only if you are using the PCIe HIP.

Compiling Source Code

To compile from your source code from the command line, type one of the following commands:

- Verilog HDL:

```
ncvlog <options> -work <library name> <design files> ←
```

- VHDL:

```
ncvhdl <options> -work <library name> <design files> ←
```

You must create a work library before compiling your design and testbench. If your design uses LPM, Altera megafunctions, or Altera primitives, you must also compile the Altera-provided functional models. The commands in [Example 4-1](#) and [Example 4-2](#) show an example of each.

Example 4-1. Compile in Verilog HDL

```
ncvlog -WORK lpm 220model.v
ncvlog -WORK altera_mf altera_mf.v
ncvhdl -WORK altera altera_primitives.v
ncvlog -WORK altera altera_primitives.v
ncvlog -WORK work toplevel.v testbench.v
```

Example 4-2. Compile in VHDL

```
ncvhdl -V93 -WORK lpm 220pack.vhd
ncvhdl -V93 -WORK lpm 220model.vhd
ncvhdl -V93 -WORK altera_mf altera_mf_components.vhd
ncvhdl -V93 -WORK altera_mf altera_mf.vhd
ncvhdl -V93 -WORK altera altera_primitives_components.vhd
ncvhdl -V93 -WORK altera altera_primitives.vhd
ncvhdl -V93 -WORK work toplevel.vhd testbench.vhd
```



For information about compiling in GUI mode, refer to [Performing a Functional Simulation with the NCSim Software](#) in Quartus II Help.

Elaborating Your Design

Before you can simulate your design, you must define the design hierarchy in a process called elaboration. The Incisive simulator elaborates your design with the language-independent `ncelab` program. The `ncelab` program constructs a design hierarchy based on the design's instantiation and configuration information, establishes signal connectivity, and computes initial values for all objects in the design. The elaborated design hierarchy is stored in a simulation snapshot, which is the representation of your design that the simulator uses to run the simulation. The snapshot is stored in the library database file, along with the other intermediate objects generated by the compiler and elaborator.

To elaborate your Verilog HDL or VHDL design from the command line, type the following command:

```
ncelab [options][<library>.<testbench module name>]
```

Example:

```
ncelab -access +rwc work.testbench_module
```

Adding the option `-access +rwc` allows signals to be viewed in the Waveform window.

If your design includes high-speed signals, you might have to add the following pulse reject options with the `ncelab` command:

```
ncelab -access +rwc work.testbench_module -PULSE_R 0 -PULSE_INT_R 0
```



For more information about the pulse reject options, refer to the *SDF Annotate Guide* from [Cadence](#).



For information about elaborating your design in GUI mode, refer to *Performing a Functional Simulation with the NCSim Software* in Quartus II Help.

Simulating Your Design

After you have compiled and elaborated your design, you can simulate it using `ncsim`. The `ncsim` program loads the file, or snapshot, generated by `ncelab` as its primary input and then loads other intermediate objects referenced by the snapshot. If you enable interactive debugging, `ncsim` can also load HDL source files and script files. The simulation output is controlled by the model or debugger. The output can include result files generated by the model, the SHM database, or the `.vcd` file.

To perform functional simulation of your Verilog HDL or VHDL design at the command line, type the following command:

```
ncsim [options][<library>.<testbench module name>]
```

Example:

```
ncsim -gui work.testbench_module
```

Adding the option `-gui` opens the SimVision window for running your simulation.



For information about performing a functional simulation in GUI mode, refer to *Performing a Functional Simulation with the NCSim Software* in Quartus II Help.

Post-Synthesis Simulation

The following sections provide detailed instructions for performing post-synthesis simulation with the Incisive platform software and output files and simulation files from the Quartus II software.



You cannot perform post-synthesis or post-fit simulation if you are targeting the Stratix V device family.



For a list of the gate-level simulation models, refer to *Altera Post-Fit Libraries* in Quartus II Help.

Quartus II Simulation Output Files

After performing synthesis with either a third-party synthesis tool or with Quartus II integrated synthesis, you must generate a simulation netlist for functional simulations.



To generate a simulation netlist for functional simulation, refer to the “Generating Post-Synthesis Netlist Files” section in the *Simulating Altera Designs* chapter in volume 3 of the *Quartus II Handbook*.

Creating Libraries

Create the following libraries for your simulation:

- Work library
- Device family library targeting your design targets using the following files in the `<path to Quartus II installation>/eda/sim_lib` directory:
 - `<device_family>_atoms.v`
 - `<device_family>_atoms.vhd`
 - `<device_family>_components.vhd`

Refer to “[Creating Libraries](#)” on page 4-4 for instructions about creating libraries.

Compiling Project Files and Libraries

Compile the project files and libraries into your work directory using the `ncvlog`, `ncvhdl` programs, or the GUI. Include the following files:

- Testbench file
- The Quartus II software functional output netlist file (`.vo` file or `.vho` file)
- Atom library file for the device family `<device family>_atoms.<v | vhd>`
- For VHDL, `<device family>_components.vhd`

Refer to “[Compiling Source Code](#)” on page 4-5 for instructions about compiling.

Elaborating Your Design

Elaborate your design using the `ncelab` program as described in “[Elaborating Your Design](#)” on page 4-6.

Simulating Your Design

Simulate your design using the `ncsim` program as described in “[Simulating Your Design](#)” on page 4-6.

Gate-Level Timing Simulation

The following sections provide detailed instructions for performing a timing simulation using the Quartus II output files, simulation libraries, and Cadence NC tools.

-  You cannot perform post-synthesis or post-fit simulations if you are targeting the Stratix V device family.
-  For a list of the gate-level simulation models, refer to [Altera Post-Fit Libraries](#) in Quartus II Help.
-  For details about how to perform gate-level timing simulation using the Quartus II software and the Cadence NC-Sim software, refer to [Performing a Timing Simulation with the NCSim Software](#) in Quartus II Help.

Generating a Gate-Level Timing Simulation Netlist

To perform a gate-level timing simulation, your design should provide the NC-Sim software with information about how the design was placed into device-specific architectural blocks. The Quartus II software provides this information in the form of a `.vo` file for Verilog HDL designs and a `.vho` file for VHDL designs. The accompanying timing information is stored in the `.sdo` file, which annotates the delay for the elements found in the `.vo` file or `.vho` file.

-  To generate the `.vo` or `.vho` file and the `.sdo` file, refer to the “[Generating Gate-Level Timing Simulation Netlist Files](#)” section in the [Simulating Altera Designs](#) chapter in volume 3 of the *Quartus II Handbook*.

Disabling Timing Violation on Registers

In certain situations, the timing violations can be ignored and you can disable the timing violation on registers. For example, timing violations that occur in internal synchronization registers used for asynchronous clock-domain crossing can be ignored and disabled.

By default, the `x_on_violation_option logic` option is **On**, which means the simulation shows “X” whenever a timing violation occurs. To disable showing the timing violation on certain registers, you can set the `x_on_violation_option logic` option to **Off** for those registers. The following command is the Quartus II Tcl command to disable timing violation on registers.

```
set_instance_assignment -name X_ON_VIOLATION_OPTION OFF -to \ <register_name>
```

This Tcl command is also stored in the `.qsf` file.

Creating Libraries

Create the following libraries for your simulation:

- Work library
- Device family libraries targeting using the following files in the `<path to Quartus II installation>/eda/sim_lib` directory:
 - `<device_family>_atoms.v`
 - `<device_family>_atoms.vhd`
 - `<device_family>_components.vhd`

For instructions about creating libraries, refer to [“Creating Libraries” on page 4-4](#).

Compiling Project Files and Libraries

Compile the project files and libraries into your work directory with the `ncvlog` or `ncvhdl` programs or the GUI. Include the following files:

- Testbench file
- The Quartus II software functional output netlist file (`.vo` file or `.vho` file)
- Atom library file for the device family `<device family>_atoms.<v | vhd>`
- For VHDL, `<device family>_components.vhd`

For instructions about compiling, refer to [“Compiling Source Code” on page 4-5](#).

Elaborating Your Design

When performing elaboration with the Quartus II-generated Verilog HDL netlist file, the `.sdo` file is read automatically. The `ncelab` executable recognizes the embedded system task `$$sdf_annotate` and automatically compiles and annotates the `.sdo` file (runs `ncsdfc` automatically).



The `.sdo` file should be located in the same directory where you perform an elaboration or simulation, because the `$$sdf_annotate` task references the `.sdo` file without using a full path. If you are starting an elaboration or simulation from a different directory, you can either comment out the `$$sdf_annotate` and annotate the `.sdo` file with the GUI, or add the full path of the `.sdo` file.

Refer to [“Elaborating Your Design” on page 4-6](#) for elaboration instructions.

VHDL netlist files do not contain system task calls to locate your `.sdf` file; therefore, you must compile the Standard `.sdo` file manually. For information about compiling the `.sdo` file, refer to [“Compiling the .sdo File \(VHDL Only\) in Command-Line Mode”](#) and [“Compiling the .sdo File \(VHDL Only\) in GUI Mode”](#).

Compiling the .sdo File (VHDL Only) in Command-Line Mode

To annotate the .sdo file timing data from the command line, perform the following steps:

1. Compile the .sdo file using the `ncsdfc` program by typing the following command at the command prompt:

```
ncsdfc <project name>_vhd.sdo -output <output name> ←
```

The `ncsdfc` program generates an `<output name>.sdf.X` compiled .sdo file.

 If you do not specify an output name, `ncsdfc` uses `<project name>.sdo.X`.

2. Specify the compiled .sdf file for the project by adding the following command to an ASCII SDF command file for the project:

```
COMPILED_SDF_FILE = "<project name>.sdf.X" SCOPE = <instance path>
```

[Example 4-3](#) shows an example of an SDF command file.

Example 4-3. SDF Command File

```
// SDF command file sdf_file
COMPILED_SDF_FILE = "lpm_ram_dp_test_vhd.sdo.X",
SCOPE = :tb,
MTM_CONTROL = "TYPICAL",
SCALE_FACTORS = "1.0:1.0:1.0",
SCALE_TYPE = "FROM_MTM";
```

After you compile the .sdf file, type the following command to elaborate the design:

```
ncelab worklib.<project name>:entity -SDF_CMD_FILE <SDF Command File> ←
```

Compiling the .sdo File (VHDL Only) in GUI Mode

-  To compile the .sdo file in GUI mode, refer to [Performing a Timing Simulation with the NCSim Software](#) in Quartus II Help.

Simulating Your Design

Simulate your design using the `ncsim` program as described in [“Simulating Your Design”](#) on page 4-6.

-  For the design examples to run gate-level timing simulation, refer to the [Cadence NC-Sim Simulation Design Example](#) web page.

Simulating Designs that Include Transceivers

If your design includes Arria®, Arria II, Cyclone IV, HardCopy IV, Stratix, Stratix II, or Stratix IV transceivers, you must compile additional library files to perform functional or gate-level timing simulations.

For high-speed simulation, you must select **ps** in the **Resolution** list for your simulator resolutions (**Design** tab of the **Start Simulation** dialog box). If you choose slower than **ps**, the high-speed simulation might fail.

-  If your design contains PCI Express® hard IP, refer to the “Simulate the Design” section in the *PCI Express Compiler User Guide*.

Functional Simulation for Stratix GX Devices

To perform a functional simulation of your design that instantiates the ALTGXB megafunction, enabling the gigabit transceiver block (GXB) on Stratix GX devices, compile the `stratixgx_mf` model file into the `altgxb` library.

-  The `stratixgx_mf` model file references the `lpm` and `sgate` libraries. You must create these libraries to perform a simulation.

Compiling Library Files for Functional Simulation in VHDL

To compile the libraries necessary for functional simulation of a VHDL design targeting a Stratix GX device, type the commands shown in [Example 4-4](#) at the NC-Sim command prompt.

Example 4-4. Compile Libraries Commands for Functional Simulation in VHDL

```
ncvhd1 -work lpm 220pack.vhd 220model.vhd
ncvhd1 -work altera_mf altera_mf_components.vhd altera_mf.vhd
ncvhd1 -work sgate sgate_pack.vhd sgate.vhd
ncvhd1 -work altgxb stratixgx_mf.vhd stratixgx_mf_components.vhd
ncsim work.<my design>
```

Compiling Library Files for Functional Simulation in Verilog HDL

To compile the libraries necessary for a functional simulation of a Verilog HDL design targeting a Stratix GX device, type the commands shown in [Example 4-5](#) at the NC-Sim command prompt.

Example 4-5. Compile Libraries Commands for Functional Simulation in Verilog HDL

```
ncvlog -work lpm 220model.v
ncvlog -work altera_mf altera_mf.v
ncvlog -work sgate sgate.v
ncvlog -work altgxb stratixgx_mf.v
ncsim work.<my design>
```

Gate-Level Timing Simulation for Stratix GX Devices

To perform a gate-level timing simulation of your design that includes a Stratix GX transceiver, compile the `stratixgx_atoms` and `stratixgx_hssi_atoms` model files into the `stratixgx` and `stratixgx_gxb` libraries, respectively.

-  You must create these libraries to perform a simulation because the `stratixgx_hssi_atoms` model file references the `lpm` and `sgate` libraries.

Compiling Library Files for Gate-Level Timing Simulation in VHDL

To compile the libraries necessary for a timing simulation of a VHDL design targeting a Stratix GX device, type the commands shown in [Example 4-6](#) at the NC-Sim command prompt.

Example 4-6. Compile Libraries Commands for Timing Simulation in VHDL

```
ncvhd1 -work lpm 220pack.vhd 220model.vhd
ncvhd1 -work altera_mf altera_mf_components.vhd altera_mf.vhd
ncvhd1 -work sgate sgate_pack.vhd sgate.vhd
ncvhd1 -work stratixgx stratixgx_atoms.vhd stratixgx_components.vhd
ncvhd1 -work stratixgx_gxb stratixgx_hssi_atoms.vhd \
stratixgx_hssi_components.vhd
ncelab work.<my design> -TIMESCALE 1ps/1ps -PULSE_R 0 -PULSE_INT_R 0
```

Compiling Library Files for Gate-Level Timing Simulation in Verilog HDL

To compile the libraries necessary for a timing simulation of a Verilog HDL design targeting a Stratix GX device, type the commands shown in [Example 4-7](#) at the NC-Sim command prompt.

Example 4-7. Compile Libraries Commands for Timing Simulation in Verilog HDL

```
ncvlog -work lpm 220model.v
ncvlog -work altera_mf altera_mf.v
ncvlog -work sgate sgate.v
ncvlog -work stratixgx stratixgx_atoms.v
ncvlog -work stratixgx_gxb stratixgx_hssi_atoms.v
ncelab work.<my design> -TIMESCALE 1ps/1ps -PULSE_R 0 -PULSE_INT_R 0
```

Functional Simulation for Stratix II GX Devices

Functional simulation of Stratix II GX devices is similar to functional simulation of Arria GX devices. [Example 4-9 on page 4-13](#) and “[Compile Libraries Commands for Functional Simulation in Verilog HDL](#)” on page 4-13 show only the functional simulation for designs that include transceivers in Stratix II GX devices. To simulate transceivers in Arria GX devices, replace the `stratixiigx_hssi` model file with the `arriagx_hssi` model file.

To perform a functional simulation of your design that instantiates the ALT2GXB megafunction, edit your `cds.lib` file so all of the libraries point to the work library, and compile the `stratixiigx_hssi` model file into the `stratixiigx_hssi` library. When compiling the library files, you can safely ignore the following warning message:

```
"Multiple logical libraries mapped to a single location"
```

[Example 4-8](#) shows the `cds.lib` file.

Example 4-8. cds.lib File

```
SOFTINCLUDE ${CDS_INST_DIR}/tools/inca/files/cdsvhdl.lib
SOFTINCLUDE ${CDS_INST_DIR}/tools/inca/files/cdsvlog.lib
DEFINE work ./ncsim_work
DEFINE stratixiigx_hssi ./ncsim_work
DEFINE stratixiigx ./ncsim_work
DEFINE lpm ./ncsim_work
DEFINE sgate ./ncsim_work
```

 The `stratixiigx_hssi_atoms` model file references the `lpm` and `sgate` libraries. You must create these libraries to perform a simulation.

Generate a functional simulation netlist by turning on **Create a simulation library for this design** in the last page of the ALT2GXB MegaWizard Plug-In Manager. The `<alt2gxb entity name>.vho` or `<alt2gxb module name>.vo` is generated in the current project directory.

 The ALT2GXB functional simulation library file generated by the Quartus II software references `stratixiigx_hssi` WYSIWYG atoms.

Compiling Library Files for Functional Simulation in VHDL

To compile the libraries necessary for functional simulation of a VHDL design targeting a Stratix II GX device, type the commands shown in [Example 4-9](#) at the NC-Sim command prompt.

Example 4-9. Compile Libraries Commands for Functional Simulation in VHDL

```
ncvhd1 -work lpm 220pack.vhd 220model.vhd
ncvhd1 -work altera_mf altera_mf_components.vhd altera_mf.vhd
ncvhd1 -work sgate sgate_pack.vhd sgate.vhd
ncvhd1 -work stratixiigx_hssi stratixiigx_hssi_components.vhd \
stratixiigx_hssi_atoms.vhd
ncvhd1 -work work <alt2gxb entity name>.vho
ncelab work.<my design>
```

Compiling Library Files for Functional Simulation in Verilog HDL

To compile the libraries necessary for functional simulation of a Verilog HDL design targeting a Stratix II GX device, type the commands shown in [Example 4-10](#) at the NC-Sim command prompt.

Example 4-10. Compile Libraries Commands for Functional Simulation in Verilog HDL

```
ncvlog -work lpm 220model.v
ncvlog -work altera_mf altera_mf.v
ncvlog -work sgate sgate.v
ncvlog -work stratixiigx_hssi stratixiigx_hssi_atoms.v
ncvlog -work work <alt2gxb module name>.vo
ncelab work.<my design>
```

Gate-Level Timing Simulation for Stratix II GX Devices

Stratix II GX functional simulation is similar to Arria GX functional simulation. [Example 4-11 on page 4-14](#) and [Example 4-12 on page 4-14](#) show only the gate-level timing simulation for designs that include transceivers in Stratix II GX. To simulate transceivers in Arria GX, replace the `stratixiigx_hssi` model file with the `arriagx_hssi` model file.

To perform a post-fit timing simulation of your design that includes the ALT2GXB megafunction, edit your `cds.lib` file so that all the libraries point to the work library and compile `stratixiigx_atoms` and `stratixiigx_hssi_atoms` into the `stratixiigx` and `stratixiigx_hssi` libraries, respectively. When compiling the library files, you can safely ignore the following warning message:

```
"Multiple logical libraries mapped to a single location"
```

For an example of a `cds.lib` file, refer to “[Functional Simulation for Stratix II GX Devices](#)” on page 4-12.



The `stratixiigx_hssi_atoms` model file references the `lpm` and `sgate` libraries. You must create these libraries to perform a simulation.

Compiling Library Files for Gate-Level Timing Simulation in VHDL

To compile the libraries necessary for timing simulation of a VHDL design targeting a Stratix II GX device, type the commands shown in [Example 4-11](#) at the NC-Sim command prompt.

Example 4-11. Compile Libraries Commands for Timing Simulation in VHDL

```
ncvhd1 -work lpm 220pack.vhd 220model.vhd
ncvhd1 -work altera_mf altera_mf_components.vhd altera_mf.vhd
ncvhd1 -work sgate sgate_pack.vhd sgate.vhd
ncvhd1 -work stratixiigx stratixiigx_atoms.vhd \
stratixiigx_components.vhd
ncvhd1 -work stratixiigx_hssi stratixiigx_hssi_components.vhd \
stratixiigx_hssi_atoms.vhd
ncvhd1 -work work <alt2gxb>.vho
ncelab work.<my design> -TIMESCALE 1ps/1ps -PULSE_R 0 -PULSE_INT_R 0
```

Compiling Library Files for Gate-Level Timing Simulation in Verilog HDL

To compile the libraries necessary for timing simulation of a Verilog HDL design targeting a Stratix II GX device, type the commands shown in [Example 4-12](#) at the NC-Sim command prompt.

Example 4-12. Compile Libraries Commands for Timing Simulation in Verilog HDL

```
ncvlog -work lpm 220model.v
ncvlog -work altera_mf altera_mf.v
ncvlog -work sgate sgate.v
ncvlog -work stratixiigx stratixiigx_atoms.v
ncvlog -work stratixiigx_hssi stratixiigx_hssi_atoms.v
ncelab work.<my design> -TIMESCALE 1ps/1ps -PULSE_R 0 -PULSE_INT_R 0
```

Functional Simulation for Stratix IV GX Devices

Functional simulation for Stratix IV devices is similar to functional simulation for Arria II, Cyclone IV, and HardCopy IV devices. [Example 4-14](#) shows only the functional simulation for designs that include transceivers in Stratix IV devices. To simulate transceivers in Arria II, Cyclone IV, and HardCopy IV devices, replace the `stratixiv_hssi` model file with the `arriaaii_hssi`, `cycloneiv_hssi`, and `hardcopyiv_hssi` model files, respectively.

To perform a functional simulation of your design that instantiates the ALTGX megafunction, edit your `cds.lib` file so that all of the libraries point to the work library, and compile the `stratixiv_hssi` model file into the `stratixiv_hssi` library.

When compiling the library files, you can safely ignore the following warning message:

```
"Multiple logical libraries mapped to a single location"
```

Example 4-13 shows the `cds.lib` file.

Example 4-13. cds.lib File

```
SOFTINCLUDE ${CDS_INST_DIR}/tools/inca/files/cdsvhdl.lib
SOFTINCLUDE ${CDS_INST_DIR}/tools/inca/files/cdsvlog.lib
DEFINE work ./ncsim_work
DEFINE stratixiv_hssi ./ncsim_work
DEFINE stratixiv ./ncsim_work
DEFINE lpm ./ncsim_work
DEFINE sgate ./ncsim_work
```

The `stratixiv_hssi_atoms` model file references the `lpm` and `sgate` libraries. You must create these libraries to perform a simulation.

Compiling Library Files for Functional Simulation in VHDL

To compile the libraries necessary for functional simulation of a VHDL design targeting a Stratix IV device, type the commands shown in Example 4-14 at the NC-Sim command prompt.

Example 4-14. Compile Libraries Commands for Functional Simulation in VHDL

```
ncvhd1 -work lpm 220pack.vhd 220model.vhd
ncvhd1 -work altera_mf altera_mf_components.vhd altera_mf.vhd
ncvhd1 -work sgate sgate_pack.vhd sgate.vhd
ncvhd1 -work stratixiv_hssi stratixiv_hssi_components.vhd \
stratixiv_hssi_atoms.vhd
ncvhd1 -work work <altgx entity name>.vhd
ncelab work.<my design>
```

Compiling Library Files for Functional Simulation in Verilog HDL

To compile the libraries necessary for functional simulation of a Verilog HDL design targeting a Stratix IV device, type the commands shown in Example 4-15 at the NC-Sim command prompt.

Example 4-15. Compile Libraries Commands for Functional Simulation in Verilog HDL

```
ncvlog -work lpm 220model.v
ncvlog -work altera_mf altera_mf.v
ncvlog -work sgate sgate.v
ncvlog -work stratixiv_hssi stratixiv_hssi_atoms.v
ncvlog -work work <altgx module name>.v
ncelab work.<my design>
```

Gate-Level Timing Simulation for Stratix IV GX Devices

Stratix IV gate-level timing simulation is similar to Arria II gate-level timing simulation.

Example 4-16 and Example 4-17 show only the gate-level timing simulation for designs that include transceivers in Stratix IV devices. To simulate transceivers in Arria II, Cyclone IV, and HardCopy IV devices, replace the `stratixiv_hssi` model file with the `arriaii_hssi`, `cycloneiv_hssi`, and `hardcopyiv_hssi` model files, respectively.

To perform a post-fit timing simulation of your design that includes the ALTGX megafunction, edit your `cds.lib` file so that all of the libraries point to the work library and compile `stratixiv_atoms` and `stratixiv_hssi_atoms` into the `stratixiv` and `stratixiv_hssi` libraries, respectively. When compiling the library files, you can safely ignore the following warning message:

```
"Multiple logical libraries mapped to a single location"
```

For an example of a `cds.lib` file, refer to [Example 4-13 on page 4-15](#).

The `stratixiv_hssi_atoms` model file references the `lpm` and `sgate` libraries. You must create these libraries to perform a simulation.

Compiling Library Files for Gate-Level Timing Simulation in VHDL

To compile the libraries necessary for timing simulation of a VHDL design targeting a Stratix IV device, type the commands shown in [Example 4-16](#) at the NC-Sim command prompt.

Example 4-16. Compile Libraries Commands for Timing Simulation in VHDL

```
ncvhd1 -work lpm 220pack.vhd 220model.vhd
ncvhd1 -work altera_mf altera_mf_components.vhd altera_mf.vhd
ncvhd1 -work sgate sgate_pack.vhd sgate.vhd
ncvhd1 -work stratixiv stratixiv_atoms.vhd \
stratixiv_components.vhd
ncvhd1 -work stratixiv_hssi stratixiv_hssi_components.vhd \
stratixiv_hssi_atoms.vhd
ncvhd1 -work work <altgx>.vho
ncsdfc <project name>_vhd.sdo
ncelab work.<my design> -TIMESCALE 1ps/1ps \
-SDF_CMD_FILE <SDF Command File> -PULSE_R 0 -PULSE_INT_R 0
```

Compiling Library Files for Gate-Level Timing Simulation in Verilog HDL

To compile the libraries necessary for timing simulation of a Verilog HDL design targeting a Stratix IV device, type the commands shown in [Example 4-17](#) at the NC-Sim command prompt.

Example 4-17. Compile Libraries Commands for Timing Simulation in Verilog HDL

```
ncvlog -work lpm 220model.v
ncvlog -work altera_mf altera_mf.v
ncvlog -work sgate sgate.v
ncvlog -work stratixiv stratixiv_atoms.v
ncvlog -work stratixiv_hssi stratixiv_hssi_atoms.v
ncvlog -work work <altgx>.vo
ncelab work.<my design> -TIMESCALE 1ps/1ps -PULSE_R 0 -PULSE_INT_R 0
```

Functional Simulation for Stratix V GX Devices

Functional simulation for Stratix V devices is similar to functional simulation for Arria II, Cyclone IV, HardCopy IV, and Stratix IV devices. To simulate transceivers in Arria II, Cyclone IV, HardCopy IV, and Stratix IV devices, you only have to replace the `stratixiv_hssi` model file with the `arriaii_hssi`, `cycloneiv_hssi`, `hardcopyiv_hssi`, and `stratiiv_hssi` model files, respectively.

The `stratixiv_hssi_atoms` model file references the `lpm` and `sgate` libraries. You must compile these libraries to perform a simulation.



The transceiver module from the MegaWizard Plug-In Manager is created in **Interfaces/Transceiver PHY**. Select **Custom PHY**.

Compiling Library Files for Functional Simulation

To compile the libraries necessary for functional simulation of a Verilog HDL or VHDL design targeting a Stratix V device, type the following command at the VCS command prompt:

[Example 4-18](#) shows the `cds.lib` file.

Example 4-18. `cds.lib` File

```
SOFTINCLUDE ${CDS_INST_DIR}/tools/inca/files/cdsvhdl.lib
SOFTINCLUDE ${CDS_INST_DIR}/tools/inca/files/cdsvlog.lib
DEFINE work ./ncsim_work
DEFINE stratixv_hssi ./ncsim_work
DEFINE stratixv ./ncsim_work
DEFINE lpm ./ncsim_work
DEFINE sgate ./ncsim_work
```

The `stratixv_hssi_atoms` model file references the `lpm` and `sgate` libraries. You must create these libraries to perform a simulation.

Compiling Library Files for Functional Simulation in VHDL

To compile the libraries necessary for functional simulation of a VHDL design targeting a Stratix V device, type the commands shown in [Example 4-19](#) at the NC-Sim command prompt.

Example 4-19. Compile Libraries Commands for Functional Simulation in VHDL

```
ncvhdl -work lpm 220pack.vhd 220model.vhd
ncvhdl -work altera_mf altera_mf_components.vhd altera_mf.vhd
ncvhdl -work sgate sgate_pack.vhd sgate.vhd
ncvhdl -work stratixv_hssi stratixv_hssi_components.vhd \
ncvlog +v2k -work stratixv_hssi \
quartus/eda/sim_lib/cadence/stratixv_hssi_atoms_ncrypt.v
stratixv_hssi_atoms.vhd
ncvhdl -work work <my_design>.vhd
ncelab work.<my design>
```

Compiling Library Files for Functional Simulation in Verilog HDL

To compile the libraries necessary for functional simulation of a Verilog HDL design targeting a Stratix V device, type the commands shown in [Example 4-20](#) at the NC-Sim command prompt.

Example 4-20. Compile Libraries Commands for Functional Simulation in Verilog HDL

```
ncvlog -work lpm 220model.v
ncvlog -work altera_mf altera_mf.v
ncvlog -work sgate sgate.v
ncvlog +v2k -work stratixv_hssi \
quartus/eda/sim_lib/cadence/stratixv_hssi_atoms_ncrypt.v
ncvlog -work stratixv_hssi stratixv_hssi_atoms.v
ncvlog -work work <my design>.v
ncelab work.<my design>
```

Pulse Reject Delays

By default, the NC-Sim software filters out all pulses that are shorter than the propagation delay between primitives. Setting the pulse reject delays options in the NC-Sim software prevents the simulation tool from filtering out these pulses. Use the following options to ensure that all signal pulses are seen in the simulation results.

-PULSE_R

Use this option when the pulses in your simulation are shorter than the delay within a gate-level primitive. The argument is the percentage of delay for pulse reject limit for the path.

-PULSE_INT_R

Use this option when the pulses in your simulation are shorter than the interconnect delay between gate-level primitives. The argument is the percentage of delay for pulse reject limit for the path.



The **-PULSE_R** and **-PULSE_INT_R** options are also used by default in the NativeLink feature for gate-level timing simulation.

The following NC-Sim software command shows the command-line syntax to perform a gate-level timing simulation with the device family library.

```
ncelab worklib.<project name>:entity -SDF_CMD_FILE <SDF Command File> \  
-TIMESCALE 1ps/1ps -PULSE_R 0 -PULSE_INT_R 0
```

Using the NativeLink Feature with NC-Sim

The NativeLink feature in the Quartus II software facilitates the seamless transfer of information between the Quartus II software and EDA tools and allows you to run NC-Sim within the Quartus II software.



For more information, refer to the “Using the NativeLink Feature” section in the *Simulating Altera Designs* chapter in volume 3 of the *Quartus II Handbook*.

Generating a Timing VCD File for the PowerPlay Power Analyzer

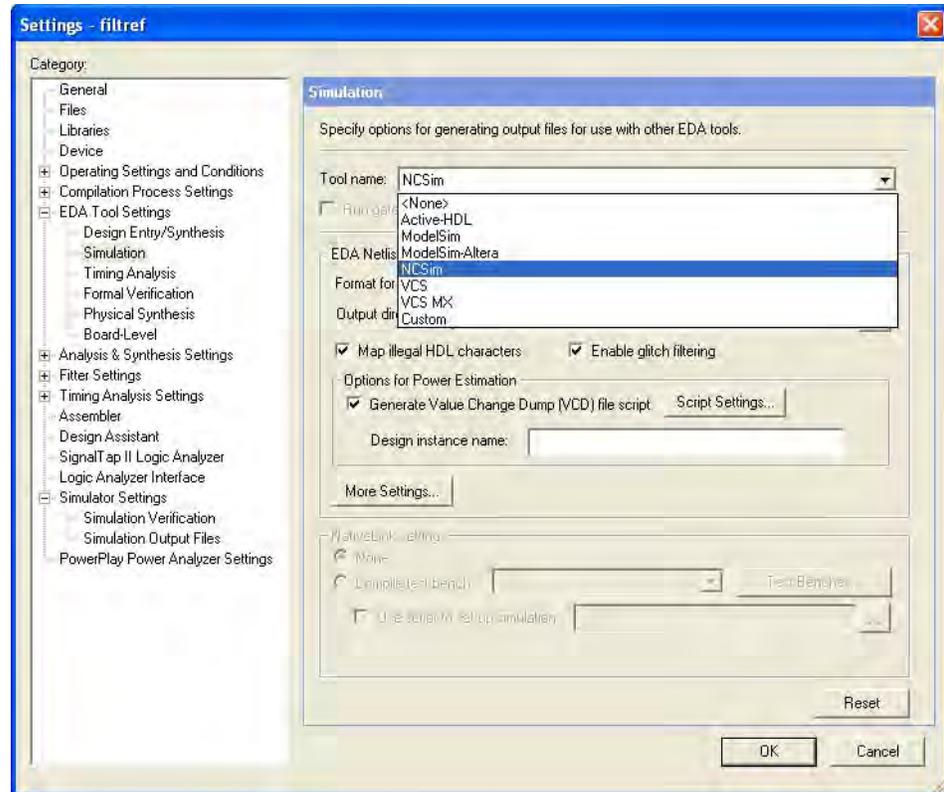
To generate a timing **.vcd** file for PowerPlay, you must first generate a VCD script in the Quartus II software and run the VCD script from the NC-Sim software to generate a timing **.vcd** file. This timing **.vcd** file can then be used by the PowerPlay Power Analyzer for power analysis. The following instructions show you how to generate a timing **.vcd** file.

Perform the following steps to generate timing VCD scripts in the Quartus II software:

1. In the Quartus II software, on the Assignments menu, click **Settings**. The **Settings** dialog box appears (Figure 4-1).
2. In the **Category** list, click the “+” icon to expand **EDA Tool Settings**.
3. Click **Simulation**.

4. In the **Tool name** list, click **NC-Sim**.
5. Turn on the **Generate Value Change Dump (VCD) File Script** option.

Figure 4-1. Simulation Settings Dialog Box



6. Click **OK**.
7. To generate the VCD script file, perform a full compilation.
Perform the following steps to generate a timing .vcd file in NC-Sim:

1. In the NC-Sim software, before simulating your design, source the `<revision_name>_dump_all_vcd_nodes.tcl` script. To source the .tcl script, use the `-input` switch while running the `nssim` command. For example:

```
nssim -input <revision_name>_dump_all_vcd_nodes.tcl <my design>
```

2. Continue to run the simulation until it finishes. Exit `nssim` and the `<revision_name>.vcd` is generated in the simulation directory.

 For more detailed information about using the timing .vcd file for power analysis, refer to the *PowerPlay Power Analysis* chapter in volume 3 of the *Quartus II Handbook*.

Viewing a Waveform from a .trn File

A .trn file is automatically generated when your simulation is done. The .trn file is not readable. It is used for generating the waveform view through SimVision.

To view a waveform from a .trn file through SimVision, perform the following steps:

1. Type `simvision` on a command line. The **Design Browser** dialog box appears.
2. On the File menu, click **Open Database**. The **Open File** dialog box appears.
3. In the **Directories** field, browse to the directory that contains your `.trn` file.
4. Double-click the `.trn` file.
5. In the **Design Browser** dialog box, select the signals that you want to observe from the Hierarchy.
6. Right-click the selected signals and click **Send to Waveform Window**.



You cannot view a waveform from a `.vcd` file in SimVision, and the `.vcd` file cannot be converted to a `.trn` file.

Scripting Support

You can run procedures and make settings described in this chapter in a Tcl script. You can also run some procedures at a command prompt.

For detailed information about scripting command options, refer to the Quartus II Command-Line and Tcl API Help browser.

To run the Help browser, type the following command at the command prompt:

```
quartus_sh --qhelp ←
```



For more information, refer to *About Quartus II Scripting* in Quartus II Help.



For more information about Tcl scripting, refer to the *Tcl Scripting* chapter in volume 2 of the *Quartus II Handbook*. For information about all settings and constraints in the Quartus II software, refer to the *Quartus II Settings File Manual*. For more information about command-line scripting, refer to the *Command-Line Scripting* chapter in volume 2 of the *Quartus II Handbook*.

Generating NC-Sim Simulation Output Files

You can generate `.vo` files and `.sdo` simulation output files with Tcl commands or at a command prompt.

For more information about generating `.vo` simulation output files and `.sdo` file simulation output files, refer to “[Quartus II Simulation Output Files](#)” on page 4–7.

Tcl Commands

The following three assignments cause a Verilog HDL netlist to be written out when you run the Quartus II netlist writer.

```
set_global_assignment -name EDA_OUTPUT_DATA_FORMAT VERILOG -section_id
eda_simulation
set_global_assignment -name EDA_TIME_SCALE "1 ps" -section_id
eda_simulation
set_global_assignment -name EDA_SIMULATION_TOOL "NC-Verilog (Verilog)"
```

The netlist has a 1 ps timing resolution for the NC-Sim simulation software.

Type the following Tcl command to run the Quartus II Netlist Writer:

```
execute_module -tool eda
```

Command Prompt

Type the following command to generate a simulation output file for the Cadence NC-Sim software simulator. Specify Verilog HDL or VHDL for the format.

```
quartus_eda <project name> --simulation --format=<verilog/vhdl> --tool=ncsim ←
```

Conclusion

The Cadence NC family of simulators work within an Altera FPGA design flow to perform functional, post-synthesis, and gate-level timing simulation, easily and accurately.

Altera provides functional models of LPM and Altera-specific megafunctions that you can compile with your testbench or design. For timing simulation, use the atom netlist file generated by Quartus II compilation.

The seamless integration of the Quartus II software and Cadence NC tools make this simulation flow an ideal method for fully verifying an FPGA design.

Document Revision History

Table 4–2 shows the revision history for this chapter.

Table 4–2. Document Revision History (Part 1 of 2)

Date	Version	Changes
December 2010	10.0.1	Changed to new document template. No change to content.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Linked to Help where appropriate ■ Minor text edits ■ Removed Referenced Documents section
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Removed NativeLink information and referenced new <i>Simulating Designs with EDA Tools</i> chapter in volume 3 of the <i>Quartus II Handbook</i> ■ Added “RTL Functional Simulation for Stratix IV Devices” and “Gate-Level Timing Simulation for Stratix IV Devices” sections ■ Minor text edits
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Removed “Compile Libraries Using the Altera Simulation Library Compiler” ■ Added “Compile Libraries Using the EDA Simulation Library Compiler” on page 4–5 ■ Added “Generate Simulation Script from EDA Netlist Writer” on page 4–35 ■ Added “Viewing a Waveform from a .trn File” on page 4–36 ■ Minor editorial updates

Table 4-2. Document Revision History (Part 2 of 2)

Date	Version	Changes
November 2008	8.1.0	<ul style="list-style-type: none"> ■ Added “Compile Libraries Using the Altera Simulation Library Compiler” on page 4-5. ■ Added information about the <code>--simlib_comp</code> utility. ■ Minor editorial updates. ■ Updated entire chapter using 8½” x 11” chapter template.
May 2008	8.0.0.	<ul style="list-style-type: none"> ■ Updated Table 4-1. ■ Updated Figure 4-1. ■ Updated “Compilation in Command-Line Mode” on page 4-9. ■ Updated “Generating a Timing Netlist with Different Timing Models” on page 4-18. ■ Added “Disable Timing Violation on Registers” on page 4-20. ■ Updated “Simulating Designs that Include Transceivers” on page 4-23. ■ Updated “Performing a Gate Level Simulation Using NativeLink” on page 4-30. ■ Added “Generating a Timing VCD File for PowerPlay” on page 4-33. ■ Added hyperlinks to referenced documents throughout the chapter. ■ Minor editorial updates.

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter describes how to use the Active-HDL and Riviera-PRO software to simulate designs that target Altera® FPGAs. This chapter provides step-by-step instructions about how to perform functional simulations, post-synthesis simulations, and gate-level timing simulations for Verilog HDL or VHDL designs.

This chapter includes the following topics:

- “Software Requirements”
- “Using Active-HDL or Riviera-PRO Software in Quartus II Design Flows”
- “Simulation Libraries” on page 5–2
- “Performing Simulation Using the Active-HDL Software (GUI Mode)” on page 5–3
- “Performing Simulation Using the Riviera-PRO Software (GUI Mode)” on page 5–5
- “Performing Simulation Using the Active-HDL or Riviera-PRO Software (Command-Line Mode)” on page 5–6
- “Compiling System Verilog Files” on page 5–11
- “Simulating Designs that Include Transceivers” on page 5–11
- “Using the NativeLink Feature in Active-HDL or Riviera-PRO Software” on page 5–18
- “Generating .vcd Files for the PowerPlay Power Analyzer” on page 5–18
- “Scripting Support” on page 5–19

Software Requirements

To simulate your design using the Active-HDL or Riviera-PRO software, you must first set up the Altera libraries. These libraries are installed with the Quartus II software.



For more information about installing the software and directories created during the Quartus II software installation, refer to the *Altera Software Installation and Licensing* manual.

Using Active-HDL or Riviera-PRO Software in Quartus II Design Flows

You can perform the following types of simulations using the Active-HDL or Riviera-PRO software:

- Functional simulation
- Post-synthesis simulation

- Gate-level timing simulation

Simulation Libraries

Simulation model libraries are required to run a simulation whether you are running a functional simulation, post-synthesis simulation, or gate-level timing simulation. In general, running a functional simulation requires the functional simulation model libraries and running a post-synthesis or gate-level timing simulation requires the gate-level timing simulation model libraries. You must compile the necessary library files before you can run the simulation.

However, there are a few exceptions where you are also required to compile gate-level timing simulation library files to perform functional simulation. For example, some of Altera megafunctions require gate-level libraries to perform a functional simulation in third-party simulators.



For each megafunction that you are using, refer to the last page in the Altera MegaWizard™ Plug-In Manager, which lists the simulation library files required to perform a functional simulation for that megafunction.

The transceiver megafunction (for example, ALTGX) also requires the gate-level libraries to perform functional simulation.

For detailed, step-by-step instructions about how to simulate the transceiver megafunction, refer to [“Simulating Designs that Include Transceivers”](#) on page 5–11.

Simulation Library Files in the Quartus II Software

In Active-HDL or Riviera-PRO software, you must compile the necessary libraries to perform functional, post-synthesis functional, or gate-level timing simulation. You can refer to these library files for any particular simulation model that you are looking for.

- ② For a list of all functional simulation library files in the Quartus II directory, refer to [Altera Functional Simulation Libraries](#) in Quartus II Help. For a list of all gate-level timing simulation and post-fit library files in the Quartus II directory, refer to [Altera Post-Fit Libraries](#) in Quartus II Help.

Disabling Timing Violation on Registers

In certain situations, the timing violation can be ignored and you can disable the timing violation on registers; for example, timing violations that occur in internal synchronization registers used for asynchronous clock-domain crossing.

By default, the `x_on_violation_option logic` option is **On**, which means the simulation shows “x” whenever a timing violation occurs. To disable showing the timing violation on certain registers, set the `x_on_violation_option logic` option to **Off** on those registers. The following command is an example of the QSF file:

```
set_instance_assignment -name X_ON_VIOLATION_OPTION OFF -to <register_name>
```

Compiling Libraries Using the EDA Simulation Library Compiler

The EDA Simulation Library Compiler is used to compile Verilog HDL and VHDL simulation libraries for all Altera devices and supported third-party simulators. You can use this tool to compile all libraries required by and gate-level timing simulation.



The `altera_mf_components.vhd` and `altera_mf.vhd` model files should be compiled into the `altera_mf` library. The `220pack.vhd` and `220model.vhd` model files should be compiled into the `lpm` library.



For more information about this tool, refer to the “EDA Simulation Library Compiler” section in the *Simulating Altera Designs* chapter in volume 3 of the *Quartus II Handbook*.

Performing Simulation Using the Active-HDL Software (GUI Mode)

Perform simulation of Verilog HDL or VHDL designs with Active-HDL software at various levels to verify designs from different aspects. There are three types of simulation:

- Functional simulation
- Post-synthesis simulation
- Gate-level timing simulation

Simulation helps you verify your designs and debug them against any errors the designs may have. The following sections provide step-by-step instructions to perform the simulation through the GUI.

For high-speed simulation, you must select **ps** in the **Resolution** list for your simulator resolutions. If you choose slower than **ps**, the high-speed simulation may fail.

Workspace creation is the mandatory first step to start working in the Active-HDL GUI. You must create a new workspace to add the simulation model files, design files, and testbench file before you can compile them.

Simulating VHDL Designs

When you simulate VHDL designs using the Active-HDL GUI, you do not have to remember the commands to compile the libraries or load and simulate the VHDL design files. You can use the Active-HDL GUI to perform the functional simulation, post-synthesis simulation, and gate-level timing simulation.

Performing Functional Simulation

Functional simulation is typically performed to verify the syntax of the code and to check the functionality of the design.



For detailed information about how to perform functional simulation in the Active-HDL software for VHDL designs, refer to *Performing a Simulation of a VHDL Design with the Active-HDL Software* in Quartus II Help.

If you are targeting a Stratix V device, compile the following files in the `quartus/eda/sim_lib/aldec` directory:

```
stratixv_atoms_ncrypt.v
stratixv_hssi_atoms_ncrypt.v
stratixv_pcie_hip_atoms_ncrypt.v
```

These files contain IEEE encrypted Verilog models suitable for VHDL/Verilog co-simulation. You need a co-simulation license from Aldec to use these models in VHDL.

 Compile these encrypted Verilog files before you compile any VHDL files. Compile `stratixv_pcie_hip_atoms_ncrypt.v` with the SystemVerilog option.

Also, compile the following files in the `quartus/eda/sim_lib` directory:

```
stratixv_atoms.vhd
stratixv_components.vhd
stratixv_hssi_components.vhd
stratixv_pcie_hip_components.vhd
stratixv_hssi_atoms.vhd
stratixv_pcie_hip_atoms.vhd
```

 The PCIe[®] files are required only if you are using the PCIe HIP.

Simulating Verilog HDL Designs

When you simulate Verilog HDL designs using the Active-HDL GUI, you do not have to remember the commands to compile the libraries or load and simulate the Verilog HDL design files. You can use the Active-HDL GUI to perform functional simulation, post-synthesis simulation, and gate-level timing simulation.

Performing Functional Simulation

Functional simulation is performed to verify the syntax of the code and to check the functionality of the design.

 For detailed information about how to perform functional simulation in the Active-HDL software for Verilog HDL designs, refer to *Performing a Simulation of a Verilog HDL Design with the Active-HDL Software* in Quartus II Help.

If you are targeting a Stratix V device, compile the following files in the `quartus/eda/sim_lib/aldec` directory:

```
stratixv_atoms_ncrypt.v
stratixv_hssi_atoms_ncrypt.v
stratixv_pcie_hip_atoms_ncrypt.v
```

These files contain IEEE encrypted Verilog models suitable for VHDL/Verilog co-simulation.

 Compile these encrypted Verilog files before you compile any VHDL files.

 Compile `stratixv_pcie_hip_atoms_ncrypt.v` with the SystemVerilog option.

Also, compile the following files in the `quartus/eda/sim_lib` directory:

`stratixv_atoms.v`
`stratixv_hssi_atoms.v`
`stratixv_pcie_hip_atoms.v`

 Compile `stratixv_pcie_hip_atoms.v` with the SystemVerilog option.

 The PCIe files are required only if you are using the PCIe HIP.

Performing Simulation Using the Riviera-PRO Software (GUI Mode)

You can use the Riviera-PRO software to perform the following types of simulation:

- Functional simulation
- Post-synthesis simulation
- Gate-level timing simulation

 For detailed information about how to perform functional simulation with the Riviera-PRO software, refer to *Performing an RTL Functional Simulation with the Riviera-PRO Software* in Quartus II Help. For detailed information about how to perform post-synthesis simulation with the Riviera-PRO software, refer to *Performing a Post-Synthesis Simulation with the Riviera-PRO Software* in Quartus II Help. For detailed information about how to perform gate-level timing simulation with the Riviera-PRO software, refer to *Performing a Gate-Level Simulation with the Riviera-PRO Software* in Quartus II Help.

If you are performing a functional simulation in Verilog HDL and targeting a Stratix V device, compile the following files in the `quartus/eda/sim_lib/aldec` directory:

`stratixv_atoms_ncrypt.v`
`stratixv_hssi_atoms_ncrypt.v`
`stratixv_pcie_hip_atoms_ncrypt.v`

These files contain IEEE encrypted Verilog models suitable for VHDL/Verilog co-simulation.

 Compile these encrypted Verilog files before you compile any VHDL files.

 Compile `stratixv_pcie_hip_atoms_ncrypt.v` with the SystemVerilog option.

Also, compile the following files in the `quartus/eda/sim_lib` directory:

`stratixv_atoms.v`
`stratixv_hssi_atoms.v`
`stratixv_pcie_hip_atoms.v`

 Compile `stratixv_pcie_hip_atoms.v` with the SystemVerilog option.

 The PCIe files are required only if you are using the PCIe HIP.

If you are performing a functional simulation in VHDL and targeting a Stratix V device, compile the following files in the `quartus/eda/sim_lib/aldec` directory:

```
stratixv_atoms_ncrypt.v
stratixv_hssi_atoms_ncrypt.v
stratixv_pcie_hip_atoms_ncrypt.v
```

These files contain IEEE encrypted Verilog models suitable for VHDL/Verilog co-simulation. You need a co-simulation license from Aldec to use these models in VHDL.

 Compile these encrypted Verilog files before you compile any VHDL files. Compile `stratixv_pcie_hip_atoms_ncrypt.v` with the `SystemVerilog` option.

Also, compile the following files in the `quartus/eda/sim_lib` directory:

```
stratixv_atoms.vhd
stratixv_components.vhd
stratixv_hssi_components.vhd
stratixv_pcie_hip_components.vhd
stratixv_hssi_atoms.vhd
stratixv_pcie_hip_atoms.vhd
```

 The PCIe files are required only if you are using the PCIe HIP.

Performing Simulation Using the Active-HDL or Riviera-PRO Software (Command-Line Mode)

Perform simulation of Verilog HDL or VHDL designs with Active-HDL or Riviera-PRO software at various levels to verify designs from different aspects. There are three categories of simulation:

- Functional simulation
- Post-synthesis simulation
- Gate-level simulation

Simulation helps you verify your design and debug it against any possible errors in the design.

For high-speed simulation, you must select **ps** in the **Resolution** list for your simulator resolutions. If you choose slower than **ps**, the high-speed simulation may fail.

In command-line mode, standalone commands, such as `vlib`, `vcom`, and `vsim`, are executed in the system shell (for example, `cygwin`). These standalone commands can be grouped into script files (`tcl`, `perl`, `windows batch`) that are run from the system shell.

Before running Active-HDL or Riviera-PRO from the command line, ensure that the **Active-HDL/bin** or **Riviera-PRO/bin** directory is located in `PATH` environment variables.

Simulating Verilog HDL Designs Using the Active-HDL or Riviera-PRO Software

If you are running the simulation commands from the Active-HDL GUI console, you must first create the workspace. Type the following commands to create and open the workspace:

```
createdesign DELAY_TEST C:/DELAY_TEST/simulation/activehdl ←
opendesign -a DELAY_TEST.adf ←
```

If you are running Riviera-PRO, you can skip the step above.

Performing Functional Simulation

Use the following commands to perform a functional simulation for Verilog HDL designs with one of the libraries (lib1) listed in *Altera Functional Simulation Libraries* in Quartus II Help:

```
# Create and compile Altera libraries
vlib <lib1> ←
vlog -v2k -dbg -work <lib1> <lib1.v> ←

vlib <lib2> ←
vlog -v2k -dbg -work <lib2> <lib2.v> ←

# Create work library and compile design files and testbench file
vlib work ←
vlog -v2k -dbg -work work <design_file1.v> <design_file2.v> <testbench
 \ file.v> ←

# Load Design
vsim +access +r -t lps -L work -L <lib1> -L <lib2> work.<testbench module
 \ name> ←

#add signals at waveform and run
add wave * ←
run ←
```

Example

```
vlib verilog_libs/lpm_ver
vlog -v2k -dbg -work lpm_ver
c:/altera/91/quartus/eda/sim_lib/220model.v

vlib verilog_libs/altera_mf_ver
vlog -v2k -dbg -work altera_mf_ver
c:/altera/91/quartus/eda/sim_lib/altera_mf.v

vlib work
vlog -v2k -dbg -work work C:/project/adder.v C:/project/adder.vt

vsim +access +r -t lps -L work -L lpm_ver -L altera_mf_ver
work.adder_vlg_vec_tst

add wave *
```

Performing Post-Synthesis Simulation

Before you run post-synthesis simulation, generate post-synthesis simulation netlist files. Refer to the “Generating Post-Synthesis Simulation Netlist Files” section in the *Simulating Altera Designs* chapter in volume 3 of the *Quartus II Handbook*.



You cannot perform post-synthesis or post-fit simulation if you are targeting the Stratix V device family.

Use the following commands to perform a post-synthesis simulation for Verilog HDL designs with one of the libraries (lib1) listed in *Altera Post-Fit Libraries* in Quartus II Help:

```
# Create and compile Altera libraries
vlib <lib1> ←
vlog -v2k -dbg -work <lib1> <lib1.v> ←

vlib <lib2> ←
vlog -v2k -dbg -work <lib2> <lib2.v> ←

# Create work library and compile EDA output netlist files and testbench
file
vlib work ←
vlog -v2k -dbg -work work <EDA_output_netlist.vo> <testbench file.v> ←

# Load Design
vsim +access +r -t lps +transport_int_delays +transport_path_delays -L
work \
-L <lib1> -L <lib2> work.<testbench module name> ←

#add signals at waveform and run
add wave * ←
run ←
```

Example

```
vlib verilog_libs/lpm_ver
vlog -v2k -dbg -work lpm_ver
c:/altera/91/quartus/eda/sim_lib/220model.v

vlib verilog_libs/altera_ver
vlog -v2k -dbg -work altera_ver
c:/altera/91/quartus/eda/sim_lib/altera_primitives.v

vlib verilog_libs/stratixiv_ver
vlog -v2k -dbg -work stratixiv_ver
c:/altera/91/quartus/eda/sim_lib/stratixiv_atoms.v

vlib work

vlog -v2k -dbg -work work C:/project/simulation/activehdl/adder.vo
C:/project/adder.vt

vsim +access +r -t lps +transport_int_delays +transport_path_delays -L
work -L lpm_ver -L altera_mf_ver work.adder_vlg_vec_tst

add wave *
run
```

Performing Gate-Level Timing Simulation

The steps for gate-level timing simulation are almost same as the steps for post-synthesis simulation. The only difference is that the SDO file must be back-annotated for gate level-timing simulation.

For Verilog HDL designs, the back-annotating process is done within the **EDA_output_netlist.vo** script. Therefore, you are not required to back-annotate the SDO file again.



You cannot perform post-synthesis or post-fit simulation if you are targeting the Stratix V device family.

Simulating VHDL Designs Using the Active-HDL or Riviera-PRO Software

If you are running the simulation commands from the Active-HDL GUI console, you first create the workspace. The following commands create and open the workspace:

```
createdesign <workspace name> <your design path> ←
opendesign -a <workspace name>.adf ←
```

If you are running Riviera-PRO, you can skip the step above.

Performing Functional Simulation

Use the following commands to perform a functional simulation for VHDL designs with one of the libraries (lib1) listed in *Altera Functional Simulation Libraries* in Quartus II Help:

```
# Create and compile Altera libraries
vlib <lib1> ←
vcom -strict93 -dbg -work <lib1> <lib1_component/pack.vhd> <lib1.vhd> ←

vlib <lib1> ←
vcom -strict93 -dbg -work <lib2> <lib2_component/pack.vhd> <lib2.vhd> ←

# Create work library and compile design files and testbench file
vlib work ←
vcom -strict93 -dbg -work work <design_file1.vhd> <design_file2.vhd>
\ <testbench_file.vhd> ←

# Load Design
vsim +access +r -t lps -L work -L <lib1> -L <lib2> work.<testbench module
\ name> ←

#add signals at waveform and run
add wave * ←
run ←
```

Example

```
vlib vhdl_libs/lpm
vcom -strict93 -dbg -work lpm
c:/altera/91/quartus/eda/sim_lib/220pack.vhd
vcom -strict93 -dbg -work lpm
c:/altera/91/quartus/eda/sim_lib/220model.vhd

vlib vhdl_libs/altera_mf
vcom -strict93 -dbg -work altera_mf
c:/altera/91/quartus/eda/sim_lib/altera_mf_components.vhd
vcom -strict93 -dbg -work altera_mf
c:/altera/91/quartus/eda/sim_lib/altera_mf.vhd

vlib work
vcom -strict93 -dbg -work work C:/project/adder.vhd
C:/project/adder.vht

vsim +access +r -t lps -L adder -L work -L lpm -L altera_mf
work.adder_vhd_vec_tst

#add signals at waveform and run
add wave *
run
```

Performing Post-Synthesis Simulation

Before you run post-synthesis simulation, generate post-synthesis simulation netlist files. Refer to the “Generating Post-Synthesis Simulation Netlist Files” section in the *Simulating Altera Designs* chapter in volume 3 of the *Quartus II Handbook*.



You cannot perform post-synthesis or post-fit simulation if you are targeting the Stratix V device family.

Use the following commands to perform a functional simulation for VHDL designs with one of the libraries (lib1) listed in *Altera Post-Fit Libraries* in Quartus II Help:

```
# Create and compile Altera libraries
vlib <lib1> ←
vcom -strict93 -dbg -work <lib1> <lib1_component/pack.vhd> <lib1.vhd> ←

vlib <lib2> ←
vcom -strict93 -dbg -work <lib2> <lib2_component/pack.vhd> <lib2.vhd> ←

# Create work library and compile EDA output netlist files and testbench
file
vlib work ←
vcom -strict93 -dbg -work work <EDA output netlist.vho> <testbench \
file.vhd> ←

# Load Design
vsim +access+r -t lps +transport_int_delays +transport_path_delays -L
work \ -L <lib1> -L <lib2> work.<testbench module name> ←

#add signals at waveform and run
add wave * ←
run ←
```

Example

```
vlib vhdl_libs/lpm
vcom -strict93 -dbg -work lpm
c:/altera/91/quartus/eda/sim_lib/220pack.vhd
vcom -strict93 -dbg -work lpm
c:/altera/91/quartus/eda/sim_lib/220model.vhd

vlib vhdl_libs/altera
vcom -strict93 -dbg -work altera
c:/altera/91/quartus/eda/sim_lib/altera_primitives_components.vhd
vcom -strict93 -dbg -work altera
c:/altera/91/quartus/eda/sim_lib/altera_primitives.vhd

vlib work
vcom -strict93 -dbg -work work
C:/project/simulation/activehdl/adder.vho C:/project/adder.vht

vsim +access +r -t lps +transport_int_delays +transport_path_delays -L
adder -L work -L lpm -L altera_mf work.adder_vhd_vec_tst

add wave *
run
```

Performing Gate-Level Timing Simulation

The steps for gate-level timing simulation are almost same as the steps for post-synthesis simulation. The only difference is that the SDO file must be back-annotated for gate level-timing simulation.

 You cannot perform post-synthesis or post-fit simulation if you are targeting the Stratix V device family.

For VHDL designs, the back-annotating process is done by adding the `-sdftyp` option.

Example

```
vsim +access +r -t lps +transport_int_delays +transport_path_delays  
-sdftyp <instance path to design>= <path to SDO file> -L adder -L work  
-L lpm -L altera_mf work.adder_vhd_vec_tst
```

Compiling System Verilog Files

If your design includes multiple System Verilog files, you must compile the System Verilog files together with a single `alog` command.

If you have Verilog files and System Verilog files in your design, it is recommended that you compile the Verilog files, and then compile only the System Verilog files in the single `alog` command.

Simulating Designs that Include Transceivers

If your design includes Arria®, Arria II, Cyclone IV, HardCopy IV, Stratix, Stratix II, or Stratix IV transceivers, you must compile additional library files to perform functional or gate-level timing simulations. The following example shows how to perform simulation on designs that include Stratix GX and Stratix II GX transceivers.

For high-speed simulation, you must select `ps` in the **Resolution** list for your simulator resolutions (**Design** tab of the **Start Simulation** dialog box). If you choose slower than `ps`, the high-speed simulation may fail.

Performing simulation with transceivers in Arria GX or Stratix II GX is very similar. The only requirement is to replace the `stratixiigx_atoms` and `stratixiigx_hssi_atoms` model files with the `arriagx_atoms` and `arriagx_hssi_atoms` model files, respectively.

 If your design contains PCI Express hard IP, refer to the “Simulate the Design” section in the *PCI Express Compiler User Guide*.

Functional Simulation for Stratix II GX Devices

Functional simulation for Stratix II GX devices is similar to functional simulation for Arria GX devices. The following example shows only the functional simulation for designs that include transceivers in Stratix II GX devices. To simulate transceivers in Arria GX devices, replace the `stratixiigx_hssi` model file with the `arriagx_hssi` model file.

To perform an functional simulation of your design that instantiates the ALT2GXB megafunction, which enables the gigabit transceiver blocks on Stratix II GX devices, you must generate a functional simulation netlist and compile the `stratixiigx_hssi` model file into the `stratixiigx_hssi` library.

 The `stratixgx_hssi_atoms` model file references the `lpm` and `sgate` libraries; you must create these libraries to perform a simulation.

To run the functional simulation, you must generate a functional simulation netlist by turning on **Generate Simulation Model** in the **Simulation Libraries** tab of the ALT2GXB MegaWizard Plug-In Manager.

The `<alt2gxb entity name>.vho` or `<alt2gxb module name>.vo` is generated in the current project directory.

The ALT2GXB functional simulation library file generated by the Quartus II software references `stratixiigx_hssi` WYSIWYG atoms.

Performing Functional Simulation in VHDL

Type the commands in [Example 5-1](#) to compile and simulate the design.

Example 5-1.

```
vcom -work lpm 220pack.vhd 220model.vhd
vcom -work altera_mf altera_mf_components.vhd altera_mf.vhd
vcom -work sgate sgate_pack.vhd sgate.vhd
vcom -work stratixiigx_hssi stratixiigx_hssi_components.vhd \
stratixiigx_hssi_atoms.vhd
vcom -work work <alt2gxb entity name>.vho
vcom -work work <my design>.vhd <my testbench>.vhd
vsim -L lpm -L altera_mf -L sgate -L stratixgx_hssi work.<my testbench>
```

Performing Functional Simulation in Verilog HDL

Type the commands in [Example 5-2](#) to compile and simulate the design.

Example 5-2.

```
vlog -work lpm_ver 220model.v
vlog -work altera_mf_ver altera_mf.v
vlog -work sgate_ver sgate.v
vlog -work stratixiigx_hssi_ver stratixiigx_hssi_atoms.v
vlog -work work <alt2gxb module name>.vo
vlog -work work <my design>.v <my testbench>.v
vsim -L lpm_ver -L altera_mf_ver -L sgate_ver -L stratixgx_hssi \
work.<my testbench>
```

Gate-Level Timing Simulation for Stratix II GX Devices

To perform a gate-level timing simulation of your design that includes a Stratix II GX transceiver, you must compile `stratixiigx_atoms` and `stratixiigx_hssi_atoms` into the `stratixiigx` and `stratixiigx_hssi` libraries, respectively.

Performing Gate-Level Timing Simulation in VHDL

Type the commands in [Example 5-3](#) to compile and simulate the design.

Example 5-3.

```
vcom -work lpm 220pack.vhd 220model.vhd
vcom -work altera_mf altera_mf_components.vhd altera_mf.vhd
vcom -work sgate sgate_pack.vhd sgate.vhd
vcom -work stratixiigx stratixiigx_atoms.vhd stratixiigx_components.vhd
vcom -work stratixiigx_hssi stratixiigx_hssi_components.vhd stratixiigx_hssi_atoms.vhd
vcom -work work <my design>.vho <my testbench>.vhd
vsim -L lpm -L altera_mf -L sgate -L stratixiigx -L stratixiigx_hssi \
work.<my testbench> -t ps -sdftyp <design instance>=<path to SDO file>.sdo \
+transport_int_delays +transport_path_delays
```

Performing Gate-Level Timing Simulation in Verilog HDL

Type the commands in [Example 5-4](#) to compile and simulate the design.

Example 5-4.

```
vlog -work lpm_ver 220model.v
vlog -work altera_mf_ver altera_mf.v
vlog -work sgate_ver sgate.v
vlog -work stratixiigx_ver stratixiigx_atoms.v
vlog -work stratixiigx_hssi_ver stratixiigx_hssi_atoms.v
vlog -work work <my design>.vo <my testbench>.v
vsim -L lpm -L altera_mf_ver -L sgate_ver -L stratixiigx_ver -L stratixiigx_hssi_ver \
work.<my testbench> -t ps +transport_int_delays +transport_path_delays
```

Functional Simulation for Stratix GX Devices

To perform a functional simulation of your design that instantiates the ALTGXB megafunction, which enables the gigabit transceiver block on Stratix GX devices, compile the **stratixgx_mf** model file into the **altgxb** library.



The **stratixgx_mf** model file references the **lpm** and **sgate** libraries. You must create these libraries to perform a simulation.

Performing Functional Simulation in VHDL

Type the commands in [Example 5-5](#) to compile and simulate the design.

Example 5-5.

```
vcom -work altera_mf altera_mf_components.vhd altera_mf.vhd
vcom -work lpm 220pack.vhd 220model.vhd
vcom -work sgate sgate_pack.vhd sgate.vhd
vcom -work altgxb stratixgx_mf.vhd stratixgx_mf_components.vhd
vcom -work work<altgxb entity name>.vhd
vsim -L lpm -L altera_mf -L sgate -L altgxb work.<my testbench>
```

Performing Functional Simulation in Verilog HDL

Type the commands in [Example 5-6](#) to compile and simulate the design.

Example 5-6.

```
vlib work
vlib lpm_ver
vlib altera_mf_ver
vlib sgate_ver
vlib altgxb_ver
vlog -work lpm_ver 220model.v
vlog -work altera_mf_ver altera_mf.v
vlog -work sgate_ver sgate.v
vlog -work altgxb_ver stratixgx_mf.v
vlog -work work <altgxb module name>.v
vsim -L lpm_ver -L altera_mf_ver -L sgate_ver -L altgxb_ver work.<my testbench>
```

Gate-Level Timing Simulation for Stratix GX Devices

Perform a gate-level timing simulation of your design that includes a Stratix GX transceiver by compiling the `stratixgx_atoms` and `stratixgx_hssi_atoms` model files into the `stratixgx` and `stratixgx_gxb` libraries, respectively.

Performing Gate-Level Timing Simulation in VHDL

Type the commands in [Example 5-7](#) to compile and simulate the design.

Example 5-7.

```
vcom -work lpm 220pack.vhd 220model.vhd
vcom -work altera_mf altera_mf_components.vhd altera_mf.vhd
vcom -work sgate sgate_pack.vhd sgate.vhd
vcom -work stratixgx stratixgx_atoms.vhd stratixgx_components.vhd
vcom -work stratixgx_gxb stratixgx_hssi_atoms.vhd \
stratixgx_hssi_components.vhd
vcom -work work <my design>.vho <my testbench>.vhd
vsim -L lpm -L altera_mf -L sgate -L stratixgx -L stratixgx_gxb work. \
<my testbench> -t ps -sdftyp <design instance>=<path to SDO file>.sdo \
+transport_int_delays +transport_path_delays
```

Performing Gate-Level Timing Simulation in Verilog HDL

Type the commands in [Example 5-8](#) to compile and simulate the design.

Example 5-8.

```
vlog -work lpm_ver 220model.v
vlog -work altera_mf_ver altera_mf.v
vlog -work sgate_ver sgate.v
vlog -work stratixgx_ver stratixgx_atoms.v
vlog -work stratixgx_gxb_ver stratixgx_hssi_atoms.v
vlog -work work <my design>.vo <my testbench>.v
vsim -L lpm_ver -L altera_mf_ver -L sgate_ver -L stratixgx_ver -L stratixgx_gxb_ver \
work.<my testbench> -t ps +transport_int_delays +transport_path_delaysr
```

Functional Simulation for Stratix IV GX Devices

Functional simulation for Stratix IV devices is similar to functional simulation for Arria II, Cyclone IV, and HardCopy IV devices.

The following example shows only the functional simulation for designs that include transceivers in Stratix IV devices. To simulate transceivers in Arria II, Cyclone IV, and HardCopy IV devices, replace the **stratixiv_hssi** model file with the **arriaii_hssi**, **cycloneiv_hssi**, and **hardcopyiv_hssi** model files, respectively.

To perform a functional simulation of your design that instantiates the ALTGX megafunction, which enables the gigabit transceiver blocks on Stratix IV devices, you must generate a functional simulation netlist and compile the **stratixiv_hssi** model file into the **stratixiv_hssi** library.



The **stratixiv_hssi** model file references the **lpm** and **sgate** libraries. You must create these libraries to perform a simulation.

Performing Functional Simulation in VHDL

Type the commands in [Example 5-9](#) to compile and simulate the design.

Example 5-9.

```
vcom -work altera_mf altera_mf_components.vhd altera_mf.vhd
vcom -work lpm 220pack.vhd 220model.vhd
vcom -work sgate sgate_pack.vhd sgate.vhd
vcom -work stratixiv_hssi stratixiv_hssi.vhd\
stratixiv_hssi_components.vhd
vcom -work work <altgxb entity name>.vhd
vsim -L lpm -L altera_mf -L sgate -L stratixiv_hssi work.<my testbench>
```

Performing Functional Simulation in Verilog HDL

Type the commands in [Example 5-10](#) to compile and simulate the design.

Example 5-10.

```
vlib work
vlib lpm_ver
vlib altera_mf_ver
vlib sgate_ver
vlib stratixiv_hssi_ver
vlog -work lpm_ver 220model.v
vlog -work altera_mf_ver altera_mf.v
vlog -work sgate_ver sgate.v
vlog -work stratixiv_hssi_ver stratixiv_hssi.v
vlog -work work <altgxb module name>.v
vsim -L lpm_ver -L altera_mf_ver -L sgate_ver \
-L stratixiv_hssi_ver work.<my testbench>
```

Gate-Level Timing Simulation for Stratix IV GX Devices

Perform a gate-level timing simulation of your design that includes a Stratix IV transceiver by compiling the **stratixiv_atoms** and **stratixiv_hssi_atoms** model files into the **stratixiv** and **stratixiv_hssi** libraries, respectively.

Performing Gate-Level Timing Simulation in VHDL

Type the commands in [Example 5-11](#) to compile and simulate the design.

Example 5-11.

```
vcom -work lpm 220pack.vhd 220model.vhd
vcom -work altera_mf altera_mf_components.vhd altera_mf.vhd
vcom -work sgate sgate_pack.vhd sgate.vhd
vcom -work stratixiv stratixiv_atoms.vhd stratixiv_components.vhd
vcom -work stratixiv_hssi stratixiv_hssi_atoms.vhd \
stratixiv_hssi_components.vhd
vcom -work work <my design>.vho <my testbench>.vhd
vsim -L lpm -L altera_mf -L sgate -L stratixiv -L stratixiv_hssi \
work.<my testbench> -t ps \
-sdftyp <design instance>=<path to SDO file>.sdo \
+transport_int_delays +transport_path_delays
```

Performing Gate-Level Timing Simulation in Verilog HDL

Type the commands in [Example 5-12](#) to compile and simulate the design.

Example 5-12.

```
vlog -work lpm_ver 220model.v
vlog -work altera_mf_ver altera_mf.v
vlog -work sgate_ver sgate.v
vlog -work stratixiv_ver stratixiv_atoms.v
vlog -work stratixiv_hssi_ver stratixiv_hssi_atoms.v
vlog -work work <my design>.vo <my testbench>.v
vsim -L lpm_ver -L altera_mf_ver -L sgate_ver -L stratixiv_ver -L
stratixiv_hssi_ver \
work.<my testbench> -t ps +transport_int_delays +transport_path_delays
```

Functional Simulation for Stratix V GX Devices

Functional simulation for Stratix V devices is similar to functional simulation for Arria II, Cyclone IV, HardCopy IV, and Stratix IV devices.

The following example shows only the functional simulation for designs that include transceivers in Stratix V devices. To simulate transceivers in Arria II, Cyclone IV, HardCopy IV, and Stratix V devices, replace the **stratixv_hssi** model file with the **arriaii_hssi**, **cycloneiv_hssi**, **hardcopyiv_hssi**, and **stratixiv_hssi** model files, respectively.



The transceiver module from the MegaWizard Plug-In Manager is created in **Interfaces/Transceiver PHY**. Select **Custom PHY**.

Performing Functional Simulation in VHDL

Type the commands in [Example 5-13](#) to compile and simulate the design.

Example 5-13.

```
vcom -work altera_mf altera_mf_components.vhd altera_mf.vhd
vcom -work lpm 220pack.vhd 220model.vhd
vcom -work sgate sgate_pack.vhd sgate.vhd
vlog +v2k -work stratixv_hssi \
quartus/eda/sim_lib/aldec/stratixv_hssi_atoms_ncrypt.v
vcom -work stratixv_hssi stratixv_hssi.vhd \
stratixv_hssi_components.vhd
vcom -work work <my design>.vhd
vsim -L lpm -L altera_mf -L sgate -L stratixv_hssi work.<my testbench>
```

Performing Functional Simulation in Verilog HDL

Type the commands in [Example 5-14](#) to compile and simulate the design.

Example 5-14.

```
vlib work
vlib lpm_ver
vlib altera_mf_ver
vlib sgate_ver
vlib stratixv_hssi_ver
vlog -work lpm_ver 220model.v
vlog -work altera_mf_ver altera_mf.v
vlog -work sgate_ver sgate.v
vlog +v2k -work stratixv_hssi \
quartus/eda/sim_lib/aldec/stratixv_hssi_atoms_ncrypt.v
vlog -work stratixv_hssi_ver stratixv_hssi.v
vlog -work work <my design>.v
vsim -L lpm_ver -L altera_mf_ver -L sgate_ver \
-L stratixv_hssi_ver work.<my testbench>
```



The `stratixv_hssi` model file references the `lpm` and `sgate` libraries. You must create these libraries to perform a simulation.



In addition to the top-level variant wrapper, `<variant>.v`, you also get a simulation files subdirectory, `<variant>_sim/`. All Verilog (`.v`) and SystemVerilog (`.sv`) files in the simulation directory must also be compiled into the simulation project.

Transport Delays

By default, the Active-HDL or Riviera-PRO software filters out all pulses that are shorter than the propagation delay between primitives. Turning on the **transport delay** options in the Active-HDL or Riviera-PRO software prevents the simulation tool from filtering out these pulses. Use the following options to ensure that all signal pulses are visible in the simulation results:

- **+transport_path_delays**

Use this option when the pulses in your simulation may be shorter than the delay within a gate-level primitive.

- **+transport_int_delays**

Use this option when the pulses in your simulation may be shorter than the interconnect delay between gate-level primitives.

 The `+transport_path_delays` and `+transport_int_delays` options are also used by default in the NativeLink feature for gate-level timing simulation.

 For more information about either of these options, refer to the Active-HDL online documentation installed with the Active-HDL software.

Example 5-15 shows an Active-HDL software command in command-line syntax to perform a gate-level timing simulation with the device family library.

Example 5-15.

```
vsim -t lps -L stratixii -sdftyp /il=filtref_vhd.sdo \
work.filtref_vhd_vec_tst +transport_int_delays +transport_path_delays
```

Using the NativeLink Feature in Active-HDL or Riviera-PRO Software

The NativeLink feature in the Quartus II software facilitates the seamless transfer of information between the Quartus II software and EDA tools and allows you to run the Active-HDL or Riviera-PRO software within the Quartus II software.

 For more information, refer to the “Using the NativeLink Feature” section in the *Simulating Altera Designs* chapter in volume 3 of the *Quartus II Handbook*.

Generating .vcd Files for the PowerPlay Power Analyzer

To generate a Value Change Dump File (.vcd) for the PowerPlay power analyzer, you must first generate a VCD script in the Quartus II software and run the VCD script from the Active-HDL software to generate a VCD file. This VCD file can then be used by PowerPlay for power analysis. The following instructions show you how to generate a VCD file.

Perform the following steps to generate VCD Scripts in the Quartus II software:

1. In the Quartus II software, on the Assignments menu, click **Settings**. The **Settings** dialog box appears.
2. In the **Category** list, select **Simulator Settings**.
3. On the **Simulator Settings** page, in the **Tool name** list, select **Active-HDL** and turn on the **Generate Value Change Dump File Script** option.
4. To generate the VCD Script file, perform a full compilation.

Perform the following steps to generate a VCD file in the Active-HDL software:

1. In the Active-HDL software, before simulating your design, source the `<revision_name>_dump_all_vcd_nodes.tcl` script. To source the TCL script, run the following command before running the `vsim` command:

```
source <revision_name>_dump_all_vcd_nodes.tcl ↵
```

2. Continue to run the simulation until the simulation is completed. Exit the Active-HDL software. If you do not exit the software, the Active-HDL software may end the writing process of the VCD files improperly, resulting in a corrupted VCD file.

- For more details about using the VCD file for power analysis, refer to the *PowerPlay Power Analysis* chapter in volume 3 of the *Quartus II Handbook*.

Scripting Support

You can run procedures and create settings described in this chapter in a Tcl script. You can also run some procedures at the command-line prompt.

- For more information about Tcl scripting, refer to the *Tcl Scripting* chapter in volume 2 of the *Quartus II Handbook*.
- For more information about command-line scripting, refer to the *Command-Line Scripting* chapter in volume 2 of the *Quartus II Handbook*.
- For detailed information about scripting command options, refer to the **Qhelp** command line and Tcl API help browser.

To start the Qhelp help browser, type the following command:

```
quartus_sh -qhelp ←
```

Generating a Post-Synthesis Simulation Netlist for Active-HDL or Riviera-PRO

You can use the Quartus II software to generate a post-synthesis simulation netlist with Tcl commands or with a command at the command-line prompt. The following examples assume you are selecting Active-HDL or Riviera-PRO (Verilog HDL output from the Quartus II software).

Tcl Commands

Use the following Tcl commands to set the output format to Verilog HDL, the simulation tool to Active-HDL or Riviera-PRO for Verilog HDL, and to generate a functional netlist:

```
set_global_assignment-name EDA_SIMULATION_TOOL "Active-HDL(Verilog)" ←  
set_global_assignment-name EDA_GENERATE_FUNCTIONAL_NETLIST ON ←
```

or

```
set_global_assignment-name EDA_SIMULATION_TOOL "Riviera-PRO (Verilog)" ←  
set_global_assignment-name EDA_GENERATE_FUNCTIONAL_NETLIST ON ←
```

Command Line

Use the following command to generate a simulation output file for the Active-HDL or Riviera-PRO software. Specify VHDL or Verilog HDL for the format:

```
quartus_eda <project name> --simulation=on --format=<format> \  
--tool=activehdl --functional ←
```

or

```
quartus_eda <project name> --simulation=on --format=<format> \  
--tool=rivierapro --functional ←
```

Generating a Gate-Level Timing Simulation Netlist for Active-HDL or Riviera-PRO

You can use the Quartus II software to generate a gate-level timing simulation netlist with Tcl commands or with a command at the command prompt.

Tcl Commands

Use one of the following Tcl commands:

```
set_global_assignment -name EDA_SIMULATION_TOOL "Active-HDL (Verilog)"  
←  
set_global_assignment -name EDA_SIMULATION_TOOL "Active-HDL (VHDL)" ←  
or  
set_global_assignment -name EDA_SIMULATION_TOOL "Riviera-PRO (Verilog)"  
←  
set_global_assignment -name EDA_SIMULATION_TOOL "Riviera-PRO (VHDL)" ←
```

Command Line

Generate a simulation output file for the Active-HDL or Riviera-PRO software by specifying VHDL or Verilog HDL for the format by typing the following command at the command prompt:

```
quartus_eda <project name> --simulation=on --format=<format> \  
--tool=activehdl ←  
or  
quartus_eda <project name> --simulation=on --format=<format> \  
--tool=rivierapro ←
```

Conclusion

Using the Active-HDL or Riviera-PRO simulation software within the Altera FPGA design flow allows you to easily and accurately perform functional simulations, post-synthesis simulations, and gate-level timing simulations on your designs. Proper verification of designs at the functional, post-synthesis, and post place-and-route stages using the Active-HDL or Riviera-PRO software helps ensure your design functions correctly and, ultimately, a quick time-to-market.

Document Revision History

Table 5-1 shows the revision history for this chapter.

Table 5-1. Document Revision History

Date	Version	Changes
December 2010	10.0.1	<ul style="list-style-type: none"> ■ Changed to new document template. No change to content.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Linked to Quartus II Help ■ Revised simulation procedures ■ Added Stratix V simulation information ■ Added Riviera-PRO support ■ Minor text edits ■ Removed Referenced Documents section
November 2000	9.1.0	<ul style="list-style-type: none"> ■ Updated Table 6-1 ■ Removed Simulation Library tables and EDA Simulation Library Compiler sections and referenced new <i>Simulating Designs with EDA Tools</i> chapter ■ Added “RTL Functional Simulation for Stratix IV Devices” and “Gate-Level Timing Simulation for Stratix IV Devices” sections ■ Minor text edits
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Removed “Compile Libraries Using the Altera Simulation Library Compiler” ■ Added “Compile Libraries Using the EDA Simulation Library Compiler” on page 5-10 ■ Added “Generate Simulation Script from EDA Netlist Writer” on page 5-51 ■ Minor editorial updates
November 2008	8.1.0	<p>Added the following sections:</p> <ul style="list-style-type: none"> ■ “Compile Libraries Using the Altera Simulation Library Compiler” on page 5-10 ■ Added steps to the procedure “Performing an RTL Simulation Using NativeLink” on page 5-45 for using the Altera Simulation Library Compilation ■ Added steps to the procedure “Performing a Gate-Level Timing Simulation Using NativeLink” on page 5-47 for using the Altera Simulation Library Compilation ■ Minor editorial updates ■ Updated entire chapter using 8½” × 11” chapter template
May 2008	8.0.0	Initial release

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

As designs become more complex, advanced timing analysis capability requirements grow. Static timing analysis is a method of analyzing, debugging, and validating the timing performance of a design. The Quartus® II software provides the features necessary to perform advanced timing analysis for today's system-on-a-programmable-chip (SOPC) designs.

Synopsys PrimeTime is an industry standard sign-off tool, used to perform static timing analysis on most ASIC designs. The Quartus II software provides a path to enable you to run PrimeTime on your Quartus II software designs, and export a netlist, timing constraints, and libraries to the PrimeTime environment.

This section explains the basic principles of static timing analysis, the advanced features supported by the Quartus II Timing Analyzer, and how you can use PrimeTime to analyze your Quartus II projects.

This section includes the following chapters:

- [Chapter 6, The Quartus II TimeQuest Timing Analyzer](#)

This chapter describes the Quartus II TimeQuest Timing Analyzer, which is a powerful ASIC-style timing analysis tool that validates the timing performance of all logic in your design using an industry-standard constraint, analysis, and reporting methodology.

- [Chapter 7, Best Practices for the Quartus II TimeQuest Timing Analyzer](#)

This chapter provides the steps to fully constrain an FPGA design with the Quartus II TimeQuest Timing Analyzer.

- [Chapter 8, Switching to the Quartus II TimeQuest Timing Analyzer](#)

This chapter describes the benefits of switching to the Quartus II TimeQuest Timing Analyzer, the differences between the Quartus II TimeQuest and Quartus II Classic Timing Analyzers, and the process you should follow to switch a design from using the Quartus II Classic Timing Analyzer to the Quartus II TimeQuest Timing Analyzer.

- [Chapter 9, Synopsys PrimeTime Support](#)

This chapter describes the PrimeTime software that uses data from either best-case or worst-case Quartus II timing models to measure timing. The PrimeTime software is controlled using a Tcl script generated by the Quartus II software that you can customize to direct the PrimeTime software to produce violation and slack reports.

The Quartus® II TimeQuest Timing Analyzer is a powerful ASIC-style timing analysis tool that validates the timing performance of all logic in your design using an industry-standard constraint, analysis, and reporting methodology. Use the TimeQuest analyzer GUI or command-line interface to constrain, analyze, and report results for all timing paths in your design.

-  For more information about interactive timing analysis, refer to the *TimeQuest Timing Analyzer Quick Start Tutorial*.
-  For more information about Altera resources available for the TimeQuest analyzer, refer to the *TimeQuest Timing Analyzer Resource Center* of the Altera website.
-  For more information about constraining circuits and reporting timing analysis results in the TimeQuest analyzer, including examples, refer to the *TimeQuest Design Examples* page of the Altera website and the *Quartus II TimeQuest Timing Analyzer Cookbook*.
-  For more information about constraining designs, refer to the *Best Practices for the Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

This chapter contains the following sections:

- “Understanding Timing Analysis with the TimeQuest Analyzer” on page 6–1
- “Clock Setup Check” on page 6–5
- “Using Tcl Commands” on page 6–23
- “Creating Clocks and Clock Constraints” on page 6–27
- “Creating I/O Constraints” on page 6–37
- “Creating Delay and Skew Constraints” on page 6–38
- “Creating Timing Exceptions” on page 6–39
- “Timing Reports” on page 6–40

Understanding Timing Analysis with the TimeQuest Analyzer

A comprehensive static timing analysis includes analysis of register-to-register, I/O, and asynchronous reset paths. The TimeQuest analyzer uses data required times, data arrival times, and clock arrival times to verify circuit performance and detect possible timing violations. The TimeQuest analyzer determines the timing relationships that must be met for the design to correctly function, and checks arrival times against required times to verify timing.

Table 6-1 describes TimeQuest analyzer terminology.

Table 6-1. TimeQuest Analyzer Terminology

Term	Definition
nodes	Most basic timing netlist unit. Used to represent ports, pins, and registers.
keepers	Ports or registers. (1)
cells	Look-up tables (LUT), registers, digital signal processing (DSP) blocks, memory blocks, input/output elements, and so on. (2)
pins	Inputs or outputs of cells.
nets	Connections between pins.
ports	Top-level module inputs or outputs; for example, device pins.
clocks	Abstract objects outside of the design.

Notes to Table 6-1:

- (1) Pins can indirectly refer to keepers. For example, a pin refers to a keeper when the value in the `-from` field of a constraint is a clock pin to a dedicated memory block. In this case, the clock pin refers to a collection of registers.
- (2) For Stratix® devices, the LUTs and registers are contained in logic elements (LE) and act as cells.

Timing Netlists and Timing Paths

The TimeQuest analyzer requires a timing netlist to perform timing analysis on any design. After you generate a timing netlist, the TimeQuest analyzer uses the data to help determine the different design elements in your design and how to optimize timing.

The Timing Netlist

Figure 6-1 shows an example design for which the TimeQuest analyzer can generate a netlist equivalent.

Figure 6-1. Sample Design

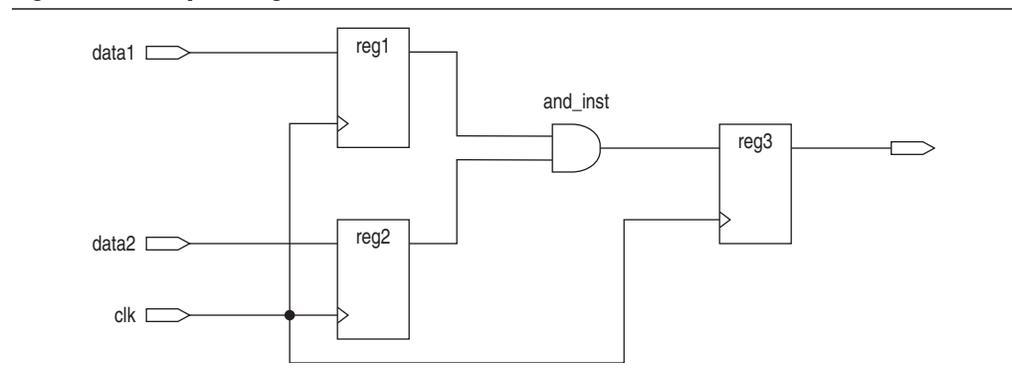
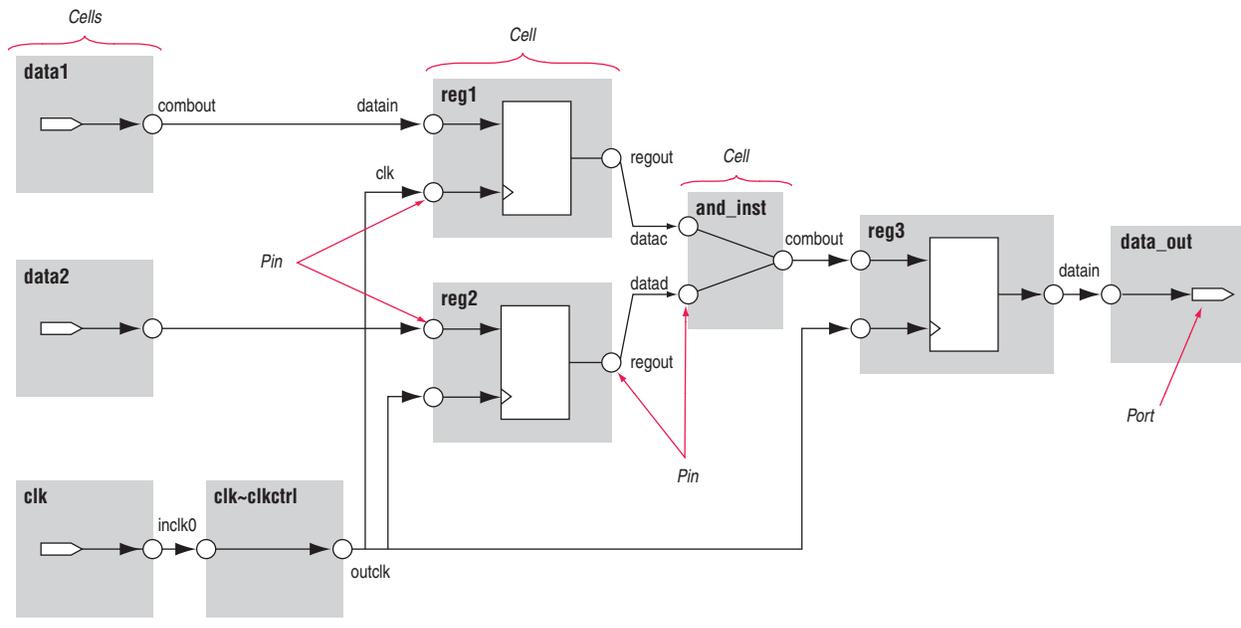


Figure 6-2 shows the timing netlist for the design, including how different design elements are divided into various cells, pins, nets, and ports.

Figure 6-2. The TimeQuest Analyzer Timing Netlist



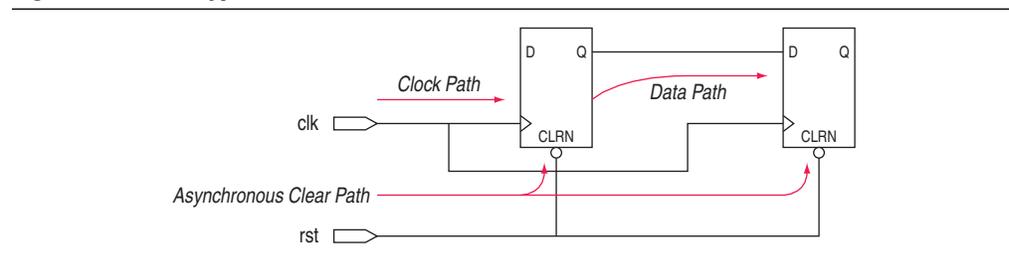
Timing Paths

Timing paths connect two design nodes, such as the output of a register to the input of another register. Timing paths play a significant role in timing analysis. Understanding the types of timing paths is important to timing closure and optimization. The TimeQuest analyzer uses the following commonly analyzed paths:

- **Edge paths**—connections from ports-to-pins, from pins-to-pins, and from pins-to-ports.
- **Clock paths**—connections from device ports or internally generated clock pins to the clock pin of a register.
- **Data paths**—connections from a port or the data output pin of a sequential element to a port or the data input pin of another sequential element.
- **Asynchronous paths**—connections from a port or sequential element to the asynchronous reset or asynchronous clear pin of another sequential element.

Figure 6-3 shows path types commonly analyzed by the TimeQuest analyzer.

Figure 6-3. Path Types



In addition to identifying various paths in a design, the TimeQuest analyzer analyzes clock characteristics to compute the worst-case requirement between any two registers in a single register-to-register path. You should constrain all clocks in your design before analyzing clock characteristics.

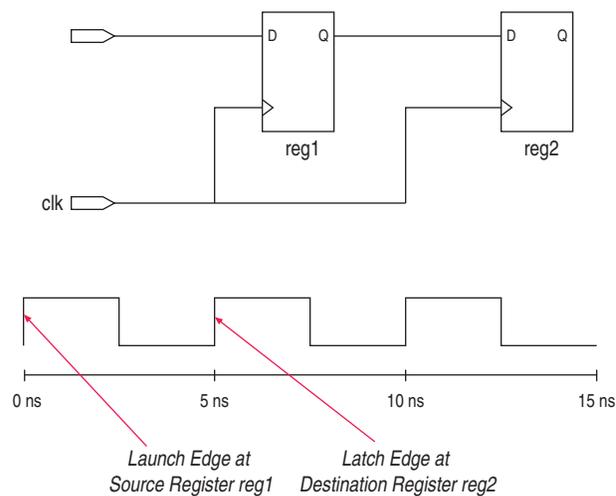
Launch and Latch Edges

Figure 6-4 shows a single-cycle system that uses consecutive clock edges to transfer and capture data, a register-to-register path, and the corresponding timing diagram for the launch and latch edges. The launch edge is an active clock edge that sends data out of a sequential element, acting as a source for the data transfer. A latch edge is the active clock edge that captures data at the data port of a sequential element, acting as a destination for the data transfer. In this example, the launch edge sends the data out of register reg1 at 0 ns, and the register reg2 latch edge captures the data at 5 ns.



The TimeQuest analyzer validates clock setup and hold requirements relative to the launch and latch edges.

Figure 6-4. Launch Edge and Latch Edge

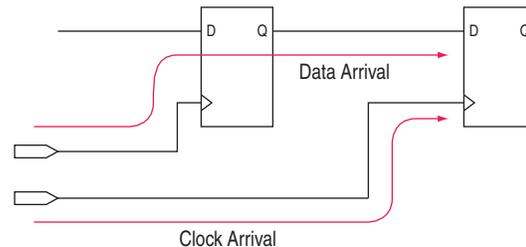


Data and Clock Arrival Times

After the TimeQuest analyzer identifies the path type, it can report data and clock arrival times at register pins. The TimeQuest analyzer calculates data arrival time by adding the delay from the clock source to the clock pin of the source register, the micro clock-to-output delay (μt_{CO}) of the source register, and the delay from the source register's Q pin to the destination register's D pin, where the μt_{CO} is the intrinsic clock-to-out for the internal registers in the FPGA.

The TimeQuest analyzer calculates clock arrival time by adding the delay from the clock source to the clock pin of the destination register. Figure 6-5 shows a data arrival path and a clock arrival path. The TimeQuest analyzer calculates the data required time by accounting for the clock arrival time and micro setup time (μt_{SU}) of the destination register, where the μt_{SU} is the intrinsic setup for the internal registers in the FPGA.

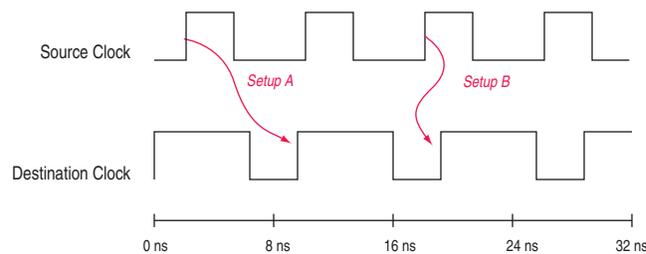
Figure 6-5. Data Arrival and Clock Arrival Paths



Clock Setup Check

To perform a clock setup check, the TimeQuest analyzer determines a setup relationship by analyzing each launch and latch edge for each register-to-register path. For each latch edge at the destination register, the TimeQuest analyzer uses the closest previous clock edge at the source register as the launch edge. Figure 6-6 shows two setup relationships, setup A and setup B. For the latch edge at 10 ns, the closest clock that acts as a launch edge is at 3 ns and is labeled setup A. For the latch edge at 20 ns, the closest clock that acts as a launch edge is 19 ns and is labeled setup B.

Figure 6-6. Setup Check



The TimeQuest analyzer reports the result of clock setup checks as slack values. Slack is the margin by which a timing requirement is met or not met. Positive slack indicates the margin by which a requirement is met; negative slack indicates the margin by which a requirement is not met. Equation 6-1 shows the TimeQuest analyzer clock setup slack time calculation for internal register-to-register paths.

Equation 6-1. Clock Setup Slack for Internal Register-to-Register paths

$$\text{Clock Setup Slack} = \text{Data Required Time} - \text{Data Arrival Time}$$

$$\text{Data Arrival Time} = \text{Launch Edge} + \text{Clock Network Delay to Source Register} + \mu t_{CO} + \text{Register-to-Register Delay}$$

$$\text{Data Required} = \text{Clock Arrival Time} - \mu t_{SU} - \text{Setup Uncertainty}$$

$$\text{Clock Arrival Time} = \text{Latch Edge} + \text{Clock Network Delay to Destination Register}$$

Equation 6-2 shows the TimeQuest analyzer clock setup slack time calculation if the data path is from an input port to an internal register.

Equation 6-2. Clock Setup Slack from Input Port to Internal Register

Clock Setup Slack = Data Required Time – Data Arrival Time

Data Arrival Time = Launch Edge + Clock Network Delay +

Input Maximum Delay + Pin-to-Register Delay

Data Required Time = Latch Edge + Clock Network Delay to Destination Register – μt_{SU}

Equation 6-3 shows the TimeQuest analyzer clock setup slack time calculation if the data path is an internal register to an output port.

Equation 6-3. Clock Setup Slack from Internal Register to Output Port

Clock Setup Slack = Data Required Time – Data Arrival Time

Data Arrival Time = Launch Edge + Clock Network Delay to Source Register +

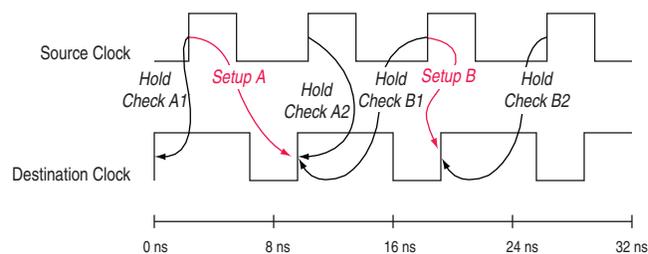
μt_{CO} + Register-to-Pin Delay

Data Required Time = Latch Edge + Clock Network Delay – Output Maximum Delay of Pin

Clock Hold Check

To perform a clock hold check, the TimeQuest analyzer determines a hold relationship for each possible setup relationship that exists for all source and destination register pairs. The TimeQuest analyzer checks all adjacent clock edges from all setup relationships to determine the hold relationships. The TimeQuest analyzer performs two hold checks for each setup relationship. The first hold check determines that the data launched by the current launch edge is not captured by the previous latch edge. The second hold check determines that the data launched by the next launch edge is not captured by the current latch edge. From the possible hold relationships, the TimeQuest analyzer selects the hold relationship that is the most restrictive. The most restrictive hold relationship is the hold relationship with the smallest difference between the latch and launch edges and determines the minimum allowable delay for the register-to-register path. Figure 6-7 shows two setup relationships, setup A and setup B, and their respective hold checks. In this example, the TimeQuest analyzer selects hold check A2 as the most restrictive hold relationship.

Figure 6-7. Hold Checks



Equation 6-4 shows the TimeQuest analyzer clock hold slack time calculation.

Equation 6-4. Clock Hold Slack

$$\begin{aligned} \text{Clock Hold Slack Time} &= \text{Data Arrival Time} - \text{Data Required Time} \\ \text{Arrival Time} &= \text{Launch Edge} + \text{Clock Network Delay to Source Register} + \mu t_{\text{CO}} + \\ &\quad \text{Register-to-Register Delay} \\ \text{Data Required Time} &= \text{Clock Arrival Time} + \mu t_{\text{H}} + \text{Hold Uncertainty} \\ \text{Clock Arrival Time} &= \text{Latch Edge} + \text{Clock Network Delay to Destination Register} \end{aligned}$$

Equation 6-5 shows the TimeQuest analyzer hold slack time calculation if the data path is from an input port to an internal register.

Equation 6-5. Clock Hold Slack from Input Port to Internal Register

$$\begin{aligned} \text{Clock Hold Slack Time} &= \text{Data Arrival Time} - \text{Data Required Time} \\ \text{Data Arrival Time} &= \text{Launch Edge} + \text{Clock Network Delay} + \\ &\quad \text{Input Minimum Delay of Pin} + \text{Pin-to-Register Delay} \\ \text{Data Required Time} &= \text{Latch Edge} + \text{Clock Network Delay to Destination Register} + \mu t_{\text{H}} \end{aligned}$$

Equation 6-6 shows the TimeQuest analyzer hold slack time calculation if the data path is from an internal register to an output port.

Equation 6-6. Clock Hold Slack from Internal Register to Output Port

$$\begin{aligned} \text{Clock Hold Slack Time} &= \text{Data Arrival Time} - \text{Data Required Time} \\ \text{Data Arrival Time} &= \text{Latch Edge} + \text{Clock Network Delay to Source Register} + \mu t_{\text{CO}} + \\ &\quad \text{Register-to-Pin Delay} \\ \text{Data Required Time} &= \text{Latch Edge} + \text{Clock Network Delay} - \text{Output Minimum Delay of Pin} \end{aligned}$$

Recovery and Removal Time

Recovery time is the minimum length of time for the deassertion of an asynchronous control signal; for example, signals such as `clear` and `preset` must be stable before the next active clock edge. The recovery slack calculation is similar to the clock setup slack calculation, but it applies to asynchronous control signals. Equation 6-7 shows the TimeQuest analyzer recovery slack time calculation if the asynchronous control signal is registered.

Equation 6-7. Recovery Slack if Asynchronous Control Signal Registered

$$\begin{aligned} \text{Recovery Slack Time} &= \text{Data Required Time} - \text{Data Arrival Time} \\ \text{Data Arrival Time} &= \text{Launch Edge} + \text{Clock Network Delay to Source Register} + \\ &\quad \mu t_{\text{CO}} + \text{Register-to-Register Delay} \\ \text{Data Required Time} &= \text{Latch Edge} + \text{Clock Network Delay to Destination Register} - \mu t_{\text{SU}} \end{aligned}$$

Equation 6-8 shows the TimeQuest analyzer recovery slack time calculation if the asynchronous control signal is not registered.

Equation 6-8. Recovery Slack if Asynchronous Control Signal not Registered

$$\begin{aligned} \text{Recovery Slack Time} &= \text{Data Required Time} - \text{Data Arrival Time} \\ \text{Data Arrival Time} &= \text{Launch Edge} + \text{Clock Network Delay} + \text{Maximum Input Delay} + \\ &\quad \text{Port-to-Register Delay} \\ \text{Data Required Time} &= \text{Latch Edge} + \text{Clock Network Delay to Destination Register Delay} - \mu t_{\text{SU}} \end{aligned}$$

 If the asynchronous reset signal is from a device I/O port, you must assign the **Input Maximum Delay** timing assignment to the asynchronous reset port for the TimeQuest analyzer to perform recovery analysis on the path.

Removal time is the minimum length of time the deassertion of an asynchronous control signal must be stable after the active clock edge. The TimeQuest analyzer removal slack calculation is similar to the clock hold slack calculation, but it applies asynchronous control signals. Equation 6-9 shows the TimeQuest analyzer removal slack time calculation if the asynchronous control signal is registered.

Equation 6-9. Removal Slack if Asynchronous Control Signal Registered

$$\begin{aligned} \text{Removal Slack Time} &= \text{Data Arrival Time} - \text{Data Required Time} \\ \text{Data Arrival Time} &= \text{Launch Edge} + \text{Clock Network Delay to Source Register} + \\ &\quad \mu t_{\text{CO}} \text{ of Source Register} + \text{Register-to-Register Delay} \\ \text{Data Required Time} &= \text{Latch Edge} + \text{Clock Network Delay to Destination Register} + \mu t_{\text{H}} \end{aligned}$$

Equation 6-10 shows the TimeQuest analyzer removal slack time calculation if the asynchronous control signal is not registered.

Equation 6-10. Removal Slack if Asynchronous Control Signal not Registered

$$\begin{aligned} \text{Removal Slack Time} &= \text{Data Arrival Time} - \text{Data Required Time} \\ \text{Data Arrival Time} &= \text{Launch Edge} + \text{Clock Network Delay} + \text{Input Minimum Delay of Pin} + \\ &\quad \text{Minimum Pin-to-Register Delay} \\ \text{Data Required Time} &= \text{Latch Edge} + \text{Clock Network Delay to Destination Register} + \mu t_{\text{H}} \end{aligned}$$

 If the asynchronous reset signal is from a device pin, you must assign the **Input Minimum Delay** timing assignment to the asynchronous reset pin for the TimeQuest analyzer to perform removal analysis on the path.

Multicycle Paths

Multicycle paths are data paths that require more than one clock cycle to latch data at the destination register. For example, a register may be required to capture data on every second or third rising clock edge. Figure 6-8 shows an example of a multicycle path between the input registers of a multiplier and an output register where the destination latches data on every other clock edge.

Figure 6-8. Multicycle Path

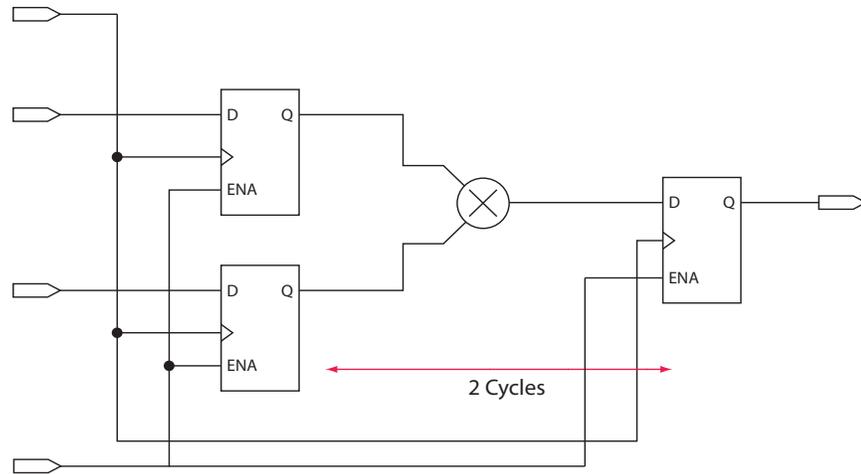
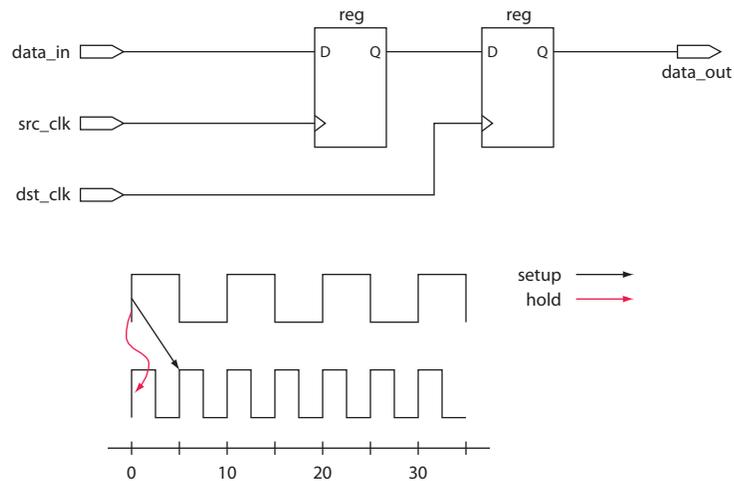


Figure 6-9 shows a register-to-register path used for the default setup and hold relationship, the respective timing diagrams for the source and destination clocks, and the default setup and hold relationships, when the source clock, `src_clk`, has a period of 10 ns and the destination clock, `dst_clk`, has a period of 5 ns. The default setup relationship is 5 ns; the default hold relationship is 0 ns.

Figure 6-9. Register-to-Register Path and Default Setup and Hold Timing Diagram



To accommodate the system requirements you can modify the default setup and hold relationships with the `set_multicycle_path` command.

Figure 6-10 shows the timing diagram after the TimeQuest analyzer uses a multicycle setup assignment of two. The command moves the latch edge time to 10 ns from the default value of 5 ns.

Figure 6-10. Modified Setup Diagram

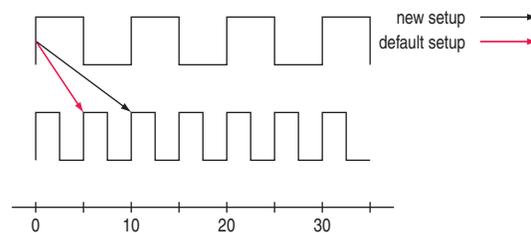


Table 6-2 shows the commands you can use to modify either the launch or latch edge times that the TimeQuest analyzer uses to determine a setup relationship or hold relationship.

Table 6-2. Commands to Modify Edge Times

Command	Description of Modification
<code>set_multicycle_path -setup -end</code>	Latch edge time of the setup relationship
<code>set_multicycle_path -setup -start</code>	Launch edge time of the setup relationship
<code>set_multicycle_path -hold -end</code>	Latch edge time of the hold relationship
<code>set_multicycle_path -hold -start</code>	Launch edge time of the hold relationship

- For more information about the `set_multicycle_path` command—including full syntax information, options, and example usage—refer to [set_multicycle_path](#) in Quartus II Help.

Metastability

Metastability problems can occur when a signal is transferred between circuitry in unrelated or asynchronous clock domains because the designer cannot guarantee that the signal will meet setup and hold time requirements. To minimize the failures due to metastability, circuit designers typically use a sequence of registers, also known as a synchronization register chain, or synchronizer, in the destination clock domain to resynchronize the data signals to the new clock domain.

The mean time between failures (MTBF) is an estimate of the average time between instances of failure due to metastability.

The TimeQuest analyzer analyzes the robustness of your design for metastability and can calculate the MTBF for synchronization register chains in your design. The MTBF of the entire design is then estimated based on the synchronization chains it contains.

In addition to reporting synchronization register chains found in the design, the Quartus II software also protects these registers from optimizations that might negatively impact MTBF, such as register duplication and logic retiming. The Quartus II software can also optimize the MTBF of your design if the MTBF is too low.

- For more information about metastability, its effects in FPGAs, and how MTBF is calculated, refer to the [Understanding Metastability in FPGAs](#) white paper. For more information about metastability analysis, reporting, and optimization features in the Quartus II software, refer to the [Managing Metastability with the Quartus II Software](#) chapter in volume 1 of the *Quartus II Handbook*.

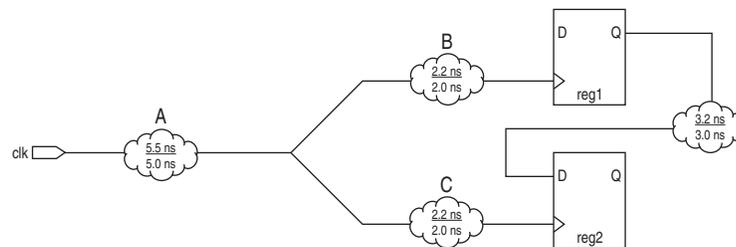
Common Clock Path Pessimism Removal

Common clock path pessimism removal accounts for the minimum and maximum delay variation associated with common clock paths during static timing analysis by adding the difference between the maximum and minimum delay value of the common clock path to the appropriate slack equation.

Minimum and maximum delay variation can occur when two different delay values are used for the same clock path. For example, in a simple setup analysis, the maximum clock path delay to the source register is used to determine the data arrival time. The minimum clock path delay to the destination register is used to determine the data required time. However, if the clock path to the source register and to the destination register share a common clock path, both the maximum delay and the minimum delay are used to model the common clock path during timing analysis. The use of both the minimum delay and maximum delay results in an overly pessimistic analysis since two different delay values, the maximum and minimum delays, cannot be used to model the same clock path.

Figure 6-11 shows a typical register-to-register path with the maximum and minimum delay values shown.

Figure 6-11. Common Clock Path

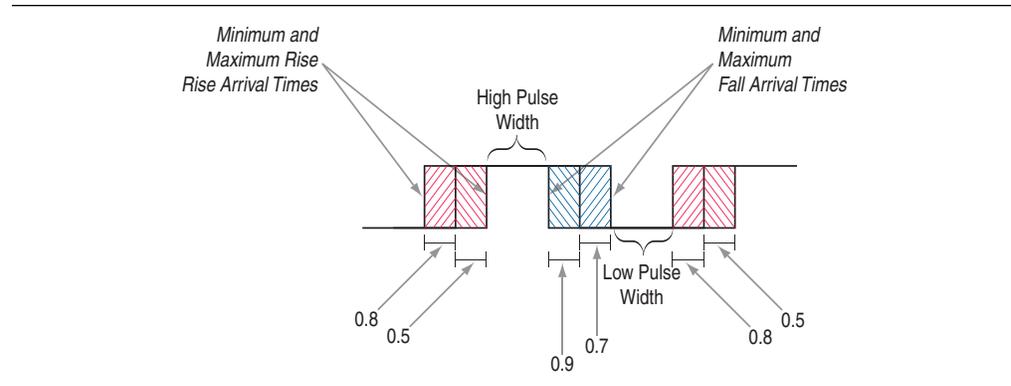


Segment A is the common clock path between reg1 and reg2. The minimum delay is 5.0 ns; the maximum delay is 5.5 ns. The difference between the maximum and minimum delay value equals the common clock path pessimism removal value; in this case, the common clock path pessimism is 0.5 ns. The TimeQuest analyzer adds the common clock path pessimism removal value to the appropriate slack equation to determine overall slack. Therefore, if the setup slack for the register-to-register path in Figure 6-11 equals 0.7 ns without common clock path pessimism removal, the slack would be 1.2 ns with common clock path pessimism removal.

You can also use common clock path pessimism removal to determine the minimum pulse width of a register. A clock signal must meet a register's minimum pulse width requirement to be recognized by the register. A minimum high time defines the minimum pulse width for a positive-edge triggered register. A minimum low time defines the minimum pulse width for a negative-edge triggered register.

Clock pulses that violate the minimum pulse width of a register prevent data from being latched at the data pin of the register. To calculate the slack of the minimum pulse width, the TimeQuest analyzer subtracts the required minimum pulse width time from the actual minimum pulse width time. The TimeQuest analyzer determines the actual minimum pulse width time by the clock requirement you specified for the clock that feeds the clock port of the register. The TimeQuest analyzer determines the required minimum pulse width time by the maximum rise, minimum rise, maximum fall, and minimum fall times. Figure 6-12 shows a diagram of the required minimum pulse width time for both the high pulse and low pulse.

Figure 6-12. Required Minimum Pulse Width



With common clock path pessimism, the minimum pulse width slack can be increased by the smallest value of either the maximum rise time minus the minimum rise time, or the maximum fall time minus the minimum fall time. For Figure 6-12, the slack value can be increased by 0.2 ns, which is the smallest value between 0.3 ns (0.8 ns - 0.5 ns) and 0.2 ns (0.9 ns - 0.7 ns).

To remove common clock path pessimism with the TimeQuest analyzer, follow these steps:

1. In the Quartus II software, on the Assignments menu, click **Settings**.
2. In the **Category** list, click **TimeQuest Timing Analyzer**.
3. Turn on **Enable common clock path pessimism removal**.
4. Click **OK**.

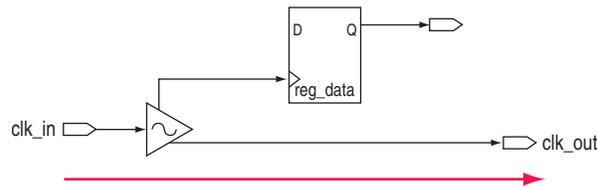
 For more information, refer to [TimeQuest Timing Analyzer Page](#) in Quartus II Help.

Clock-As-Data Analysis

The majority of FPGA designs contain simple connections between any two nodes known as either a data path or a clock path. A data path is a connection between the output of a synchronous element to the input of another synchronous element. A clock is a connection to the clock pin of a synchronous element. However, for more complex FPGA designs, such as designs that use source-synchronous interfaces, this simplified view is no longer sufficient.

The connection between the input clock port and output clock port can be treated either as a clock path or a data path. Figure 6-13 shows a design where the path from port `clk_in` to port `clk_out` is both a clock and a data path. The clock path is from the port `clk_in` to the register `reg_data` clock pin. The data path is from port `clk_in` to the port `clk_out`.

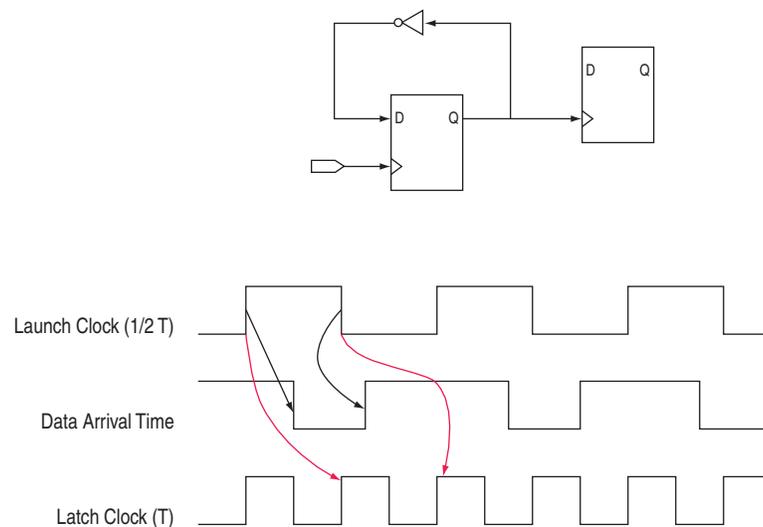
Figure 6-13. Simplified Source Synchronous Output



With clock-as-data analysis, the TimeQuest analyzer provides a more accurate analysis of the path based on user constraints. For the clock path analysis, any phase shift associated with the phase-locked loop (PLL) is taken into consideration. For the data path analysis, any phase shift associated with the PLL is taken into consideration rather than ignored.

The clock-as-data analysis also applies to internally generated clock dividers. Figure 6-14 shows an internally generated clock divider. A source-synchronous interface contains a clock signal that travels in parallel with data signals. The clock and data pair originates or terminates at the same device.

Figure 6-14. Clock Divider (Note 1)



Note to Figure 6-14:

- (1) In this figure, the inverter feedback path is analyzed during timing analysis. The output of the divider register is used to determine the launch time and the clock port of the register is used to determine the latch time.

Multicorner Analysis

The TimeQuest analyzer performs multicorner timing analysis to verify your design under a variety of operating conditions—such as voltage, process, and temperature—while performing static timing analysis.

To change the operating conditions or speed grade of the device used for static timing analysis, use the `set_operating_conditions` command.

If you specify an operating condition Tcl object, the `-model`, `speed`, `-temperature`, and `-voltage` options are optional. If you do not specify an operating condition Tcl object, the `-model` option is required; the `-speed`, `-temperature`, and `-voltage` options are optional.



To obtain a list of available operating conditions for the target device, use the `get_available_operating_conditions -all` command.



For more information about the `set_operating_conditions` and `get_available_operating_conditions` commands—including full syntax information, options, and example usage—refer to [set_operating_conditions](#) and [get_available_operating_conditions](#) in Quartus II Help.

To ensure that no violations occur under various conditions during the device operation, perform static timing analysis under all available operating conditions. [Table 6-3](#) shows the operating conditions for the slow and fast timing models for device families that support only slow and fast operating conditions.

Table 6-3. Operating Conditions for Slow and Fast Models

Model	Speed Grade	Voltage	Temperature
Slow	Slowest speed grade in device density	V_{cc} minimum supply (1)	Maximum T_J (1)
Fast	Fastest speed grade in device density	V_{cc} maximum supply (1)	Minimum T_J (1)

Note to Table 6-3:

(1) Refer to the DC & Switching Characteristics chapter of the applicable device Handbook for V_{cc} and T_J values

[Example 6-1](#) shows how to set the operating conditions for a Stratix III design to the slow timing model, with a voltage of 1100 mV, and temperature of 85° C.

Example 6-1. Setting Operating Conditions with Individual Values

```
set_operating_conditions -model slow -temperature 85 -voltage 1100
```

[Example 6-2](#) shows how to set the operating conditions in [Example 6-1](#) with a Tcl object.

Example 6-2. Setting Operating Conditions with a Tcl Object

```
set_operating_conditions 4_slow_1100mv_85c
```

Example 6-3 shows how to use the `set_operating_conditions` command to generate different reports for various operating conditions.

Example 6-3. Script Excerpt for Analysis of Various Operating Conditions

```
#Specify initial operating conditions
set_operating_conditions -model slow -speed 3 -grade c -temperature 85
-voltage 1100

#Update the timing netlist with the initial conditions
update_timing_netlist

#Perform reporting

#Change initial operating conditions. Use a temperature of 0C
set_operating_conditions -model slow -speed 3 -grade c -temperature 0
-voltage 1100

#Update the timing netlist with the new operating condition
update_timing_netlist

#Perform reporting

#Change initial operating conditions. Use a temperature of 0C and a
model of fast
set_operating_conditions -model fast -speed 3 -grade c -temperature 0
-voltage 1100

#Update the timing netlist with the new operating condition
update_timing_netlist

#Perform reporting
```

Getting Started with the TimeQuest Analyzer

This section provides a brief overview of the design steps necessary to set up your project for timing and analysis and the steps to perform to constrain your design with the TimeQuest analyzer.

Running the TimeQuest Analyzer

You can run the TimeQuest analyzer in the following ways:

- Directly from the Quartus II software GUI
- As a stand-alone GUI application
- At a system command prompt

For more information about prerequisite steps to perform before opening the TimeQuest analyzer, refer to [“Recommended Flows” on page 6-19](#).

To run the TimeQuest analyzer from the Quartus II software, on the Tools menu, click **TimeQuest Timing Analyzer**.

To run the TimeQuest analyzer as a stand-alone application, type the following command at the command prompt:

```
quartus_staw ←
```

- ❓ For more information about the TimeQuest analyzer GUI, refer to [About TimeQuest Timing Analysis](#) in Quartus II Help.

To run the TimeQuest analyzer in command-line mode for easy integration with scripted design flows, type the following command at a system command prompt:

```
quartus_sta ←
```

For more information about using Tcl commands to constrain and analyze your design, refer to [“Using Tcl Commands” on page 6-23](#).

[Table 6-4](#) shows a summary of the command-line options available in command-line mode.

Table 6-4. Summary of Command-Line Options (Part 1 of 2)

Command-Line Option	Description
-h --help	Provides help information on <code>quartus_sta</code> .
-t <script file> --script=<script file>	Sources the <script file>.
-s --shell	Enters shell mode.
--tcl_eval <tcl command>	Evaluates the Tcl command <tcl command>.
--do_report_timing	For all clocks in the design, run the following commands: report_timing -npaths 1 -to_clock \$clock report_timing -setup -npaths 1 -to_clock \$clock report_timing -hold -npaths 1 -to_clock \$clock report_timing -recovery -npaths 1 -to_clock \$clock report_timing -removal -npaths 1 -to_clock \$clock
--force_dat	Forces the delay annotator to annotate the new delays from the recently compiled design to the compiler database.
--lower_priority	Lowers the computing priority of the <code>quartus_sta</code> process.
--post_map	Uses the post-map database results.
--qsf2sdc	Converts assignments from the Quartus II Settings File (.qsf) format to the .sdc format.

Table 6-4. Summary of Command-Line Options (Part 2 of 2)

Command-Line Option	Description
--sdc=<SDC file>	Specifies the .sdc to read.
--report_script=<script>	Specifies a custom report script to call.
--speed=<value>	Specifies the device speed grade used for timing analysis.
--tq2hc	Generate temporary files to convert the TimeQuest analyzer .sdc file(s) to a PrimeTime .sdc that can be used by the HardCopy® Design Center.
--tq2pt	Generates temporary files to convert the TimeQuest Timing Analyzer .sdc file(s) to a PrimeTime .sdc.
-f <argument file>	Specifies a file containing additional command-line arguments.
-c <revision name> --rev=<revision_name>	Specifies which revision and its associated .qsf to use.
--multicorner	Specifies that all slack summary reports be generated for both slow- and fast-corners.
--multicorner[=on off]	Turns off the multicorner timing analysis.
--voltage=<value_in_mV>	Specifies the device voltage, in mV, used in timing analysis.
--temperature= <value_in_C>	Specifies the device temperature in degrees Celsius, used in timing analysis.
--parallel [=<num_processors>]	Specifies the number of computer processors to use on a multiprocessor system.
--64bit	Enables 64-bit version of the executable.

Locating Timing Paths in Other Tools

You can locate paths and elements from the TimeQuest analyzer to other tools in the Quartus II software. Use the **Locate Path** command in the TimeQuest analyzer GUI or the `locate` command-line command to locate paths.

- ❓ For more information about locating paths from the TimeQuest analyzer, refer to [Viewing Timing Analysis Results](#) and [locate](#) in Quartus II Help.

Example 6-4 shows how to locate ten paths from TimeQuest analyzer to the Chip Planner and locate all data ports in the Technology Map Viewer.

Example 6-4. Locating from the TimeQuest Analyzer

```
# Locate in the Chip Planner all of the nodes in the longest ten paths
locate [get_path -npaths 10] -chip

# locate all ports that begin with data to the Technology Map Viewer
locate [get_ports data*] -tmv
```

Recommended Flows

Figure 6-15 shows the recommended design flow for setting up your design for timing analysis with the TimeQuest analyzer.

Figure 6-15. Design Flow with the TimeQuest Timing Analyzer

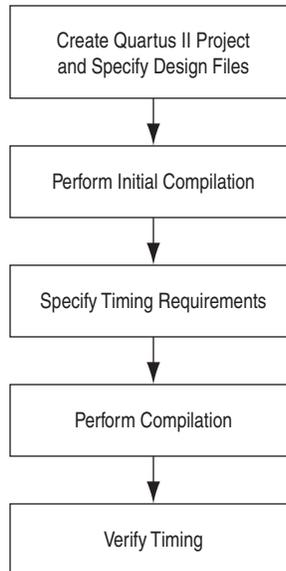
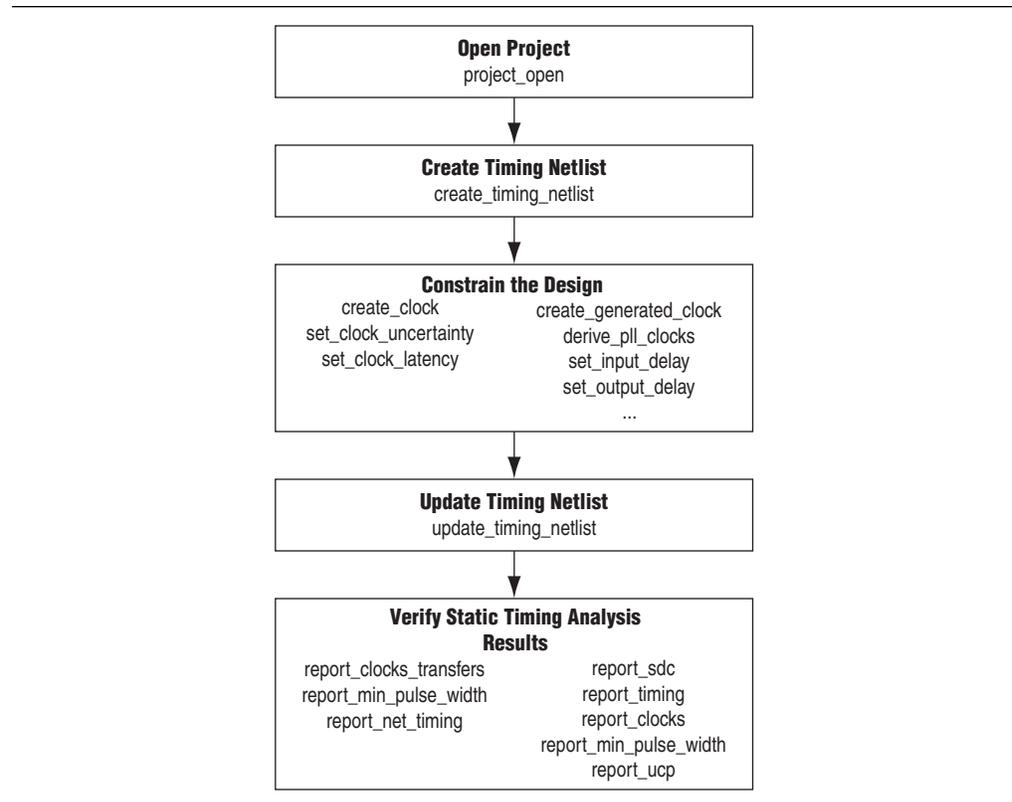


Figure 6–16 shows the recommended flow for constraining and analyzing your design with the TimeQuest analyzer. Included are the corresponding Tcl commands for each step.

Figure 6–16. The TimeQuest Timing Analyzer Flow



Creating and Setting Up your Design

You must first create your project in the Quartus II software. Make sure to include all the necessary device files, including any existing Synopsys Design Constraints (.sdc) files that contain timing constraints for your design.

-  If you previously created and specified an .sdc for your project, you should perform a full compilation to create a post-fit database.
-  For more information, refer to *Managing Files in a Project* in Quartus II Help.

Performing an Initial Compilation

After your project is set up, you must compile your design to create an initial design database before you specify timing constraints. You can either perform Analysis and Synthesis to create a post-map database, or perform a full compilation to create a post-fit database. Creating a post-map database reduces processing time and is sufficient for creating initial timing constraints. The type of database you create determines the type of timing netlist generated by the TimeQuest analyzer, a post-map netlist if you perform Analysis and Synthesis or a post-fit netlist if you perform a full compilation.

-  If you want to create a post-map database and your design uses incremental compilation, you must merge your design partitions before performing Analysis and Synthesis.
-  For more information, refer to *Setting up and Running Analysis and Synthesis* and *Setting up and Running a Compilation* in Quartus II Help.

Specifying Timing Requirements

Before running timing analysis with the TimeQuest analyzer, you must first create a timing netlist, and then you can specify initial timing constraints that describe the clock characteristics, timing exceptions, and signal transition arrival and required times. You can use the TimeQuest Timing Analyzer Wizard to enter basic, initial constraints for your design, or you can specify timing constraints with the dialog boxes in the TimeQuest analyzer or with a Tcl script. The timing constraints you create with the TimeQuest analyzer are saved in the **.sdc** format.

-  For more information, refer to *Specifying Timing Constraints and Exceptions, Running a Timing Analysis*, and *TimeQuest Timing Analyzer Wizard* in Quartus II Help.
-  When you create timing constraints with the TimeQuest analyzer GUI, the **.sdc** is not automatically updated. To write your constraints to an **.sdc**, use the `write_sdc` command in command-line mode, or the **Write SDC File** command in the TimeQuest analyzer GUI. Writing constraints to an existing **.sdc** overwrites the existing file. After editing timing constraints in your design, if you want to save the constraints to an **.sdc**, you should create a new file rather than overwriting the existing file.

You can also use an existing **.sdc** rather than creating new timing constraints. To use timing constraints from an existing **.sdc** and any SDC timing constraints embedded in your HDL files, use the `read_sdc` command in command-line mode, or the **Read SDC File** in the TimeQuest analyzer GUI.

-  For more information, refer to *read_sdc*, *write_sdc*, *Write SDC File Command*, and *Read SDC File Command* in Quartus II Help.

The **.sdc** should contain only SDC commands; commands to manipulate the timing netlist or control the compilation flow should be run as part of a separate Tcl script. After you create timing constraints, you must update the timing netlist. The TimeQuest analyzer applies all constraints to the netlist for verification and removes any invalid or false paths in the design from verification.

-  The constraints in the **.sdc** are read in the order in which they are declared. You must first declare a constraint before making any references to it. For example, if a generated clock references a base clock, the base clock constraint must be declared before the generated clock constraint.

Fitting and Timing Analysis with .sdc Files

You can specify the same or different .sdc files for the Quartus II Fitter to use during the place-and-route process, and the TimeQuest analyzer for static timing analysis. Using different .sdc files allows you to have one set of constraints for the place-and-route process and another set of constraints for final timing sign-off in the TimeQuest analyzer.

To specify an .sdc for the Fitter, you must add the .sdc to your project. The Fitter optimizes your design based on the requirements in the .sdc.

- For more information, refer to *Managing Files in a Project* in Quartus II Help.

Performing a Full Compilation

After creating initial timing constraints, you must fully compile your design. During full compilation the Fitter optimizes the placement of logic to meet your constraints. When compilation is complete, you can open the TimeQuest analyzer to verify timing results and to generate summary, clock setup and clock hold, recovery, and removal reports for all defined clocks in the design.

Verifying Timing

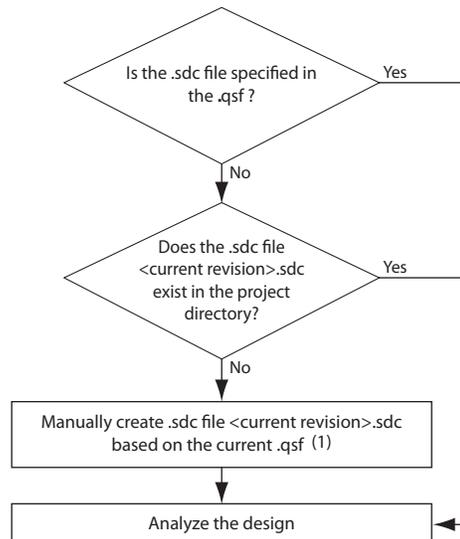
During timing analysis, the TimeQuest analyzer analyzes the timing paths in the design, calculates the propagation delay along each path, checks for timing constraint violations, and reports timing results as slack. If the TimeQuest analyzer reports any timing violations, you can view precise timing information about specific paths by customize the reporting, and then constrain those paths to correct the violations. As you verify timing, you might encounter failures along critical paths. If you encounter failures along critical paths, use the timing reports to analyze your design and determine how best to optimize your design. If you modify, remove, or add constraints, you should perform a full compilation. This iterative process allows you to resolve your timing violations in the design.

SDC File Precedence

The Fitter and the TimeQuest analyzer read the .sdc files from the files list in the .qsf in the order they are listed, from top to bottom.

Figure 6-17 shows the order in which the Quartus II software searches for an `.sdc`.

Figure 6-17. Synopsys Design Constraints File Order of Precedence



Note to Figure 6-17:

(1) This flow occurs only in the TimeQuest analyzer and not during compilation in the Quartus II software. The TimeQuest analyzer has the ability to automate the conversion of the `.qsf` timing assignments to SDC if no `.sdc` exists when you open the TimeQuest analyzer.

 If you type the `read_sdc` command at the command line without any arguments, the TimeQuest analyzer follows precedence order shown in Figure 6-17.

 For more information, refer to *Viewing Timing Analysis Results* in Quartus II Help.

Using Tcl Commands

You can use Tcl commands from the Quartus II software Tcl Application Programming Interface (API) to constrain and analyze your design, and to collect information about your design. This chapter focuses on completing timing analysis tasks with Tcl commands, however; you can perform many of the same functions in the TimeQuest analyzer GUI. You can use standard SDC Tcl commands and SDC extension commands in addition to TimeQuest analyzer Tcl commands. The following Tcl packages are available in the Quartus II software:

- `::quartus::sta`
- `::quartus::sdc`
- `::quartus::sdc_ext`

 For more information about TimeQuest analyzer Tcl commands and a complete list of commands, refer to `::quartus::sta` in Quartus II Help. For more information about standard SDC commands and a complete list of commands, refer to `::quartus::sdc` in Quartus II Help. For more information about Altera extensions of SDC commands and a complete list of commands, refer to `::quartus::sdc_ext` in Quartus II Help.

Wildcard Characters

To apply constraints to many nodes in a design, use the “*” and “?” wildcard characters. The “*” wildcard character matches any string; the “?” wildcard character matches any single character.

If you make an assignment to node `reg*`, the TimeQuest analyzer searches for and applies the assignment to all design nodes that match the prefix `reg` with any number of following characters, such as `reg`, `reg1`, `reg[2]`, `regbank`, and `reg12bank`.

If you make an assignment to a node specified as `reg?`, the TimeQuest analyzer searches and applies the assignment to all design nodes that match the prefix `reg` and any single character following; for example, `reg1`, `rega`, and `reg4`.

Collection Commands

The commands in the Tcl API for the TimeQuest analyzer often return port, pin, cell, or node names in a data set called a collection. In your Tcl scripts you can iterate over the values in collections to analyze or modify constraints in your design. Collection commands are part of the `::quartus::sta`, `::quartus::sdc`, and `::quartus::sdc_ext` Tcl packages.

Adding and Removing Collection Items

Filters used with collection commands limit collection items identified by the command. For example, if a design contains registers named `src0`, `src1`, `src2`, and `dst0`, the collection command `[get_registers src*]` identifies registers `src0`, `src1`, and `src2`, but not register `dst0`. To identify register `dst0`, you must use an additional command, `[get_registers dst*]`.

To overcome this limitation when using filters, use the `add_to_collection` and `remove_from_collection` commands. The `add_to_collection` command allows you to add additional items to an existing collection. [Example 6-5](#) shows the `add_to_collection` command and arguments.

Example 6-5. `add_to_collection` Command

```
add_to_collection <first collection> <second collection>
```

The `remove_from_collection` command allows you to remove items from an existing collection. [Example 6-6](#) shows the `remove_from_collection` command and arguments.

Example 6-6. `remove_from_collection` Command

```
remove_from_collection <first collection> <second collection>
```

Example 6-7 shows examples of how to add elements to collections.

Example 6-7. Adding Items to a Collection

```
#Setting up initial collection of registers
set regs1 [get_registers a*]

#Setting up initial collection of keepers
set kprs1 [get_keepers b*]

#Creating a new set of registers of $regs1 and $kprs1
set regs_union [add_to_collection $kprs1 $regs1]

#OR
# Creating a new set of registers of $regs1 and b*
# Note that the new collection appends only registers with name b*
# not all keepers
set regs_union [add_to_collection $regs1 b*]
```

- ② For more information about the `add_to_collection` and `remove_from_collection` commands—including full syntax information, options, and example usage—refer to [add_to_collection](#) and [remove_from_collection](#) in Quartus II Help.

Refining Collections with Wildcards

The collection commands `get_cells` and `get_pins` have options that allow you to refine searches that include wildcard characters.

Table 6-5 shows examples of search strings that use options to refine the search and wildcards. The examples in Table 6-5 use the following cells and pin names:

- foo
- foo|bar
- foo|dataa
- foo|datab
- foo|bar|datac
- foo|bar|datad

Table 6-5. Sample Search Strings and Search Results

Search String	Search Result
get_pins * dataa	foo dataa
get_pins * datac	<empty>
get_pins * * datac	foo bar datac
get_pins foo* *	foo dataa, foo datab
get_pins -hierarchical * * datac	<empty> (1)
get_pins -hierarchical foo *	foo dataa, foo datab
get_pins -hierarchical * datac	foo bar datac
get_pins -hierarchical foo * datac	<empty> (1)
get_pins -compatibility * datac	foo bar datac
get_pins -compatibility * * datac	foo bar datac

Note to Table 6-5:

(1) The search result is <empty> because of the additional *|*| in the search string.

Removing Constraints and Exceptions

When you use the TimeQuest analyzer interactively, it is sometimes necessary to remove a constraint or exception. Use the following commands to remove constraints and exceptions:

- `remove_clock`
- `remove_clock_groups`
- `remove_clock_latency`
- `remove_clock_uncertainty`
- `remove_input_delay`
- `remove_output_delay`
- `remove_annotated_delay`
- `reset_design`
- `reset_timing_derate`

- ❓ For more information—including a complete list of commands and full syntax information, options, and example usage—refer to `::quartus::sdc` and `::quartus::sdc_ext` in Quartus II Help.

Creating Clocks and Clock Constraints

To ensure accurate static timing analysis results, you must specify all clocks and any associated clock characteristics in your design. The TimeQuest analyzer supports SDC commands that accommodate various clocking schemes and clock characteristics.

Creating Clocks

To create a clock at any register, port, or pin, use the `create_clock` command. You can create each clock with unique characteristics. Clocks defined with the `create_clock` command have a default source latency value of zero. The TimeQuest analyzer automatically computes the clock's network latency for non-virtual clocks.

- ❓ For more information about the `create_clock` command—including full syntax information, options, and example usage—refer to `create_clock` in Quartus II Help.

Example 6–8 shows how to create a 10 ns clock with a 50% duty cycle that is phase shifted by 90 degrees applied to port `clk_sys`.

Example 6–8. 100MHz Shifted by 90 Degrees Clock Creation

```
create_clock -period 10 -waveform { 2.5 7.5 } [get_ports clk_sys]
```

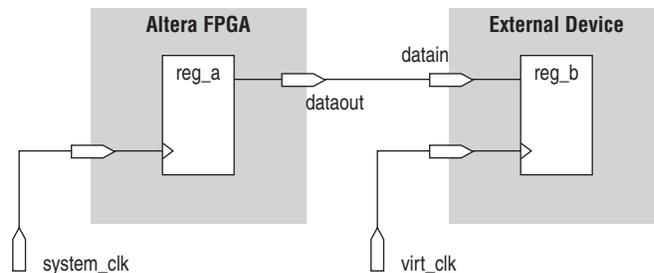
Creating Virtual Clocks

To create virtual clocks, use the `create_clock` command with no value specified for the `<targets>` option. A virtual clock is a clock that does not have a real source in the design or that does not interact directly with the design. For example, if a clock feeds only an external device's clock port and not a clock port in the design, and the external device then feeds, or is fed by, a port in the design, it is considered a virtual clock.

 Use virtual clocks for `set_input_delay` and `set_output_delay` constraints.

Figure 6-18 shows a design where a virtual clock is required for the TimeQuest analyzer to properly analyze the relationship between the external register and the registers in the design. Because the oscillator, `virt_clk`, does not interact with the Altera device, but acts as the clock source for the external register, you must declare the clock as a virtual clock. After you create the virtual clock, you can perform a register-to-register analysis between the register in the Altera device and the register in the external device.

Figure 6-18. Virtual Clock Board Topology



Example 6-9 shows how to create a 10 ns virtual clock named `virt_clk` with a 50% duty cycle where the first rising edge occurs at 0 ns. The virtual clock is then used as the clock source for an output delay constraint.

Example 6-9. Virtual Clock Example 1

```
#create base clock for the design
create_clock -period 5 [get_ports system_clk]

#create the virtual clock for the external register
create_clock -period 10 -name virt_clk -waveform { 0 5 }

#set the output delay referencing the virtual clock
set_output_delay -clock virt_clk -max 1.5 [get_ports dataout]
```

Creating Multifrequency Clocks

To create multifrequency clocks, use the `create_clock` command with the `-add` option to add more than one clock to a clock node. You should create a multifrequency clock if your design has more than one clock source feeding a single clock node in the design. The additional clock may act as a low-power clock, with a lower frequency than the primary clock.

Example 6-10 shows how to create a 10 ns clock applied to clock port `clk`, and then add an additional 15 ns clock to the same clock port. The TimeQuest analyzer uses both clocks when it performs timing analysis.

Example 6-10. Multifrequency Clock Example

```
create_clock -period 10 -name clock_primary -waveform { 0 5 } \
    [get_ports clk]
create_clock -period 15 -name clock_secondary -waveform { 0 7.5 } \
    [get_ports clk] -add
```

Creating Generated Clocks

To create generated clocks, use the `create_generated_clock` command. The TimeQuest analyzer considers clock dividers, ripple clocks, or circuits that modify or change the characteristics of the incoming or master clock as generated clocks. You should define as generated clocks the output of circuits that modify or change the characteristics of the incoming or master clock. Defining the output of the circuits as generated clocks allows the TimeQuest analyzer to analyze these clocks and account for any associate network latency. Source latencies are based on clock network delays from the master clock, but not necessarily the master pin. You can use the `set_clock_latency -source` command to override source latency.

- For more information about the `create_generated_clock` command—including full syntax information, options, and example usage—refer to [create_generated_clock](#) in Quartus II Help.

Figure 6-19 shows how to generate an inverted clock based on a 10 ns clock.

Figure 6-19. Generating an Inverted Clock

```
create_clock -period 10 [get_ports clk]
create_generated_clock -divide_by 1 -invert -source [get_registers clk] \
    [get_registers gen|clkreg]
```

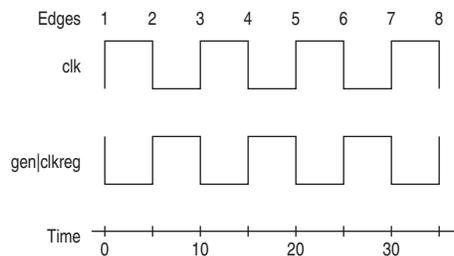


Figure 6–20 shows how to modify the generated clock by defining and shifting the edges.

Figure 6–20. Edge Shifting a Generated Clock

```
create_clock -period 10 -waveform { 0 5 } [get_ports clk]

# Creates a divide-by-two clock
create_generated_clock -source [get_ports clk] -edges {1 3 5 } [get_registers \
clkdivA|clkreg]

# Creates a divide-by-two clock independent of the master clocks' duty cycle (now 50%)
create_generated_clock -source [get_ports clk] -edges {1 1 5} -edge_shift { 0 2.5 0 } \
[get_registers clkdivB|clkreg]
```

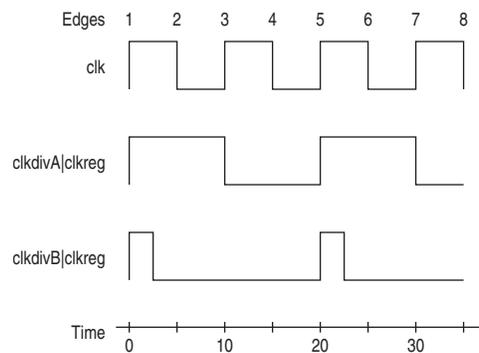
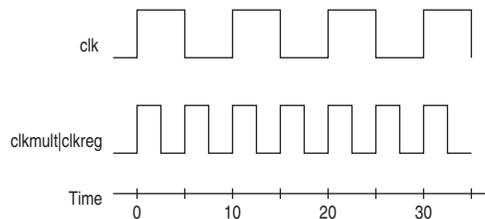


Figure 6–21 shows the effect of the applying a multiplication factor to the generated clock.

Figure 6–21. Multiplying a Generated Clock

```
create_clock -period 10 -waveform { 0 5 } [get_ports clk]

# Creates a multiply-by-two clock
create_generated_clock -source [get_ports clk] -multiply_by 2 [get_registers \
clkmult|clkreg]
```



Automatically Detecting Clocks and Creating Default Clock Constraints

To automatically create clocks for all clock nodes in your design, use the `derive_clocks` command. The `derive_clocks` command is equivalent to using the `create_clock` command for each register or port feeding the clock pin of a register. The `derive_clocks` command creates clocks on ports or registers to ensure every register in the design has a clock, and it applies one period to all base clocks in your design.

To provide a complete clock analysis, if there are no clock constraints in your design, the TimeQuest analyzer automatically creates default clock constraints for all detected unconstrained clock nodes. The TimeQuest analyzer automatically creates clocks only when all synchronous elements have no associated clocks. For example, the TimeQuest analyzer does not create a default clock constraint if your design contains two clocks and you assigned constraints to one of the clocks. However, if you did not assign constraints to either clock, then the TimeQuest analyzer creates a default clock constraint.

Example 6–11 shows how the TimeQuest analyzer creates a base clock with a 1 GHz requirement for unconstrained clock nodes.

Example 6–11. Create Base Clock for Unconstrained Clock Nodes

```
derive_clocks -period 1
```

-  To achieve a thorough and realistic analysis of your design's timing requirements, you should make individual clock constraints for all clocks in your design. Do not use the `derive_clocks` command for final timing sign-off; instead, you should create clocks for all clock sources with the `create_clock` and `create_generated_clock` commands.
-  For more information about the `derive_clocks` command—including full syntax information, options, and example usage—refer to [derive_clocks](#) in Quartus II Help.

Deriving PLL Clocks

Because you should create clocks for all clock nodes, you must create generated clocks for all PLL outputs in your design. Use the `derive_pll_clocks` command to direct the TimeQuest analyzer to automatically create the appropriate generated clocks, or use the `create_generated_clock` command to manually create the generated clocks.

-  For more information about the `derive_pll_clocks` command—including full syntax information, options, and example usage—refer to [derive_pll_clocks](#) in Quartus II Help.

Use the `derive_pll_clocks` command to direct the TimeQuest analyzer to automatically search the timing netlist for all unconstrained PLL output clocks. The `derive_pll_clocks` command calls the `create_generated_clock` command to create generated clocks on the outputs of the PLL. The source for the `create_generated_clock` command is the input clock pin of the PLL.

You must create manually a base clock for the input clock port of the PLL. If you do not define a clock for the input clock node of the PLL, no clocks are reported for the PLL outputs and the TimeQuest analyzer issues a warning message when the timing netlist is updated.

-  You can use the `-create_base_clocks` option to create the input clocks for the PLL inputs automatically.

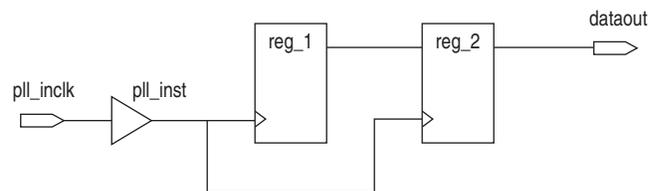
You can include the `derive_pll_clocks` command in your `.sdc`, which allows the `derive_pll_clocks` command to automatically detect any changes to the PLL. Each time the TimeQuest analyzer reads your `.sdc`, it generates the appropriate `create_generated_clocks` command for the PLL output clock pin. If you use the `write_sdc -expand` command after the `derive_pll_clocks` command, the new `.sdc`

contains the individual `create_generated_clock` commands for the PLL output clock pins and not the `derive_pll_clocks` command. Any changes to the properties of the PLL are not automatically reflected in the new `.sdc`. You must manually update the `create_generated_clock` commands in the new `.sdc` written by the `derive_pll_clocks` command to reflect the changes to the PLL.

-  Writing constraints to an existing `.sdc` overwrites the existing file. After editing timing constraints in your design, if you want to save the constraints to an `.sdc`, you should create a new file rather than overwriting the existing file.
-  The `derive_pll_clocks` command also constrains any LVDS transmitters or LVDS receivers in the design by adding the appropriate multicycle constraints to account for any deserialization factors.

Figure 6-22 shows a simple PLL design with a register-to-register path.

Figure 6-22. Simple PLL Design



Example 6-12 shows the messages generated by the TimeQuest analyzer when you use the `derive_pll_clocks` command to automatically constrain the PLL for the design shown in Figure 6-22.

Example 6-12. `derive_pll_clocks` Command Messages

```

Info:
Info: Deriving PLL Clocks:
Info: create_generated_clock -source
pll_inst|altpll_component|pll|inclk[0] -divide_by 2 -name
pll_inst|altpll_component|pll|clk[0]
pll_inst|altpll_component|pll|clk[0]
Info:
  
```

The input clock pin of the PLL is the node

`pll_inst|altpll_component|pll|inclk[0]` used for the `-source` option. The name of the output clock of the PLL is the PLL output clock node, `pll_inst|altpll_component|pll|clk[0]`.

If the PLL is in clock switchover mode, multiple clocks are created for the output clock of the PLL; one for the primary input clock (for example, `inclk[0]`), and one for the secondary input clock (for example, `inclk[1]`). You should create exclusive clock groups for the primary and secondary output clocks.

For more information about creating exclusive clock groups, refer to “Creating Clock Groups” on page 6-33.

Before you can generate any reports for this design, you must create a base clock for the PLL input clock port. You do not have to generate the base clock on the input clock pin of the PLL, `pll_inst|altpll_component|pll|inclk[0]`. The clock created on the PLL input clock port propagates to all fan-outs of the clock port, including the PLL input clock pin. [Example 6-13](#) shows how to create a base clock for the PLL input clock port.

Example 6-13. Create Base Clock for PLL input Clock Port

```
create_clock -period 5 [get_ports pll_inclk]
```

Creating Clock Groups

To specify clocks in your design that are exclusive or asynchronous, use the `set_clock_groups` command.

- ① For more information about the `set_clock_groups` command—including full syntax information, options, and example usage—refer to [set_clock_groups](#) in Quartus II Help.

Exclusive Clock Groups

Use the `-exclusive` option to declare that two clocks are mutually exclusive. You may want to declare clocks as mutually exclusive when multiple clocks are created on the same node or for multiplexed clocks. For example, a port can be clocked by either a 25-MHz or a 50-MHz clock. To constrain this port, you should create two clocks on the port, and then create clock groups to declare that they cannot coexist in the design at the same time. Declaring the clocks as mutually exclusive eliminates any clock transfers that may be derived between the 25-MHz clock and the 50-MHz clock.

[Example 6-14](#) shows how to create mutually exclusive clock groups.

Example 6-14. Create Mutually Exclusive Clock Groups

```
create_clock -period 40 -name clk_A [get_ports {port_A}]  
create_clock -add -period 20 -name clk_B [get_ports {port_A}]  
set_clock_groups -exclusive -group {clk_A} -group {clk_B}
```

A group is defined with the `-group` option. The TimeQuest analyzer excludes the timing paths between clocks for each of the separate groups.

Asynchronous Clock Groups

Use the `-asynchronous` option to create asynchronous clock groups. Clocks contained within an asynchronous clock group are considered asynchronous to clocks in other clock groups; however, any clocks within a clock group are considered synchronous to each other.

-  The TimeQuest analyzer assumes all clocks are related unless constrained otherwise.

For example, if your design has three clocks, `clk_A`, `clk_B`, and `clk_C`, and you establish that `clk_A` and `clk_B` are related to each other, but clock `clk_C` operates completely asynchronously, you can set up clock groups to define the clock behavior. [Example 6-15](#) shows how to create a clock group containing clocks `clk_A` and `clk_B` and a second unrelated clock group containing `clk_C`.

Example 6-15. Create Asynchronous Clock Groups

```
set_clock_groups -asynchronous -group {clk_A clk_B} -group {clk_C}
```

Alternatively, in this example, you can create a clock group containing only `clk_C` to ensure that `clk_A` and `clk_B` are synchronous with each other and asynchronous with `clk_C`. Because `clk_C` is the only clock in the constraint, it is asynchronous with every other clock in the design.

Accounting for Clock Effect Characteristics

The clocks you create with the TimeQuest analyzer are ideal clocks that do not account for any board effects. You can account for clock effect characteristics with clock latency and clock uncertainty.

Clock Latency

There are two forms of clock latency, clock source latency and clock network latency. Source latency is the propagation delay from the origin of the clock to the clock definition point (for example, a clock port). Network latency is the propagation delay from a clock definition point to a register's clock pin. The total latency at a register's clock pin is the sum of the source and network latencies in the clock path.

To specify source latency to any clock ports in your design, use the `set_clock_latency` command.



The TimeQuest analyzer automatically computes network latencies; therefore, you only can create source latency with the `set_clock_latency` command. You must use the `-source` option.



For more information about the `set_clock_latency` command—including full syntax information, options, and example usage—refer to [set_clock_latency](#) in Quartus II Help.

Clock Uncertainty

To specify clock uncertainty, or skew, for clocks or clock-to-clock transfers, use the `set_clock_uncertainty` command. You can specify the uncertainty separately for setup and hold, and you can specify separate rising and falling clock transitions. The TimeQuest analyzer subtracts setup uncertainty from the data required time for each applicable path and adds the hold uncertainty to the data required time for each applicable path.

To automatically apply interclock, intraclock, and I/O interface uncertainties, use the `derive_clock_uncertainty` command. The TimeQuest analyzer automatically applies clock uncertainties to clock-to-clock transfers in the design, and calculates both setup and hold uncertainties for each clock-to-clock transfer.

Any clock uncertainty constraints applied to source and destination clock pairs with the `set_clock_uncertainty` command have a higher precedence than the clock uncertainties derived with the `derive_clock_uncertainty` command for the same source and destination clock pairs. For example, if you use the `set_clock_uncertainty` command to set clock uncertainty between `clka` and `clkb`, the TimeQuest analyzer ignores the values for the clock transfer calculated with the `derive_clock_uncertainty` command. The TimeQuest analyzer reports the values calculated with the `derive_clock_uncertainty` command even if they are not used.

To automatically remove previous clock uncertainty assignments, use the `-overwrite` option. To manually remove previous clock uncertainty assignments, use the `remove_clock_uncertainty` command.

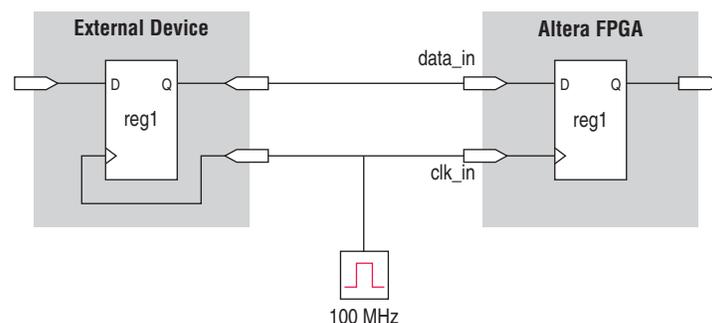
- ④ For more information about the `set_clock_uncertainty`, `derive_clock_uncertainty`, and `remove_clock_uncertainty` commands—including full syntax information, options, and example usage—refer to [set_clock_uncertainty](#), [derive_clock_uncertainty](#), and [remove_clock_uncertainty](#) in Quartus II Help.

I/O Interface Uncertainty

To specify I/O interface uncertainty, you must create a virtual clock and constrain the input and output ports with the `set_input_delay` and `set_output_delay` commands that reference the virtual clock. You must use a virtual clock to prevent the `derive_clock_uncertainty` command from applying clock uncertainties for either intraclock or interclock transfers to an I/O interface clock transfer when the `set_input_delay` or `set_output_delay` commands reference a clock port or PLL output. If you do not reference a virtual clock with the `set_input_delay` or `set_output_delay` commands, the `derive_clock_uncertainty` command calculates intraclock or interclock uncertainty values for the I/O interface.

Create the virtual clock with the same properties as the original clock that is driving the I/O port. [Figure 6-23](#) shows a typical input I/O interface with clock specifications.

Figure 6-23. I/O Interface Clock Specifications



Example 6-16 shows the SDC commands to constrain the I/O interface shown in Figure 6-23.

Example 6-16. SDC Commands to Constrain the I/O Interface

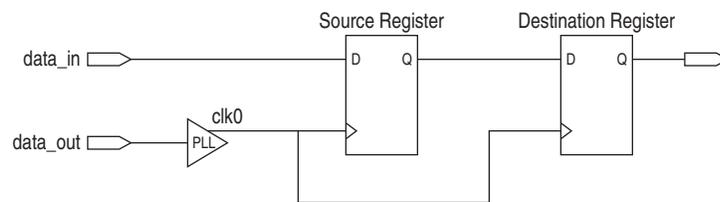
```
# Create the base clock for the clock port
create_clock -period 10 -name clk_in [get_ports clk_in]

# Create a virtual clock with the same properties of the base clock
# driving the source register
create_clock -period 10 -name virt_clk_in

# Create the input delay referencing the virtual clock and not the base
# clock
# DO NOT use set_input_delay -clock clk_in <delay_value>
# [get_ports data_in]
set_input_delay -clock virt_clk_in <delay value> [get_ports data_in]
```

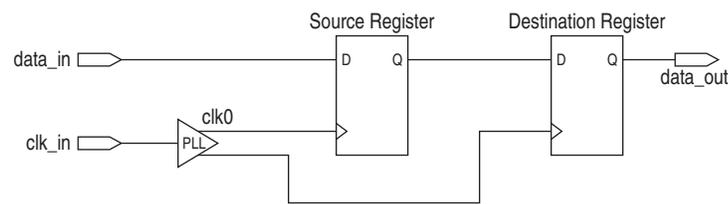
Intraclock transfers occur when the register-to-register transfer takes place in the core of the FPGA and the source and destination clocks come from the same PLL output pin or clock port. Figure 6-24 shows an intraclock transfer.

Figure 6-24. Intraclock Transfer



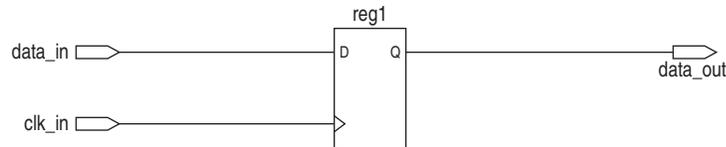
Interlock transfers occur when a register-to-register transfer takes place in the core of the FPGA and the source and destination clocks come from a different PLL output pin or clock port. Figure 6-25 shows an interlock transfer.

Figure 6-25. Interlock Transfer



I/O interface clock transfers occur when data transfers from an I/O port to the core of the FPGA or from the core of the FPGA to the I/O port. Figure 6-26 shows an I/O interface clock transfer.

Figure 6-26. I/O Interface Clock Transfer



Creating I/O Constraints

To specify any external device or board timing parameters, use input and output delay constraints. When you apply these constraints, the TimeQuest analyzer performs static timing analysis on the entire system.

To specify input delay constraints to ports in the design and the data arrival time at a port with respect to a given clock, use the `set_input_delay` command. To specify output delay constraints to ports in the design and the data required time at a port with respect to a given clock, use the `set_output_delay` command.

Figure 6-27 shows an input delay path.

Figure 6-27. Input Delay

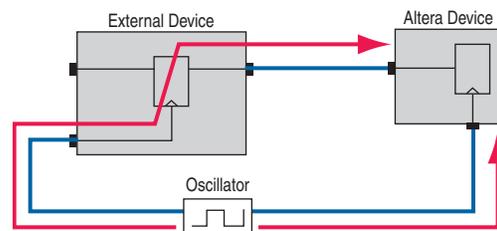
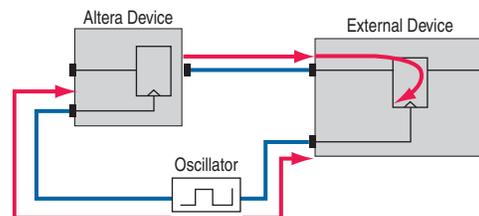


Figure 6-28 shows an output delay path.

Figure 6-28. Output Delay



By default, a set of input or output delays, that is a `-min` and `-max` pair or a `-rise` and `-fall` pair, is allowed for only one `-clock`, `-clock_fall`, and `-reference_pin` combination. Specifying an input delay on the same port for a different `-clock`, `-clock_fall`, or `-reference_pin` setting removes any previously set input delays, unless you specify the `-add_delay` option. When you specify the `-add_delay` option, the TimeQuest analyzer uses the worst-case values.

-  The `-min` and `-max` and `-rise` and `-fall` options are mutually exclusive.
-  For more information about the `set_input_delay` and `set_output_delay` commands—including full syntax information, options, and example usage—refer to [set_input_delay](#) and [set_output_delay](#) in Quartus II Help.

Creating Delay and Skew Constraints

The TimeQuest analyzer supports the Synopsys Design Constraint format for constraining timing for the ports in your design. These constraints allow the TimeQuest analyzer to perform a system static timing analysis that includes not only the FPGA internal timing, but also any external device timing and board timing parameters.

Net Delay

To perform minimum or maximum analysis across nets and report the net delays, use the `set_net_delay` command in conjunction with the `report_net_delay` command. You can use the `set_net_delay` and `report_net_delay` commands to verify timing-critical delays for high-speed interfaces.

-  For more information about the `set_net_delay` and `report_net_delay` commands—including full syntax information, options, and example usage—refer to [set_net_delay](#) and [report_net_delay](#) in Quartus II Help.

Advanced I/O Timing and Board Trace Model Delay

The TimeQuest analyzer can use advanced I/O timing and board trace model assignments to model I/O buffer delays in your design.

If you change any advanced I/O timing settings or board trace model assignments, recompile your design before you analyze timing, or use the `-force_dat` option to force delay annotation when you create a timing netlist. [Example 6-17](#) shows how to force delay annotation when creating a timing netlist.

Example 6-17. Forcing Delay Annotation

```
create_timing_netlist -force_dat ←
```

-  For more information about using advanced I/O timing, refer to [Using Advanced I/O Timing](#) in Quartus II Help.
-  For more information about advanced I/O timing, refer to the [I/O Management](#) chapter in volume 2 of the *Quartus II Handbook*.

Maximum Skew

To specify the maximum path-based skew requirements for registers and ports in the design and report the results of maximum skew analysis, use the `set_max_skew` command in conjunction with the `report_max_skew` command.

By default, the `set_max_skew` command excludes any input or output delay constraints.

- For more information about the `set_max_skew` and `report_max_skew` commands—including full syntax information, options, and example usage—refer to [set_max_skew](#) [report_max_skew](#) in Quartus II Help.

Creating Timing Exceptions

Timing exceptions modify the default analysis performed by the TimeQuest analyzer. You can create several different timing exceptions with the TimeQuest analyzer to adjust the timing of your design.

Precedence

If a conflict of node names occurs between timing exceptions, the following order of precedence applies:

1. False path
2. Minimum delays and maximum delays
3. Multicycle path

The false path timing exception has the highest precedence. Within each category, assignments to individual nodes have precedence over assignments to clocks. Finally, the remaining precedence for additional conflicts is order-dependent, such that the assignments most recently created overwrite, or partially overwrite, earlier assignments.

False Paths

To specify false paths in your design, use the `set_false_path` command. False paths are paths that can be ignored during timing analysis.

If you specify a false path between two timing notes, the false path applies only to the path between the two nodes. If you specify a false path to a clock, the false path applies to all paths where the source node (`-from`) or destination node (`-to`) is clocked by the clock.

- For more information about the `set_false_path` command—including full syntax information, options, and example usage—refer to [set_false_path](#) in Quartus II Help.

Minimum and Maximum Delays

To specify an absolute minimum or maximum delay for a path, use the `set_min_delay` command or the `set_max_delay` commands, respectively.

If the source or destination node is clocked, the TimeQuest analyzer takes into account the clock paths, allowing more or less delay on the data path. If the source or destination node has an input or output delay, that delay is also included in the minimum or maximum delay check.

If you specify a minimum or maximum delay between timing nodes, the delay applies only to the path between the two nodes. If you specify a minimum or maximum delay for a clock, the delay applies to all paths where the source node (-from) or destination node (-to) is clocked by the clock.

You can create a minimum or maximum delay exception for an output port that does not have an output delay constraint. You cannot report timing for the paths associated with the output port; however, the TimeQuest analyzer reports any slack for the path in the setup summary and hold summary reports. Because there is no clock associated with the output port, no clock is reported for timing paths associated with the output port.

-  To report timing with clock filters for output paths with minimum and maximum delay constraints, you can set the output delay for the output port with a value of zero. You can use an existing clock from the design or a virtual clock as the clock reference.
-  For more information about the `set_min_delay` and `set_max_delay` commands—including full syntax information, options, and example usage—refer to [set_min_delay](#) and [set_max_delay](#) in Quartus II Help.

Delay Annotation

To modify the default delay values used during timing analysis, use the `set_annotated_delay` and `set_timing_derate` commands. You must update the timing netlist to see the results of these commands

To specify different operating conditions in a single `.sdc`, rather than having multiple `.sdc` files that specify different operating conditions, use the `set_annotated_delay` command with the `-operating_conditions` option.

-  For more information about the `set_annotated_delay` and `set_timing_derate` commands—including full syntax information, options, and example usage—refer to [set_annotated_delay](#) and [set_timing_derate](#) in Quartus II Help.

Timing Reports

The TimeQuest analyzer provides real-time static timing analysis result reports. The TimeQuest analyzer does not automatically generate reports; you must create each report individually in the TimeQuest analyzer GUI or with command-line commands. You can customize in which report to display specific timing information, excluding fields that are not required.

Table 6-6 shows some of the different command-line commands you can use to generate reports in the TimeQuest analyzer and the equivalent reports shown in the TimeQuest analyzer GUI.

Table 6-6. TimeQuest Analyzer Reports

Command-Line Command	Report
report_timing	Timing report
report_exceptions	Exceptions report
report_clock_transfers	Clock Transfers report
report_min_pulse_width	Minimum Pulse Width report
report_ucp	Unconstrained Paths report

- For more information—including a complete list of commands to generate timing reports and full syntax information, options, and example usage—refer to `::quartus::sta` in Quartus II Help.

During compilation, the Quartus II software generates timing reports that contain basic timing information. You can configure various options for the TimeQuest analyzer reports generated during compilation.

- For more information about the options you can set to customize the reports, refer to *TimeQuest Timing Analyzer Page* in Quartus II Help.

You can also use the `TIMEQUEST_REPORT_WORST_CASE_TIMING_PATHS` assignment to generate a report of the worst-case timing paths for each clock domain. This report contains worst-case timing data for setup, hold, recovery, removal, and minimum pulse width checks.

Use the `TIMEQUEST_REPORT_NUM_WORST_CASE_TIMING_PATHS` assignment to specify the number of paths to report for each clock domain.

Example 6-18 shows an example of how to use the `TIMEQUEST_REPORT_WORST_CASE_TIMING_PATHS` and `TIMEQUEST_REPORT_NUM_WORST_CASE_TIMING_PATHS` assignments in the `.qsf` to generate reports.

Example 6-18. Generating Worst-Case Timing Reports

```
#Enable Worst-Case Timing Report
set_global_assignment -name TIMEQUEST_REPORT_WORST_CASE_TIMING_PATHS ON
#Report 10 paths per clock domain
set_global_assignment -name TIMEQUEST_REPORT_NUM_WORST_CASE_TIMING_PATHS 10
```

Document Revision History

Table 6-7 shows the revision history for this chapter.

Table 6-7. Document Revision History

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Changed to new document template. ■ Revised and reorganized entire chapter. ■ Linked to Quartus II Help.
July 2010	10.0.0	Updated to link to content on SDC commands and the TimeQuest analyzer GUI in Quartus II Help.
November 2009	9.1.0	Updated for the Quartus II software version 9.1, including: <ul style="list-style-type: none"> ■ Added information about commands for adding and removing items from collections ■ Added information about the <code>set_timing_derate</code> and <code>report_skew</code> commands ■ Added information about worst-case timing reporting ■ Minor editorial updates
November 2008	8.1.0	Updated for the Quartus II software version 8.1, including: <ul style="list-style-type: none"> ■ Added the following sections: <ul style="list-style-type: none"> ■ “set_net_delay” on page 7-42 ■ “Annotated Delay” on page 7-49 ■ “report_net_delay” on page 7-66 ■ Updated the descriptions of the <code>-append</code> and <code>-file <name></code> options in tables throughout the chapter ■ Updated entire chapter using 8½” × 11” chapter template ■ Minor editorial updates

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

Timing constraints and exceptions are vital to all designs that target FPGAs because they allow designers to specify requirements and verify timing of their systems or FPGAs. Constraints and exceptions allow the Fitter to spend more time fitting the critical paths in your design and reduce the amount of time spent on noncritical parts of the design.

This chapter provides recommendations on how to fully constrain an FPGA design with the Quartus® II TimeQuest Timing Analyzer. The sections are ordered in the recommended flow for applying timing constraints and exceptions in the TimeQuest analyzer.

-  For more information about interactive timing analysis, refer to the *TimeQuest Timing Analyzer Quick Start Tutorial*.
-  For more information about Altera resources available for the TimeQuest analyzer, refer to the *TimeQuest Timing Analyzer Resource Center* of the Altera website.
-  For more information about constraining circuits and reporting timing analysis results in the TimeQuest analyzer, including examples, refer to the *TimeQuest Design Examples* page of the Altera website and the *Quartus II TimeQuest Timing Analyzer Cookbook*.
-  For more information about constraints and exceptions supported by the TimeQuest analyzer, refer to the *Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.
-  For more information about TimeQuest analyzer Tcl commands and a complete list of commands, refer to `::quartus::sta` in Quartus II Help. For more information about standard SDC commands and a complete list of commands, refer to `::quartus::sdc` in Quartus II Help. For more information about Altera extensions of SDC commands and a complete list of commands, refer to `::quartus::sdc_ext` in Quartus II Help.

This chapter contains the following sections:

- “Creating Clock Requirements” on page 7–2
- “Creating I/O Requirements” on page 7–5
- “Creating Timing Exceptions” on page 7–7

Creating Clock Requirements

The TimeQuest analyzer supports the following types of clocks:

- Base clocks
- Derived clocks
- Virtual clocks

Clocks are used to specify register-to-register requirements for synchronous transfers and guide the Fitter optimization algorithms to achieve the best possible placement for your design.

Specify clock constraints first in Synopsys Design Constraint Files (.sdc) because other constraints may reference previously defined clocks. The TimeQuest analyzer reads SDC constraints and exceptions from top to bottom in the file.



For more information about creating clocks and clock constraints, refer to the *Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

Base Clocks

Base clocks are the primary input clocks to the FPGA. Unlike clocks from phase-locked loops (PLLs) that are derived within the FPGA, base clocks are usually generated in off-chip oscillators or forwarded from an external device. Define base clocks first because derived clocks and other constraints can reference the base clocks.

Use the `create_clock` command to constrain all primary input clocks. The target for the `create_clock` command is usually an FPGA device pin. To specify the FPGA device pin as the target, use the `get_ports` command. [Example 7-1](#) shows how to specify a 100-MHz requirement on a `clk_sys` input clock port.

Example 7-1. `create_clock` Command

```
create_clock -period 10 -name clk_sys [get_ports clk_sys]
```

You can apply multiple clocks on the same clock node with the `-add` option. [Example 7-2](#) shows how to specify that two oscillators drive the same clock port on the FPGA.

Example 7-2. Two Oscillators Driving the Same Clock Port

```
create_clock -period 10 -name clk_100 [get_ports clk_sys]
create_clock -period 5 -name clk_200 [get_ports clk_sys] -add
```

- ❓ For more information about the `create_clock` and `get_ports` commands—including full syntax information, options, and example usage—refer to [create_clock](#) and [get_ports](#) in Quartus II Help.

Derived Clocks

Derived clocks are generated in the FPGA when you modify the properties, including phase, frequency, offset, and duty cycle, of a source clock signal. Derived clocks, which are PLLs or register clock dividers, are constrained after all base clocks are constrained in the `.sdc`. Derived clocks capture all clock delays and clock latency where the derived clock target is defined, ensuring that all base clock properties are accounted for in the derived clock.

You can use the `create_generated_clock` command to constrain all generated clocks in your design. The source of the `create_generated_clock` command should be a node in your design and not a previously constrained clock.

[Example 7-3](#) shows a divide-by-two clock divider.

Example 7-3. Clock Divider

```
create_clock -period 10 -name clk_sys [get_ports clk_sys]
create_generated_clock -name clk_div_2 -divide_by 2 -source
[get_pins reg|clk] [get_pins reg|regout]
```

Use the `create_generated_clock` command with the `-source` option when you specify multiple clock constraints for the same pin in a design. When you use the `create_generated_clock` command, the `-source` option should refer to the nearest clock pin of the specified target. Do not use the clock port as the source for a generated clock, because multiple source clocks can feed the clock port. In [Example 7-3](#), the `-source` option assigns the clock pin of the register as the source for the generated clock instead of the clock port `clk` feeding the register's `reg` clock pin.

The TimeQuest analyzer provides the `derive_pll_clocks` command to automatically generate derived clocks for all PLL clock outputs. The properties of the generated clocks on the PLL outputs match the properties defined for the PLL.

- For more information about the `create_generated_clock` and `derive_pll_clocks` commands—including for full syntax information, options, and example usage—refer to [create_generate_clock](#) and [derive_pll_clocks](#) in Quartus II Help.

Virtual Clocks

A virtual clock does not have a real source in your design and does not interact directly with your design. You can create virtual clocks with the `create_clock` command, with no targets specified. [Example 7-4](#) shows how to create a 10 ns virtual clock.

Example 7-4. Create Virtual Clock

```
create_clock -period 10 -name my_virt_clk
```

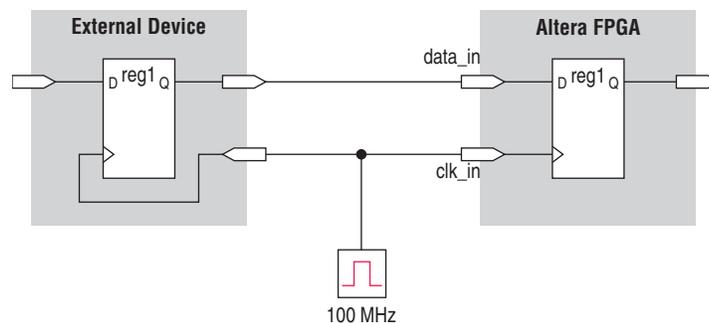
If you use the `derive_clock_uncertainty` command for your design, use virtual clocks with the `set_input_delay` and `set_output_delay` commands. It is important to use virtual clocks to allow the TimeQuest analyzer to calculate clock uncertainty separately for I/O interfaces and internal register-to-register paths.

If an FPGA interfaces with an external device, and both the FPGA and external device have different clock sources, you should model the clock source for the external device with a virtual clock.

- ② For more information about the `set_input_delay`, `set_output_delay`, and `derive_clock_uncertainty` commands—including full syntax information, options, and example usage—refer to [set_input_delay](#), [set_output_delay](#), and [derive_clock_uncertainty](#) in Quartus II Help.

Figure 7-1 shows a typical I/O interface that contains an FPGA that interfaces with an external device and also shows a virtual clock. Example 7-5 shows how to create an equivalent virtual clock for each clock in your design that feeds an input or output port.

Figure 7-1. Design with Virtual Clock



Example 7-5. Commands to Create Clocks

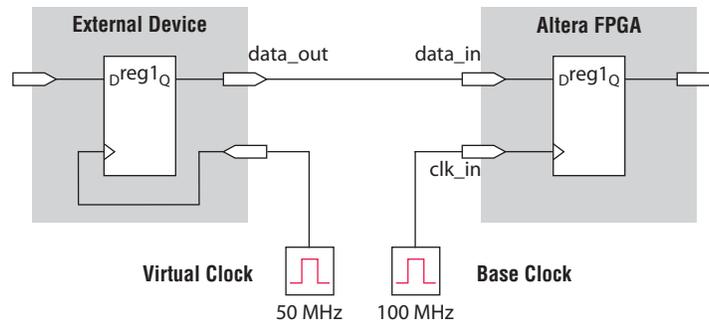
```
# Create the base clock for the clock port
create_clock -period 10 -name clk_in [get_ports clk_in]

# Create a virtual clock with the same properties of the base clock
# driving the source register
create_clock -period 10 -name virt_clk_in

# Create the input delay referencing the virtual clock and not the base
# clock
# DO NOT use set_input_delay -clock clk_in <delay_value>
# [get_ports data_in]
set_input_delay -clock virt_clk_in <delay value> [get_ports data_in]
```

Figure 7-2 shows a design that contains an FPGA that interfaces with an external device and also shows the base clock and the virtual clock. Example 7-6 shows how to create virtual clock for the design.

Figure 7-2. Design with Base Clock and Virtual Clock



Example 7-6. Commands to Create Clocks

```
#create base clock for the design
create_clock -period 10 -name clk_in [get_ports clk_in]

#create the virtual clock for the external register
create_clock -period 20 -name virt_clk -waveform {0 10}

#set the output delay referencing the virtual clock
set_output_delay -clock virt_clk -max 1.5 [get_ports data_out]
```

Creating I/O Requirements

You should specify timing requirements, including internal and external timing requirements, before you fully analyze a design. With external timing requirements specified, the TimeQuest analyzer verifies the I/O interface, or periphery of the FPGA, against any system specification. The TimeQuest analyzer supports input and output external delay modeling.

You should specify I/O requirements after you constrain all clocks in your design. When specifying I/O requirements, reference a virtual clock in the constraints.

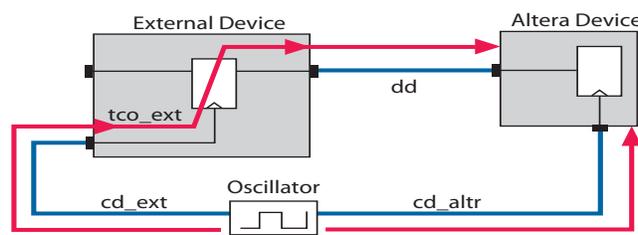
 For more information about specifying I/O interface requirements and uncertainty, refer to the *Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

Input Constraints

Input constraints allow you to specify all the external delays feeding into the FPGA. Specify input requirements for all input ports in your design.

You can use the `set_input_delay` command to specify external input delay requirements. Use the `-clock` option to reference a virtual clock. Using a virtual clock allows the TimeQuest analyzer to correctly derive clock uncertainties for interclock and intraclock transfers. The virtual clock defines the launching clock for the input port. The TimeQuest analyzer automatically determines the latching clock inside the device that captures the input data, because all clocks in the device are defined. [Figure 7-3](#) shows an example of an input delay referencing a virtual clock.

Figure 7-3. Input Delay



[Equation 7-1](#) shows a typical input delay calculation.

Equation 7-1. Input Delay Calculation

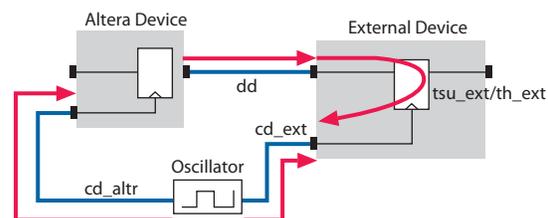
$$\begin{aligned} \text{input delay}_{\text{MAX}} &= (\text{cd_ext}_{\text{MIN}} - \text{cd_altr}_{\text{MAX}}) + \text{tco_ext}_{\text{MAX}} + \text{dd}_{\text{MAX}} \\ \text{input delay}_{\text{MIN}} &= (\text{cd_ext}_{\text{MAX}} - \text{cd_altr}_{\text{MIN}}) + \text{tco_ext}_{\text{MIN}} + \text{dd}_{\text{MIN}} \end{aligned}$$

Output Constraints

Output constraints allow you to specify all external delays from the FPGA for all output ports in your design.

You can use the `set_output_delay` command to specify external output delay requirements. Use the `-clock` option to reference a virtual clock. The virtual clock defines the latching clock for the output port. The TimeQuest analyzer automatically determines the launching clock inside the device that launches the output data, because all clocks in the device are defined. [Figure 7-4](#) shows an example of an output delay referencing a virtual clock.

Figure 7-4. Output Delay



Equation 7-2 shows a typical output delay calculation.

Equation 7-2. output Delay Calculation

$$\begin{aligned}\text{input delay}_{\text{MAX}} &= \text{dd}_{\text{MAX}} - \text{tsu}_{\text{ext}} + (\text{cd}_{\text{altr}}_{\text{MIN}} - \text{cd}_{\text{ext}}_{\text{MAX}}) \\ \text{input delay}_{\text{MIN}} &= \text{dd}_{\text{MIN}} - \text{th}_{\text{ext}} + (\text{cd}_{\text{altr}}_{\text{MAX}} - \text{cd}_{\text{ext}}_{\text{MIN}})\end{aligned}$$

- ❓ For more information about the `set_input_delay` and `set_output_delay` commands—including full syntax information, options, and example usage—refer to [set_input_delay](#) and [set_output_delay](#) in Quartus II Help.

Creating Timing Exceptions

Timing exceptions in the TimeQuest analyzer provide a way to modify the default timing analysis behavior to match the analysis required by your design. Specify timing exceptions after clocks and input and output delay constraints because timing exceptions can modify the default analysis. The TimeQuest analyzer supports the following three major categories of timing exceptions:

- False paths
- Minimum and maximum delays
- Multicycle paths

 For more information about creating timing exceptions, refer to the [Quartus II TimeQuest Timing Analyzer](#) chapter in volume 3 of the *Quartus II Handbook*.

False Paths

Specifying a false path in your design removes the path from timing analysis. Use the `set_false_path` command to specify false paths in your design. You can specify either a point-to-point or clock-to-clock path as a false path. For example, a path you should specify as false path is a static configuration register that is written once during power-up initialization, but does not change state again. Although signals from static configuration registers often cross clock domains, you may not want to make false path exceptions to a clock-to-clock path, because some data may transfer across clock domains. However, you can selectively make false path exceptions from the static configuration register to all endpoints.

[Example 7-7](#) shows how to make false path exceptions from all registers beginning with A to all registers beginning with B.

Example 7-7. False Path

```
set_false_path -from [get_pins A*] -to [get_pins B*]
```

The TimeQuest analyzer assumes all clocks are related unless you specify otherwise. You can use clock groups to make false path exceptions for clock-to-clock timing relationships in your design. Clock groups are a more efficient way to make false path exceptions between clocks, compared to writing multiple `set_false_path` exceptions between every clock transfer you want to eliminate.

-  The TimeQuest analyzer runs more efficiently with clock groups rather than individual false path assignments.

Use the `set_clock_groups` command to collect groups of signals related to each other, and use the `-asynchronous` option to specify that each group of clocks is asynchronous with each other. If you apply multiple clocks to the same port, use the `set_clock_groups` command with the `-exclusive` option to place the clocks into separate groups and declare that the clocks are mutually exclusive. The clocks cannot physically exist in your design at the same time.

-  For more information about the `set_clock_groups` and `set_false_path` commands—including full syntax information, options, and example usage—refer to [set_clock_groups](#) and [set_false_path](#) in Quartus II Help.

Minimum and Maximum Delays

Specifying minimum and maximum delay constraints in your design creates a bounded minimum and maximum path delay. Use the `set_min_delay` and `set_max_delay` commands to create constraints for asynchronous signals that do not have a specific clock relationship in your design, but require a minimum and maximum path delay. You can create minimum and maximum delay exceptions for port-to-port paths through the FPGA without a register stage in the path. If you use minimum and maximum delay exceptions to constrain the path delay, specify both the minimum and maximum delay of the path; do not constrain only the minimum or maximum value.

You can also use the `set_net_delay` command to specify the minimum delay, maximum delay, or skew for any edge in your design when no clock relationships are defined or required.

-  For more information about the `set_min_delay`, `set_max_delay`, and `set_net_delay` commands—including full syntax information, options, and example usage—refer to [set_min_delay](#), [set_max_delay](#), and [set_net_delay](#) in Quartus II Help.

Creating Multicycle Exceptions

By default, the TimeQuest analyzer uses single-cycle path analysis. The TimeQuest Timing Analyzer examines all register-to-register paths and performs setup and hold check analysis on those paths. The setup and hold check analysis evaluates the launch and latch edge relationships. When analyzing a path, the TimeQuest analyzer determines the setup launch and latch edge times by finding the closest two active edges in the respective waveforms. When analyzing setup and hold relationships, the TimeQuest analyzer analyzes the path against two timing conditions for every possible setup relationship, not just the worst-case setup relationship; therefore, the hold launch and latch times may be unrelated to the setup launch and latch edges. The TimeQuest analyzer does not report negative setup or hold relationships. When either a negative setup or a negative hold relationship is calculated, the TimeQuest analyzer moves both the launch and latch edges such that the setup and hold relationship becomes positive.

Multicycle exceptions relax the timing requirements for a register-to-register path, allowing the Fitter to optimally place and route a design in an FPGA. Multicycle exceptions also can reduce compilation time and increase the quality of results.

For example, if your design contains a long combinational path in which the latching register does not require data stability on every clock edge, but only on every second clock edge, you can assign a multicycle exception to the path. The multicycle path is dependent on the endpoint register's use of the clock signal. [Example 7-8](#) shows how to create a multicycle path for the combinational path where the data is stable at the endpoint every two clock cycles of the endpoint latch clock.

Example 7-8. Multicycle Path

```
set_multicycle_path -setup 2
```

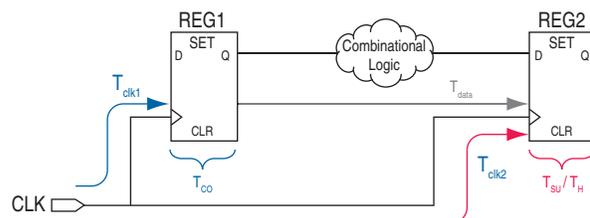
If you specify a multicycle path, define both the setup and hold multicycle relationships. For the preceding example, setting data at the endpoint can take two clock cycles and the minimum hold time relationship is defined with a multicycle exception as well. Use the `set_multicycle_path` command with the `-hold` option to define the hold relationship. The value of the `-hold` option is $(N - 1)$, where N is equal to the multicycle setup assignment value for a register-to-register path in the same clock domain. However, if data crosses different clock domains, the phase and period of the launch and latch clock may change the default multicycle setup and hold values. If you use multicycle paths that cross different clock domains, you must carefully examine the timing paths in the TimeQuest analyzer before and after applying the multicycle exception to determine if the launch and latch clock edges function as you intend.

- For more information about the `set_multicycle_path` command, including full syntax information, options, and example usage, refer to [set_multicycle_path](#) in Quartus II Help.

Multicycle Clock Setup Check and Hold Check Analysis

You can modify the setup and hold relationship when you apply a multicycle exception to a register-to-register path. [Figure 7-5](#) shows a register-to-register path with various timing parameters labeled.

Figure 7-5. Register-to-Register Path



Multicycle constraints adjust setup or hold relationships by the specified number of clock cycles based on the source (-start) or destination (-end) clock. An end multicycle setup constraint of two extends the worst-case setup latch edge by one destination clock period.

Hold multicycle constraints are based on the hold position—the default hold position is zero. An end multicycle hold constraint of one effectively subtracts one destination clock period from the default hold latch edge.

To specify the multicycle constraints in your design, use the `set_multicycle_path` command.

- ④ For more information about the `set_multicycle_path` command, including full syntax information, options, and example usage, refer to [set_multicycle_path](#) in Quartus II Help.

If you specify a multicycle constraint between timing nodes, the multicycle constraint applies only to the path between the two nodes. If you specify a multicycle constraint for a clock, the multicycle constraint applies to all paths where the source node (-from) or destination node (-to) is clocked by the clock.

Table 7-1 summarizes commonly used multicycle assignments.

Table 7-1. Common Multicycle Assignments

Type	Clock	Timing Check	SDC Command
End Multicycle Setup	Destination	Setup	<code>set_multicycle_path -end -setup</code>
End Multicycle Hold	Destination	Hold	<code>set_multicycle_path -end -hold</code>
Start Multicycle Setup	Source	Setup	<code>set_multicycle_path -start -setup</code>
Start Multicycle Hold	Source	Hold	<code>set_multicycle_path -start -hold</code>

Multicycle Clock Setup

The setup relationship is defined as the number of clock periods between the latch edge and the launch edge. By default, the TimeQuest analyzer performs a single-cycle path analysis, which results in the setup relationship being equal to one clock period (latch edge – launch edge). When analyzing a path, the TimeQuest analyzer determines the setup launch and latch edge times by finding the closest two active edges in the respective waveforms. Applying a multicycle setup assignment, increases—or relaxes—the setup relationship by the multicycle setup value.

For every register-to-register path, the TimeQuest analyzer calculates the setup slack for the path. Equation 7-3 shows the setup slack calculation.

Equation 7-3. Setup Slack (1) (2) (3) (4) (5) (6)

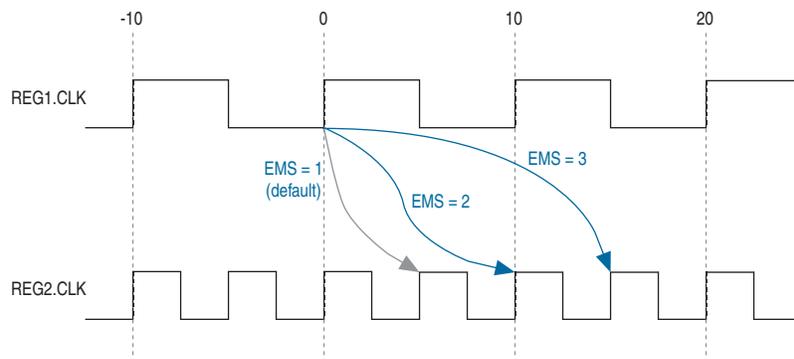
$$\begin{aligned} \text{setup slack} &= \text{data required time} - \text{data arrival time} \\ &= (\text{latch edge} + t_{\text{clk2}} - t_{\text{SU}}) - (\text{launch edge} + t_{\text{clk1}} + t_{\text{CO}} + t_{\text{data}}) \\ &= (\text{latch edge} - \text{launch edge}) + (t_{\text{clk2}} - t_{\text{clk1}}) - (t_{\text{CO}} + t_{\text{data}} + t_{\text{SU}}) \end{aligned}$$

Notes to Equation 7-3:

- (1) t_{clk1} = the propagation delay from clock source to clock input on source register
- (2) t_{clk2} = the propagation delay from clock source to clock input on destination register
- (3) t_{data} = the propagation delay from source register to data input on destination register
- (4) t_{CO} = the clock to output delay of source register
- (5) t_{SU} = the setup requirement of destination register
- (6) t_{H} = the hold requirement of destination register

An end multicycle setup assignment modifies the latch edge of the destination clock by moving the latch edge the specified number of clock periods to the right of the determined default latch edge. Figure 7-6 shows various values of the end multicycle setup assignment and the resulting latch edge.

Figure 7-6. End Multicycle Setup Values



A start multicycle setup assignment modifies the launch edge of the source clock by moving the launch edge the specified number of clock periods to the left of the determined default launch edge. Figure 7-7 shows various values of the start multicycle setup assignment and the resulting launch edge.

Figure 7-7. Start Multicycle Setup Values

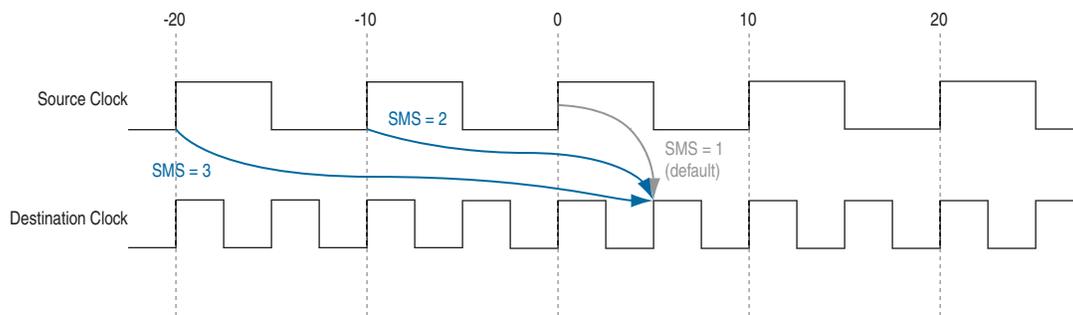
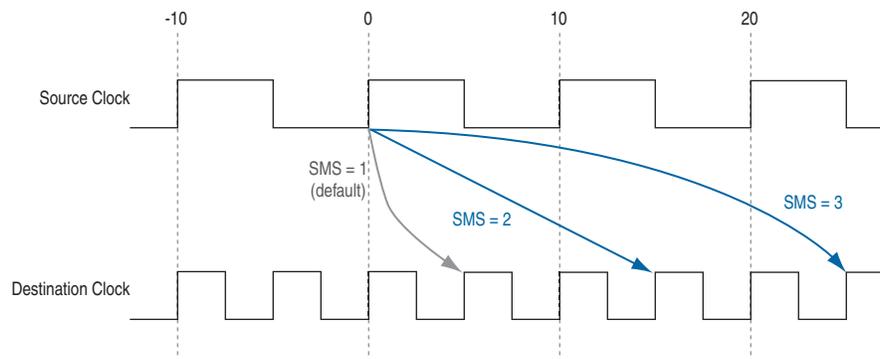


Figure 7-8 shows the setup relationship reported by the TimeQuest analyzer for the negative setup relationship shown in Figure 7-7.

Figure 7-8. Start Multicycle Setup Values Reported by the TimeQuest Analyzer



Multicycle Clock Hold

The hold relationship is defined as the number of clock periods between the launch edge and the latch edge. By default, the TimeQuest analyzer performs a single-cycle path analysis, which results in the hold relationship being equal to one clock period (launch edge – latch edge). When analyzing a path, the TimeQuest analyzer performs two hold checks. The first hold check determines that the data launched by the current launch edge is not captured by the previous latch edge. The second hold check determines that the data launched by the next launch edge is not captured by the current latch edge. The TimeQuest analyzer reports only the most restrictive hold check. Equation 7-4 shows the calculation that the TimeQuest analyzer performs to determine the hold check.

Equation 7-4. Hold Check

hold check 1 = current launch edge – previous latch edge

hold check 2 = next launch edge – current latch edge

 If a hold check overlaps a setup check, the hold check is ignored.

Applying a multicycle hold assignment, increases—or relaxes—the hold slack equation by the number of specified clock cycles.

For every register-to-register path, the TimeQuest analyzer calculates the hold slack for the path. Equation 7-5 shows the hold slack calculation.

Equation 7-5. Hold Slack (1) (2) (3) (4) (5) (6)

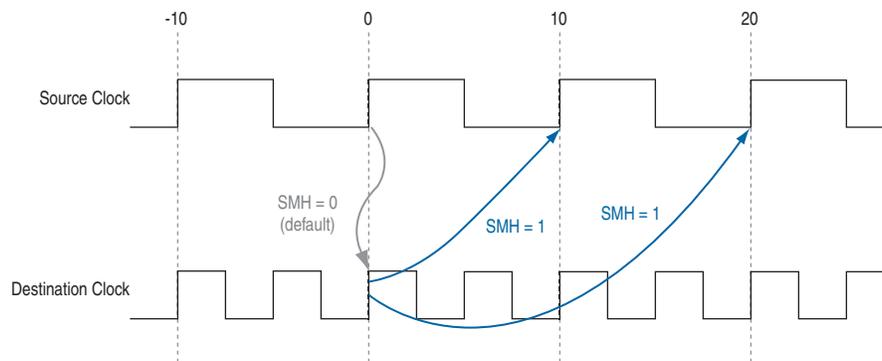
$$\begin{aligned} \text{hold slack} &= \text{data arrival time} - \text{data required time} \\ &= (\text{launch edge} + t_{\text{clk1}} + t_{\text{CO}} + t_{\text{data}}) - (\text{latch edge} + t_{\text{clk2}} - t_{\text{H}}) \\ &= (\text{launch edge} - \text{latch edge}) - (t_{\text{clk2}} - t_{\text{clk1}}) + (t_{\text{CO}} + t_{\text{data}} - t_{\text{H}}) \end{aligned}$$

Notes to Equation 7-5:

- (1) t_{clk1} = the propagation delay from clock source to clock input on source register
- (2) t_{clk2} = the propagation delay from clock source to clock input on destination register
- (3) t_{data} = the propagation delay from source register to data input on destination register
- (4) t_{CO} = the clock to output delay of source register
- (5) t_{SU} = the setup requirement of destination register
- (6) t_{H} = the hold requirement of destination register

A start multicycle hold assignment modifies the launch edge of the destination clock by moving the latch edge the specified number of clock periods to the right of the determined default launch edge. Figure 7-9 shows various values of the start multicycle hold assignment and the resulting launch edge.

Figure 7-9. Start Multicycle Hold Values



An end multicycle hold assignment modifies the latch edge of the destination clock by moving the latch edge the specific ed number of clock periods to the left of the determined default latch edge. Figure 7-10 shows various values of the end multicycle hold assignment and the resulting latch edge.

Figure 7-10. End Multicycle Hold Values

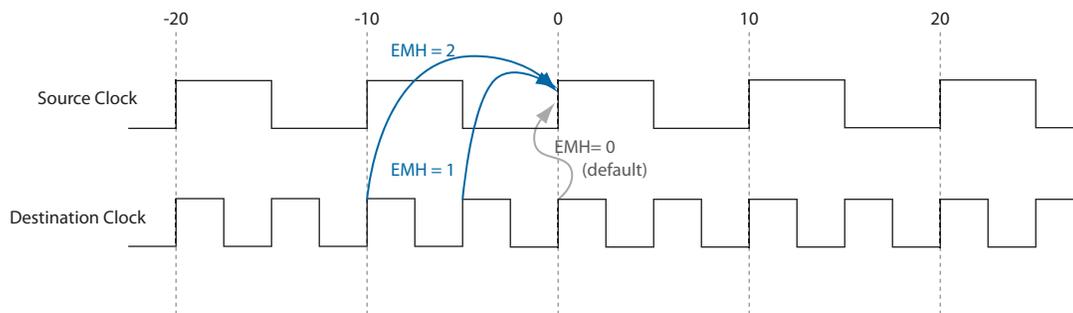
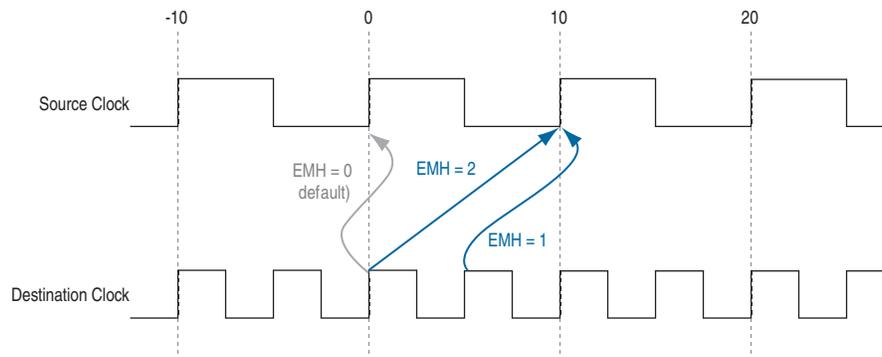


Figure 7-11 shows the hold relationship reported by the TimeQuest analyzer for the negative hold relationship shown in Figure 7-10.

Figure 7-11. End Multicycle Hold Values Reported by the TimeQuest Analyzer



Examples of Basic Multicycle Exceptions

This section describes the following examples of combinations of multicycle exceptions:

- “Default Settings” on page 7-15
- “End Multicycle Setup = 2 and End Multicycle Hold = 0” on page 7-17
- “End Multicycle Setup = 1 and End Multicycle Hold = 1” on page 7-20
- “End Multicycle Setup = 2 and End Multicycle Hold = 1” on page 7-23
- “Start Multicycle Setup = 2 and Start Multicycle Hold = 0” on page 7-26
- “Start Multicycle Setup = 1 and Start Multicycle Hold = 1” on page 7-29
- “Start Multicycle Setup = 2 and Start Multicycle Hold = 1” on page 7-32

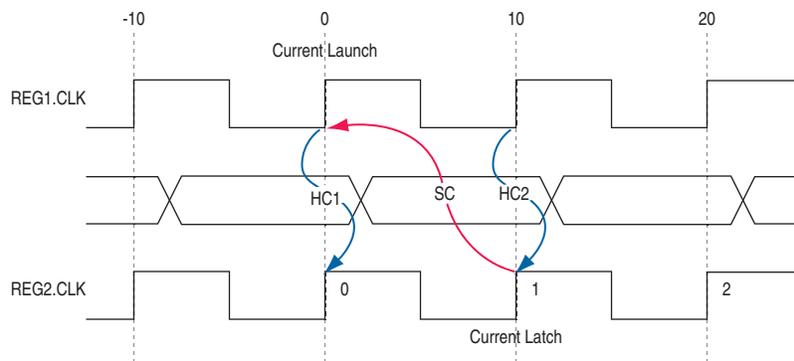
Each example explains how the multicycle exceptions affect the default setup and hold analysis in the TimeQuest analyzer. The multicycle exceptions are applied to a simple register-to-register circuit. Both the source and destination clocks are set to 10 ns.

Default Settings

By default, the TimeQuest analyzer performs a single-cycle analysis to determine the setup and hold checks. Also, by default, the TimeQuest Timing Analyzer sets the end multicycle setup assignment value to one and the end multicycle hold assignment value to zero.

Figure 7-12 shows the source and the destination timing waveform for the source register and destination register, respectively.

Figure 7-12. Default Timing Diagram



Equation 7-6 shows the calculation that the TimeQuest analyzer performs to determine the setup check.

Equation 7-6. Setup Check

$$\begin{aligned}
 \text{setup check} &= \text{current latch edge} - \text{closest previous launch edge} \\
 &= 10\text{ns} - 0\text{ns} \\
 &= 10\text{ns}
 \end{aligned}$$

The most restrictive setup relationship with the default single-cycle analysis, that is, a setup relationship with an end multicycle setup assignment of one, is 10 ns.

Figure 7-13 shows the setup report for the default setup in the TimeQuest analyzer with the launch and latch edges highlighted.

Figure 7-13. Setup Report

Path #1: Setup slack is 9.077

Data Arrival Path						
	Total	Incr	RF	Type	Fanout	Element
1:	0.000	0.000				launch edge time
2	2.522	2.522	R			clock network delay
3	2.606	0.084		uTco	1	src
4	2.606	0.000	RR	CELL	1	srcfq
5	2.864	0.258	RR	IC	1	dst~feeder\dataf
6	2.960	0.096	RR	CELL	1	dst~feeder/combout
7	2.960	0.000	RR	IC	1	dstld
8	3.065	0.105	RR	CELL	1	dst

Path #1: Setup slack is 9.077

Data Required Path						
	Total	Incr	RF	Type	Fanout	Element
1:	10.000	10.000				latch edge time
2	12.248	2.248	R			clock network delay
3	12.142	-0.106		uTsu	1	dst

Property	Value
1 From Node	src
2 To Node	dst
3 Launch Clock	clk_src
4 Latch Clock	clk_dst
5 Data Arrival Time	3.065
6 Data Required Time	12.142
7 Slack	9.077

Equation 7-7 shows the calculation that the TimeQuest analyzer performs to determine the hold check. Both hold checks are equivalent.

Equation 7-7. Hold Check

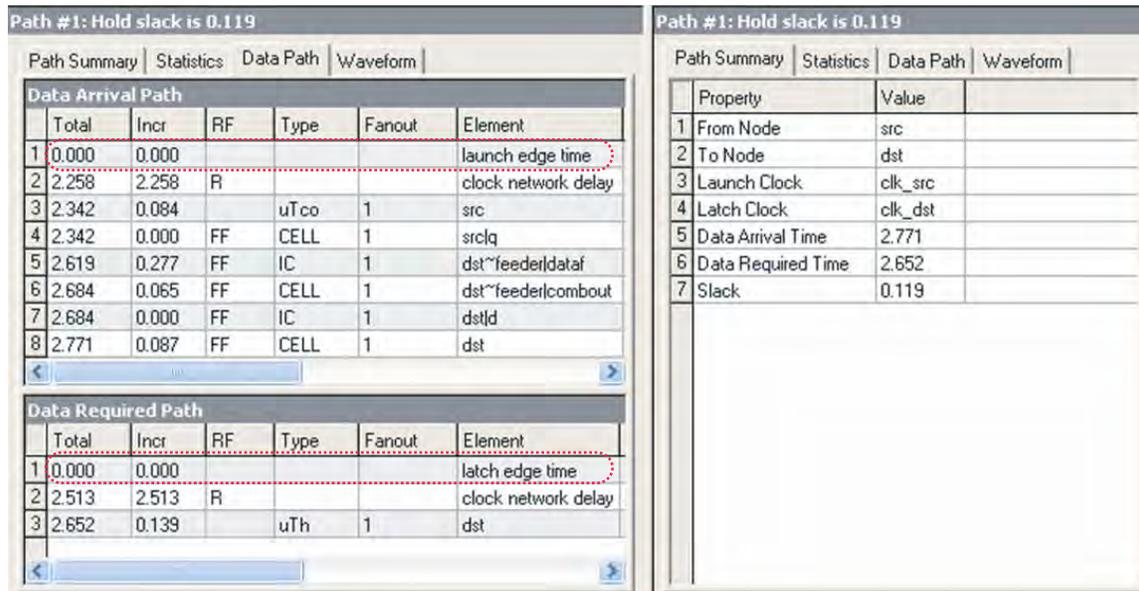
$$\begin{aligned}
 \text{hold check 1} &= \text{next launch edge} - \text{current latch edge} \\
 &= 0\text{ns} - 0\text{ns} \\
 &= 0\text{ns}
 \end{aligned}$$

$$\begin{aligned}
 \text{hold check 2} &= 10\text{ ns} - 10\text{ ns} \\
 &= 0\text{ ns}
 \end{aligned}$$

The most restrictive hold relationship with the default single-cycle analysis, that a hold relationship with an end multicycle hold assignment of zero, is 0 ns.

Figure 7-14 shows the hold report for the default setup in the TimeQuest analyzer with the launch and latch edges highlighted.

Figure 7-14. Hold Report



End Multicycle Setup = 2 and End Multicycle Hold = 0

In this example, the end multicycle setup assignment value is two, and the end multicycle hold assignment value is zero. Example 7-9 shows the multicycle exceptions applied to the register-to-register design for this example.

Example 7-9. Multicycle Exceptions

```
set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_\
dst] -setup -end 2
```

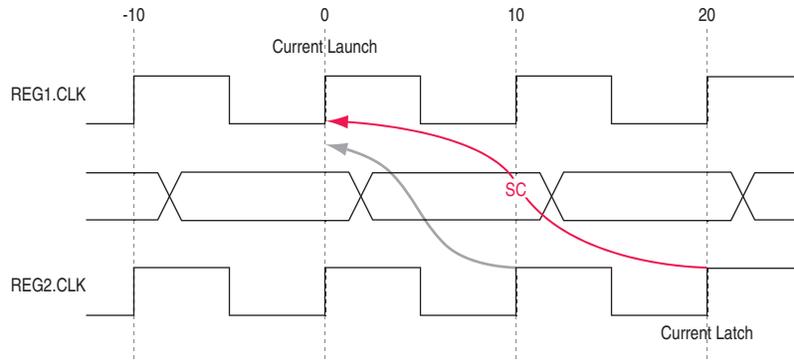


An end multicycle hold value is not required because the default end multicycle hold value is zero.

In this example, the setup relationship is relaxed by a full clock period by moving the latch edge to the next latch edge. The hold analysis is unchanged from the default settings.

Figure 7-15 shows the setup timing diagram. The latch edge is a clock cycle later than in the default single-cycle analysis.

Figure 7-15. Setup Timing Diagram



Equation 7-8 shows the calculation that the TimeQuest analyzer performs to determine the setup check.

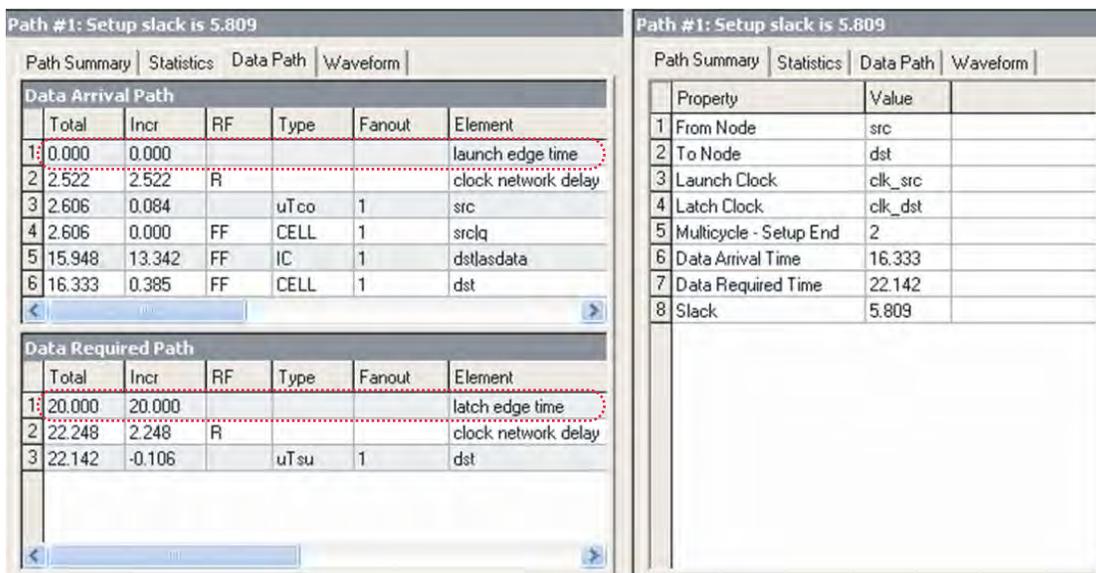
Equation 7-8. Setup Check

$$\begin{aligned}
 \text{setup check} &= \text{current latch edge} - \text{closest previous launch edge} \\
 &= 20 \text{ ns} - 0 \text{ ns} \\
 &= 20 \text{ ns}
 \end{aligned}$$

The most restrictive setup relationship with an end multicycle setup assignment of two is 20 ns.

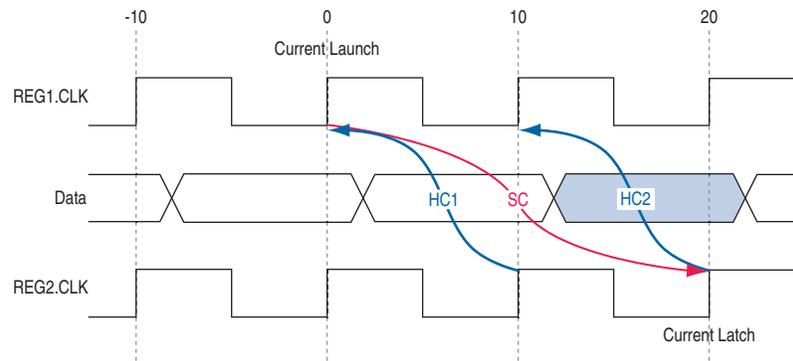
Figure 7-16 shows the setup report in the TimeQuest analyzer with the launch and latch edges highlighted.

Figure 7-16. Setup Report



Because the multicycle hold latch and launch edges are the same as the results of hold analysis with the default settings, the multicycle hold analysis in this example is equivalent to the single-cycle hold analysis. Figure 7-17 shows the timing diagram for the hold checks for this example. The hold checks are relative to the setup check. Usually, the TimeQuest analyzer performs hold checks on every possible setup check, not only on the most restrictive setup check edges.

Figure 7-17. Hold Timing Diagram



Equation 7-9 shows the calculation that the TimeQuest analyzer performs to determine the hold check. Both hold checks are equivalent.

Equation 7-9. Hold Check

$$\begin{aligned}
 \text{hold check 1} &= \text{current launch edge} - \text{previous latch edge} \\
 &= 0 \text{ ns} - 10 \text{ ns} \\
 &= -10 \text{ ns}
 \end{aligned}$$

$$\begin{aligned}
 \text{hold check 2} &= \text{next launch edge} - \text{current latch edge} \\
 &= 10 \text{ ns} - 20 \text{ ns} \\
 &= -10 \text{ ns}
 \end{aligned}$$

The most restrictive hold relationship with an end multicycle setup assignment value of two and an end multicycle hold assignment value of zero is 10 ns.

Figure 7-18 shows the hold report for this example in the TimeQuest analyzer with the launch and latch edges highlighted.

Figure 7-18. Hold Report

Path #1: Hold slack is 3.196						
Data Arrival Path						
	Total	Incr	RF	Type	Fanout	Element
1:	0.000	0.000				launch edge time
2:	2.258	2.258	R			clock network delay
3:	2.342	0.084		uTco	1	src
4:	2.342	0.000	RR	CELL	1	srcdq
5:	15.606	13.264	RR	IC	1	dstlasdata
6:	15.848	0.242	RR	CELL	1	dst

Path #1: Hold slack is 3.196						
Data Required Path						
	Total	Incr	RF	Type	Fanout	Element
1:	10.000	10.000				latch edge time
2:	12.513	2.513	R			clock network delay
3:	12.652	0.139		uTh	1	dst

Path Summary		
Property	Value	
1 From Node	src	
2 To Node	dst	
3 Launch Clock	clk_src	
4 Latch Clock	clk_dst	
5 Multicycle - Setup End	2	
6 Data Arrival Time	15.848	
7 Data Required Time	12.652	
8 Slack	3.196	

End Multicycle Setup = 1 and End Multicycle Hold = 1

In this example, the end multicycle setup assignment value is one, and the end multicycle hold assignment value is one. Example 7-10 shows the multicycle exceptions applied to the register-to-register design for this example.

Example 7-10. Multicycle Exceptions

```
set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_\
dst] -hold -end 1
```

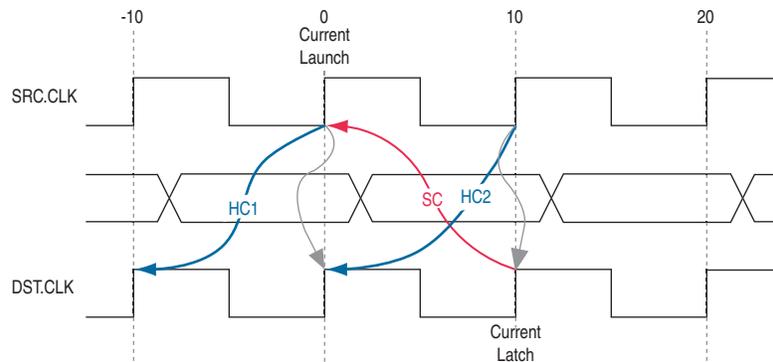


An end multicycle setup value is not required because the default end multicycle setup value is one.

In this example, the hold relationship is relaxed by one clock period by moving the latch edge to the previous latch edge. The setup analysis is unchanged from the default settings.

Figure 7-19 shows the setup timing diagram. The latch edge is the same as the default single-cycle analysis.

Figure 7-19. Setup Timing Diagram



Equation 7-10 shows the calculation that the TimeQuest analyzer performs to determine the setup check.

Equation 7-10. Setup Check

$$\begin{aligned} \text{setup check} &= \text{current latch edge} - \text{closest previous launch edge} \\ &= 10 \text{ ns} - 0 \text{ ns} \\ &= 10 \text{ ns} \end{aligned}$$

The most restrictive setup relationship with an end multicycle setup assignment of one is 10 ns.

Figure 7-20 shows the setup report in the TimeQuest analyzer with the launch and latch edges highlighted.

Figure 7-20. Setup Report

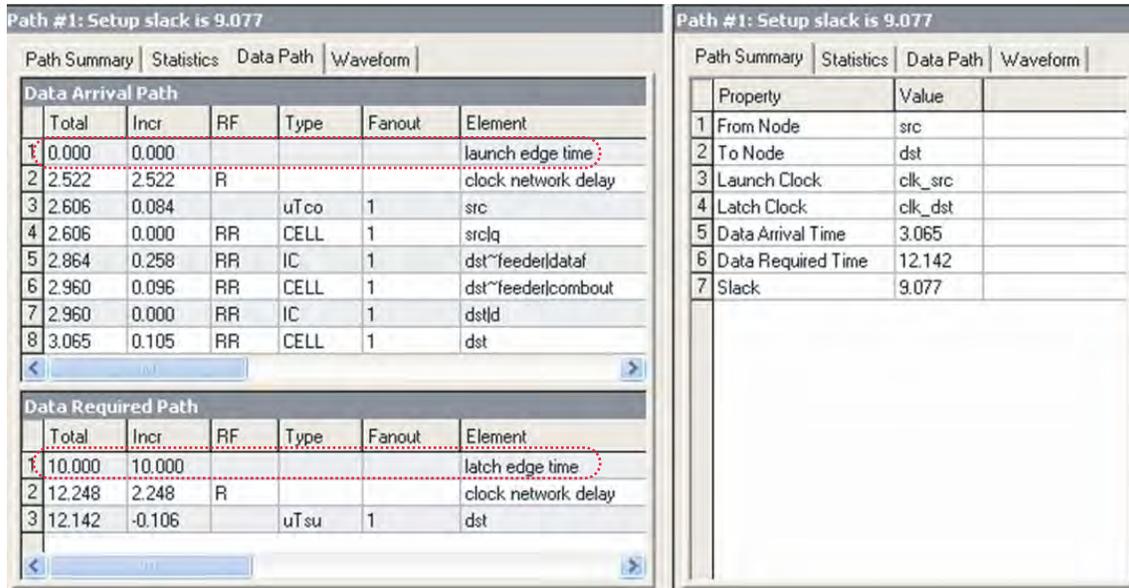
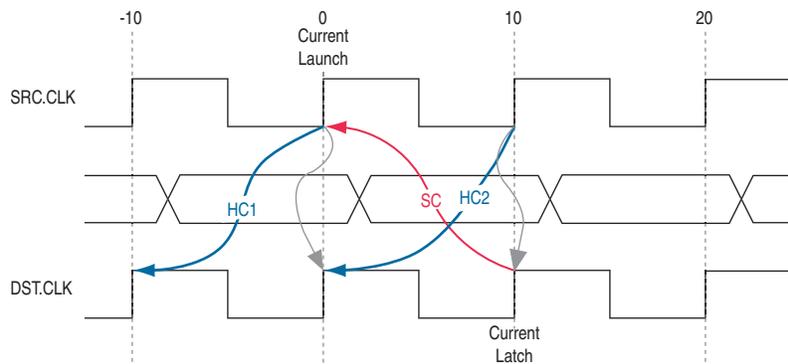


Figure 7-21 shows the timing diagram for the hold checks for this example. The hold checks are relative to the setup check.

Figure 7-21. Hold Timing Diagram



Equation 7-11 shows the calculation that the TimeQuest analyzer performs to determine the hold check. Both hold checks are equivalent.

Equation 7-11. Hold Check

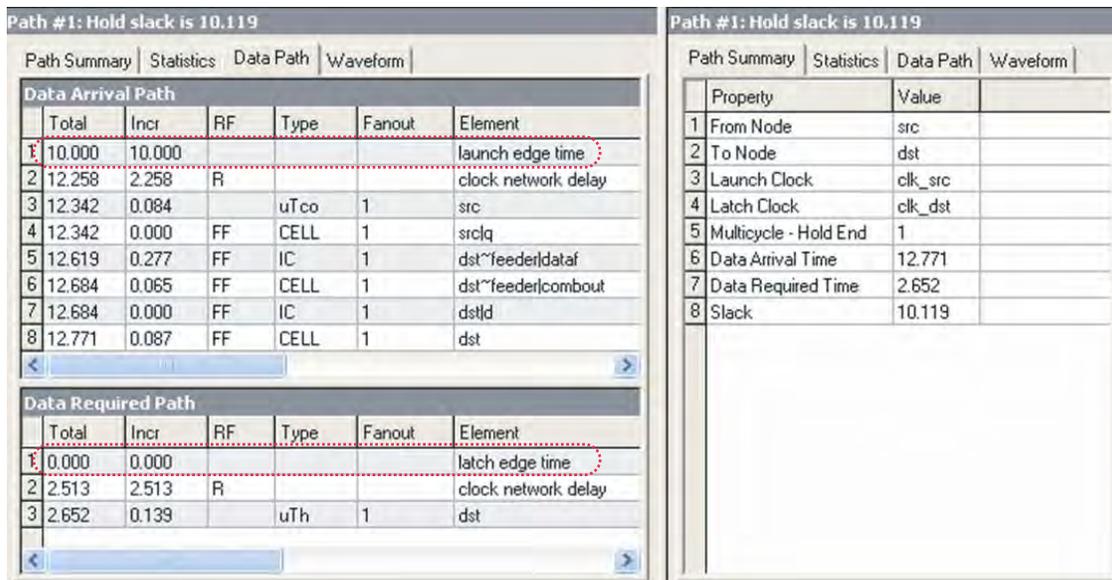
$$\begin{aligned} \text{hold check 1} &= \text{current launch edge} - \text{previous latch edge} \\ &= 0 \text{ ns} - (-10 \text{ ns}) \\ &= 10 \text{ ns} \end{aligned}$$

$$\begin{aligned} \text{hold check 2} &= \text{next launch edge} - \text{current latch edge} \\ &= 10 \text{ ns} - 0 \text{ ns} \\ &= 10 \text{ ns} \end{aligned}$$

The most restrictive hold relationship with an end multicycle setup assignment value of one and an end multicycle hold assignment value of one is 10 ns.

Figure 7-22 shows the hold report for this example in the TimeQuest analyzer with the launch and latch edges highlighted.

Figure 7-22. Hold Report



End Multicycle Setup = 2 and End Multicycle Hold = 1

In this example, the end multicycle setup assignment value is two, and the end multicycle hold assignment value is one. Example 7-11 shows the multicycle exceptions applied to the register-to-register design for this example.

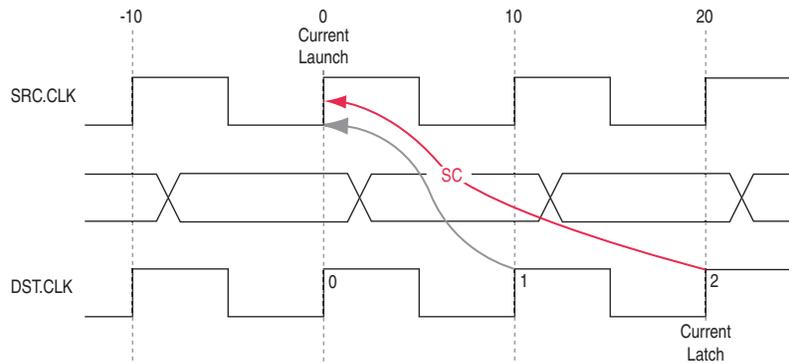
Example 7-11. Multicycle Exceptions

```
set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_dst]
-setup -end 2
set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_dst]
-hold -end 1
```

In this example, the setup relationship is relaxed by two clock periods by moving the latch edge to the left two clock periods. The hold relationship is relaxed by a full period by moving the latch edge to the previous latch edge.

Figure 7-23 shows the setup timing diagram.

Figure 7-23. Setup Timing Diagram



Equation 7-12 shows the calculation that the TimeQuest analyzer performs to determine the setup check.

Equation 7-12. Setup Check

$$\begin{aligned}
 \text{setup check} &= \text{current latch edge} - \text{closest previous launch edge} \\
 &= 20 \text{ ns} - 0 \text{ ns} \\
 &= 20 \text{ ns}
 \end{aligned}$$

The most restrictive hold relationship with an end multicycle setup assignment value of two is 20 ns.

Figure 7-24 shows the setup report for this example in the TimeQuest analyzer with the launch and latch edges highlighted.

Figure 7-24. Setup Report

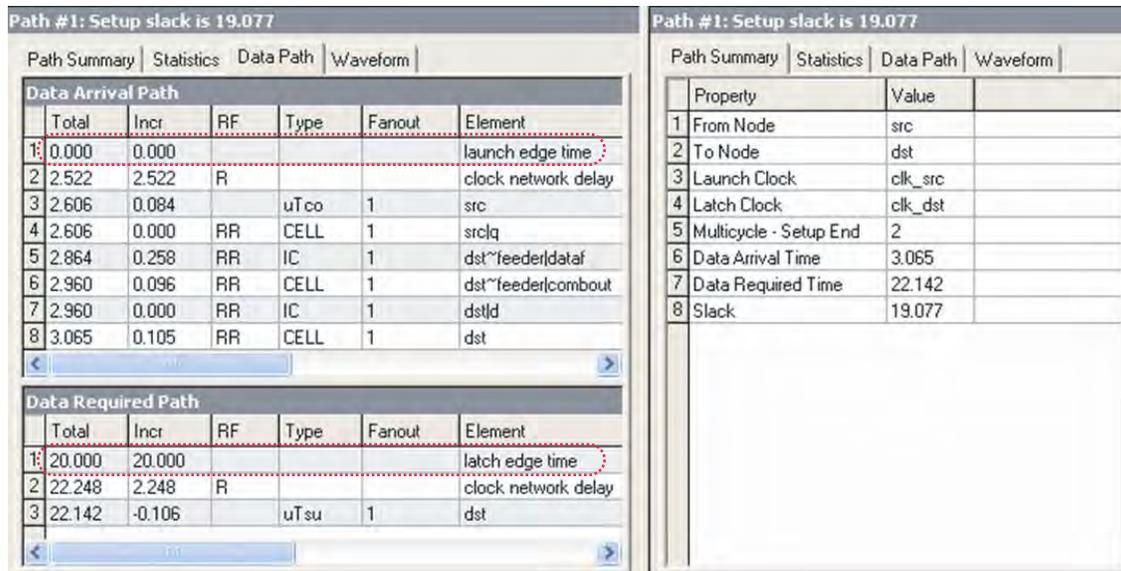
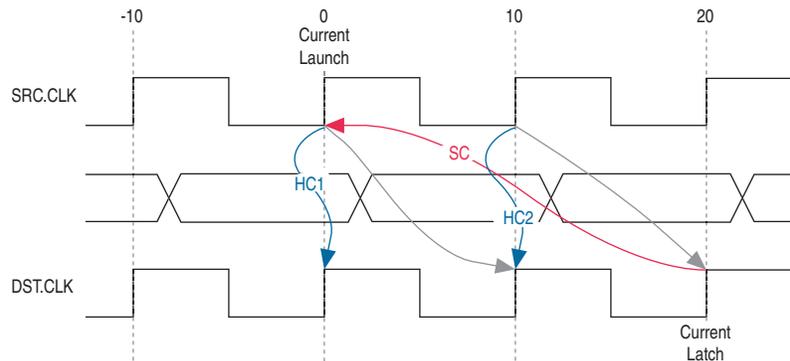


Figure 7-25 shows the timing diagram for the hold checks for this example. The hold checks are relative to the setup check.

Figure 7-25. Hold Timing Diagram



Equation 7-13 shows the calculation that the TimeQuest analyzer performs to determine the hold check. Both hold checks are equivalent.

Equation 7-13. Hold Check

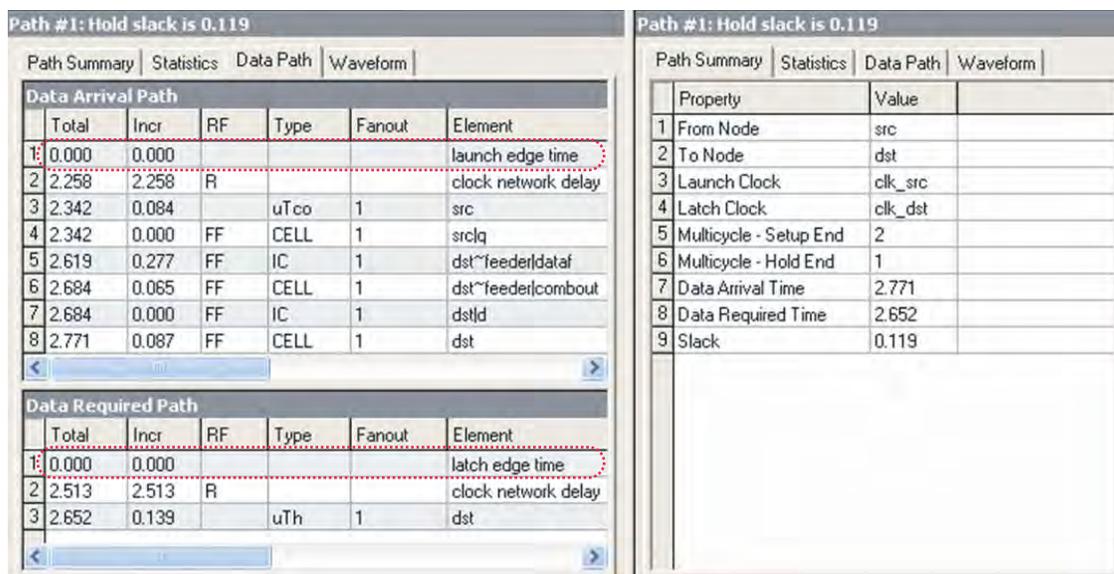
$$\begin{aligned} \text{hold check 1} &= \text{current launch edge} - \text{previous latch edge} \\ &= 0 \text{ ns} - 0 \text{ ns} \\ &= 0 \text{ ns} \end{aligned}$$

$$\begin{aligned} \text{hold check 2} &= \text{next launch edge} - \text{current latch edge} \\ &= 10 \text{ ns} - 10 \text{ ns} \\ &= 0 \text{ ns} \end{aligned}$$

The most restrictive hold relationship with an end multicycle setup assignment value of two and an end multicycle hold assignment value of one is 0 ns.

Figure 7-26 shows the hold report for this example in the TimeQuest analyzer with the launch and latch edges highlighted.

Figure 7-26. Hold Report



Start Multicycle Setup = 2 and Start Multicycle Hold = 0

In this example, the start multicycle setup assignment value is two, and the start multicycle hold assignment value is zero. Example 7-12 shows the multicycle exceptions applied to the register-to-register design for this example.

Example 7-12. Multicycle Exceptions

```
set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_\
dst] -setup -start 2
```

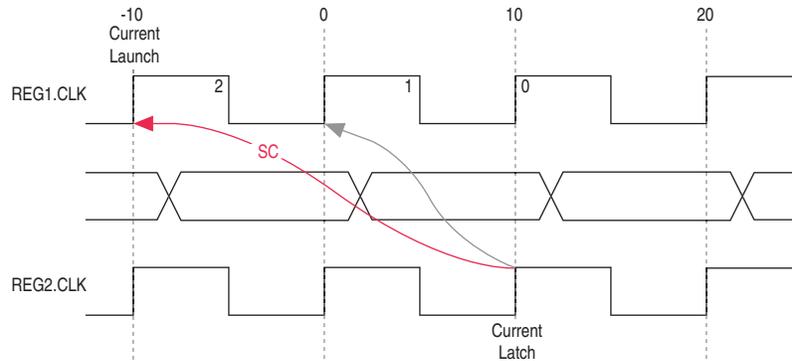


A start multicycle hold value is not required because the default start multicycle hold value is zero.

In this example, the setup relationship is relaxed by moving the latch edge to the left two clock periods. Hold analysis is unchanged from the default settings.

Figure 7-27 shows the setup timing diagram.

Figure 7-27. Setup Timing Diagram



Equation 7-14 shows the calculation that the TimeQuest analyzer performs to determine the setup check.

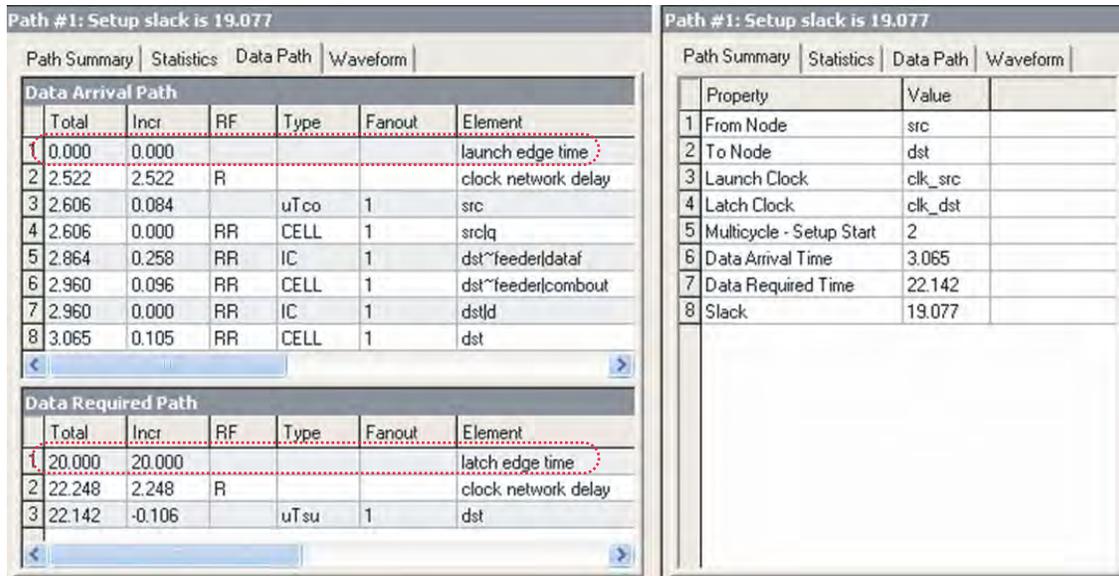
Equation 7-14. Setup Check

$$\begin{aligned} \text{setup check} &= \text{current latch edge} - \text{closest previous launch edge} \\ &= 10 \text{ ns} - (-10 \text{ ns}) \\ &= 20 \text{ ns} \end{aligned}$$

The most restrictive hold relationship with a start multicycle setup assignment value of two is 20 ns.

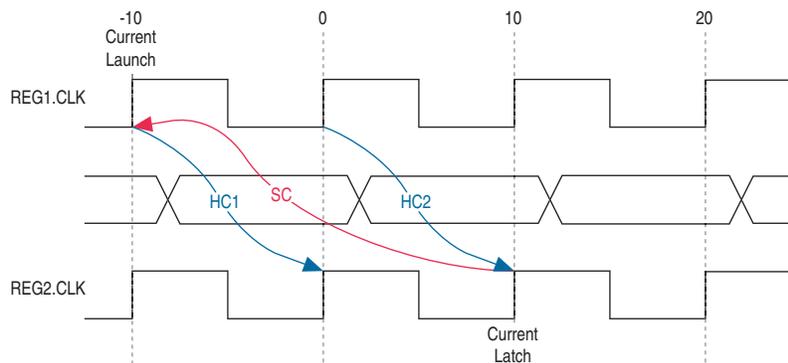
Figure 7-28 shows the setup report for this example in the TimeQuest analyzer with the launch and latch edges highlighted.

Figure 7-28. Setup Report



Because the multicycle hold latch and launch edges are the same as the results of hold analysis with the default settings, the multicycle hold analysis in this example is equivalent to the single-cycle hold analysis. Figure 7-29 shows the timing diagram for the hold checks for this example.

Figure 7-29. Hold Timing Diagram



Equation 7-15 shows the calculation that the TimeQuest analyzer performs to determine the hold check. Both hold checks are equivalent.

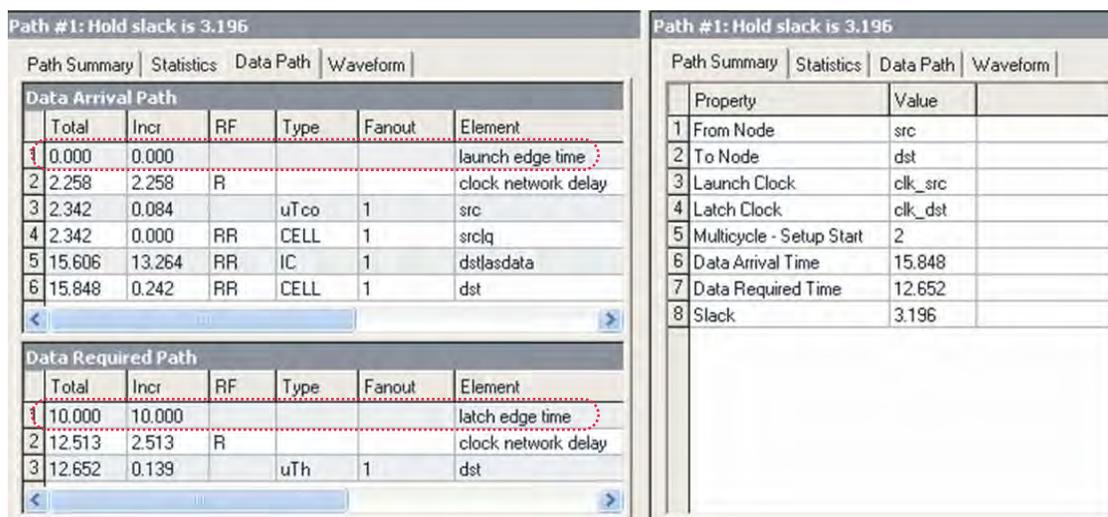
Equation 7-15. Hold Check

$$\begin{aligned} \text{hold check 1} &= \text{current launch edge} - \text{previous latch edge} \\ &= 0 \text{ ns} - 10 \text{ ns} \\ &= -10 \text{ ns} \\ \\ \text{hold check 2} &= \text{next launch edge} - \text{current latch edge} \\ &= 10 \text{ ns} - 0 \text{ ns} \\ &= 10 \text{ ns} \end{aligned}$$

The most restrictive hold relationship with a start multicycle setup assignment value of two and a start multicycle hold assignment value of zero is 10 ns.

Figure 7-30 shows the hold report for this example in the TimeQuest analyzer with the launch and latch edges highlighted.

Figure 7-30. Hold Report



Start Multicycle Setup = 1 and Start Multicycle Hold = 1

In this example, the start multicycle setup assignment value is one, and the start multicycle hold assignment value is one. Example 7-13 shows the multicycle exceptions applied to the register-to-register design for this example.

Example 7-13. Multicycle Exceptions

```
set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_\
dst] -hold -start 1
```

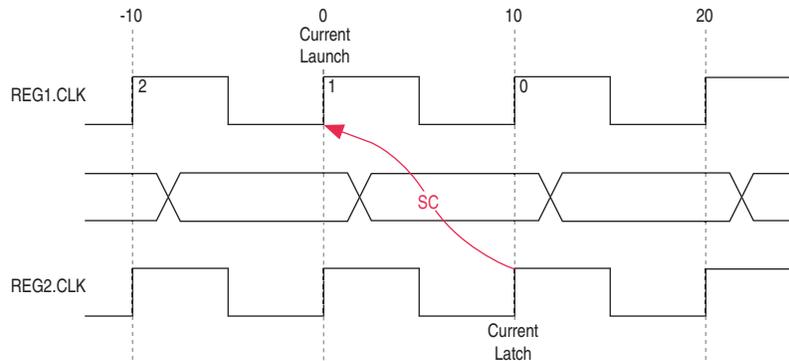


A start multicycle setup value is not required, because the default start multicycle hold value is one.

In this example, the hold relationship is relaxed by one clock period by moving the launch edge to the left.

Figure 7-31 shows the setup timing diagram.

Figure 7-31. Setup Timing Diagram



Equation 7-16 shows the calculation that the TimeQuest analyzer performs to determine the setup check.

Equation 7-16. Setup Check

$$\begin{aligned} \text{setup check} &= \text{current latch edge} - \text{closest previous launch edge} \\ &= 10 \text{ ns} - 0 \text{ ns} \\ &= 10 \text{ ns} \end{aligned}$$

The most restrictive setup relationship with a start multicycle setup assignment of one is 10 ns.

Figure 7-32 shows the setup report in the TimeQuest analyzer with the launch and latch edges highlighted.

Figure 7-32. Setup Report

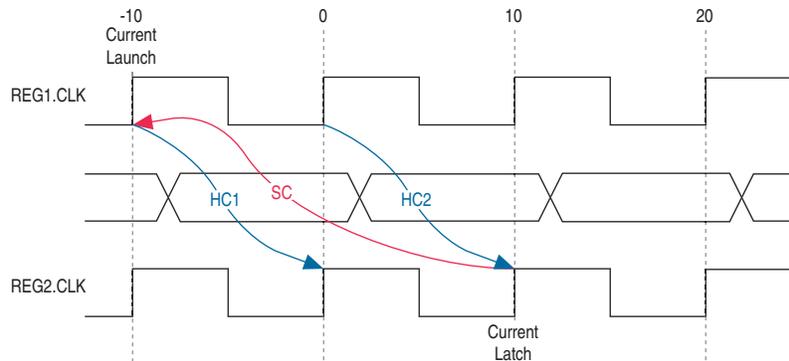
Path #1: Setup slack is 8.891							
Data Arrival Path		Total	Incr	RF	Type	Fanout	Element
1		0.000	0.000				launch edge time
2		2.584	2.584	R			clock network delay
3		2.668	0.084		uTco	1	src
4		2.668	0.000	FF	CELL	1	src1q
5		2.944	0.276	FF	IC	1	dst1asdata
6		3.329	0.385	FF	CELL	1	dst

Path #1: Setup slack is 8.891							
Data Required Path		Total	Incr	RF	Type	Fanout	Element
1		10.000	10.000				latch edge time
2		12.326	2.326	R			clock network delay
3		12.220	-0.106		uTsu	1	dst

Path #1: Setup slack is 8.891		
Property	Value	
1 From Node	src	
2 To Node	dst	
3 Launch Clock	clk_src	
4 Latch Clock	clk_dst	
5 Data Arrival Time	3.329	
6 Data Required Time	12.220	
7 Slack	8.891	

Figure 7-33 shows the timing diagram for the hold checks for this example. The hold checks are relative to the setup check.

Figure 7-33. Hold Timing Diagram



Equation 7-17 shows the calculation that the TimeQuest analyzer performs to determine the hold check. Both hold checks are equivalent.

Equation 7-17. Hold Check

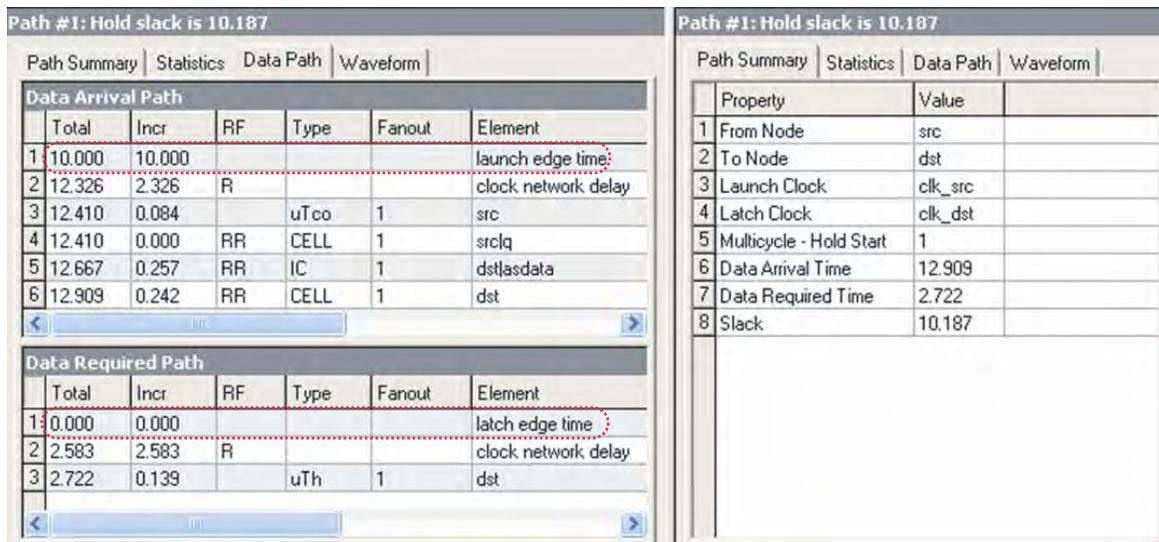
$$\begin{aligned}
 \text{hold check 1} &= \text{current launch edge} - \text{previous latch edge} \\
 &= 10 \text{ ns} - 0 \text{ ns} \\
 &= 10 \text{ ns}
 \end{aligned}$$

$$\begin{aligned}
 \text{hold check 2} &= \text{next launch edge} - \text{current latch edge} \\
 &= 20 \text{ ns} - 10 \text{ ns} \\
 &= 10 \text{ ns}
 \end{aligned}$$

The most restrictive hold relationship with a start multicycle setup assignment value of one and a start multicycle hold assignment value of one is 10 ns.

Figure 7-34 shows the hold report for this example in the TimeQuest analyzer with the launch and latch edges highlighted.

Figure 7-34. Hold Report



Start Multicycle Setup = 2 and Start Multicycle Hold = 1

In this example, the start multicycle setup assignment value is two, and the start multicycle hold assignment value is one. Example 7-14 shows the multicycle exceptions applied to the register-to-register design for this example.

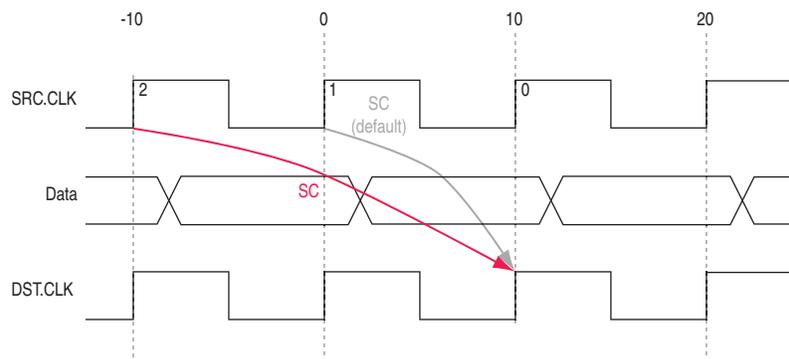
Example 7-14. Multicycle Exceptions

```
set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_dst]
-setup -start 2
set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_dst]
-hold -start 1
```

In this example, the setup relationship is relaxed by two clock periods by moving the launch edge to the left two clock periods.

Figure 7-35 shows the setup timing diagram.

Figure 7-35. Setup Timing Diagram



Equation 7-18 shows the calculation that the TimeQuest analyzer performs to determine the setup check.

Equation 7-18. Setup Check

$$\begin{aligned}
 \text{setup check} &= \text{current latch edge} - \text{closest previous launch edge} \\
 &= 10 \text{ ns} - (-10 \text{ ns}) \\
 &= 20 \text{ ns}
 \end{aligned}$$

The most restrictive hold relationship with a start multicycle setup assignment value of two is 20 ns.

Figure 7-36 shows the setup report for this example in the TimeQuest analyzer with the launch and latch edges highlighted.

Figure 7-36. Setup Report

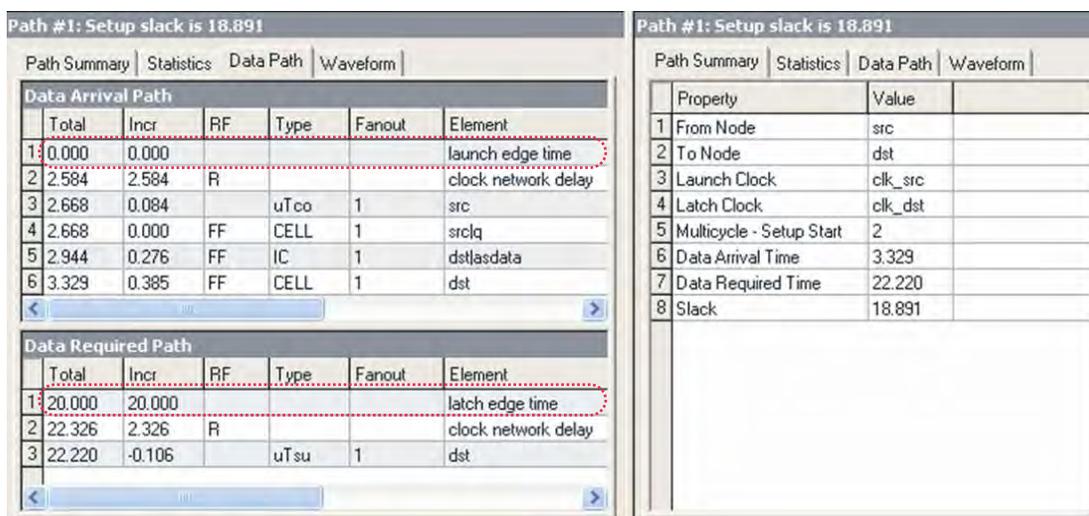
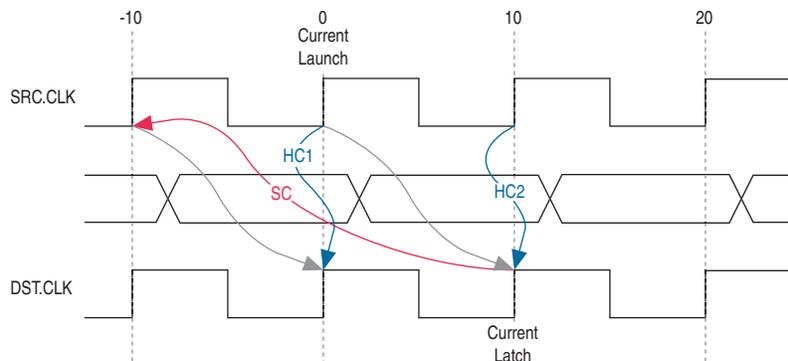


Figure 7-37 shows the timing diagram for the hold checks for this example. The hold checks are relative to the setup check.

Figure 7-37. Hold Timing Diagram



Equation 7-19 shows the calculation that the TimeQuest analyzer performs to determine the hold check. Both hold checks are equivalent.

Equation 7-19. Hold Check

$$\begin{aligned} \text{hold check 1} &= \text{current launch edge} - \text{previous latch edge} \\ &= 0 \text{ ns} - 0 \text{ ns} \\ &= 0 \text{ ns} \end{aligned}$$

$$\begin{aligned} \text{hold check 2} &= \text{next launch edge} - \text{current latch edge} \\ &= 10 \text{ ns} - 10 \text{ ns} \\ &= 0 \text{ ns} \end{aligned}$$

The most restrictive hold relationship with a start multicycle setup assignment value of two and a start multicycle hold assignment value of one is 0 ns.

Figure 7-38 shows the hold report for this example in the TimeQuest analyzer with the launch and latch edges highlighted.

Figure 7-38. Hold Report

Path #1: Hold slack is 0.187

Data Arrival Path						
	Total	Incr	RF	Type	Fanout	Element
1	0.000	0.000				launch edge time
2	2.326	2.326	R			clock network delay
3	2.410	0.084		uTco	1	src
4	2.410	0.000	RR	CELL	1	srclq
5	2.667	0.257	RR	IC	1	dstlasdata
6	2.909	0.242	RR	CELL	1	dst

Data Required Path						
	Total	Incr	RF	Type	Fanout	Element
1	0.000	0.000				latch edge time
2	2.583	2.583	R			clock network delay
3	2.722	0.139		uTh	1	dst

Path #1: Hold slack is 0.187

Path Summary		
Property	Value	
1 From Node	src	
2 To Node	dst	
3 Launch Clock	clk_src	
4 Latch Clock	clk_dst	
5 Multicycle - Setup Start	2	
6 Multicycle - Hold Start	1	
7 Data Arrival Time	2.909	
8 Data Required Time	2.722	
9 Slack	0.187	

Application of Multicycle Exceptions

This section shows the following examples of applications of multicycle exceptions:

- “Same Frequency Clocks with Destination Clock Offset” on page 7-35
- “The Destination Clock Frequency is a Multiple of the Source Clock Frequency” on page 7-37
- “The Destination Clock Frequency is a Multiple of the Source Clock Frequency with an Offset” on page 7-40
- “The Source Clock Frequency is a Multiple of the Destination Clock Frequency” on page 7-42
- “The Source Clock Frequency is a Multiple of the Destination Clock Frequency with an Offset” on page 7-45

Each example explains how the multicycle exceptions affect the default setup and hold analysis in the TimeQuest analyzer. All of the examples are between related clock domains. If your design contains related clocks, such as PLL clocks, and paths between related clock domains, you can apply multicycle constraints.

Same Frequency Clocks with Destination Clock Offset

In this example, the source and destination clocks have the same frequency, but the destination clock is offset with a positive phase shift. Both the source and destination clocks have a period of 10 ns. The destination clock has a positive phase shift of 2 ns with respect to the source clock. Figure 7-39 shows an example of a design with same frequency clocks and a destination clock offset.

Figure 7-39. Same Frequency Clocks with Destination Clock Offset

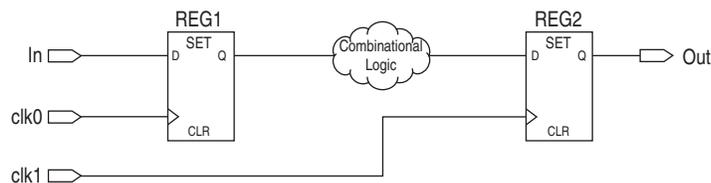
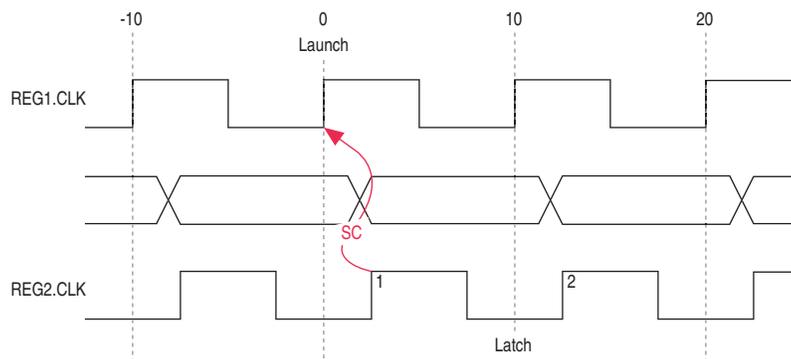


Figure 7-40 shows the timing diagram for default setup check analysis performed by the TimeQuest analyzer.

Figure 7-40. Setup Timing Diagram



Equation 7-20 shows the calculation that the TimeQuest analyzer performs to determine the setup check.

Equation 7-20. Setup Check

$$\begin{aligned} \text{setup check} &= \text{current latch edge} - \text{closest previous launch edge} \\ &= 2 \text{ ns} - 0 \text{ ns} \\ &= 2 \text{ ns} \end{aligned}$$

The setup relationship shown in Figure 7-40 is too pessimistic and is not the setup relationship required for typical designs. To correct the default analysis, you must use an end multicycle setup exception of two. Example 7-15 shows the multicycle exception used to correct the default analysis in this example.

Example 7-15. Multicycle Exceptions

```
set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_dst]
-setup -end 2
```

Figure 7-41 shows the timing diagram for the preferred setup relationship for this example.

Figure 7-41. Preferred Setup Relationship

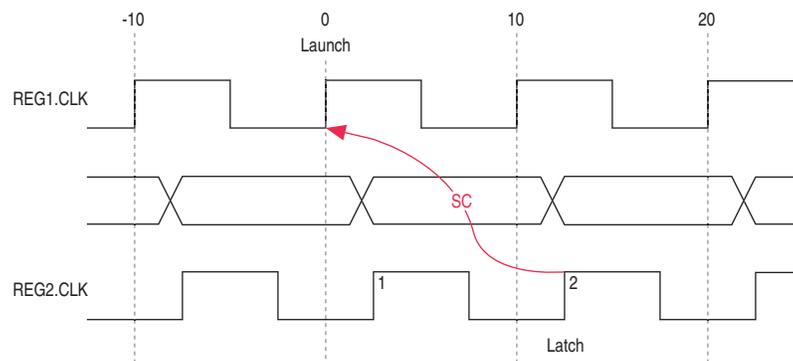
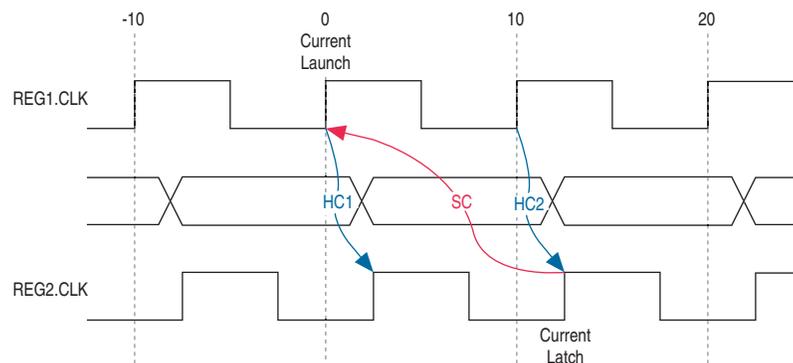


Figure 7-42 shows the timing diagram for default hold check analysis performed by the TimeQuest analyzer with an end multicycle setup value of two.

Figure 7-42. Default Hold Check



Equation 7-21 shows the calculation that the TimeQuest analyzer performs to determine the hold check.

Equation 7-21. Hold Check

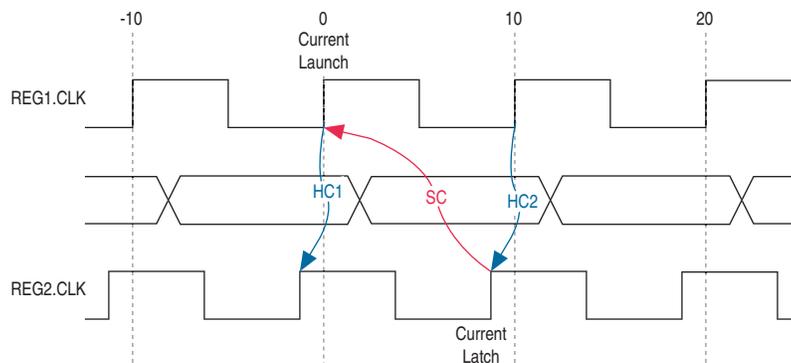
$$\begin{aligned} \text{hold check 1} &= \text{current launch edge} - \text{previous latch edge} \\ &= 0 \text{ ns} - 2 \text{ ns} \\ &= -2 \text{ ns} \end{aligned}$$

$$\begin{aligned} \text{hold check 2} &= \text{next launch edge} - \text{current latch edge} \\ &= 10 \text{ ns} - 12 \text{ ns} \\ &= 2 \text{ ns} \end{aligned}$$

In this example, the default hold analysis returns the preferred hold requirements and no multicycle hold exceptions are required.

Figure 7-43 shows the associated setup and hold analysis if the phase shift is -2 ns. In this example, the default hold analysis is correct for the negative phase shift of 2 ns, and no multicycle exceptions are required.

Figure 7-43. Negative Phase Shift



The Destination Clock Frequency is a Multiple of the Source Clock Frequency

In this example, the destination clock frequency value of 5 ns is an integer multiple of the source clock frequency of 10 ns. The destination clock frequency can be an integer multiple of the source clock frequency when a PLL is used to generate both clocks with a phase shift applied to the destination clock. Figure 7-44 shows an example of a design where the destination clock frequency is a multiple of the source clock frequency.

Figure 7-44. Destination Clock is Multiple of Source Clock

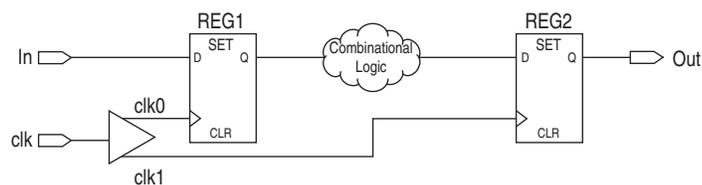
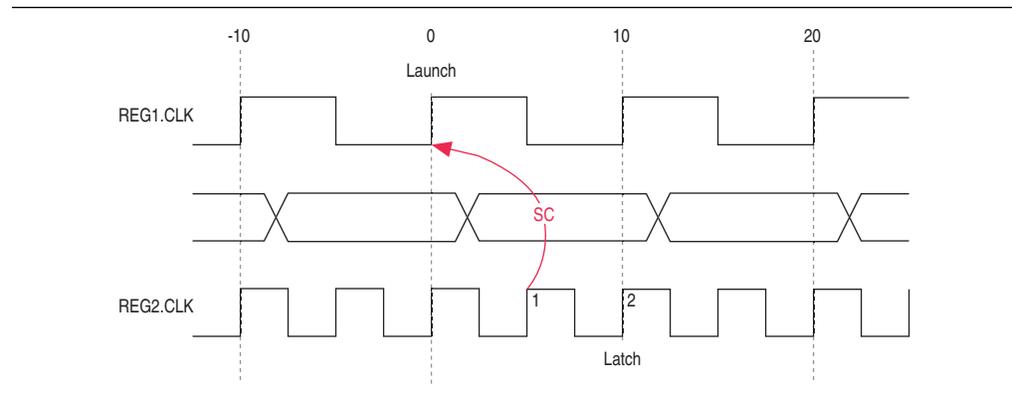


Figure 7-45 shows the timing diagram for default setup check analysis performed by the TimeQuest analyzer.

Figure 7-45. Setup Timing Diagram



Equation 7-22 shows the calculation that the TimeQuest analyzer performs to determine the setup check.

Equation 7-22. Setup Check

$$\begin{aligned} \text{setup check} &= \text{current latch edge} - \text{closest previous launch edge} \\ &= 5 \text{ ns} - 0 \text{ ns} \\ &= 5 \text{ ns} \end{aligned}$$

The setup relationship shown in Figure 7-45 demonstrates that the data does not need to be captured at edge one, but can be captured at edge two; therefore, you can relax the setup requirement. To correct the default analysis, you must shift the latch edge by one clock period with an end multicycle setup exception of two. Example 7-16 shows the multicycle exception used to correct the default analysis in this example.

Example 7-16. Multicycle Exceptions

```
set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_dst]
-setup -end 2
```

Figure 7-46 shows the timing diagram for the preferred setup relationship for this example.

Figure 7-46. Preferred Setup Analysis

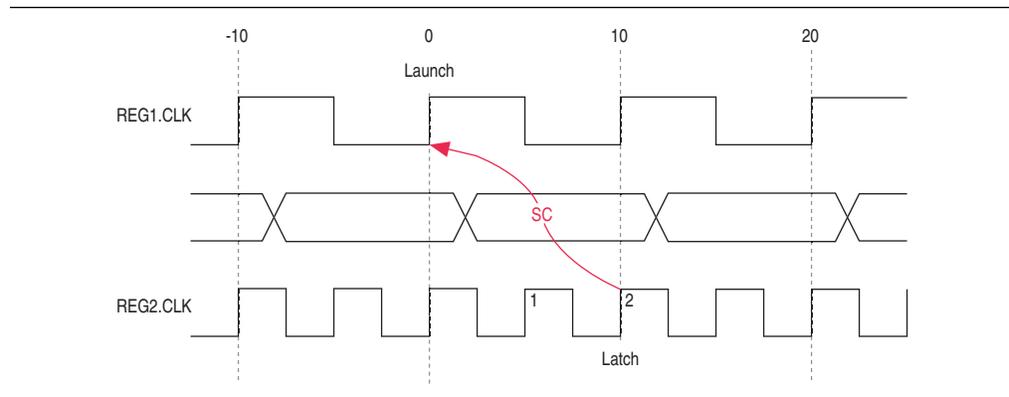
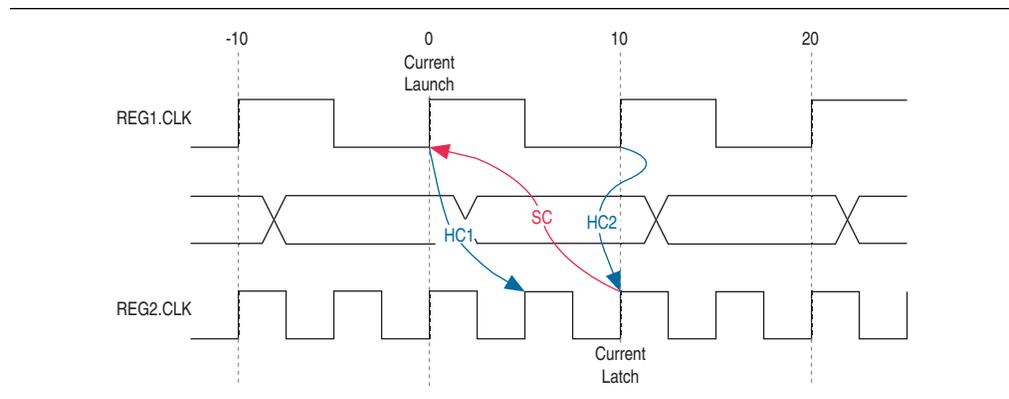


Figure 7-47 shows the timing diagram for default hold check analysis performed by the TimeQuest analyzer with an end multicycle setup value of two.

Figure 7-47. Default Hold Check



Equation 7-23 shows the calculation that the TimeQuest analyzer performs to determine the hold check.

Equation 7-23. Hold Check

$$\begin{aligned}
 \text{hold check 1} &= \text{current launch edge} - \text{previous latch edge} \\
 &= 0 \text{ ns} - 5 \text{ ns} \\
 &= -5 \text{ ns}
 \end{aligned}$$

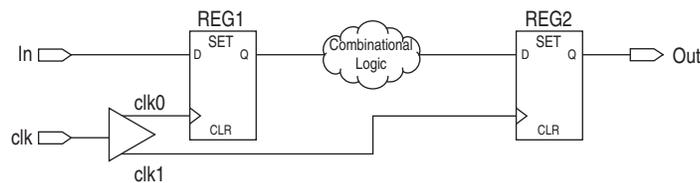
$$\begin{aligned}
 \text{hold check 2} &= \text{next launch edge} - \text{current latch edge} \\
 &= 10 \text{ ns} - 10 \text{ ns} \\
 &= 0 \text{ ns}
 \end{aligned}$$

In this example, hold check one is too restrictive. The data is launched by the edge at 0 ns and should check against the data captured by the previous latch edge at 0 ns, which does not occur in hold check one. To correct the default analysis, you must use an end multicycle hold exception of one.

The Destination Clock Frequency is a Multiple of the Source Clock Frequency with an Offset

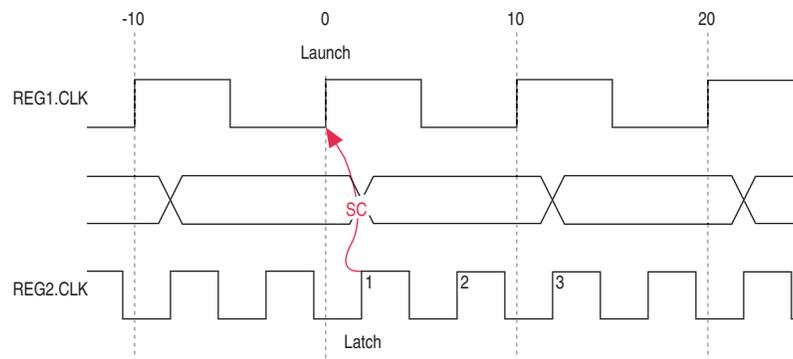
This example is a combination of the previous two examples. The destination clock frequency is an integer multiple of the source clock frequency and the destination clock has a positive phase shift. The destination clock frequency is 5 ns and the source clock frequency is 10 ns. The destination clock also has a positive offset of 2 ns with respect to the source clock. The destination clock frequency can be an integer multiple of the source clock frequency with an offset when a PLL is used to generate both clocks with a phase shift applied to the destination clock. [Figure 7-48](#) shows an example of a design in which the destination clock frequency is a multiple of the source clock frequency with an offset.

Figure 7-48. Destination Clock is Multiple of Source Clock with Offset



[Figure 7-49](#) shows the timing diagram for default setup check analysis performed by the TimeQuest analyzer.

Figure 7-49. Setup Timing Diagram



[Equation 7-24](#) shows the calculation that the TimeQuest analyzer performs to determine the setup check.

Equation 7-24. Setup Check

$$\begin{aligned} \text{setup check} &= \text{current latch edge} - \text{closest previous launch edge} \\ &= 2 \text{ ns} - 0 \text{ ns} \\ &= 2 \text{ ns} \end{aligned}$$

The setup relationship shown in [Figure 7-49](#) demonstrates that the data does not need to be captured at edge one, but can be captured at edge two; therefore, you can relax the setup requirement. To correct the default analysis, you must shift the latch edge by one clock period with an end multicycle setup exception of three.

Example 7-17 shows the multicycle exception used to correct the default analysis in this example.

Example 7-17. Multicycle Exceptions

```
set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_dst]
-setup -end 3
```

Figure 7-50 shows the timing diagram for the preferred setup relationship for this example.

Figure 7-50. Preferred Setup Analysis

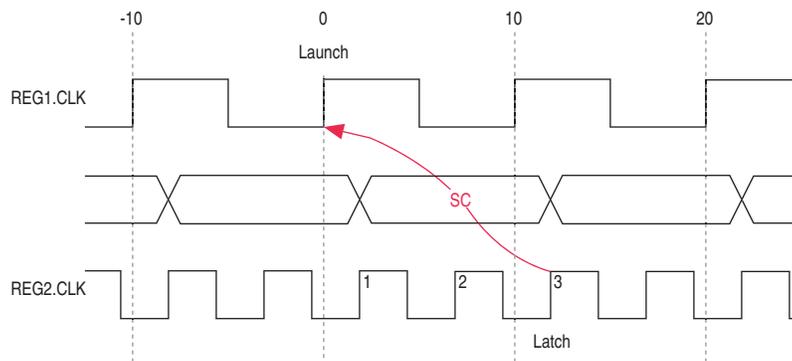
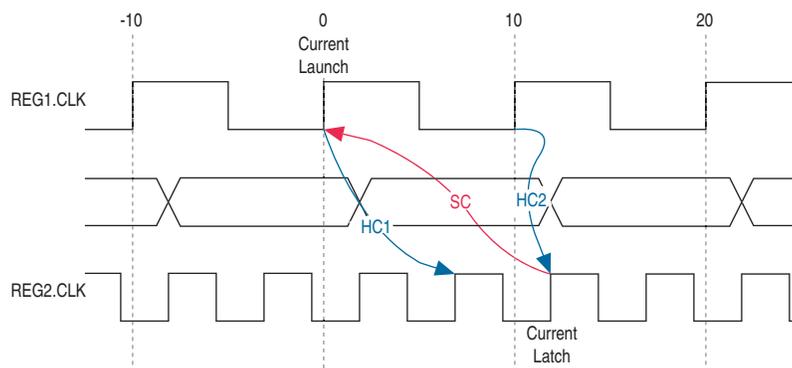


Figure 7-51 shows the timing diagram for default hold check analysis performed by the TimeQuest analyzer with an end multicycle setup value of three.

Figure 7-51. Default Hold Check



Equation 7-25 shows the calculation that the TimeQuest analyzer performs to determine the hold check.

Equation 7-25. Hold Check

$$\begin{aligned} \text{hold check 1} &= \text{current launch edge} - \text{previous latch edge} \\ &= 0 \text{ ns} - 5 \text{ ns} \\ &= -5 \text{ ns} \end{aligned}$$

$$\begin{aligned} \text{hold check 2} &= \text{next launch edge} - \text{current latch edge} \\ &= 10 \text{ ns} - 10 \text{ ns} \\ &= 0 \text{ ns} \end{aligned}$$

In this example, hold check one is too restrictive. The data is launched by the edge at 0 ns and should check against the data captured by the previous latch edge at 2 ns, which does not occur in hold check one. To correct the default analysis, you must use an end multicycle hold exception of one.

The Source Clock Frequency is a Multiple of the Destination Clock Frequency

In this example, the source clock frequency value of 5 ns is an integer multiple of the destination clock frequency of 10 ns. The source clock frequency can be an integer multiple of the destination clock frequency when a PLL is used to generate both clocks and different multiplication and division factors are used. Figure 7-52 shows an example of a design where the source clock frequency is a multiple of the destination clock frequency.

Figure 7-52. Source Clock Frequency is Multiple of Destination Clock Frequency

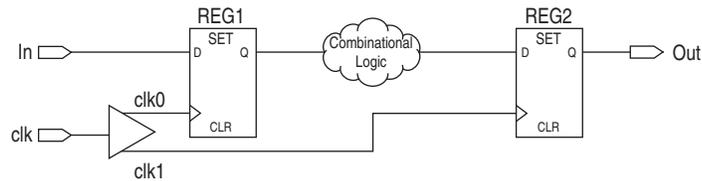
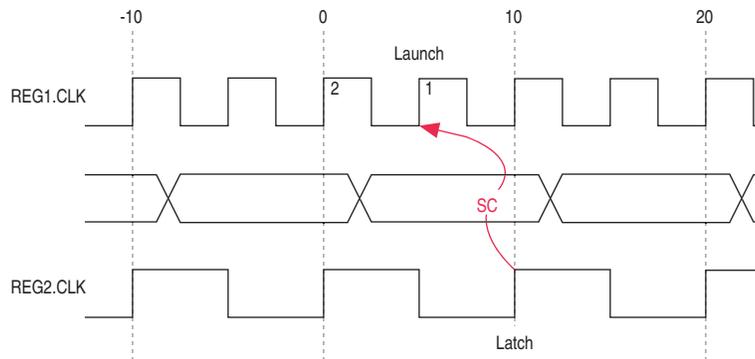


Figure 7-53 shows the timing diagram for default setup check analysis performed by the TimeQuest analyzer.

Figure 7-53. Default Setup Check Analysis



Equation 7-26 shows the calculation that the TimeQuest analyzer performs to determine the setup check.

Equation 7-26. Setup Check

$$\begin{aligned}
 \text{setup check} &= \text{current latch edge} - \text{closest previous launch edge} \\
 &= 10 \text{ ns} - 5 \text{ ns} \\
 &= 5 \text{ ns}
 \end{aligned}$$

The setup relationship shown in Figure 7-53 demonstrates that the data launched at edge one does not need to be captured, and the data launched at edge two must be captured; therefore, you can relax the setup requirement. To correct the default analysis, you must shift the launch edge by one clock period with a start multicycle setup exception of two.

Example 7-18 shows the multicycle exception used to correct the default analysis in this example.

Example 7-18. Multicycle Exceptions

```

set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_dst]
-setup -start 2
    
```

Figure 7-54 shows the timing diagram for the preferred setup relationship for this example.

Figure 7-54. Preferred Setup Check Analysis

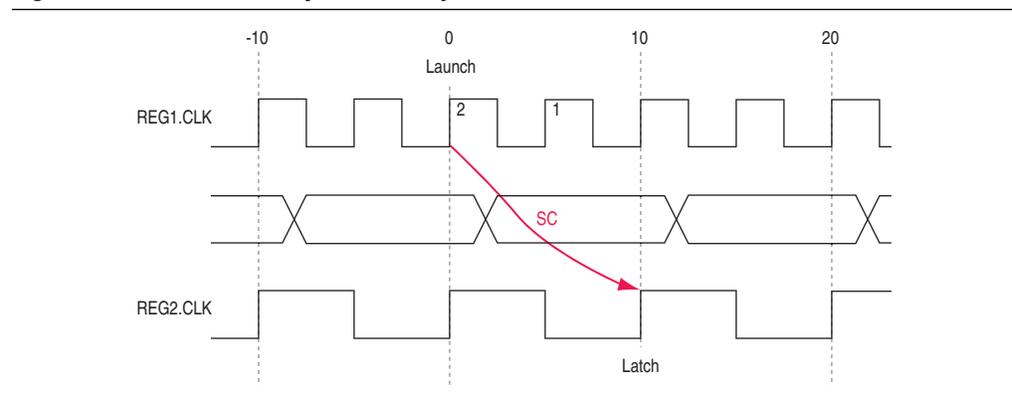
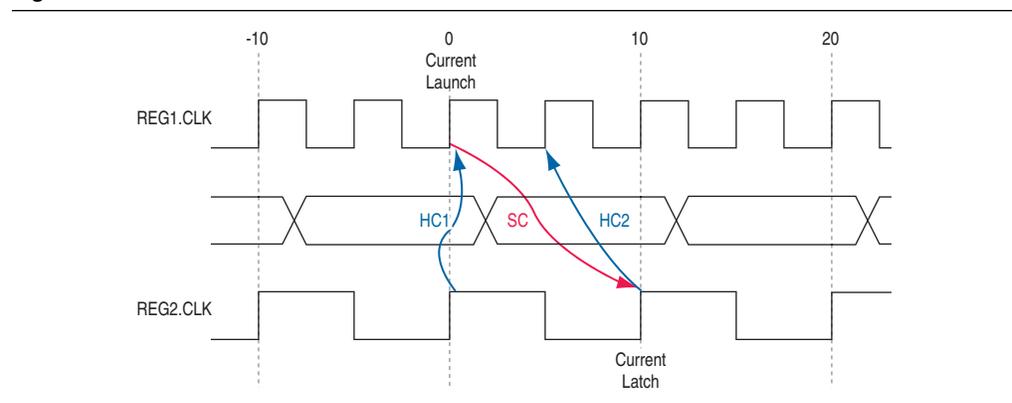


Figure 7-55 shows the timing diagram for default hold check analysis performed by the TimeQuest analyzer with a start multicycle setup value of two.

Figure 7-55. Default Hold Check



Equation 7-27 shows the calculation that the TimeQuest analyzer performs to determine the hold check.

Equation 7-27. Hold Check

$$\begin{aligned} \text{hold check 1} &= \text{current launch edge} - \text{previous latch edge} \\ &= 0 \text{ ns} - 0 \text{ ns} \\ &= 0 \text{ ns} \end{aligned}$$

$$\begin{aligned} \text{hold check 2} &= \text{next launch edge} - \text{current latch edge} \\ &= 5 \text{ ns} - 10 \text{ ns} \\ &= -5 \text{ ns} \end{aligned}$$

In this example, hold check two is too restrictive. The data is launched next by the edge at 10 ns and should check against the data captured by the current latch edge at 10 ns, which does not occur in hold check two. To correct the default analysis, you must use a start multicycle hold exception of one.

The Source Clock Frequency is a Multiple of the Destination Clock Frequency with an Offset

In this example, the source clock frequency is an integer multiple of the destination clock frequency and the destination clock has a positive phase offset. The source clock frequency is 5 ns and destination clock frequency is 10 ns. The destination clock also has a positive offset of 2 ns with respect to the source clock. The source clock frequency can be an integer multiple of the destination clock frequency with an offset when a PLL is used to generate both clocks, different multiplication and division factors are used, and a phase shift applied to the destination clock. Figure 7-56 shows an example of a design where the source clock frequency is a multiple of the destination clock frequency with an offset.

Figure 7-56. Source Clock Frequency is Multiple of Destination Clock Frequency with Offset

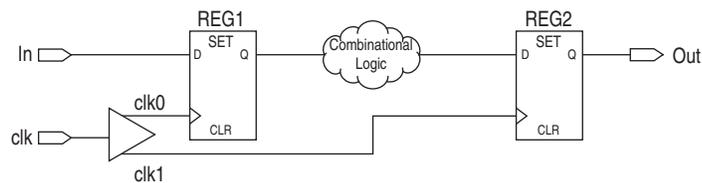
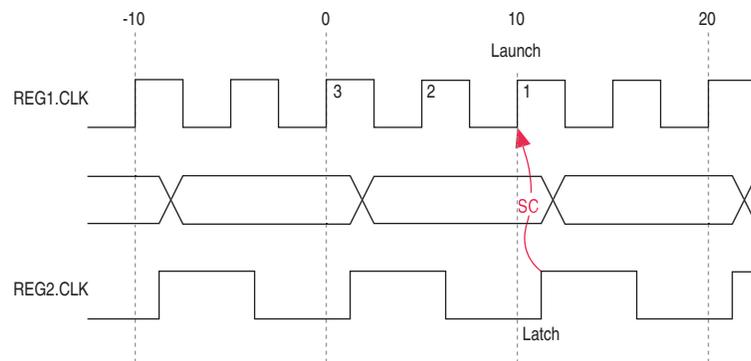


Figure 7-57 shows the timing diagram for default setup check analysis performed by the TimeQuest analyzer.

Figure 7-57. Setup Timing Diagram



Equation 7-28 shows the calculation that the TimeQuest analyzer performs to determine the setup check.

Equation 7-28. setup Check

$$\begin{aligned}
 \text{setup check} &= \text{current latch edge} - \text{closest previous launch edge} \\
 &= 12 \text{ ns} - 10 \text{ ns} \\
 &= 2 \text{ ns}
 \end{aligned}$$

The setup relationship shown in Figure 7-57 demonstrates that the data is not launched at edge one, and the data that is launched at edge three must be captured; therefore, you can relax the setup requirement. To correct the default analysis, you must shift the launch edge by two clock periods with a start multicycle setup exception of three.

Example 7-19 shows the multicycle exception used to correct the default analysis in this example.

Example 7-19. Multicycle Exceptions

```
set_multicycle_path -from [get_clocks clk_src] -to [get_clocks clk_dst]
-setup -start 3
```

Figure 7-58 shows the timing diagram for the preferred setup relationship for this example.

Figure 7-58. Preferred Setup Check Analysis

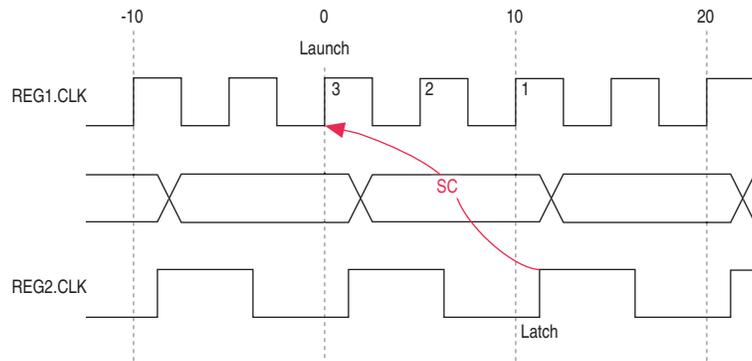
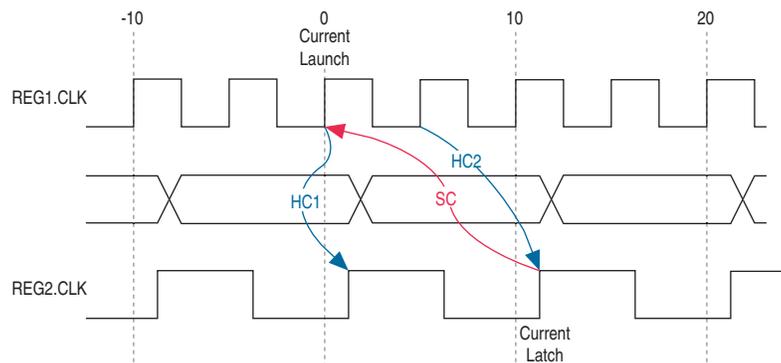


Figure 7-59 shows the timing diagram for default hold check analysis performed by the TimeQuest analyzer with a start multicycle setup value of three.

Figure 7-59. Default Hold Check Analysis



Equation 7-29 shows the calculation that the TimeQuest analyzer performs to determine the hold check.

Equation 7-29. Hold Check

$$\begin{aligned} \text{hold check 1} &= \text{current launch edge} - \text{previous latch edge} \\ &= 0 \text{ ns} - 2 \text{ ns} \\ &= -2 \text{ ns} \end{aligned}$$

$$\begin{aligned} \text{hold check 2} &= \text{next launch edge} - \text{current latch edge} \\ &= 5 \text{ ns} - 12 \text{ ns} \\ &= -7 \text{ ns} \end{aligned}$$

In this example, hold check two is too restrictive. The data is launched next by the edge at 10 ns and should check against the data captured by the current latch edge at 12 ns, which does not occur in hold check two. To correct the default analysis, you must use a start multicycle hold exception of one.

Document Revision History

Table 7-2 shows the revision history for this chapter.

Table 7-2. Document Revision History

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Changed to new document template. ■ Edited document and linked to other resources. ■ Added information that previously appeared in <i>Application Note 481: Applying Multicycle Exceptions in the TimeQuest Timing Analyzer</i>.
July 2010	10.0.0	Minor editorial changes to document. No updates to technical content.
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Added example to “Derived Clocks” on page 7-3. ■ Updated “Base Clocks” on page 7-2.
March 2009	9.0.0	Initial release.

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter describes the benefits of switching to the Quartus® II TimeQuest Timing Analyzer, the differences between the TimeQuest analyzer and the Classic Timing Analyzer, and the process you should follow to switch a design from using the Classic Timing Analyzer to the TimeQuest analyzer.

Benefits of Switching to the TimeQuest Analyzer

Increasing design complexity requires a timing analysis tool with greater capabilities and flexibility. The TimeQuest analyzer offers the following benefits:

- Industry-standard Synopsys Design Constraint (SDC) support increases productivity.
- Simple, flexible reporting uses industry-standard terminology and makes timing sign-off faster.

These features ease constraint and analysis of modern, complex designs. SDC constraints support complex clocking schemes and high-speed interfaces. An example includes designs that have multiplexed clocks, regardless of whether they are switched on or off chip. Designs with source-synchronous interfaces, such as double data rate (DDR) memory interfaces, are much simpler to constrain and analyze with the TimeQuest analyzer.

There are three main differences between the Classic Timing Analyzer and the TimeQuest analyzer. Unlike the Classic Timing Analyzer, the TimeQuest analyzer has the following benefits:

- All clocks are related by default.
- The default hold multicycle value is zero.
- You must constrain all ports and ripple clocks.

This chapter contains the following sections:

- [“Switching Your Design” on page 8-2](#)
- [“Differences Between the Quartus II Timing Analyzers” on page 8-3](#)
- [“Timing Assignment Conversion” on page 8-24](#)
- [“Conversion Utility” on page 8-43](#)
- [“Notes” on page 8-53](#)

Switching Your Design

To switch a design from the Classic Timing Analyzer to the TimeQuest analyzer, perform the steps:

1. Open your compiled design in the Quartus II software.
2. Create a Synopsys Design Constraints File (.sdc) that contains timing constraints.
3. Perform timing analysis with the TimeQuest analyzer and examine the reports.

Open Your Compiled Design

To begin, open a design you previously compiled with the Quartus II software.

Create an SDC Constraints

Create SDC Constraints Manually

If you are familiar with SDC commands and syntax, you can create an .sdc with any text editor and skip to [“Start the TimeQuest Analyzer” on page 8-3](#). Name the .sdc `<revision>.sdc`, where `<revision>` is the current revision of your project, and save it in your project directory.



For more information about TimeQuest analyzer SDC constraints, refer to the [Quartus II TimeQuest Timing Analyzer](#) chapter in volume 3 of the *Quartus II Handbook*.

Create SDC Constraints from Existing Timing Assignments

You can use the TimeQuest analyzer conversion utility to help you convert the timing assignments in an existing Quartus II Settings File (.qsf) to corresponding SDC constraints. To run the TimeQuest analyzer conversion utility, on the Constraints menu, click **Generate SDC File from QSF**. You can also run the conversion utility by typing the following command at a system command prompt:

```
quartus_sta --qsf2sdc <project name> ↵
```

The .sdc created by the conversion utility is named `<revision>.sdc`.

After conversion, review the .sdc to ensure it is correct and complete, and make changes if necessary. Refer to [“Constraint File Priority” on page 8-7](#) for recommendations.

The conversion utility cannot convert some types of Classic Timing Analyzer assignments if no corresponding SDC constraint exist, or if there is more than one potentially equivalent SDC constraint. If the conversion utility cannot convert your assignment, manually convert any ambiguous assignments. Correct conversion requires knowledge of the intended function of your design. To convert your assignments, use the guidelines in [“Timing Assignment Conversion” on page 8-24](#).

Start the TimeQuest Analyzer

To open the TimeQuest analyzer GUI, on the Tools menu, click **TimeQuest Timing Analyzer**. The TimeQuest GUI automatically opens the project open in the Quartus II software GUI.

To open the TimeQuest analyzer GUI from a system command prompt, type the following command:

```
quartus_staw ↵
```

To start the TimeQuest analyzer Tcl shell type the following command:

```
quartus_sta -s ↵
```

Differences Between the Quartus II Timing Analyzers

The differences between the TimeQuest analyzer and the Classic Timing Analyzer are described in the following sections:

- “Terminology” on page 8-3
- “Constraints” on page 8-5
- “Clocks” on page 8-10
- “Hold Multicycle” on page 8-18
- “Fitter Performance and Behavior” on page 8-19
- “Reporting” on page 8-20
- “Timing Assignment Conversion” on page 8-24

Terminology

This section introduces the industry-standard SDC terminology used by the TimeQuest analyzer.

Netlists

The TimeQuest analyzer uses SDC naming conventions for netlists. Netlists consist of cells, pins, nets, ports, and clocks.

- Cells are instances of fundamental hardware elements in Altera® FPGAs, such as logic elements, look-up tables (LUTs), and registers.
- Pins are inputs and outputs of cells.
- Nets are connections between output pins and input pins.
- Ports are top-level module inputs and outputs (device inputs and outputs).
- Clocks are abstract objects outside the netlist.



The terminology of pins and ports is opposite to that of the Classic Timing Analyzer. In the Classic Timing Analyzer, ports are inputs and outputs of cells, and pins are top-level module inputs and outputs (device inputs and outputs).

Figure 8-1 shows a sample design, and Figure 8-2 shows the TimeQuest analyzer netlist representation of the design. Netlist elements in Figure 8-2 are labeled to illustrate the SDC terminology.

Figure 8-1. Sample Design

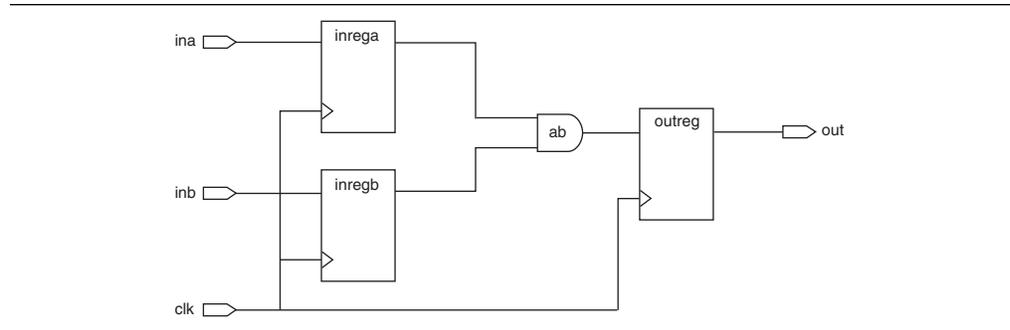
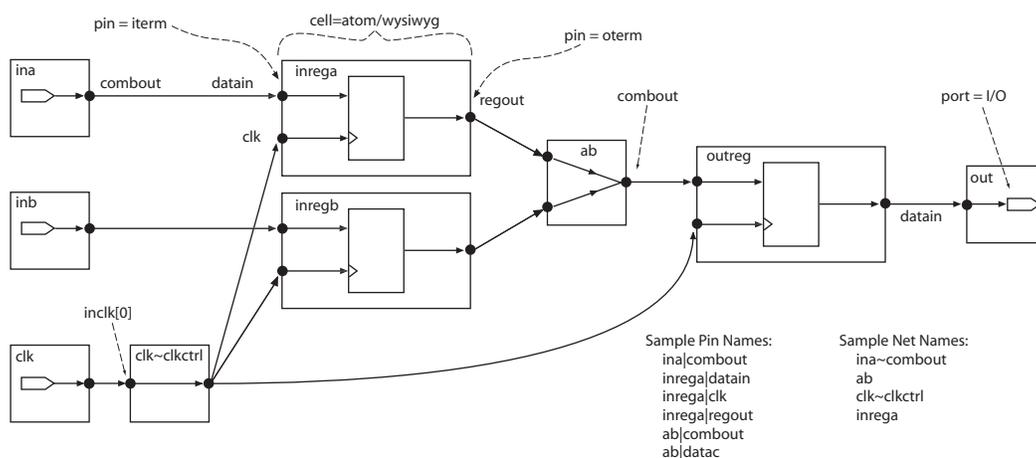


Figure 8-2. TimeQuest Analyzer Netlist



Collections

In addition to standard SDC collections, the TimeQuest analyzer supports the following Altera-specific collection types:

- **Keepers**—Noncombinational nodes in a netlist
- **Nodes**—Nodes can be combinational, registers, latches, or ports (device inputs and outputs)
- **Registers**—Registers or latches in the netlist

You can use the `get_keepers`, `get_nodes`, or `get_registers` commands to access these collections.

 For more information about TimeQuest analyzer terminology, refer to the *Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

Constraints

The Classic Timing Analyzer and TimeQuest analyzer store constraints in different files, support different methods for constraint entry, and prioritize constraints differently. The following sections explain these differences.

Constraint Files

The TimeQuest analyzer stores all SDC timing constraints in **.sdc** files. The Classic Timing Analyzer stores all timing assignments in the **.qsf** for your project.

When you use the TimeQuest analyzer, your **.qsf** contains all assignments and settings except for SDC constraints. The TimeQuest analyzer ignores the timing assignments in your **.qsf** except when the conversion utility converts Classic Timing Analyzer timing assignments to TimeQuest analyzer SDC constraints. There is no automatic process that keeps timing constraints synchronized between your **.qsf** and **.sdc** files. You must manually keep the constraints synchronized, if so desired.

Constraint Entry

With the Classic Timing Analyzer, you enter timing assignments with the **Settings** dialog box, the Assignment Editor, or with commands in Tcl scripts. With the TimeQuest analyzer, you cannot use the Assignment Editor to enter SDC constraints. You must use one of the following methods to enter TimeQuest analyzer constraints:

- Enter constraints at the Tcl prompt in the TimeQuest analyzer
- Enter constraints in an **.sdc** with a text editor or SDC editor
- Use the constraint entry commands on the Constraints menu in the TimeQuest analyzer GUI

You can enter timing assignments for the Classic Timing Analyzer even if no timing netlist exists for your design. The TimeQuest analyzer requires that a netlist exist for interactive constraint entry. Each TimeQuest analyzer constraint is a Tcl command evaluated in real-time, if entered directly in the Tcl console. As part of this evaluation, the TimeQuest analyzer validates all names. To do this, SDC commands can only be evaluated after a netlist is created. An **.sdc** can be created at any time using the TimeQuest analyzer or any other text editor, but a netlist is required before an **.sdc** can be sourced. You must create a timing netlist in the TimeQuest analyzer before you can enter constraints with either of the following interactive methods:

- At the Tcl console of the TimeQuest analyzer
- With commands on the Constraints menu in the TimeQuest analyzer GUI

If you enter constraints with a text editor separate from the TimeQuest analyzer, no timing netlist is required.

To create a timing netlist in the TimeQuest analyzer, use the `create_timing_netlist` command, or double-click **Create Timing Netlist** in the Tasks pane of the TimeQuest analyzer GUI.

Time Units

Enter time values in default time units of nanoseconds (ns) with up to three decimal places. Note that the TimeQuest analyzer does not display the default time unit when it displays time values.

You can specify a different default time unit with the `set_time_format -unit <default time unit>` command, or specify another unit when you enter a time value, for example, `300ps`.



Specifying time units after a value is not part of the standard SDC format. Unit specification is a TimeQuest analyzer extension.

You can specify clock constraints with period or frequency in the TimeQuest analyzer. For example, you can use any of the following constraints:

- `create_clock -period 10.000`
(assuming default units and decimal places)
- `create_clock -period "100 MHz"`
- `create_clock -period "10 ns"`

MegaCore Functions

If you change any MegaCore function settings and regenerate the core after you convert your timing assignments to SDC constraints, you must manually update the SDC constraints or reconvert your assignments. You must update or reconvert, because changes to MegaCore function settings can affect timing assignments embedded in the hardware description language files of the function. The timing assignments are not converted automatically when the core settings change.



You should make a backup copy of your `.sdc` before reconverting assignments. If you make changes to the `.sdc`, you can manually copy the updated MegaCore timing constraints to your `.sdc`.

Bus Name Format

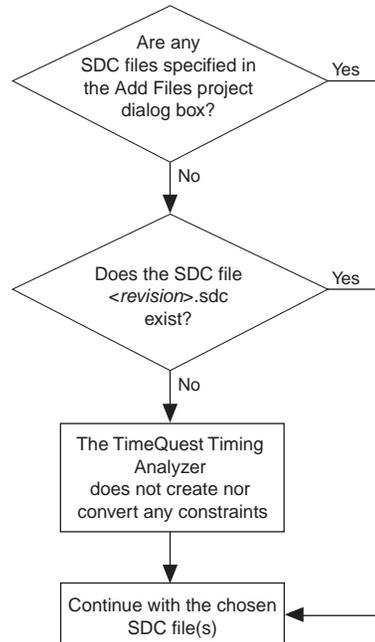
In the Classic Timing Analyzer, you can make a timing assignment to all bits in a bus with the bus name (or the bus name followed by an asterisk enclosed in square brackets) as the target. For example, to make an assignment to all bits of a bus called `address`, use `address` or `address[*]` as the target of the assignment.

In the TimeQuest analyzer, you must use the bus name followed by square brackets enclosing an asterisk, as follows: `address[*]` to make all timing assignments to all bits within a bus.

Constraint File Priority

Figure 8-3 shows the priority order in which the TimeQuest analyzer searches for .sdc files.

Figure 8-3. .sdc Search Order



If you specify constraints in multiple .sdc files, or if you use a single .sdc with a name other than <revision>.sdc, you must add the files to your project so the TimeQuest analyzer can find them.

- For more information, refer to *Managing Files in a Project* in Quartus II Help.

You can also add .sdc files to your project with the following Tcl command in your .qsf, repeated once for each .sdc:

```
set_global_assignment -name SDC_FILE <.sdc file name>
```

The TimeQuest analyzer reads constraint files from the files list in the order they are listed.

If you use an .sdc created by the conversion utility, you should place it first in the list of files. When conflicting constraints apply to the same node, the last constraint has the highest priority. Therefore, .sdc files with your additions or changes should be listed after the .sdc created by the conversion utility, so your constraints have higher priority.

When you process your design, and the TimeQuest analyzer cannot find an `.sdc` it attempts to meet a default 1 GHz constraint on all clocks in your design. The Quartus II software may prompt you to run the conversion utility if your design does not reference any `.sdc` files.

You must review the `.sdc` as you would when manually running the conversion utility. Refer to “[Reviewing Conversion Results](#)” on page 8–50 for information about reviewing the converted constraints.

Constraint Priority

The Classic Timing Analyzer prioritizes assignments based on the specificity of the nodes to which they are assigned. The more specific an assignment is, the higher its priority. The TimeQuest analyzer simplifies these precedence rules. When overlaps occur in the nodes to which the constraints apply, constraints at the bottom of the file take priority over constraints at the top of the file.

As an example, in the Classic Timing Analyzer, point-to-point multicycle assignments have higher priority than single-point multicycle assignments. The two assignments in [Example 8–1](#) result in a multicycle assignment of two between `A_reg` and all nodes beginning with `B`, including `B_reg`. The single-point assignment does not apply to paths from `A_reg` to `B_reg`, because the specific point-to-point assignment takes priority over the general single-point assignment.

Example 8–1. Classic Timing Analyzer Multicycle Assignments

```
set_instance_assignment -name MULTICYCLE -from A_reg -to B* 2
set_instance_assignment -name MULTICYCLE -to B_reg 3
```

[Example 8–2](#) shows SDC versions of the preceding Classic Timing Analyzer timing assignments. However, the TimeQuest analyzer evaluates the constraints from top to bottom, regardless of whether the assignments are point-to-point assignments or single-point assignments, so the path from `A_reg` to `B_reg` receives a multicycle exception of three because it is ordered second.

Example 8–2. TimeQuest Analyzer Multicycle Exceptions

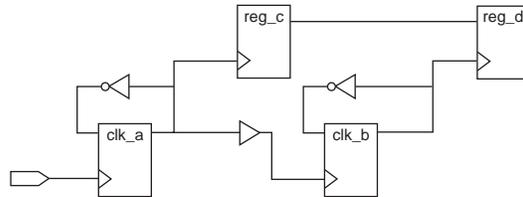
```
set_multicycle_path -from [get_keepers A_reg] -to [get_keepers B*] 2
set_multicycle_path -to [get_keepers B_reg] 3
```

Ambiguous Constraints

Because of capabilities in the TimeQuest analyzer, some Classic Timing Analyzer assignments can be ambiguous after conversion by the conversion utility. These assignments require manual updating based on your knowledge of your design.

Figure 8-4 shows a ripple clock circuit. The explanation that follows shows an ambiguous constraint for that circuit, and how to edit the constraint to remove the ambiguity in the TimeQuest analyzer.

Figure 8-4. Ripple Clock Circuit



In the Classic Timing Analyzer, the following QSF multicycle assignment from `clk_a` to `clk_b` with a value of two applies to paths transferring data from the `clk_a` domain to the `clk_b` domain:

```
set_instance_assignment -name MULTICYCLE -from clk_a -to clk_b 2
```

In Figure 8-4, this assignment applies to the path from `reg_c` to `reg_d`. In the TimeQuest analyzer, the use of the clock node names in the following equivalent multicycle exception is ambiguous:

```
set_multicycle_path -setup -from clk_a -to clk_b 2
```

The exception could apply to the path between `clk_a` and `clk_b`, or it could apply to paths from one ripple clock domain to the other ripple clock domain (`reg_c` to `reg_d`).

The TimeQuest analyzer exceptions shown in Example 8-3 are not ambiguous because they use collections to explicitly specify the targets of the exception.

Example 8-3. Unambiguous TimeQuest Analyzer Exceptions

```
# Applies to path from reg_c to reg_d
set_multicycle_path -setup -from [get_clocks clk_a] \
  -to [get_clocks clk_b] 2
# Applies to path from clk_a to clk_b
set_multicycle_path -setup -from [get_registers clk_a] \
  -to [get_registers clk_b] 2
```

Clocks

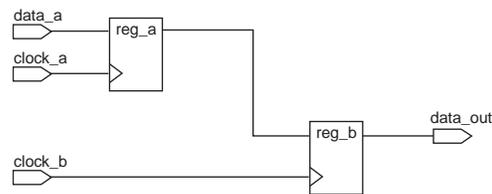
The Classic Timing Analyzer and TimeQuest analyzer detect, analyze, and report clocks differently. The following sections describe these differences.

Related and Unrelated Clocks

In the TimeQuest analyzer, all clocks are related by default, and you must add assignments to indicate unrelated clocks. However, in the Classic Timing Analyzer, all base clocks are unrelated by default. All derived clocks of a base clock are related to each other, but are unrelated to other base clocks and their derived clocks.

Figure 8-5 shows a circuit with a path between two clock domains. The TimeQuest analyzer analyzes the path from `reg_a` to `reg_b` because all clocks are related by default. The Classic Timing Analyzer does not analyze the path from `reg_a` to `reg_b` by default.

Figure 8-5. Cross Clock Domain Path



To make clocks unrelated in the TimeQuest analyzer, use the `set_clock_groups` command with the `-exclusive` option. The TimeQuest analyzer does not analyze paths between the two clock domains. For example, the following command renders `clock_a` and `clock_b` unrelated:

```
set_clock_groups -exclusive -group {clock_a} -group {clock_b}
```

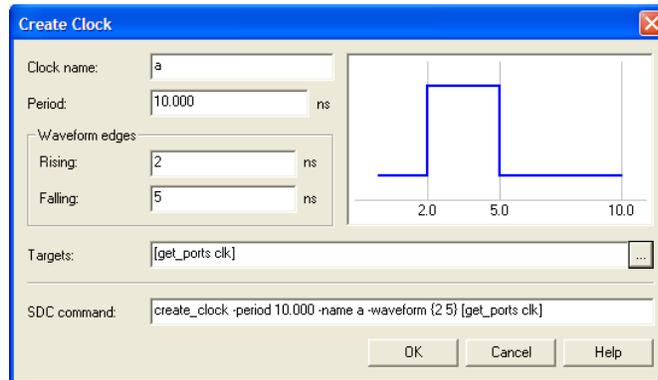
Clock Offset

In the TimeQuest analyzer, clocks can have nonzero values for the rising edge of the waveform, a feature that the Classic Timing Analyzer does not support. To specify an offset for your clock, use the `create_clock` command with the `-waveform` option to specify the rising and falling edge times, as shown in this example:

```
-waveform {<rising edge time> <falling edge time>}
```

Figure 8-6 shows a clock constraint with an offset in the TimeQuest analyzer GUI.

Figure 8-6. Create Clock Dialog Box



Clock offset affects setup and hold relationships. Launch and latch edges are evaluated after offsets are applied. Depending on the offset, the setup relationship can be the offset value, or the difference between the period and offset. You should use the clock latency constraint, instead of clock offset to emulate latency. Refer to “[Offset and Latency Example](#)” on page 8-11 for an example that illustrates the different effects of offset and latency.

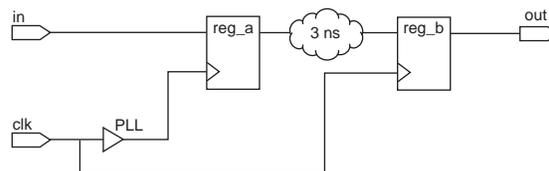
Clock Latency

The TimeQuest analyzer does not ignore early clock latency and late clock latency differences when the clock source is the same, as the Classic Timing Analyzer does. When you specify latencies, you should take common clock path pessimism into account and use uncertainty to control pessimism differences for clock-to-clock data transfers. Unlike clock offset, clock latency affects skew, and launch and latch edges are evaluated before latencies are applied, so the setup relationship is always equal to the period.

Offset and Latency Example

Figure 8-7 shows a simple register-to-register circuit used to illustrate the different effects of offset and latency. The examples show why you should not use clock offset to emulate clock latency.

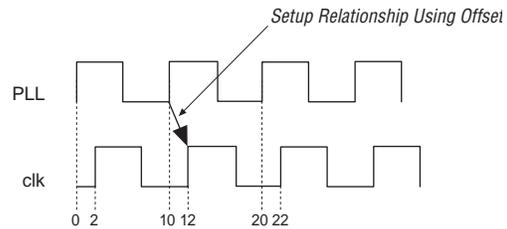
Figure 8-7. Offset and Latency Example Circuit



The period for `clk` is 10 ns, and the period for the phase-locked loop (PLL) output is 10 ns. The PLL compensation value is -2 ns. The network delay from the PLL to `reg_a` equals the network delay from `clk` to `reg_b`. Finally, the delay from `reg_a` to `reg_b` is 3 ns.

Clock Offset Scenario

Treat the PLL compensation value of -2 ns as a clock offset of -2 ns with a clock skew of 0 ns. Launch and latch edges are evaluated after offsets are applied, so the setup relationship is 2 ns (Figure 8-8).

Figure 8-8. Setup Relationship Using Offset

Equation 8-1 shows how to calculate the slack value for the path from reg_a to reg_b.

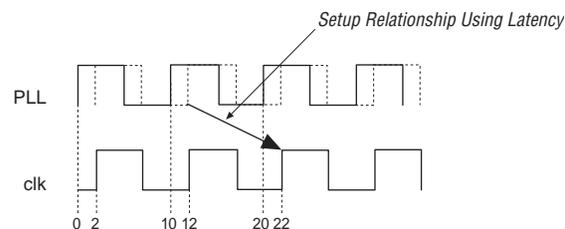
Equation 8-1.

$$\begin{aligned} \text{slack} &= \text{clock arrival} - \text{data arrival} \\ \text{slack} &= \text{setup relationship} + \text{clock skew} - \text{reg_to_reg delay} \\ \text{slack} &= 2\text{ns} + 0\text{ns} - 3\text{ns} \\ \text{slack} &= -1\text{ns} \end{aligned}$$

The negative slack requires a multicycle assignment with a value of two and a hold multicycle assignment with a value of one to correct. With these assignments from reg_a to reg_b, the setup relationship is then 12 ns, resulting in a slack of 9 ns.

Clock Latency Scenario

Treat the PLL compensation value of -2 ns as latency with a clock skew of 2 ns. Because launch and latch edges are evaluated before latencies are applied, the setup relationship is 10 ns (the period of clk and the PLL) (Figure 8-9).

Figure 8-9. Setup Relationship Using Latency

Equation 8-2 shows how to calculate the slack value for the path from reg_a to reg_b.

Equation 8-2.

$$\begin{aligned} \text{slack} &= \text{clock arrival} - \text{data arrival} \\ \text{slack} &= \text{setup relationship} + \text{clock skew} - \text{reg_to_reg delay} \\ \text{slack} &= 10\text{ns} + 2\text{ns} - 3\text{ns} \\ \text{slack} &= 9\text{ns} \end{aligned}$$

The slack of 9 ns is identical to the slack computed in the previous example, but because this example uses latency instead of offset, no multicycle assignment is required.

Clock Uncertainty

The Classic Timing Analyzer ignores **Clock Setup Uncertainty** and **Clock Hold Uncertainty** assignments when you specify a setup or hold relationship between two clocks. However, the TimeQuest analyzer does not ignore clock uncertainty when you specify a setup or hold relationship between two clocks. Figure 8-10 and Figure 8-11 illustrate the different behavior between the TimeQuest analyzer and the Classic Timing Analyzer.

In both figures, the constraints are identical. There is a 20-ns period for clk_a and clk_b. There is a setup relationship (a set_max_delay exception in the TimeQuest analyzer) of 7 ns from clk_a to clk_b, and a clock setup uncertainty constraint of 1 ns from clk_a to clk_b. The actual setup relationship in the TimeQuest analyzer is 1 ns less than in the Classic Timing Analyzer because of the way they analyze clock uncertainty.

Figure 8-10. Classic Timing Analyzer Behavior

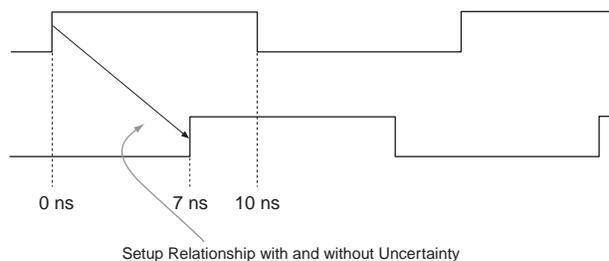
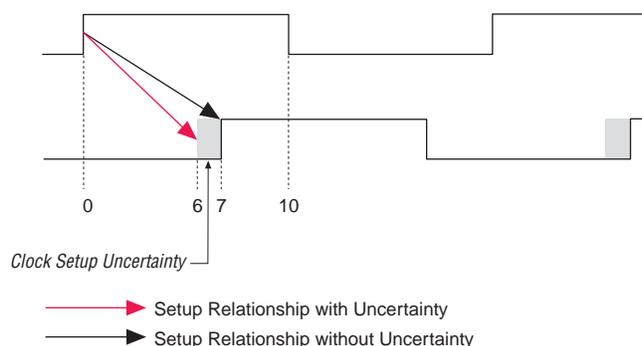


Figure 8-11. TimeQuest Analyzer Behavior



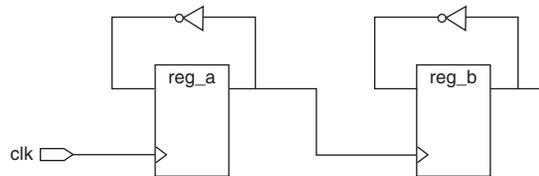
Derived and Generated Clocks

Generated clocks in the TimeQuest analyzer are different than derived clocks in the Classic Timing Analyzer. In the Classic Timing Analyzer, the source of a derived clock must be a base clock. However, in the TimeQuest analyzer, the source of a generated clock can be any other clock in the design (including virtual clocks), or any node to which a clock propagates through the clock network. Because generated clocks are related through the clock network, you can specify generated clocks for isolated modules, such as IP, without knowing the details of the clocks outside of the module.

In the TimeQuest analyzer, you can specify generated clocks relative to specific edges and edge shifts of a master clock, a feature that the Classic Timing Analyzer does not support.

Figure 8-12 shows a simple ripple clock that you should constrain with generated clocks in the TimeQuest analyzer.

Figure 8-12. Generated Clocks Circuit



The TimeQuest analyzer constraints shown in Example 8-4 constrain the clocks in the circuit above. Note that the source of each generated clock can be the input pin of the register itself, not the name of another clock.

Example 8-4. Generated Clock Constraints

```
create_clock -period 10 -name clk clk
create_generated_clock -divide_by 2 -source reg_a|CLK -name reg_a reg_a
create_generated_clock -divide_by 2 -source reg_b|CLK -name reg_b reg_b
```

Automatic Clock Detection

The Classic Timing Analyzer and TimeQuest analyzer identify clock sources of registers that do not have a defined clock source. The Classic Timing Analyzer traces back along the clock network, through registers and logic, until it reaches a top-level port in your design. The TimeQuest analyzer also traces back along the clock network, but it stops at registers.

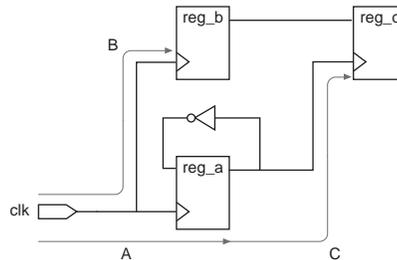
You can use two SDC commands in the TimeQuest analyzer to automatically detect and create clocks for unconstrained clock sources:

- **derive_clocks**—creates clocks on sources of clock pins that do not already have at least one clock sourcing the clock pin
- **derive_pll_clocks**—identifies PLLs and creates generated clocks on the clock output pins

derive_clocks Command

Figure 8-13 shows a circuit with a divide-by-two register and indicates the clock network delays as A, B, and C. The following explanation describes how the Classic Timing Analyzer and TimeQuest analyzer detect the clocks in Figure 8-13.

Figure 8-13. Example Circuit for derive_clocks



The Classic Timing Analyzer detects that `clk` is the clock source for registers `reg_a`, `reg_b`, and `reg_c`. It detects that `clk` is the clock source for `reg_c` because it traces back along the clock network for `reg_c` through `reg_a`, until it finds the `clk` port. The Classic Timing Analyzer computes the clock arrival time for `reg_c` as $A + C$.

The `derive_clocks` command in the TimeQuest analyzer creates two base clocks, one on the `clk` port and one on `reg_a`, because the command does not trace through registers on the clock network. The clock arrival time for `reg_c` is C because the clock starts at `reg_a`.

To make the TimeQuest analyzer compute the same clock arrival time ($A + C$) as the Classic Timing Analyzer for `reg_c`, make the following modifications to the clock constraints created by the `derive_clocks` command:

1. Change the base clock named `reg_a` to a generated clock
2. Make the source the clock pin of `reg_a` (`reg_a | clk`) or the port `clk`
3. Add a `-divide_by 2` option

These modifications cause the clock arrival times to `reg_c` to match between the Classic Timing Analyzer and the TimeQuest analyzer. However, the clock for `reg_c` is shown as `reg_a` instead of `clk`, and the launch and latch edges may change for some paths due to the `-divide_by` option.

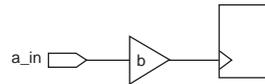
You can use the `derive_clocks` command at the beginning of your design cycle when you do not know all of the clock constraints for your design, but you should not use it during timing sign-off. Instead, you should constrain each clock in your design with the `create_clock` or `create_generated_clocks` commands.

The `derive_clocks` command detects clocks in your design using the following rules:

1. An input clock port is detected as a clock only if there are no other clocks feeding the destination registers.
 - a. Input clock ports are not detected automatically if they feed only other base clocks.
 - b. If other clocks feed the port's register destinations, the port is assumed to be an enable or data port for a gated clock.
 - c. When no clocks are defined, and multiple clocks feed a destination register, the auto-detected clock is selected arbitrarily.
2. All ripple clocks (registers in a clock path) are detected as clocks automatically using the same rules for input clock ports. If both an input port and a register feed register clock pins, the input port is selected as the clock.

The following examples show how the `derive_clocks` command detects clocks in the circuit shown in [Figure 8-14](#).

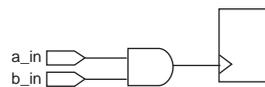
Figure 8-14. Detectable Clock



- If you do not make any clock settings, and then you use the `derive_clocks` command, it detects `a_in` as a clock according to rule 1, because there are no other clocks feeding the register.
- If you create a clock with `b` as its target, and then you run the `derive_clocks` command, it does not detect `a_in` as a clock according to rule 1a, because `a_in` feeds only another clock.

The following examples show how the `derive_clocks` command detects clocks in the circuit shown in [Figure 8-15](#).

Figure 8-15. Two Detectable Clocks



- If you do not make any clock settings and then you use the `derive_clocks` command, it selects a clock arbitrarily, according to rule 1c.
- If you create a clock with `a_in` as its target and then you use the `derive_clocks` command, it does not detect `b_in` as a clock according to rule 1b, because another clock (`a_in`) feeds the register.

derive_pll_clocks Command

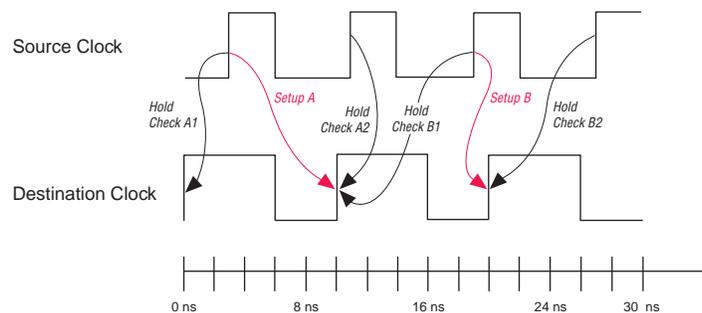
The `derive_pll_clocks` command names the generated clocks according to the names of the PLL output pins by default, and you cannot change these generated clock names. If you want to use your own clock names, you must use the `create_generated_clock` command for each PLL output clock and specify the names with the `-name` option.

If you use the PLL clock-switchover feature, the `derive_pll_clocks` command creates additional generated clocks on each output clock pin of the PLL based on the secondary clock input to the PLL. This may require the `set_clock_groups` or `set_false_path` commands to cut the primary and secondary clock outputs. For information about how to make clocks unrelated, refer to “[Related and Unrelated Clocks](#)” on page 8-10.

Hold Relationship

The TimeQuest analyzer and Classic Timing Analyzer choose the worst-case hold relationship differently. Refer to [Figure 8-16](#) for sample waveforms to illustrate the different effects.

Figure 8-16. Worst-Case Hold



The Classic Timing Analyzer first identifies the worst-case setup relationship. The worst-case setup relationship is Setup B. Then the Classic Timing Analyzer chooses the worst-case hold relationship (Hold Check B1 or Hold Check B2) for that specific setup relationship, Setup B. The Classic Timing Analyzer chooses Hold Check B2 for the worst-case hold relationship.

However, the TimeQuest analyzer calculates worst-case hold relationships for all possible setup relationships and chooses the absolute worst-case hold relationship. The TimeQuest analyzer checks two hold relationships for every setup relationship:

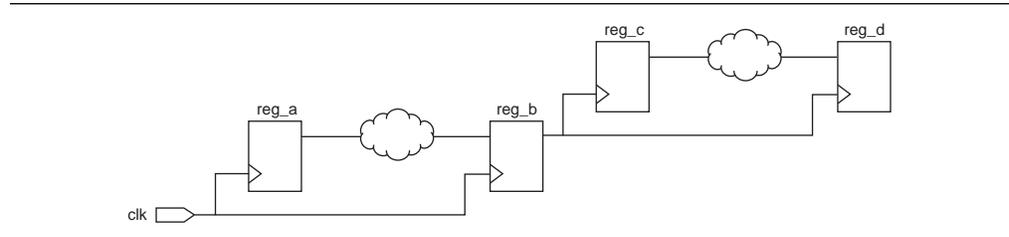
- Data launched by the current launch edge not captured by the previous latch edge (Hold Check A1 and Hold Check B1)
- Data launched by the next launch edge not captured by the current latch edge (Hold Check A2 and Hold Check B2)

The TimeQuest analyzer chooses Hold Check A2 as the absolute worst-case hold relationship.

Clock Objects

The Classic Timing Analyzer treats nodes with clock settings assigned to them as special objects in the timing netlist. Any node in the timing netlist with a clock setting is treated as a clock object, regardless of its actual type, such as a register. When a register has a clock setting assigned to it, the Classic Timing Analyzer does not analyze register-to-register paths beginning or ending at that register. Figure 8-17 shows a circuit that illustrates this situation.

Figure 8-17. Clock Objects



With no clock assignments on any of the registers, the Classic Timing Analyzer analyzes timing on the path from reg_a to reg_b, and from reg_c to reg_d. If you make a clock setting assignment to reg_b, reg_b is no longer considered a register node in the netlist, and the path from reg_a to reg_b is no longer analyzed.

In the TimeQuest analyzer, clocks are abstract objects that are associated with nodes in the timing netlist. The TimeQuest analyzer analyzes the path from reg_a to reg_b even if there is a clock assigned to reg_b.

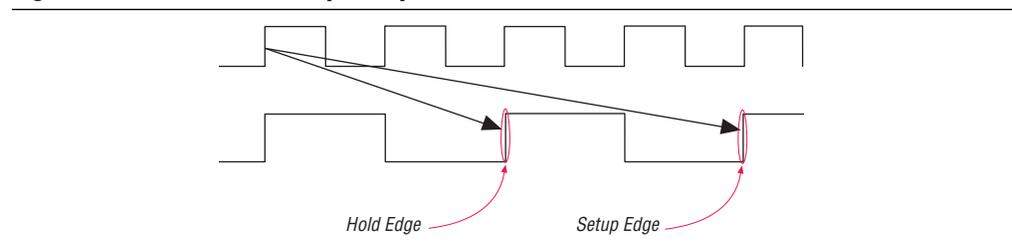
Hold Multicycle

The hold multicycle value numbering scheme is different in the Classic Timing Analyzer and TimeQuest analyzer. Also, you can choose between two values for the default hold multicycle value in the Classic Timing Analyzer but you cannot change the default value in the TimeQuest analyzer. The hold multicycle value specifies which clock edge is used for hold analysis when you change the latch edge with a multicycle assignment.

In the Classic Timing Analyzer, the hold multicycle value is based on one, and is the number of clock cycles away from the setup edge. In the TimeQuest analyzer, the hold multicycle value is based on zero, and is the number of clock cycles away from the default hold edge. In the TimeQuest analyzer, the default hold edge is one edge before or after the setup edges. Subtract one from any hold multicycle value in the Classic Timing Analyzer to compute the equivalent value for the TimeQuest analyzer.

Figure 8-18 shows simple waveforms for a cross-clock domain transfer with the indicated setup and hold edges.

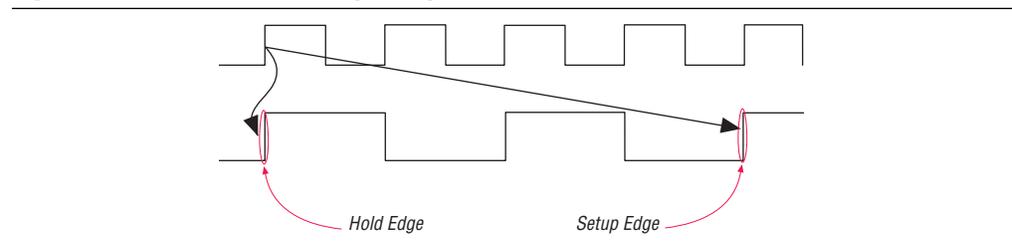
Figure 8-18. First Relationship Example



In the TimeQuest analyzer, only a multicycle exception of two is required to constrain the design for the indicated setup and hold relationships.

Figure 8-19 shows simple waveforms for a different cross-clock domain transfer with indicated setup and hold edges. The following explanation shows what exceptions to apply to achieve the desired setup and hold relationships.

Figure 8-19. Second Relationship Example



In the TimeQuest analyzer, you must use the following two exceptions:

- A multicycle exception of two
- A hold multicycle exception of one, because the hold edge is one edge behind the default hold edge, which is one edge after the setup edge.

You should always add a hold multicycle assignment for every multicycle assignment to ensure the correct exceptions are applied regardless of the timing analyzer you use.

Fitter Performance and Behavior

When you analyze a design with the TimeQuest analyzer, the Fitter memory use and compilation time may increase as compared to the Classic Timing Analyzer, however the timing analysis time may decrease.

The behavior for one value of the **Optimize hold time** Fitter assignment differs between the TimeQuest analyzer and the Classic Timing Analyzer. In the TimeQuest analyzer, the **I/O Paths and Minimum TPD Paths** setting and the **All Paths** setting are equivalent, whereas in the Classic Timing Analyzer the settings directed the Fitter to optimize hold times differently.

Reporting

The TimeQuest analyzer provides a more flexible and powerful interface for reporting timing analysis results than the Classic Timing Analyzer. Although the interface and constraints are more flexible and powerful, both analyzers use the same device timing models, except for device families that support rise/fall analysis. The Classic Timing Analyzer does not support rise/fall analysis, but the TimeQuest analyzer does. Therefore, you may see slightly different delays on identical paths in device families that support rise/fall analysis if you analyze timing in both analyzers.

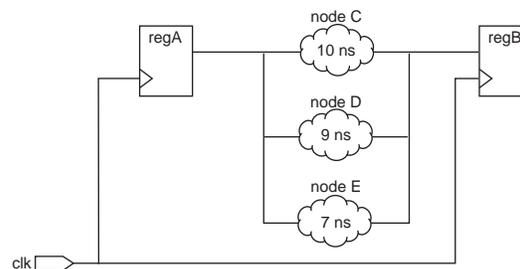
Both analyzers report identical delays along identically constrained paths in your design. The TimeQuest analyzer allows you to constrain some paths that you could not constrain with the Classic Timing Analyzer. Differently constrained paths result in different reported values, but for identical paths in your design that are constrained identically, the delays are exactly the same. Both timing analyzers use the same timing models.

Paths and Pairs

In reporting, the most significant difference between the two analyzers is that the TimeQuest analyzer reports paths, while the Classic Timing Analyzer reports pairs. Path reporting means that the analyzer separately reports every path between two registers. Pair reporting means that the analyzer reports only the worst-case path between two registers. One benefit of path reporting over pair reporting is that you can more easily identify common points in failing paths that may be good targets for optimization.

If your design does not meet timing constraints, this reporting difference can give the impression that there are many more timing failures when you use the TimeQuest analyzer. Figure 8-20 shows a sample circuit followed by a description of the differences between path and pair reporting.

Figure 8-20. Failing Paths



There is an 8-ns period constraint on the `clk` pin, resulting in two paths that fail timing: `regA → C → regB` and `regA → D → regB`. The Classic Timing Analyzer reports only worst-case path `regA → C → regB`. The TimeQuest analyzer reports both failing paths `regA → C → regB` and `regA → D → regB`. It also reports path `regA → E → regB` with positive slack.

Default Reports

The TimeQuest analyzer generates only a small number of reports by default, as compared to the Classic Timing Analyzer, which generates every available report by default. With the TimeQuest analyzer, you generate reports on demand.

Netlist Names

The Classic Timing Analyzer uses register names in reporting, but the TimeQuest analyzer uses register pin names (with the exception of port names of the top-level module). Buried nodes or register names are used when necessary.

[Example 8-5](#) shows how register names are used in Classic Timing Analyzer reports.

Example 8-5. Netlist Names in the Classic Timing Analyzer

```
Info: + Shortest register to register delay is 0.538 ns
      Info: 1: + IC(0.000 ns) + CELL(0.000 ns) = 0.000 ns; Loc. =
LCFF_X1_Y5_N1;
Fanout = 1; REG Node = 'inst'
      Info: 2: + IC(0.305 ns) + CELL(0.149 ns) = 0.454 ns; Loc. =
LCCOMB_X1_Y5_N20; Fanout = 1; COMB Node = 'inst3~feeder'
      Info: 3: + IC(0.000 ns) + CELL(0.084 ns) = 0.538 ns; Loc. =
LCFF_X1_Y5_N21; Fanout = 1; REG Node = 'inst3'
      Info: Total cell delay = 0.233 ns ( 43.31 % )
      Info: Total interconnect delay = 0.305 ns ( 56.69 % )
```

[Example 8-6](#) shows the same information as presented in a TimeQuest analyzer report. In this example, register pin names are used in place of register names.

Example 8-6. Netlist Names in the TimeQuest Analyzer

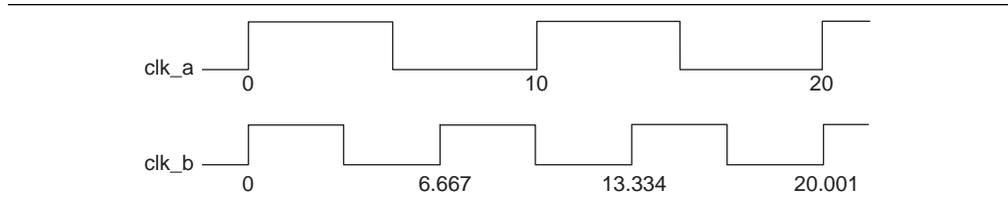
```
Info:      3.788      0.250      uTco      inst
Info:      3.788      0.000 RR      CELL      inst|regout
Info:      4.093      0.305 RR      IC      inst3~feeder|datad
Info:      4.242      0.149 RR      CELL      inst3~feeder|combout
Info:      4.242      0.000 RR      IC      inst3|datain
Info:      4.326      0.084 RR      CELL      inst3
```

Non-Integer Clock Periods

In some cases when related clock periods are not integer multiples of each other, a lack of precision in clock period definition in the TimeQuest analyzer can result in reported setup or hold relationships of a few picoseconds. In addition, launch and latch times for the relationships can be very large. If you experience this, use the `set_max_delay` and `set_min_delay` commands to specify the correct relationships. The Classic Timing Analyzer can maintain additional information about clock frequency that mitigates the lack of precision in clock period definition.

When the clock period cannot be expressed as an integer in terms of picoseconds, you have the situation shown in [Figure 8-21](#). This figure shows two clocks: `clk_a` has a 10 ns period, and `clk_b` has a 6.667 ns period.

Figure 8-21. Very Small Setup Relationship



There is a 1 ps setup relationship at 20 ns because you cannot specify the 6.667 ns period beyond picosecond precision. You should apply the maximum and minimum delay exceptions shown in [Example 8-7](#) between the two clocks to specify the correct relationships.

Example 8-7. Minimum and Maximum Delay Exceptions

```
set_max_delay -from [get_clocks clk_a] -to [get_clocks clk_b] 3.333
set_min_delay -from [get_clocks clk_a] -to [get_clocks clk_b] 0
```

Other Features

The TimeQuest analyzer reports time values without units. By default, the units are nanoseconds, and three decimal places are displayed. You can change the default time unit and decimal places with the `set_time_format` command.

When you use the TimeQuest analyzer in a Tcl shell, output is ASCII-formatted, and columns are aligned for easy reading on 80-column consoles. [Example 8-8](#) shows sample output from the TimeQuest analyzer `report_timing` command.

Example 8-8. ASCII-Formatted TimeQuest Analyzer Report

```
tcl> report_timing -from inst -to inst5
Info: Report Timing: Found 1 setup paths (0 violated). Worst case slack
is 3.634
Info: -from [get_keepers inst]
Info: -to [get_keepers inst5]
Info: Path #1: Slack is 3.634
Info:
=====
Info: From Node      : inst
Info: To Node        : inst5
Info: Launch Clock   : clk_a
Info: Latch Clock    : clk_b
Info:
Info: Data Arrival Path:
Info:
Info: Total (ns)  Incr (ns)      Type  Node
Info: =====  =====  ==  ====
=====
Info:          0.000      0.000           launch edge time
Info:          2.347      2.347    R           clock network delay
Info:          2.597      0.250    uTco    inst
Info:          2.597      0.000  RR    CELL  inst|regout
Info:          3.088      0.491  RR     IC   inst6|datac
Info:          3.359      0.271  RR    CELL  inst6|combout
Info:          3.359      0.000  RR     IC   inst5|datain
Info:          3.443      0.084  RR    CELL  inst5
Info:
Info: Data Required Path:
Info:
Info: Total (ns)  Incr (ns)      Type  Node
Info: =====  =====  ==  ====
=====
Info:          4.000      4.000           latch edge time
Info:          7.041      3.041    R           clock network delay
Info:          7.077      0.036    uTsu    inst5
Info:
Info: Data Arrival Time   :      3.443
Info: Data Required Time  :      7.077
Info: Slack               :      3.634
Info:
=====
Info:
1 3.634
```

Timing Assignment Conversion

This section describes Classic Timing Analyzer QSF timing assignments and the equivalent TimeQuest analyzer constraints. In some cases, there is more than one potentially equivalent SDC constraint. In these cases, correct conversion requires knowledge of the intended function of your design.

This section includes the following topics:

- “Clock Enable Multicycle” on page 8-28
- “Clock Latency” on page 8-25
- “Clock Uncertainty” on page 8-25
- “Cut Timing Path” on page 8-40
- “Default Required f_{MAX} Assignment” on page 8-26
- “Hold Relationship” on page 8-25
- “Input and Output Delay” on page 8-29
- “Inverted Clock” on page 8-25
- “Maximum Clock Arrival Skew” on page 8-41
- “Maximum Data Arrival Skew” on page 8-41
- “Maximum Delay” on page 8-40
- “Minimum Delay” on page 8-41
- “Minimum t_{CO} Requirement” on page 8-36
- “Minimum t_{PD} Requirement” on page 8-40
- “Multicycle” on page 8-27
- “Not a Clock” on page 8-26
- “Setup Relationship” on page 8-24
- “ t_{CO} Requirement” on page 8-34
- “ t_H Requirement” on page 8-32
- “ t_{PD} Requirement” on page 8-38
- “ t_{SU} Requirement” on page 8-30
- “Virtual Clock Reference” on page 8-26

Setup Relationship

The **Setup Relationship** assignment overrides the setup relationship between two clocks. By default, the Classic Timing Analyzer automatically calculates the setup relationship based on your clock settings. The QSF variable for the **Setup Relationship** assignment is `SETUP_RELATIONSHIP`. In the TimeQuest analyzer, use the `set_max_delay` command to specify the maximum setup relationship for a path.

The setup relationship value is the time between latch and launch edges before the TimeQuest analyzer accounts for clock latency, source μt_{CO} , or destination μt_{SU} .

Hold Relationship

The **Hold Relationship** assignment overrides the hold relationship between two clocks. By default, the Classic Timing Analyzer automatically calculates the hold relationship based on your clock settings. The QSF variable for the **Hold Relationship** assignment is `HOLD_RELATIONSHIP`. In the TimeQuest analyzer, use the `set_min_delay` command to specify the minimum hold relationship for a path.

Clock Latency

Table 8–1 shows the equivalent SDC constraints for the **Early Clock Latency** and **Late Clock Latency** Classic Timing Analyzer assignments.

Table 8–1. Classic Timing Analyzer Assignments and SDC Equivalent Constraints for Clock Latency Assignments

Classic Timing Analyzer Assignment		SDC Constraint
Assignment Name	QSF Variable	
Early Clock Latency	<code>EARLY_CLOCK_LATENCY</code>	<code>set_clock_latency -source -late</code>
Late Clock Latency	<code>LATE_CLOCK_LATENCY</code>	<code>set_clock_latency -source -early</code>

For more information about clock latency support in the TimeQuest analyzer, refer to “Clock Latency” on page 8–11.

Clock Uncertainty

Table 8–2 shows the equivalent SDC constraints for the **Clock Setup Uncertainty** and **Clock Hold Uncertainty** Classic Timing Analyzer assignments.

Table 8–2. Classic Timing Analyzer Assignments and SDC Equivalent Constraints for Clock Uncertainty Assignments

Classic Timing Analyzer Timing Assignment		SDC Constraint
Assignment Name	QSF Variable	
Clock Setup Uncertainty	<code>CLOCK_SETUP_UNCERTAINTY</code>	<code>set_clock_uncertainty -setup</code>
Clock Hold Uncertainty	<code>CLOCK_HOLD_UNCERTAINTY</code>	<code>set_clock_uncertainty -hold</code>

Inverted Clock

The Classic Timing Analyzer detects inverted clocks automatically when the clock inversion occurs at the input of the LCELL that contains the register specified in the assignment. You must make an **Inverted Clock** assignment in all other situations for Classic Timing Analyzer analysis. The QSF variable for the **Inverted Clock** assignment is `INVERTED_CLOCK`. The TimeQuest analyzer detects inverted clocks automatically, regardless of the type of inversion circuit, in designs that target device families that support unateness. For designs that target any other device family, you must create a generated clock with the `-invert` option on the output of the cell that inverts the clock.

Not a Clock

The **Not a Clock** assignment directs the Classic Timing Analyzer to identify specified node as not a clock source when it would normally be detected as a clock because of a global f_{MAX} requirement. The QSF variable for the **Not a Clock** assignment is `NOT_A_CLOCK`. This assignment is not supported in the TimeQuest analyzer and there is no equivalent constraint. Appropriate clock constraints are created in the TimeQuest analyzer only.

Default Required f_{MAX} Assignment

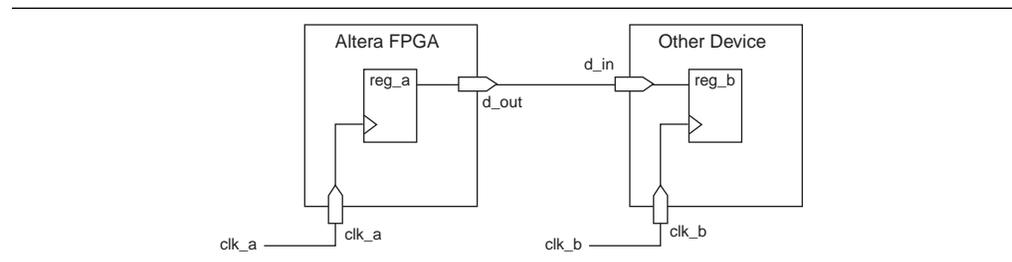
The **Default Required f_{MAX}** assignment allows you to specify an f_{MAX} requirement for the Classic Timing Analyzer for all unconstrained clocks in your design. The QSF variable for the **Default Required f_{MAX}** assignment is `FMAX_REQUIREMENT`. You can use the `derive_clocks` command to create clocks on sources of clock pins in your design that do not already have clocks assigned to them. You should constrain each individual clock in your design with the `create_clock` or `created_generated_clock` command, instead of the `derive_clocks` command. Refer to “[Automatic Clock Detection](#)” on page 8-14 to learn why you should constrain individual clocks in your design.

Virtual Clock Reference

The **Virtual Clock Reference** assignment allows you to define timing characteristics of a reference clock outside the FPGA. The QSF variable for the **Virtual Clock Reference** assignment is `VIRTUAL_CLOCK_REFERENCE`. The TimeQuest analyzer supports virtual clocks by default, while the Classic Timing Analyzer requires the **Virtual Clock Reference** assignment to indicate that a clock setting is for a virtual clock. To create a virtual clock in the TimeQuest analyzer, use the `create_clock` or `create_generated_clock` commands with the `-name` option and no targets.

[Figure 8-22](#) shows a circuit that requires a virtual clock, and the following example shows how to constrain the circuit. The circuit shows data transfer between an Altera FPGA and another device, and the clocks for the two devices are not related. You can constrain the path with an output delay assignment, but that assignment requires a virtual clock that defines the clock characteristics of the destination device.

Figure 8-22. Virtual Clock Circuit



Assume the circuit has the following assignments in the Classic Timing Analyzer:

- Clock period of 10 ns on `system_clk` (clock for the Altera FPGA)
- Clock period of 8 ns on `virt_clk` (clock for the other device)
- The **Virtual Clock Reference** setting for `virt_clk` is **On** (indicates that `virt_clk` is a virtual clock)
- The **Output Maximum Delay** setting of 5 ns on `dataout` with respect to `virt_clk` (constrains the path between the two devices)

The SDC commands shown in [Example 8-9](#) constrain the circuit the same way.

Example 8-9. SDC Constraints

```
# Clock for the Altera FPGA
create_clock -period 10 -name system_clk [get_ports system_clk]
# Virtual clock for the other device, with no targets
create_clock -period 8 -name virt_clk
# Constrains the path between the two devices
set_output_delay -clock virt_clk 5 [get_ports dataout]
```

Clock Settings

The Classic Timing Analyzer includes a variety of assignments to describe clock settings. These include duty cycle, f_{MAX} , offset, and others. In the TimeQuest analyzer, use the `create_clock` and `create_generated_clock` commands to constrain clocks.

Multicycle

[Table 8-3](#) shows the equivalent SDC exceptions for each of these Classic Timing Analyzer timing assignments.

Table 8-3. Classic Timing Analyzer Multicycle Assignments and SDC Equivalent Exceptions

Classic Timing Analyzer Multicycle Assignment		SDC Exception
Assignment Name	QSF Variable	
Multicycle (1)	MULTICYCLE	<code>set_multicycle_path -setup -end</code>
Source Multicycle (2)	SRC_MULTICYCLE	<code>set_multicycle_path -setup -start</code>
Multicycle Hold (3)	HOLD_MULTICYCLE	<code>set_multicycle_path -hold -end</code>
Source Multicycle Hold	SRC_HOLD_MULTICYCLE	<code>set_multicycle_path -hold -start</code>

Notes to Table 8-3:

- (1) A multicycle assignment is also known as a “destination multicycle setup” assignment.
- (2) A source multicycle assignment is also known as a “source multicycle setup” assignment.
- (3) A multicycle hold is also known as a “destination multicycle hold” assignment.

The default value and numbering scheme for the hold multicycle value is different in the Classic Timing Analyzer and TimeQuest analyzer. Refer to “Hold Multicycle” on page 8-18 for more information about the difference between the default value and numbering scheme for the hold multicycle value in the Classic Timing Analyzer and TimeQuest analyzer.



For more information about how to convert the hold multicycle value, refer to the *Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

Clock Enable Multicycle

The Classic Timing Analyzer supports the following clock enable multicycle assignments:

- Clock Enable Multicycle
- Clock Enable Source Multicycle
- Clock Enable Multicycle Hold
- Clock Enable Source Multicycle Hold

Corresponding types of multicycle assignments are applied to all registers enabled by the targets of the specified clock. The TimeQuest analyzer supports clock-enabled multicycle constraints with the `get_fanouts` command. Use the `get_fanouts` command to create a collection of nodes that have a common source signal, such as a clock enable.

I/O Constraints

FPGA I/O timing assignments have typically been made with FPGA-centric t_{SU} and t_{CO} requirements for the Classic Timing Analyzer. However, the Classic Timing Analyzer also supports input and output delay assignments to accommodate industry-standard, system-centric timing constraints. Where possible, you should use system-centric constraints to constrain your designs for the TimeQuest analyzer. Table 8-4 includes Classic Timing Analyzer I/O assignments, the equivalent FPGA-centric SDC constraints, and recommended system-centric SDC constraints.

For setup checks (t_{SU} and t_{CO}), $\langle latch - launch \rangle$ equals the clock period for same-clock transfers. For hold checks (t_H and Minimum t_{CO}), $\langle latch - launch \rangle$ equals 0 for same-clock transfers. Conversions from Classic Timing Analyzer assignments to `set_input_delay` and `set_output_delay` constraints work well only when the source and destination registers' clocks are the same (same clock and polarity). If the source and destination registers' clocks are different, the conversion may not be straightforward and you should take extra care when converting to `set_input_delay` and `set_output_delay` constraints.

Table 8-4. Classic Timing Analyzer and TimeQuest Analyzer Equivalent I/O Constraints

Classic Timing Analyzer Assignment	FPGA-centric SDC	System-centric SDC
t_{SU} Requirement	<code>set_max_delay <t_{SU} requirement></code>	<code>set_input_delay -max <latch - launch -- t_{SU} requirement></code>
t_H Requirement	<code>set_min_delay - <t_H requirement> (1)</code>	<code>set_input_delay -min <latch -- launch + t_H requirement></code>
t_{CO} Requirement	<code>set_max_delay <t_{CO} requirement></code>	<code>set_output_delay -max <latch -- launch - t_{CO} requirement></code>
Minimum t_{CO} Requirement	<code>set_min_delay <minimum t_{CO} requirement></code>	<code>set_output_delay -min <latch -- launch - minimum t_{CO} requirement></code>
t_{PD} Requirement	<code>set_max_delay <t_{PD} requirement></code>	(2)
Minimum t_{PD} Requirement	<code>set_min_delay <minimum t_{PD} requirement></code>	(2)

Notes to Table 8-4:

- (1) Refer to “ t_H Requirement” on page 8-32 for an explanation about why this exception uses the negative t_H requirement.
- (2) The input and output delays can be used for t_{PD} paths, such that they will be analyzed as a system f_{MAX} path. This is a feature unique to the TimeQuest analyzer.

Input and Output Delay

Table 8-5 shows the equivalent SDC exceptions for Classic Timing Analyzer timing assignments.

Table 8-5. Classic Timing Analyzer Assignments and SDC Equivalent Exceptions

Classic Timing Analyzer Assignment		SDC Exception
Assignment Name	QSF Variable	
Input Maximum Delay	INPUT_MAX_DELAY	<code>set_input_delay -max</code>
Input Minimum Delay	INPUT_MIN_DELAY	<code>set_input_delay -min</code>
Output Maximum Delay	OUTPUT_MAX_DELAY	<code>set_output_delay -max</code>
Output Minimum Delay	OUTPUT_MIN_DELAY	<code>set_output_delay -min</code>

In some circumstances, you may receive the following warning message when you update the SDC netlist:

```
Warning: For set_input_delay/set_output_delay, port "<port>" does not have delay for flag (<rise|fall>, <min|max>)
```

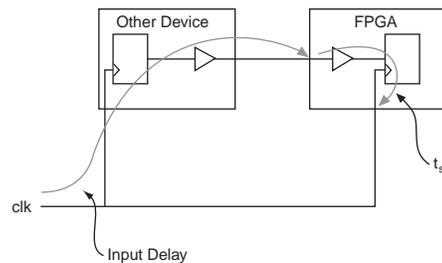

Equation 8-3 shows the derivation of this conversion.

Equation 8-3.

$$\begin{aligned} \text{required} - \text{arrival} &> 0 \\ \text{required} &= \text{latch} + \text{board.dstclk} + \text{dst.clk} - \text{dst.utsu} \\ \text{arrival} &= \text{launch} + \text{board.srcclk} + \text{src.clk} + \text{src.utco} + \text{src.out} + \text{srctodst} + \text{dst.in} \\ \text{input_delay} &= \text{board.srcclk} + \text{src.clk} + \text{src.utcu} + \text{src.out} + \text{srctodst} - \text{board.dstclk} \\ \text{required} &= \text{latch} + \text{dst.clk} - \text{dst.utsu} \\ \text{arrival} &= \text{launch} + \text{input_delay} + \text{dst.in} \\ (\text{latch} + \text{dst.clk} - \text{dst.utsu}) - (\text{launch} + \text{input_delay} + \text{dst.in}) &> 0 \\ t_{su} \text{ requirement} - \text{actual } t_{su} &> 0 \\ \text{actual } t_{su} &= \text{dst.in} + \text{dst.utsu} - \text{dst.clk} \\ t_{su} \text{ requirement} - (\text{dst.in} + \text{dst.utsu} - \text{dst.clk}) &> 0 \\ t_{su} \text{ requirement} &= \text{latch} - \text{launch} - \text{input_delay} \\ \text{input_delay} &= \text{latch} - \text{launch} - t_{su} \text{ requirement} \end{aligned}$$

The delay value is the difference between the period of the clock source of the register and the t_{su} Requirement value, as shown in Figure 8-24.

Figure 8-24. t_{su} Requirement



The delay value for the set_max_delay command is the t_{su} Requirement value. Equation 8-4 shows the derivation of this conversion.

Equation 8-4.

$$\begin{aligned} \text{required} - \text{arrival} &> 0 \\ \text{required} &= \text{latch} + \text{board.dstclk} + \text{dst.clk} - \text{dst.utsu} \\ \text{arrival} &= \text{launch} + \text{board.srcclk} + \text{src.clk} + \text{src.utco} + \text{src.out} + \text{srctodst} + \text{dst.in} \\ \text{max_delay} &= \text{latch} + \text{board.dstclk} + - \text{launch} - \text{board.srcclk} - \text{src.clk} - \text{src.out} - \text{srctodst} \\ \text{required} &= \text{max_delay} + \text{dst.clk} - \text{dst.utsu} \\ \text{arrival} &= \text{dst.in} \\ (\text{max_delay} + \text{dst.clk} - \text{dst.utsu}) - (\text{dst.in}) &> 0 \\ t_{su} \text{ requirement} - t_{su} &> 0 \\ \text{actual } t_{su} &= \text{dst.in} + \text{dst.utsu} - \text{dst.clk} \\ t_{su} \text{ requirement} - (\text{dst.in} + \text{dst.utsu} - \text{dst.clk}) &> 0 \\ \text{set_max_delay} &= t_{su} \text{ requirement} \end{aligned}$$

Table 8-6 shows the different ways you can make t_{SU} assignments in the Classic Timing Analyzer, and the corresponding options for the `set_max_delay` exception.

Table 8-6. t_{SU} Requirement and `set_max_delay` Equivalence

t_{SU} Requirement Options	<code>set_max_delay</code> Options
-to <pin>	-from [get_ports <pin>] -to [get_registers *]
-to <clock>	-from [get_ports *] -to [get_clocks <clock>]
-to <register>	-from [get_ports *] -to [get_registers <register>]
-from <pin> -to <register>	-from [get_ports <pin>] -to [get_registers <register>]
-from <clock> -to <pin>	-from [get_ports <pin>] -to [get_clocks <clock>] (1)

Note to Table 8-6:

- (1) If the pin in this assignment feeds registers clocked by the same clock, it is equivalent to the first option, -to <pin>. If the pin feeds registers clocked by different clocks, use `set_input_delay` to constrain the paths properly.

To convert a global t_{SU} assignment to an equivalent SDC exception, use the command shown in Example 8-10.

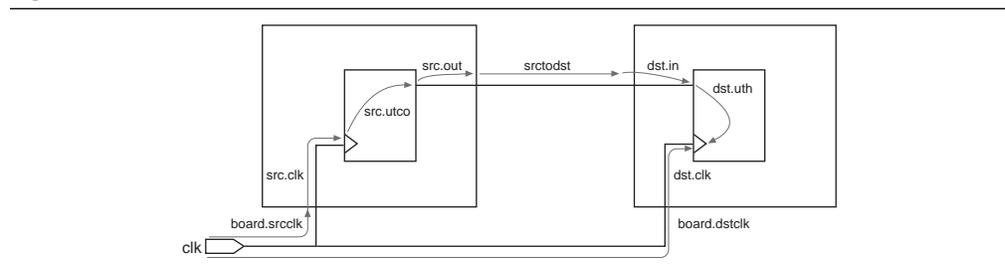
Example 8-10. Converting a Global t_{SU} Assignment to an Equivalent SDC Exception

```
set_max_delay -from [all_inputs] -to [all_registers] < $t_{SU}$  value>
```

t_H Requirement

The t_H Requirement specifies the maximum acceptable clock hold time for the input (data) pin. The QSF variable for the t_H Requirement assignment is `TH_REQUIREMENT`. You can convert the t_H Requirement assignment to the `set_min_delay` command, or the `set_input_delay` command with the `-min` option. The delay value for the `set_input_delay` command is $\langle \text{latch} - \text{launch} + t_H \text{ requirement} \rangle$. Refer to the labeled paths in Figure 8-25 to understand the names in Equation 8-5 and Equation 8-6.

Figure 8-25. Path Names



Equation 8-5 shows the derivation of this calculation.

Equation 8-5.

$$\begin{aligned}
 & \text{arrival} - \text{required} > 0 \\
 & \text{arrival} = \text{launch} + \text{board.srcclk} + \text{src.clk} + \text{src.utco} + \text{src.out} + \text{srctodst} + \text{dst.in} \\
 & \text{required} = \text{latch} + \text{board.dstclk} + \text{dst.clk} + \text{dst.uth} \\
 & \text{input_delay} = \text{board.srcclk} + \text{src.clk} + \text{srcutcu} + \text{src.out} + \text{srctodst} - \text{board.dstclk} \\
 & \text{arrival} = \text{launch} + \text{input_delay} + \text{dst.in} \\
 & \text{required} = \text{latch} + \text{dst.clk} + \text{dst.uth} \\
 & (\text{launch} + \text{input_delay} + \text{dst.in}) - (\text{latch} + \text{dst.clk} + \text{dst.uth}) > 0 \\
 & t_H \text{ requirement} - \text{actual } t_H > 0 \\
 & \text{actual } t_H = \text{dst.clk} + \text{dst.uth} - \text{dst.in} \\
 & t_H \text{ requirement} - (\text{dst.clk} + \text{dst.uth} - \text{dst.in}) > 0 \\
 & t_H \text{ requirement} = \text{launch} - \text{latch} + \text{input_delay} \\
 & \text{input_delay} = \text{latch} - \text{launch} + t_H \text{ requirement}
 \end{aligned}$$

The delay value for the `set_min_delay` command is the **t_H Requirement** value. Equation 8-6 shows the derivation of this conversion.

Equation 8-6.

$$\begin{aligned}
 & \text{arrival} - \text{required} > 0 \\
 & \text{arrival} = \text{dst.in} \\
 & \text{required} = \text{min_delay} + \text{dst.clk} + \text{dst.uth} \\
 & \text{dst.in} - (\text{min_delay} + \text{dst.clk} + \text{dst.uth}) \\
 & t_H \text{ requirement} - \text{actual } t_H > 0 \\
 & \text{actual } t_H = \text{dst.clk} + \text{dst.uth} - \text{dst.in} \\
 & t_H \text{ requirement} - (\text{dst.clk} + \text{dst.uth} - \text{dst.in}) > 0 \\
 & \text{set_min_delay} = -t_H \text{ requirement}
 \end{aligned}$$

Table 8-7 shows the different ways you can make t_H assignments in the Classic Timing Analyzer, and the corresponding options for the `set_min_delay` command.

Table 8-7. t_H Requirement and `set_min_delay` Equivalence

t_H Requirement Options	<code>set_min_delay</code> Options
-to <pin>	-from [get_ports <pin>] -to [get_registers *]
-to <clock>	-from [get_ports *] -to [get_clocks <clock>]
-to <register>	-from [get_ports *] -to [get_registers <register>]
-from <pin> -to <register>	-from [get_ports <pin>] -to [get_registers <register>]
-from <clock> -to <pin>	-from [get_ports <pin>] -to [get_clocks <clock>] (1)

Note to Table 8-7:

- (1) If the pin in this assignment feeds registers clocked by the same clock, it is equivalent to the first option, -to <pin>. If the pin feeds registers clocked by different clocks, use `set_input_delay` to constrain the paths properly. Refer to “Input and Output Delay” on page 8-29 for additional information.

To convert a global t_H assignment to an equivalent SDC exception, use the command shown in Example 8-11.

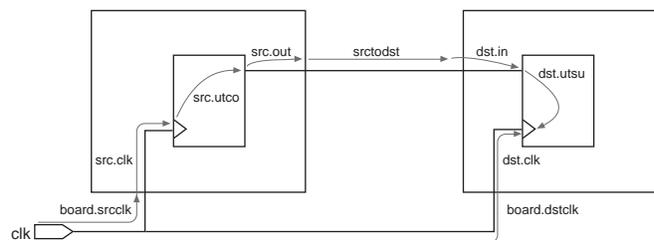
Example 8-11. Converting a Global t_H Assignment to an Equivalent SDC Exception

```
set_min_delay -from [all_inputs] -to [all registers] <negative  $t_H$  value>
```

t_{CO} Requirement

The t_{CO} Requirement assignment specifies the maximum acceptable clock to output delay to the output pin. The QSF variable for the t_{CO} Requirement assignment is `TCO_REQUIREMENT`. You can convert the t_{CO} Requirement assignment to the `set_max_delay` command or the `set_output_delay` command with the `-max` option. The delay value for the `set_output_delay` command is `<latch - launch - t_{CO} requirement>`. Refer to the labeled paths in Figure 8-26 to understand the names in Equation 8-7 and Equation 8-8.

Figure 8-26. Path Names



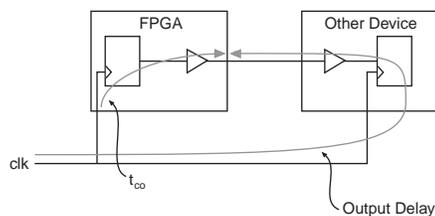
Equation 8-7 shows the derivation of this conversion.

Equation 8-7.

$$\begin{aligned} \text{required} - \text{arrival} &> 0 \\ \text{required} &= \text{latch} - \text{output_delay} \\ \text{arrival} &= \text{launch} + \text{src.clk} + \text{src.utco} + \text{src.out} \\ \text{output_delay} &= \text{srctodst} + \text{dst.in} + \text{dst.utsu} - \text{dst.clk} - \text{board.dstc.k} + \text{board.srcclk} \\ (\text{latch} - \text{output_delay}) - (\text{launch} + \text{src.clk} + \text{src.utco} + \text{src.out}) &> 0 \\ t_{\text{co}} \text{ requirement} - \text{actual } t_{\text{co}} &> 0 \\ \text{actual } t_{\text{co}} &= \text{launch} + \text{src.clk} + \text{src.utco} + \text{src.out} \\ t_{\text{co}} \text{ requirement} - (\text{src.clk} + \text{src.utco} + \text{src.out}) &> 0 \\ t_{\text{co}} \text{ requirement} &= \text{latch} - \text{launch} - \text{output_delay} \\ \text{output_delay} &= \text{latch} - \text{launch} - t_{\text{co}} \text{ requirement} \end{aligned}$$

The delay value is the difference between the period of the clock source of the register and the t_{CO} Requirement value, as illustrated in Figure 8-27.

Figure 8-27. t_{CO} Requirement



The delay value for the `set_max_delay` command is the t_{CO} Requirement value. Equation 8-8 shows the derivation of this conversion.

Equation 8-8.

$$\begin{aligned} \text{required} - \text{arrival} &> 0 \\ \text{required} &= \text{set_max_delay} \\ \text{arrival} &= \text{src.clk} + \text{src.utco} + \text{src.out} \\ \text{set_max_delay} - (\text{src.clk} + \text{src.utco} + \text{src.out}) &> 0 \\ t_{\text{co}} \text{ requirement} - \text{actual } t_{\text{co}} &> 0 \\ \text{actual } t_{\text{co}} &= \text{src.clk} + \text{src.utco} + \text{src.out} \\ t_{\text{co}} \text{ requirement} - (\text{src.clk} + \text{src.utco} + \text{src.out}) &> 0 \\ \text{set_max_delay} &= t_{\text{co}} \text{ requirement} \end{aligned}$$

Table 8-8 shows the different ways you can make t_{CO} assignments in the Classic Timing Analyzer, and the corresponding options for the `set_max_delay` exception.

Table 8-8. t_{CO} Requirement and `set_max_delay` Equivalence

t_{CO} Requirement Options	<code>set_max_delay</code> Options
-to <pin>	-from [get_registers *] -to [get_ports <pin>]
-to <clock>	-from [get_clocks <clock>] -to [get_ports *]
-to <register>	-from [get_registers <register>] -to [get_ports *]
-from <register> -to <pin>	-from [get_registers <register>] -to [get_ports <pin>]
-from <clock> -to <pin>	-from [get_clocks <clock>] -to [get_ports <pin>] (1)

Note to Table 8-8:

- (1) If the pin in this assignment feeds registers clocked by the same clock, it is equivalent to the first option, -to <pin>. If the pin feeds registers clocked by different clocks, you should use `set_output_delay` to constrain the paths properly.

To convert a global t_{CO} assignment to an equivalent SDC exception, use the command in Example 8-12.

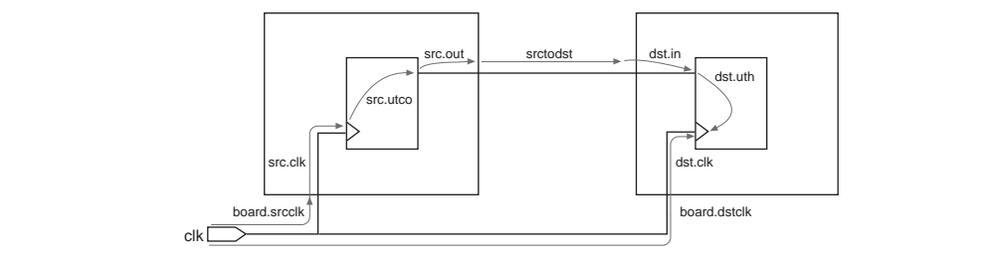
Example 8-12. Converting a Global t_{CO} Assignment to an Equivalent SDC Exception

```
set_max_delay -from [all registers] -to [all_outputs] < $t_{CO}$  value>
```

Minimum t_{CO} Requirement

The **Minimum t_{CO} Requirement** assignment specifies the minimum acceptable clock to output delay to the output pin. The QSF variable for the **Minimum t_{CO} Requirement** assignment is `MIN_TCO_REQUIREMENT`. You can convert the **Minimum t_{CO} Requirement** assignment to the `set_min_delay` command or the `set_output_delay` command with the `-min` option. The delay value for the `set_output_delay` command is <latch – launch – minimum t_{CO} requirement>. Refer to the labeled paths in Figure 8-28 to understand the names in Equation 8-9 and Equation 8-10.

Figure 8-28. Path Names



Equation 8–9 shows the derivation of this conversion.

Equation 8–9.

$$\begin{aligned}
 & \text{arrival} + \text{required} > 0 \\
 & \text{arrival} = \text{launch} + \text{src.clk} + \text{src.utco} + \text{src.out} \\
 & \text{required} = \text{latch} - \text{output_delay} \\
 & \text{output_delay} = \text{srctodst} + \text{dst.in} - \text{dst.uth} - \text{dst.clk} - \text{board.dstclk} + \text{board.srcclk} \\
 & (\text{launch} + \text{src.clk} + \text{src.utco} + \text{src.out}) - (\text{latch} - \text{output_delay}) > 0 \\
 & \text{minimum } t_{\text{co}} - \text{minimum } t_{\text{co}} \text{ requirement} > 0 \\
 & \text{minimum } t_{\text{co}} = \text{launch} + \text{src.clk} + \text{src.utco} + \text{src.out} \\
 & (\text{launch} + \text{src.clk} + \text{src.utco} + \text{src.out}) - \text{minimum } t_{\text{co}} \text{ requirement} > 0 \\
 & \text{minimum } t_{\text{co}} \text{ requirement} = \text{latch} - \text{launch} - \text{output_delay} \\
 & \text{output_delay} = \text{latch} - \text{launch} - \text{minimum } t_{\text{co}} \text{ requirement}
 \end{aligned}$$

The delay value for the `set_min_delay` command is the **Minimum t_{co} Requirement**. Equation 8–10 shows the derivation of this conversion.

Equation 8–10.

$$\begin{aligned}
 & \text{arrival} - \text{required} > 0 \\
 & \text{arrival} = \text{src.clk} + \text{src.utco} + \text{src.out} \\
 & \text{required} = \text{min_delay} \\
 & (\text{src.clk} + \text{src.utco} + \text{src.out}) - (\text{set_min_delay}) > 0 \\
 & \text{minimum } t_{\text{co}} - \text{minimum } t_{\text{co}} \text{ requirement} > 0 \\
 & \text{minimum } t_{\text{co}} = \text{src.clk} + \text{src.utco} + \text{src.out} \\
 & (\text{src.clk} + \text{src.utco} + \text{src.out}) - \text{minimum } t_{\text{co}} \text{ requirement} > 0 \\
 & \text{set_min_delay} = \text{minimum } t_{\text{co}} \text{ requirement}
 \end{aligned}$$

Table 8-9 shows the different ways you can make minimum t_{CO} assignments in the Classic Timing Analyzer, and the corresponding options for the `set_min_delay` exception.

Table 8-9. Minimum t_{CO} Requirement and `set_min_delay` Equivalence

Minimum t_{CO} Requirement Options	<code>set_min_delay</code> Options
-to <pin>	-from [get_registers *] -to [get_ports <pin>]
-to <clock>	-from [get_clocks <clock>] -to [get_ports *]
-to <register>	-from [get_registers <register>] -to [get_ports *]
-from <register> -to <pin>	-from [get_registers <register>] -to [get_ports <pin>]
-from <clock> -to <pin>	-from [get_clocks <clock>] -to [get_ports <pin>] (1)

Note to Table 8-9:

- (1) If the pin in this assignment feeds registers clocked by the same clock, it is equivalent to the first option, -to <pin>. If the pin feeds registers clocked by different clocks, use `set_output_delay` to constrain the paths properly.

To convert a global **Minimum t_{CO} Requirement** to an equivalent SDC exception, use the command shown in Example 8-13.

Example 8-13. Converting a Global Minimum t_{CO} Requirement to an Equivalent SDC Exception

```
set_min_delay -from [all_registers] -to [all_outputs] <minimum tCO value>
```

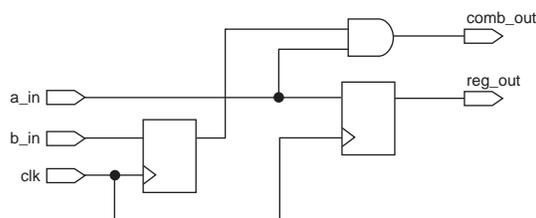
t_{PD} Requirement

The **t_{PD} Requirement** assignment specifies the maximum acceptable input to nonregistered output delay, that is, the time required for a signal from an input pin to propagate through combinational logic and appear at an output pin. The QSF variable for the **t_{PD} Requirement** assignment is `TPD_REQUIREMENT`. You can use the `set_max_delay` command in the TimeQuest analyzer as an equivalent constraint as long as you account for input and output delays. The **t_{PD} Requirement** assignment does not take into account input and output delays, but the `set_max_delay` exception does, so you must modify the `set_max_delay` value to take into account input and output delays.

Combinational Path Delay Scenario

Figure 8-29 shows a circuit to illustrate the following example of converting a **t_{PD} Requirement** to a `set_max_delay` constraint.

Figure 8-29. t_{PD} Example



Assume the circuit has the following assignments in the Classic Timing Analyzer:

- Clock period of 10 ns
- t_{PD} **Requirement** from a_in to comb_out of 10 ns
- **Input Max Delay** on a_in relative to clk of 2 ns
- **Output Max Delay** on comb_out relative to clk of 2 ns

The path from a_in to comb_out is not affected by the input and output delays. The slack is equal to the $\langle t_{PD} \text{ Requirement from a_in to comb_out} \rangle - \langle \text{path delay from a_in to comb_out} \rangle$.

Assume the circuit has the SDC constraints shown in [Example 8-14](#) in the TimeQuest analyzer.

Example 8-14. SDC Constraints

```
create_clock -period 10 -name clk [get_ports clk]
set_max_delay -from a_in -to comb_out 10
set_input_delay -clk clk 2 [get_ports a_in]
set_output_delay -clk clk 2 [get_ports comb_out]
```

The path from a_in to comb_out is affected by the input and output delays. The slack is equal to:

$\langle \text{set_max_delay value from a_in to comb_out} \rangle - \langle \text{input delay} \rangle - \langle \text{output delay} \rangle - \langle \text{path delay from a_in to comb_out} \rangle$

To convert a global t_{PD} **Requirement** assignment to an equivalent SDC exception, use the command shown in [Example 8-15](#). You should add the input and output delays to the value of your converted t_{PD} **Requirement** (**set_max_delay** exception value) to achieve an equivalent SDC exception.

Example 8-15. Converting a Global t_{PD} Requirement Assignment to an Equivalent SDC Exception

```
set_max_delay -from [all_inputs] -to [all_outputs] <value>
```

Minimum t_{PD} Requirement

The **Minimum t_{PD} Requirement** assignment specifies the minimum acceptable input to nonregistered output delay, that is, the minimum time required for a signal from an input pin to propagate through combinational logic and appear at an output pin. The QSF variable for the **Minimum t_{PD} Requirement** assignment is `MIN_TPD_REQUIREMENT`. You can use the `set_min_delay` command in the TimeQuest analyzer as an equivalent constraint as long as you account for input and output delays. The **Minimum t_{PD} Requirement** assignment does not take into account input and output delays, but the `set_min_delay` exception does.

Refer to “[Combinational Path Delay Scenario](#)” on page 8-38 to see how input and output delays affect minimum and maximum delay exceptions.

To convert a global **Minimum t_{PD} Requirement** assignment to an equivalent SDC exception, use the command shown in [Example 8-16](#).

Example 8-16. Converting a Global Minimum t_{PD} Requirement Assignment to an Equivalent SDC Exception

```
set_min_delay -from [all_inputs] -to [all_outputs] <value>
```

Cut Timing Path

The **Cut Timing Path** assignment in the Classic Timing Analyzer is equivalent to the `set_false_path` command in the TimeQuest analyzer. The QSF variable for the **Cut Timing Path** assignment is `CUT`.

Maximum Delay

The **Maximum Delay** assignment specifies the maximum required delay for the following types of paths:

- Pins to registers
- Registers to registers
- Registers to pins

The QSF variable for the **Maximum Delay** assignment is `MAX_DELAY`. This overwrites the requirement computed from the clock setup relationship and clock skew. There is no equivalent constraint in the TimeQuest analyzer.



The **Maximum Delay** assignment for the Classic Timing Analyzer is not related to the `set_max_delay` exception in the TimeQuest analyzer.

Minimum Delay

The **Minimum Delay** assignment specifies the minimum required delay for the following types of paths:

- Pins to registers
- Registers to registers
- Registers to pins

The QSF variable for the **Minimum Delay** assignment is `MIN_DELAY`. This overwrites the requirement computed from the clock hold relationship and clock skew. There is no equivalent constraint in the TimeQuest analyzer.



The **Minimum Delay** assignment for the Classic Timing Analyzer is not related to the `set_min_delay` exception in the TimeQuest analyzer.

Maximum Clock Arrival Skew

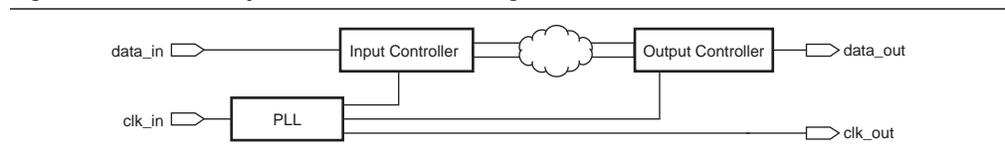
The **Maximum Clock Arrival Skew** assignment specifies the maximum clock skew between a set of registers. The QSF variable for the **Maximum Clock Arrival Skew** assignment is `MAX_CLOCK_ARRIVAL_SKEW`. In the Classic Timing Analyzer, this assignment is specified between a clock node name and a set of registers. **Maximum Clock Arrival Skew** is not supported in the TimeQuest analyzer.

Maximum Data Arrival Skew

The **Maximum Data Arrival Skew** assignment specifies the maximum data arrival skew between a set of registers, pins, or both. The QSF variable for the **Maximum Data Arrival Skew** assignment is `MAX_DATA_ARRIVAL_SKEW`. In this case, the data arrival delay represents the t_{CO} from the clock to the given register, pin, or both. This assignment is specified between a clock node and a set of registers, pins, or both.

The TimeQuest analyzer does not support a constraint to specify maximum data arrival skew, but you can specify setup and hold times relative to a clock port to constrain an interface like this. [Figure 8-30](#) shows a simplified source-synchronous interface used in the following example.

Figure 8-30. Source-Synchronous Interface Diagram

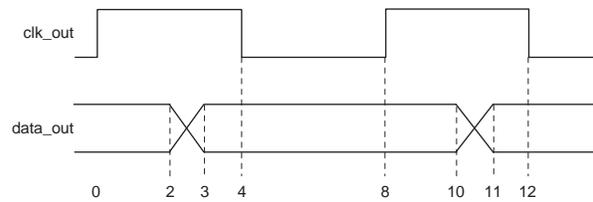


Constraining Skew on an Output Bus

This example constrains the interface so that all bits of the `data_out` bus go off-chip between 2 ns and 3 ns after the `clk_out` signal. Assume that `clk_in` and `clk_out` have a period of 8 ns.

The following equations and example show how to create timing requirements that satisfy the timing relationships shown in [Figure 8-31](#).

Figure 8-31. Source-Synchronous Timing Diagram



[Equation 8-11](#) shows how to compute the value for the `set_output_delay -min` command that creates the 2 ns hold requirement on the destination. For hold requirement calculations in which source and destination clocks are the same, $\langle latch \rangle - \langle launch \rangle = 0$.

Equation 8-11.

$$\begin{aligned} latch - launch &= 0\text{ns} \\ \text{output delay} &= latch - launch - 2\text{ns} \\ \text{output delay} &= -2\text{ns} \end{aligned}$$

[Equation 8-12](#) shows how to compute the value for the `set_output_delay` command that creates the 3 ns setup requirement on the destination. For setup requirement calculations in which source and destination clocks are the same, $\langle latch \rangle - \langle launch \rangle = \text{clock period}$.

Equation 8-12.

$$\begin{aligned} latch - launch &= 8\text{ns} \\ \text{output delay} &= latch - launch - 3\text{ns} \\ \text{output delay} &= 5\text{ns} \end{aligned}$$

Refer to [“I/O Constraints” on page 8-28](#) for an explanation of the above equations.

[Example 8-17](#) shows the three constraints together.

Example 8-17. Constraining a DDR Interface

```
set period 8.000
create_clock -period $period \
  -name clk_in \
  [get_ports data_out*]
derive_pll_clocks
set_output_delay -add_delay \
  -clock ddr_pll_1_inst|altpll_component|pll|CLK[0] \
  -reference_pin [get_ports clk_out] \
  -min -2.000 \
  [get_ports data_out*]
set_output_delay -add_delay \
  -clock ddr_pll_1_inst|altpll_component|pll|CLK[0] \
  -reference_pin [get_ports clk_out] \
  -max [expr $period - 3.000] \
  [get_ports data_out*]
```

Conversion Utility

The TimeQuest analyzer includes a conversion utility to help you convert Classic Timing Analyzer assignments in a .qsf to SDC constraints in an .sdc. The utility can use information from your project report database (in the \ddb folder), if it exists, so you should compile your design before the conversion.



The conversion utility ignores all disabled QSF assignments. Disabled assignments show **No** in the **Enabled** column of the Assignment Editor, and include the -disable option in the .qsf.

Refer to “[Create SDC Constraints from Existing Timing Assignments](#)” on page 8-2 to learn how to run the conversion utility.

Unsupported Global Assignments

The conversion utility checks whether any of the global timing assignments in [Table 8-10](#) exist in your project. Any global assignments not supported by the conversion utility are ignored during the conversion. Refer to the indicated page for information about each assignment, and how to manually convert these global assignments to SDC commands.

Table 8-10. Global Timing Assignments

Assignment Name	QSF Variable	More Information
t_{SU} Requirement	TSU_REQUIREMENT	page 8-30
t_H Requirement	TH_REQUIREMENT	page 8-32
t_{CO} Requirement	TCO_REQUIREMENT	page 8-34
Minimum t_{CO} Requirement	MIN_TCO_REQUIREMENT	page 8-36
t_{PD} Requirement	TPD_REQUIREMENT	page 8-38
Minimum t_{PD} Requirement	MIN_TPD_REQUIREMENT	page 8-40

Recommended Global Assignments

When any unsupported assignments have been identified, the conversion utility checks the global assignments in [Table 8-11](#) to ensure they match the specified values.

Table 8-11. Recommended Global Assignments

Classic Timing Analyzer Assignment Name	QSF Variable	Value
Cut off clear and preset signal paths	CUT_OFF_CLEAR_AND_PRESET_PATHS	ON
Cut off feedback from I/O pins	CUT_OFF_IO_PIN_FEEDBACK	ON
Cut off read during write signal paths	CUT_OFF_READ_DURING_WRITE_PATHS	ON
Analyze latches as synchronous elements	ANALYZE_LATCHES_AS_SYNCHRONOUS_ELEMENTS	ON
Enable Clock Latency	ENABLE_CLOCK_LATENCY	ON
Display Entity Name	PROJECT_SHOW_ENTITY_NAME	ON

The following assignments are checked to ensure the functionality of the Classic Timing Analyzer with the specified values corresponding to the behavior of the TimeQuest analyzer.

- **Cut off clear and preset signal paths**—The TimeQuest analyzer does not support this functionality. You should use Recovery and Removal analysis instead to analyze register control paths. The Classic Timing Analyzer does not support this option.
- **Cut off feedback from I/O pins**—The TimeQuest analyzer does not match the functionality of the Classic Timing Analyzer when this assignment is set to OFF.
- **Cut off read during write signal paths**—The TimeQuest analyzer does not match the functionality of the Classic Timing Analyzer when this assignment is set to OFF.
- **Analyze latches as synchronous elements**—The TimeQuest analyzer analyzes latches as synchronous elements by default and does not match the functionality of the Classic Timing Analyzer when this assignment is set to OFF. The Classic Timing Analyzer analyzes latches as synchronous elements by default.
- **Enable Clock Latency**—The TimeQuest analyzer includes clock latency in its calculations. The TimeQuest analyzer does not match the functionality of the Classic Timing Analyzer when this assignment is set to OFF. Latency on a clock can be viewed as a simple delay on the clock path, and affects clock skew. This is in contrast to an offset, which alters the setup and hold relationship between two clocks. Refer to [“Offset and Latency Example” on page 8-11](#) for an example of the different effects of offset and latency. When you turn on **Enable Clock Latency** in the Classic Timing Analyzer, it affects the following aspects of timing analysis:
 - **Early Clock Latency** and **Late Clock Latency** assignments are honored
 - The compensation delay of a PLL is analyzed as latency
 - For clock settings where you do not specify an offset, the automatically computed offset is treated as latency
- **Display Entity Name**—Any entity-specific assignments are ignored in the TimeQuest analyzer because they do not include the entity name when this option is set to ON.

If your design meets timing requirements in the Classic Timing Analyzer without all of the settings recommended in [Table 8-11 on page 8-43](#), you should perform one of the following actions:

- Change the settings and reconstrain and reverify as necessary.
or
- Add or modify SDC constraints as appropriate because analysis in the TimeQuest analyzer may be different after conversion.

Clock Conversion

Next, the conversion utility adds the `derive_pll_clocks` command to the `.sdc`. This command creates generated clocks on all PLL outputs in your design each time the `.sdc` is read. The command does not add a clock on the FPGA port that drives the PLL input.

The conversion utility includes the `derive_pll_clocks -use_net_name` command in the `.sdc` it creates. The `-use_net_name` option overrides the default clock naming behavior (the PLL pin name) so the clock name is the same as the net name in the Classic Timing Analyzer.

Including the `-use_net_name` option ensures that the conversion utility converts clock-to-clock exceptions properly. If you remove the `-use_net_name` option, you must also edit references to the clock names in other SDC commands so they match the PLL pin names.

If your design includes a global f_{MAX} assignment, the assignment is converted to a `derive_clocks` command. The behavior of a global f_{MAX} assignment is different from the behavior of clocks created with the `derive_clocks` command, and you should use the `report_clocks` command when you review conversion results to evaluate the clock settings. Refer to [“Automatic Clock Detection” on page 8-14](#) for an explanation of the differences. As soon as you know the appropriate clock settings, you should use the `create_clock` or `create_generated_clock` command instead of the `derive_clocks` command.



The conversion utility adds a `post_message` command before the `derive_clocks` command to remind you that the clocks are derived automatically. The TimeQuest analyzer displays the reminder the first time it reads the `.sdc`. Remove or comment out the `post_message` command to prevent the message from displaying.

Next, the conversion utility identifies and converts clock settings in the `.qsf`. If a project database exists, the utility also identifies and converts any additional clocks in the report file that are not in the `.qsf`, such as PLL base clocks.



If you change the PLL input frequency, you must modify the SDC constraint manually.

The conversion utility ignores clock offsets on generated clocks. Refer to [“Clock Offset” on page 8-10](#) for information about how to use offset values in the TimeQuest analyzer.

Instance Assignment Conversion

Next, the conversion utility converts the instance assignments shown in [Table 8-12](#). Refer to the indicated page for information about each assignment.

Table 8-12. Instance Timing Assignments

Assignment Name	QSF Variable	More Information
Late Clock Latency	LATE_CLOCK_LATENCY	page 8-25
Early Clock Latency	EARLY_CLOCK_LATENCY	
Clock Setup Uncertainty	CLOCK_SETUP_UNCERTAINTY	page 8-25
Clock Hold Uncertainty	CLOCK_HOLD_UNCERTAINTY	
Multicycle (1)	MULTICYCLE	page 8-27
Source Multicycle (2)	SRC_MULTICYCLE	
Multicycle Hold (3)	HOLD_MULTICYCLE	
Source Multicycle Hold	SRC_HOLD_MULTICYCLE	page 8-28
Clock Enable Multicycle	CLOCK_ENABLE_MULTICYCLE	
Clock Enable Source Multicycle	CLOCK_ENABLE_SOURCE_MULTICYCLE	
Clock Enable Multicycle Hold	CLOCK_ENABLE_MULTICYCLE_HOLD	
Clock Enable Source Multicycle Hold	CLOCK_ENABLE_SOURCE_MULTICYCLE_HOLD	page 8-40
Cut Timing Path	CUT	
Input Maximum Delay	INPUT_MAX_DELAY	page 8-29
Input Minimum Delay	INPUT_MIN_DELAY	
Output Maximum Delay	OUTPUT_MAX_DELAY	
Output Minimum Delay	OUTPUT_MIN_DELAY	

Notes to Table 8-12:

- (1) A multicycle assignment can also be known as a “destination multicycle setup” assignment.
- (2) A source multicycle assignment can also be known as a “source multicycle setup” assignment.
- (3) A multicycle hold can also be known as a “destination multicycle hold” assignment.

Depending on input and output delay assignments, you may receive a warning message when the `.sdc` is read. The message warns that the `set_input_delay` command, `set_output_delay` command, or both are missing the `-max` option, `-min` option, or both. Refer to [“Input and Output Delay” on page 8-29](#) for an explanation of why the warning occurs and how to avoid it.

The conversion utility automatically adds multicycle hold exceptions for each multicycle setup assignment. The value of each multicycle hold exception depends on the **Default hold multicycle** assignment value in your project. If the value is **One**, the conversion utility uses a value of 0 (zero) for each multicycle hold exception it adds. If the value is **Same as multicycle**, the conversion utility uses a value one less than the corresponding multicycle setup assignment value for each multicycle hold exception it adds. Refer to [“Hold Multicycle” on page 8-18](#) for more information on hold multicycle differences between the Classic Timing Analyzer and the TimeQuest analyzer.

Next, the conversion utility converts the instance assignments shown in [Table 8-13](#). Refer to the indicated page for information about each assignment. If the t_{PD} and minimum t_{PD} assignment targets also have input or output delays that apply to them, the t_{PD} and minimum t_{PD} conversion values may be incorrect. This is described in more detail on the indicated pages for the appropriate assignments.

Table 8-13. Instance Timing Assignments

Assignment Name	QSF Variable	More Information
t_{PD} Requirement (1)	TPD_REQUIREMENT	page 8-38
Minimum t_{PD} Requirement (1)	MIN_TPD_REQUIREMENT	page 8-40
Setup Relationship	SETUP_RELATIONSHIP	page 8-24
Hold Relationship	HOLD_RELATIONSHIP	page 8-25

Note to [Table 8-13](#):

- (1) Refer to " [\$t_{PD}\$ and Minimum \$t_{PD}\$ Requirement Conversion](#)" on [page 8-55](#) for more information about how the conversion utility converts single-point t_{PD} and minimum t_{PD} assignments.

The conversion utility converts Classic Timing Analyzer I/O timing assignments to FPGA-centric SDC constraints. [Table 8-14](#) includes Classic Timing Analyzer timing assignments, the equivalent FPGA-centric SDC constraints, and recommended system-centric SDC constraints.

Table 8-14. Classic Timing Analyzer and TimeQuest Analyzer Equivalent Constraints

Classic Timing Analyzer Assignment	FPGA-Centric SDC	System-Centric SDC	More Information
t_{SU} Requirement (1)	set_max_delay	set_input_delay -max	page 8-30
t_H Requirement (1)	set_min_delay	set_input_delay -min	page 8-32
t_{CO} Requirement (2)	set_max_delay	set_output_delay -max	page 8-34
Minimum t_{CO} Requirement (2)	set_min_delay	set_output_delay -min	page 8-36

Notes to [Table 8-14](#):

- (1) Refer to " [\$t_{PD}\$ and Minimum \$t_{PD}\$ Requirement Conversion](#)" on [page 8-55](#) for more information about how the conversion utility converts this type of assignment.
- (2) Refer to " [\$t_{CO}\$ Requirement Conversion](#)" on [page 8-49](#) for more information about how the conversion utility converts this type of assignment.

The conversion utility can convert Classic Timing Analyzer I/O timing assignments only to the FPGA-centric constraints without additional information about your design. Making system-centric constraints requires information about the external circuitry interfacing with the FPGA such as external clocks, clock latency, board delay, and clocking exceptions. You cannot convert Classic Timing Analyzer timing assignments to system-centric constraints without that information. If you use the conversion utility, you can modify the SDC constraints to change the FPGA-centric constraints to system-centric constraints as appropriate.

PLL Phase Shift Conversion

The conversion utility does not account for PLL phase shifts when it converts values of the following FPGA-centric I/O timing assignments:

- **t_{SU} Requirement**
- **t_H Requirement**
- **t_{CO} Requirement**
- **Minimum t_{CO} Requirement**

If any of your paths go through PLLs with a phase shift, you must correct the converted values for those paths according to the formula in [Equation 8-13](#):

Equation 8-13.

$$\langle \text{correct value} \rangle = \langle \text{converted value} \rangle - \frac{(\langle \text{pll output period} \rangle \times \langle \text{phase shift} \rangle)}{360}$$

[Example 8-18](#) shows the incorrect conversion result for a t_{CO} assignment and how to correct it. For the example, assume the PLL output frequency is 200 MHz (period is 5 ns), the phase shift is 90 degrees, the t_{CO} **Requirement** value is 1 ns, and it is made to data[0]. The .qsf contains the following assignment:

Example 8-18. Assignment

```
set_instance_assignment -name TCO_REQUIREMENT -to data[0] 1.0
```

The conversion utility generates the SDC command shown in [Example 8-19](#).

Example 8-19. SDC Command

```
set_max_delay -from [get_registers *] -to [get_ports data[0]] 1.0
```

To correct the value, use the formula and values above, as shown in the following equation:

$$1.0 - \frac{(5 \times 90)}{360} = -0.25$$

Then, change the value so the SDC command so that it looks like [Example 8-20](#).

Example 8-20. SDC Command with Correct Values

```
set_max_delay -from [get_registers *] -to [get_ports data[0]] -0.25
```

t_{CO} Requirement Conversion

The conversion utility uses a special process to convert **t_{CO} Requirement** and **Minimum t_{CO} Requirement** assignments. In addition to the `set_max_delay` or `set_min_delay` commands, the conversion utility adds a `set_output_delay` constraint relative to a virtual clock named N/C (Not a Clock). It also creates the virtual clock named N/C with a period of 10 ns. Adding the virtual clock allows you to report timing on the output paths. Without the virtual clock N/C, the clock used for reporting would be blank. [Example 8-21](#) shows how the conversion utility converts a **t_{CO} Requirement** assignment of 5.0 ns to `data[0]`.

Example 8-21. Converting a t_{CO} Requirement Assignment of 5.0 ns to data[0]

```
set_max_delay -from [get_registers *] -to [get_ports data[0]] 5.0
set_output_delay -clock "N/C" 0 [get_ports data[0]]
```

Entity-Specific Assignments

Next, the conversion utility converts the entity-specific assignments listed in [Table 8-15](#) that exist in the **Timing Analyzer Settings** report panel. This usually occurs if you have any timing assignments in your Verilog HDL or VHDL source, which can include MegaCore function files. These entity-specific assignments cannot be automatically converted unless your project is compiled and a `\db` directory exists.

Table 8-15. Entity-Specific Timing Assignments

Classic Timing Analyzer Assignment	QSF Variable	More Information
Multicycle	MULTICYCLE	page 8-27
Source Multicycle	SRC_MULTICYCLE	
Multicycle Hold	HOLD_MULTICYCLE	
Source Multicycle Hold	SRC_HOLD_MULTICYCLE	
Setup Relationship	SETUP_RELATIONSHIP	page 8-24
Hold Relationship	HOLD_RELATIONSHIP	page 8-25
Cut Timing Path	CUT	page 8-40

 You must manually convert all other entity-specific timing assignments.

Paths Between Unrelated Clock Domains

The conversion utility can create exceptions that cut paths between unrelated clock domains, which matches the default behavior of the Classic Timing Analyzer. When **Cut paths between unrelated clock domains** is turned on, the conversion utility creates clock groups with the `set_clock_groups` command and uses the `-exclusive` option to cut paths between the clock groups.

Unsupported Instance Assignments

Finally, the conversion utility checks for the unsupported instance assignments listed in [Table 8-16](#) and warns you if any exist. Refer to the indicated page for information about each assignment.



You can manually convert some of the assignments to SDC constraints.

Table 8-16. Instance Timing Assignments

Assignment Name	QSF Variable	More Information
Inverted Clock	INVERTED_CLOCK	page 8-25
Maximum Clock Arrival Skew	MAX_CLOCK_ARRIVAL_SKEW	page 8-41
Maximum Data Arrival Skew	MAX_DATA_ARRIVAL_SKEW	page 8-41
Maximum Delay	MAX_DELAY	page 8-40
Minimum Delay	MIN_DELAY	page 8-41
Virtual Clock Reference	VIRTUAL_CLOCK_REFERENCE	page 8-26

Reviewing Conversion Results

You must review the messages that are generated during the conversion process, and review the `.sdc` file for correctness and completeness. Warning and critical warning messages identify significant differences between the Classic Timing Analyzer and TimeQuest analyzer behaviors. In some cases, warning messages indicate that the conversion utility ignored assignments because it could not determine the intended functionality of your design. You must add to or modify the SDC constraints as necessary based on your knowledge of the design.

The conversion utility creates an `.sdc` with the same name as your current revision, `<revision>.sdc`, and it overwrites any existing `<revision>.sdc`. If you use the conversion utility to create an `.sdc`, you should make additions or corrections in a separate `.sdc`, or a copy of the `.sdc` created by the conversion utility. That way, you can rerun the conversion utility later without overwriting your additions and changes. If you have constraints in multiple `.sdc` files, refer to [“Constraint File Priority” on page 8-7](#) to learn how to add constraints to your project.

Warning Messages

The conversion utility may generate any of the following warning messages. Refer to the information provided for each warning message to learn what to do in that instance.

Ignored QSF Variable `<assignment>`

The conversion utility ignored the specified assignment. Determine whether an equivalent constraint is necessary and manually add one if appropriate. Refer to [“Timing Assignment Conversion” on page 8-24](#) for information about conversions for all QSF timing assignments.

Global <name> = <value>

The conversion utility ignored the global assignment <name>. Manually add an equivalent constraint if appropriate. Refer to [“Unsupported Global Assignments” on page 8-43](#) for information about conversions for these assignments.

QSF: Expected <name> to be set to <expected value> but it is set to <actual value>

The behavior of the TimeQuest analyzer is closest to the Classic Timing Analyzer when the value for the specified assignment is the expected value. Because the actual assignment value is not the expected value in your project, the TimeQuest analyzer analysis may be different from the Classic Timing Analyzer analysis. Refer to [“Recommended Global Assignments” on page 8-43](#) for an explanation about the indicated QSF variable names.

QSF: Found Global f_{MAX} Requirement. Translation will be done using derive_clocks

Your design includes a global f_{MAX} requirement, and the requirement is converted to the derive_clocks command. Refer to [“Default Required \$f_{MAX}\$ Assignment” on page 8-26](#) for information about how to convert to an SDC constraint.

TAN Report Database not found. HDL based assignments will not be migrated

You did not analyze your design with the Classic Timing Analyzer before running the conversion utility. As a result, the conversion utility did not convert any timing assignments in your HDL source code to SDC constraints. If you have timing assignments in your HDL source code, you must find and convert them manually, or analyze your design with the Classic Timing Analyzer and rerun the conversion utility.

Ignored Entity Assignment (Entity <entity>): <variable> = <value> -from <from> -to <to>

The conversion utility ignored the specified entity assignment because the utility cannot automatically convert the assignment. [Table 8-15 on page 8-49](#) lists the entity-specific assignments the script can convert automatically.

Refer to [“Timing Assignment Conversion” on page 8-24](#) for information about how to convert the entity assignment manually.

Ignoring OFFSET_FROM_BASE_CLOCK assignment for clock <clock>

In some cases, this assignment is used to work around a limitation in how the Classic Timing Analyzer handles some forms of clock inversion. The conversion script ignores the assignment because it cannot determine whether the assignment is used as a workaround. Review your clock setting and add the assignment in your .sdc if appropriate. Refer to [“Clock Offset” on page 8-10](#) for more information about converting clock offset.

Clock <clock> has no FMAX_REQUIREMENT - No clock was generated

The conversion utility did not convert the clock named <clock> because it has no f_{MAX} requirement. You should add a clock constraint with an appropriate period to your .sdc.

No Clock Settings defined in .qsf

If your .qsf has no clock settings, ignore this message. You must add clock constraints in your .sdc manually.

Clocks

Ensure that the conversion utility converted all clock assignments correctly. Run the `report_clocks` command, or double-click **Report Clocks** in the Tasks pane in the TimeQuest analyzer GUI. Make sure that the right number of clocks is reported. If any clock constraints are missing, you must add them manually with the appropriate SDC commands (`create_clock` or `create_generated_clock`). Confirm that each option for each clock is correct.

The TimeQuest analyzer can create more clocks, such as:

- `derive_clocks` selecting ripple clocks
- `derive_pll_clocks`, adding
 - Extra clocks for PLL switchover
 - Extra clocks for LVDS pulse-generated clocks (`~load_reg`)

Clock Transfers

After you confirm that all clock assignments are correct, run the `report_clock_transfers` command, or double-click **Report Clock Transfers** in the Tasks pane in the TimeQuest analyzer GUI. The command generates a summary table with the number of paths between each clock domain. If the number of cross-clock domain paths seems high, remember that all clock domains are related in the TimeQuest analyzer. You must cut unrelated clock domains. Refer to [“Related and Unrelated Clocks” on page 8–10](#) for information about how to cut unrelated clock domains.

Path Details

If you have unexpected differences between the Classic Timing Analyzer and the TimeQuest analyzer on some paths, follow these steps to identify the cause of the difference.

1. Report timing on the path in the TimeQuest analyzer.
2. Compare slack values.
3. Compare source and destination clocks.
4. Compare the launch/latch times in the TimeQuest analyzer to the setup/hold relationship in the Classic Timing Analyzer. The times are absolute in the TimeQuest analyzer and relative in the Classic Timing Analyzer.
5. Compare clock latency values.

Unconstrained Paths

Next, run the `report_ucp` command, or double-click **Report Unconstrained Paths** in the Tasks pane in the TimeQuest analyzer GUI. This command generates a series of reports that detail any unconstrained paths in your design. If your design was completely constrained in the Classic Timing Analyzer, but there are unconstrained paths in the TimeQuest analyzer, some assignments may not have been converted properly. Also, some of the assignments could be ambiguous. The TimeQuest analyzer analyzes more paths than the Classic Timing Analyzer does, so any unconstrained paths might be paths you could not constrain in the Classic Timing Analyzer.

Bus Names

If your design includes Classic Timing Analyzer timing assignments to buses, and the bus names do not include square brackets enclosing an asterisk, such as: `address[*]`, you should review the SDC constraints to ensure the conversion is correct. Refer to [“Bus Name Format” on page 8-6](#) for more information.

Other

Review the notes listed in [“Conversion Utility” on page 8-55](#).

Rerunning the Conversion Utility

You can force the conversion utility to run even if it can find an `.sdc` according to the priority described in [“Constraint File Priority” on page 8-7](#). Any method described in [“Create SDC Constraints from Existing Timing Assignments” on page 8-2](#) forces the conversion utility to run even if it can find an `.sdc`.

Notes

This section describes notes for the TimeQuest analyzer.

Output Pin Load Assignments

The TimeQuest analyzer takes **Output Pin Load** values into account when it analyzes your design. If you change **Output Pin Load** assignments and do not recompile before you analyze timing, you must use the `-force_dat` option when you create the timing netlist. Type the following command at the Tcl console of the TimeQuest analyzer:

```
create_timing_netlist -force_dat ←
```

If you change **Output Pin Load** assignments and recompile before you analyze timing, do not use the `-force_dat` option when you create the timing netlist. You can create the timing netlist with the `create_timing_netlist` command, or with the **Create Timing Netlist** task in the Tasks pane.

Also note that the SDC `set_output_load` command is not supported, so you must make all output load assignments in the `.qsf`.

Constraint Target Types

The TimeQuest analyzer supports mixed exception types; you can apply an exception of any clock/node combination.

DDR Constraints with the DDR Timing Wizard

The DDR Timing Wizard creates an **.sdc** that contains constraints for a DDR interface. You can use that **.sdc** with the TimeQuest analyzer to analyze only the DDR interface part of a design.

You should use the **.sdc** created by DDR Timing Wizard for constraining a DDR interface in the TimeQuest analyzer. Additionally, your **.qsf** should not contain timing assignments for the DDR interface if you plan to use the conversion utility to create an **.sdc**. You should run the conversion utility before you use DDR Timing Wizard, and you should choose not to apply assignments to the **.qsf**.

However, if you used DDR Timing Wizard and chose to apply assignments to a **.qsf**, before you used the conversion utility, you should remove the DDR Timing Wizard-generated QSF timing assignments and rerun the conversion utility. The conversion utility creates some incompatible SDC constraints from the DDR Timing Wizard QSF assignments.

Unsupported SDC Features

Some SDC commands and features are not supported by the current version of the TimeQuest analyzer, including the following commands and features:

- The `get_designs` command, because the Quartus II software supports a single design, so this command is not necessary
- True latch analysis with time-borrowing feature; it can, however, convert latches to negative-edge-triggered registers
- The case analysis feature
- Loads, clock transitions, input transitions, and other features

Constraint Passing and Optimization

The Quartus II software can read constraints generated by other EDA software, and write constraints to be read by other EDA software.

Other synthesis software can generate constraints that target the **.qsf**. If you change timing constraints in synthesis software after creating an **.sdc** for the TimeQuest analyzer, you must update the SDC constraints. You can use the conversion utility, or update the **.sdc** manually.

If you use physical synthesis with the TimeQuest analyzer, the design may have lower performance.

Clock Network Delay Reporting

The TimeQuest analyzer reports delay on the clock network by node-to-node segments; the Classic Timing Analyzer reports delay on the clock network as a single number, rather than node-to-node segments

Project Management

If you use the `project_open` Tcl command in the TimeQuest analyzer to open a project compiled with an earlier version of the Quartus II software, the TimeQuest analyzer overwrites the compilation results (`\db` folder) without warning. Opening a project any other way results in a warning, and you can choose not to open the project.

Conversion Utility

This section describes the notes for the QSF assignment to SDC constraint conversion utility.

t_{PD} and Minimum t_{PD} Requirement Conversion

The conversion utility treats the targets of single-point t_{PD} and minimum t_{PD} assignments as device outputs. It does not correctly convert targets of single-point t_{PD} and minimum t_{PD} assignments that are device inputs. The following QSF assignment applies to an a device input named `d_in`:

```
set_instance_assignment -name TPD_REQUIREMENT -to d_in "3 ns"
```

The conversion utility creates the following SDC command, regardless of whether `d_in` is a device input or device output:

```
set_max_delay "3 ns" -from [get_ports *] -to [get_ports d_in]
```

You must update any incorrect SDC constraints manually.

-  For more detailed information about the features and capabilities of the TimeQuest analyzer, refer to the *Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

Document Revision History

Table 8-17. Document Revision History

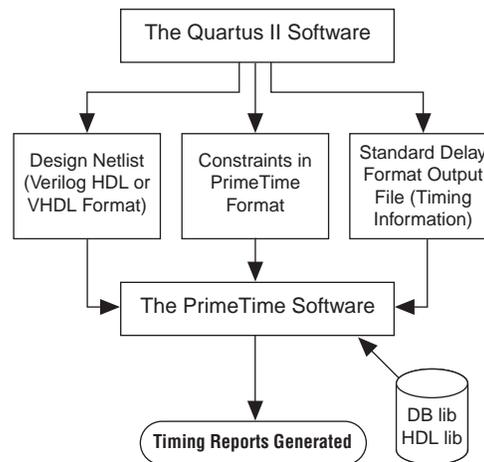
Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Changed to new document template. ■ Removed deprecated Classic Timing Analyzer features. ■ Minor updates to content.
July 2010	10.0.0	Minor updates to content.
November 2009	9.1.0	No change to content.
March 2009	9.0.0	This was chapter 8 in version 8.1.
November 2008	8.1.0	Changed to 8-1/2 x 11 page size. No change to content.

-  For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

-  Take an [online survey](#) to provide feedback about this handbook chapter.

PrimeTime is the Synopsys stand-alone full chip, gate-level static timing analyzer. The Quartus® II software makes it easy for designers to analyze their Quartus II projects using the PrimeTime software. The Quartus II software exports a netlist, design constraints (in the PrimeTime format), and libraries to the PrimeTime software environment. Figure 9–1 shows the PrimeTime flow diagram.

Figure 9–1. PrimeTime Software Flow Diagram



This chapter contains the following sections:

- “Quartus II Settings for Generating the PrimeTime Software Files”
- “Files Generated for the PrimeTime Software Environment” on page 9–2
- “Running the PrimeTime Software” on page 9–6
- “PrimeTime Timing Reports” on page 9–7
- “Static Timing Analyzer Differences” on page 9–18

Quartus II Settings for Generating the PrimeTime Software Files

To set up the Quartus II software to generate files for the PrimeTime software, perform the following steps:

1. In the Quartus II software, on the Assignments menu, click **Settings**, and then click **EDA Tool Settings**.
2. In the **Category** list, under **EDA Tool Settings**, select **Timing Analysis**.
3. In the **Tool name** list, select **PrimeTime**, and in the **Format for output netlist** list, select either **Verilog HDL** or **VHDL**.

© 2010 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, HARDCOPY, MAX, MEGACORE, NIOS, QUARTUS and STRATIX are Reg. U.S. Pat. & Tm. Off. and/or trademarks of Altera Corporation in the U.S. and other countries. All other trademarks and service marks are the property of their respective holders as described at www.altera.com/common/legal.html. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera’s standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

When you compile your project after making these settings, the Quartus II software runs the EDA Netlist Writer to create three files for the PrimeTime software. These files are saved in the `<revision_name>/timing/primetime` directory by default, where `<revision_name>` is the name of your Quartus II software revision. If it is not, you have used the wrong variable name.

Files Generated for the PrimeTime Software Environment

The Quartus II software generates a flattened netlist, a Standard Delay Output File (.sdo), and a Tcl script that prepares the PrimeTime software for timing analysis of the Quartus II project. These files are saved in the `<project_directory>/timing/primetime` directory.

The Quartus II software uses the EDA Netlist Writer to generate PrimeTime files based on either the Classic Timing Analyzer or the TimeQuest Timing Analyzer static timing analysis results. When you run the EDA Netlist Writer, the PrimeTime .sdo files are based on delays generated by the currently selected timing analysis tool in the Quartus II software.

To specify the timing analyzer, on the Assignments menu, click **Settings**. The **Settings** dialog box appears. Under **Category**, click **Timing Analysis Settings**. Select the timing analyzer of your choice.

 For more information about specifying the Quartus II timing analyzers, refer to either the *Quartus II Classic Timing Analyzer* or the *Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*. Also, refer to the *Switching to the Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook* to help you decide which timing analyzer is most appropriate for your design.

The Netlist

Depending on whether **Verilog HDL** or **VHDL** is selected as the **Format for output netlist** option, in the **Tool name** list on the **Timing Analysis** page of the **Settings** dialog box, the netlist is written and saved as either `<project name>.vo` or `<project name>.vho`, respectively. This file contains the flattened netlist representing the entire design.

 When you select the TimeQuest analyzer, only a Verilog HDL PrimeTime netlist can be generated.

The .sdo File

The Quartus II software saves the .sdo file as either `<revision_name>_v.sdo` or `<revision_name>_vhd.sdo`, depending on whether you select **Verilog HDL** or **VHDL** in the **Tool name** list on the **Timing Analysis** page of the **Settings** dialog box.

This file contains the timing information for each timing path between any two nodes in the design.

When you enable the Classic Timing Analyzer, the slow-corner (worst-case) timing models are used by default when generating the .sdo file. To generate the .sdo file using the fast-corner (best-case) timing models, perform the following steps:

1. In the Quartus II software, on the Processing menu, point to **Start** and click **Start Classic Timing Analyzer (Fast Timing Model)**.
2. After the fast-corner timing analysis is complete, on the Processing menu, point to **Start** and click **Start EDA Netlist Writer** to create a `<revision_name>_v_fast.sdo` or `<revision_name>_vhd_fast.sdo` file, which contains the best-case delay values for each timing path.

 If you are running a best-case timing analysis, the Quartus II software generates a Tcl script similar to the following: `<revision_name>_pt_v_fast.tcl`.

When the TimeQuest analyzer is run with the fast-corner netlist, or when the **Optimize fast-corner timing** check box is selected in the **Fitter Settings** dialog box, the fast-corner Synopsys Design Constraints File (.sdc) file is generated.

After the EDA Netlist Writer has finished, two .sdc files are created: `<revision_name>_v.sdo` (slow corner) and `<revision_name>_v_fast.sdo` (fast corner).

Generating Multiple Operating Conditions with the TimeQuest Analyzer

You can specify different operating conditions to the EDA Netlist Writer for PrimeTime analysis. The different operating conditions are reflected in the .sdo file generated by the EDA Netlist Writer.

 From the TimeQuest analyzer console pane, use the command `get_available_operating_conditions` to obtain a list of available operating conditions for the target device.

The following steps show how to generate the .sdo files for the three different operating conditions for a Stratix III design. Enter each command at the command prompt.

 The `--tq2pt` option for `quartus_sta` is required only if the project does not specify that the PrimeTime tool is to be used as the timing analysis tool.

1. Generate the first slow-corner model at the operating conditions: slow, 1100 mV, and 85° C.

```
quartus_sta --model=slow --voltage=1100 --temperature=85 <project name>
```
2. Generate the fast-corner model at the operating conditions: fast, 1100 mV, and 0° C.

```
quartus_sta --model=fast --voltage=1100 --temperature=0  
--tq2pt <project name>
```
3. Generate the PrimeTime output files for the corners specified above. The output files are generated in the **primetime_two_corner_files** directory.

```
quartus_eda --timing_analysis --tool=primetime  
--format=verilog  
--output_directory=primetime_two_corner_files  
--write_settings_files=off <project name>
```

4. Generate the second slow-corner model at the operating conditions: slow, 1100 mV, and 0° C.

```
quartus_sta --model=slow --voltage=1100 --temperature=0
--tq2pt <project name>
```

5. Generate the PrimeTime output files for the second slow corner. The output files are generated in the **primetime_one_slow_corner_files** directory.

```
quartus_eda --timing_analysis --tool=primetime
--format=verilog
--output_directory=primetime_one_slow_corner_files
--write_settings_files=off $revision
```

To summarize, the previous steps generate the following files for the three operating conditions:

- First slow corner (slow, 1100 mV, 85° C):
 .vo file—primetime_two_corner_files/<project name>.vo
 .sdo file—primetime_two_corner_files/<project name>_v.sdo
- Fast corner (fast, 1100 mV, 0° C):
 .vo file—primetime_two_corner_files/<project name>.vo
 .sdo file—primetime_two_corner_files/<project name>_v_fast.sdo
- Second slow corner (slow, 1100 mV, 0° C):
 .vo file—primetime_one_slow_corner_files/<project name>.vo
 .sdo file—primetime_one_slow_corner_files/<project name>_v.sdo



The **primetime_one_slow_corner_files** directory may also have files for fast corner. These files can be ignored because they were already generated in the **primetime_two_corner_files** directory.

The Tcl Script

The Tcl script generated by the Quartus II software contains information required by the PrimeTime software to analyze the timing and set up your post-fit design. This script specifies the search path and the names of the PrimeTime database library files provided with the Quartus II software. The `search_path` and `link_path` variables are defined at the beginning of the Tcl file. The `link_path` variable is a space-delimited list that contains the names of all database files used by the PrimeTime software.

Depending on whether you select **Verilog HDL** or **VHDL** in the **Format for output netlist** list on the **Timing Analysis** page of the **Settings** dialog box, when the Classic Timing Analyzer is enabled, the EDA Netlist Writer generates and saves the script as either `<revision_name>_pt_v.tcl` or `<revision_name>_pt_vhd.tcl`.

To access the **EDA Settings** dialog box, perform the following:

1. On the Assignments menu, click **Settings**, and then click **EDA Tool Settings**
2. Expand **EDA Tool Settings** under the **Category** list.

In the dialog box, you can specify VHDL or Verilog HDL for the format of the output netlist.

 The script also directs the PrimeTime software to use the `<device family>_all_pt.v` or `<device family>_all_pt.vhd` file, which contains the Verilog HDL or VHDL description of library cells for the targeted device family.

Example 9-1 shows the `search_path` and `link_path` variables defined in the Tcl script:

Example 9-1. Sample PrimeTime Setup Script

```
set quartus_root "altera/quartus/"
set search_path [list . [format "%s%s" $quartus_root "eda/synopsys/primetime/lib"] ]

set link_path [list * stratixii_lcell_comb_lib.db stratixii_lcell_ff_lib.db
stratixii_asynch_io_lib.db stratixii_io_register_lib.db stratixii_termination_lib.db
bb2_lib.db stratixii_ram_internal_lib.db stratixii_memory_register_lib.db
stratixii_memory_addr_register_lib.db stratixii_mac_out_internal_lib.db
stratixii_mac_mult_internal_lib.db stratixii_mac_register_lib.db
stratixii_lvds_receiver_lib.db stratixii_lvds_transmitter_lib.db
stratixii_asmiblock_lib.db stratixii_crcblock_lib.db stratixii_jtag_lib.db
stratixii_rublock_lib.db stratixii_pll_lib.db stratixii_dll_lib.db alt_vt1.db]

read_vhdl -vhdl_compiler stratixii_all_pt.vhd
```

The EDA Netlist Writer converts any Classic Timing Analyzer timing assignments to the PrimeTime software constraints and exceptions when it generates the PrimeTime files. The converted constraints are saved to the Tcl script. The Tcl script also includes a PrimeTime software command that reads the `.sdo` file generated by the Quartus II software. You can place additional commands in the Tcl script to analyze or report on timing paths.

Table 9-1 shows some examples of timing assignments converted by the Quartus II software for the PrimeTime software. For example, the `set_input_delay -max` command sets the input delay on an input pin.

Table 9-1. Equivalent Quartus II and PrimeTime Software Constraints

Quartus II Equivalent	PrimeTime Constraint
Clock defined on input pin, clock of 10 ns period and 50% duty cycle	<code>create_clock -period 10.000 -waveform {0 5.000} \ [get_ports clk] -name clk</code>
Input maximum delay of 1 ns on input pin, din	<code>set_input_delay -max -add_delay 1.000 -clock \ [get_clocks clk] [get_ports din]</code>
Input minimum delay of 1 ns on input pin, din	<code>set_input_delay -min -add_delay 1.000 -clock \ [get_clocks clk] [get_ports din]</code>
Output maximum delay of 3 ns on output pin, out	<code>set_output_delay -max -add_delay 3.000 -clock \ [get_clocks clk] [get_ports out]</code>

When the TimeQuest analyzer is turned on, the EDA Netlist Writer generates and saves the script as `<revision_name>.pt.tcl`.

The EDA Netlist Writer converts all TimeQuest analyzer `.sdc` constraints and exceptions into compatible PrimeTime software constraints and exceptions when it generates the PrimeTime files. The constraints and exceptions are saved to the `<revision_name>.constraints.sdc` file.

Generated File Summary

The files that are generated by the EDA Netlist Writer for the PrimeTime software depend on the Quartus II timing analysis tool you select.

Table 9-2 shows the files that are generated for the PrimeTime software when the Classic Timing Analyzer is selected.

Table 9-2. Classic Timing Analyzer-Generated PrimeTime Files

File	Description
<revision_name>.vho <revision_name>.vo	The PrimeTime software output netlist. Either a VHDL Output File (.vho) or a Verilog Output File (.vo) is generated, depending on the output netlist language set.
<revision_name>_vhd.sdo <revision_name>_v.sdo	The PrimeTime software standard delay file. Either a VHDL Standard Delay Output File (vhd.sdo) or a Verilog Standard Delay Output File (v.sdo) is generated, depending on the output netlist language set.
<revision_name>_pt_vhd.tcl <revision_name>_pt_v.tcl	PrimeTime setup and constraint script. Either a VHDL Tcl Script File (vhd.tcl) or a Verilog Tcl Script File (v.tcl) is generated, depending on the output netlist language set.

Table 9-3 shows the files that are generated for the PrimeTime software when the TimeQuest analyzer is selected. The EDA Netlist Writer supports the output netlist format only when the TimeQuest analyzer is enabled.

Table 9-3. TimeQuest Timing Analyzer-Generated PrimeTime Files

File	Description
<revision_name>.vo	The PrimeTime software output netlist. When the TimeQuest analyzer is enabled, only PrimeTime (Verilog HDL) is supported.
<revision_name>_v.sdo <revision_name>_v_fast.sdo	The PrimeTime software standard delay file. When the TimeQuest analyzer is enabled, only PrimeTime (Verilog HDL) is supported.
<revision_name>.pt.tcl	PrimeTime setup and constraint script. When the TimeQuest analyzer is enabled, only PrimeTime (Verilog HDL) is supported.
<revision_name>.collections.sdc	Contains the mapping from the TimeQuest analyzer netlist to the PrimeTime netlist.
<revision_name>.constraints.sdc	Contains the converted TimeQuest analyzer constraints for the PrimeTime software.

Running the PrimeTime Software

The PrimeTime software runs only on UNIX operating systems. If the Quartus II output files for the PrimeTime software were generated by running the Quartus II software on a PC/Windows-based system, follow these steps to run the PrimeTime software using Quartus II output files:

1. Install the PrimeTime libraries on a UNIX system by installing the Quartus II software on UNIX.

The PrimeTime libraries are located in the <Quartus II installation directory>/eda/synopsys/primetime/lib directory.

2. Copy the Quartus II output files to the appropriate UNIX directory. You may need to run a PC to UNIX program, such as dos2unix, to remove any control characters.

3. Modify the Quartus II path in Tcl scripts to point to the PrimeTime libraries using the first line of [Example 9-1](#):

```
set quartus_root "altera/quartus/" set search_path [list . [format  
"%s%s" $quartus_root "eda/synopsys/primetime/lib" ] ]
```

Analyzing Quartus II Projects

The PrimeTime software is controlled with Tcl scripts and can be run through `pt_shell`. You can run the `<revision_name>_pt_v.tcl` script file. For example, type the following at a UNIX system command prompt:

```
pt_shell -f <revision_name>_pt_v.tcl ←
```

When the TimeQuest analyzer is selected, type the following at a UNIX system command prompt:

```
pt_shell -f <revision_name>.pt.tcl ←
```

After all Tcl commands in the script are interpreted, the PrimeTime software returns control to the `pt_shell` prompt, which allows you to use other commands.

Other `pt_shell` Commands

You can run additional `pt_shell` commands at the `pt_shell` prompt, including the `man` program. For example, to read documentation about the `report_timing` command, type the following at the `pt_shell` prompt:

```
man report_timing ←
```

You can list all commands available in `pt_shell` by typing the following at the `pt_shell` prompt:

```
help ←
```

Type `quit` ← at the `pt_shell` prompt to close `pt_shell`.



You can also run `pt_shell` without a script file by typing `pt_shell` ← at the UNIX command line prompt.

PrimeTime Timing Reports

This section describes PrimeTime timing reports.

Sample PrimeTime Software Timing Report

After running the script, the PrimeTime software generates a timing report. If the timing constraints are not met, `Violated` is displayed at the end of the timing report. The timing report also gives the negative slack.

The PrimeTime software timing report is similar to the sample shown in [Example 9-2](#). The starting point in this report is a register clocked by clock signal, `clock`, and the endpoint is another register, `inst3-I.lereg`.

Example 9-2. Hold Path Report in PrimeTime

```

Startpoint: inst2~I.lereg
  (rising edge-triggered flip-flop clocked by clock)
Endpoint: inst3~I.lereg
  (rising edge-triggered flip-flop clocked by clock)
Path Group: clock
Path Type: min

```

Point	Incr	Path

clock clock (rise edge)	0.000	0.000
clock network delay (propagated)	3.166	3.166
inst2~I.lereg.clk (stratix_lcell_register)	0.000	3.166r
inst2~I.lereg.regout (stratix_lcell_register) <-	0.176*	3.342r
inst2~I.regout (stratix_lcell)	0.000*	3.342r
inst3~I.datac (stratix_lcell)	0.000*	3.342r
inst3~I.lereg.datac (stratix_lcell_register)	3.413*	6.755r
data arrival time		6.755
clock clock (rise edge)	0.000	0.000
clock network delay (propagated)	3.002	3.002
inst3~I.lereg.clk (stratix_lcell_register)		3.002r
library hold time	0.100*	3.102
data required time		3.102

data required time		3.102
data arrival time		-6.755

slack (MET)		3.653

Comparing Timing Reports from the Classic Timing Analyzer and the PrimeTime Software

Both the Classic Timing Analyzer and the TimeQuest analyzer generate a static timing analysis report for every successful design compilation. The timing report lists all of the timing paths in your design that were analyzed, and indicates whether these paths have met or violated their timing requirements. Violations are reported only if timing constraints were specified.

The TimeQuest analyzer and PrimeTime use an equivalent set of equations when reporting the static timing analysis results for a design. However, the Classic Timing Analyzer uses slightly different reporting equations when reporting the static timing analysis results for a design. This section describes the differences between the Classic Timing Analyzer and the PrimeTime software.

The timing report generated by the Classic Timing Analyzer differs from the report generated by the PrimeTime software. Both tools provide the same data, but the data is presented in different formats. The following sections show how the PrimeTime software reports the following slack values differently from the Classic Timing Analyzer report:

- [“Clock Setup Relationship and Slack” on page 9-9](#)
- [“Clock Hold Relationship and Slack” on page 9-12](#)
- [“Input Delay and Output Delay Relationships and Slack” on page 9-16](#)

Clock Setup Relationship and Slack

The Classic Timing Analyzer performs a setup check that ensures that the data launched by source registers is latched correctly at the destination registers. The Classic Timing Analyzer does this by determining the data arrival time and clock arrival time at the destination registers, and compares this data with the setup time delay of the destination register. Equation 9-1 expresses the inequality that is used for a setup check. The data arrival time includes the longest path from the clock to the source register, the clock-to-out micro delay of the source register, and the longest path from the source register to the destination register. The clock arrival time is the shortest delay from the clock to the destination register.

Equation 9-1.

$$\text{Clock Arrival} - \text{Data Arrival} \geq t_{su}$$

Slack is the margin by which a timing requirement is met or not met. Positive slack indicates the margin by which a requirement is met. Negative slack indicates the margin by which a requirement is not met. The Classic Timing Analyzer determines the clock setup slack, as shown in Equation 9-2:

Equation 9-2.

$$\text{Clock Setup Slack} = \text{Largest Register-to-Register Requirement} - \text{Longest Register-to-Register Delay}$$

 The longest register-to-register delay in the previous equation is equal to the register-to-register data delay.

Equation 9-3.

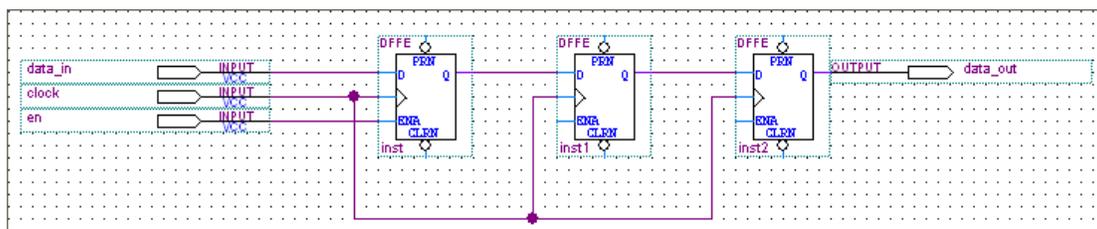
$$\begin{aligned} \text{Largest Register-to-Register Requirement} = \\ \text{Setup Relationship between Source and Destination} + \text{Largest Clock Skew} - \\ \text{Micro } t_{co} \text{ of Destination Register} - \text{Micro } t_{su} \text{ of Destination Register} \end{aligned}$$

$$\text{Setup Relationship between Source and Destination} = \text{Latch Edge} - \text{Launch Edge}$$

$$\text{Clock Skew} = \text{Shortest Clock Path to Destination} - \text{Longest Clock Path to Source}$$

Figure 9-2 shows a simple three-register design.

Figure 9-2. Simple Three-Register Design



The Classic Timing Analyzer generates a report for the design, as shown in Figure 9-3.

Figure 9-3. Timing Analyzer Report from Figure 9-2

	Slack	Actual fmax (period)	From	To	From Clock	To Clock	Required Setup Relationship	Required Longest P2P Time	Actual Longest P2P Time
1	4.237 ns	265.75 MHz (period = 3.763 ns)	inst2	inst3	clock	clock	8.000 ns	7.650 ns	3.413 ns
2	4.741 ns	306.84 MHz (period = 3.259 ns)	inst	inst2	clock	clock	8.000 ns	7.980 ns	3.239 ns

Equation 9-1, Equation 9-2, and Equation 9-3 are similar to those found in other static timing analysis tools, such as the PrimeTime software. Equation 9-4 through Equation 9-7, used by the PrimeTime software, are essentially the same as those used by the Classic Timing Analyzer, but they are rearranged.

Equation 9-4.

$$\text{Slack} = \text{Data Required} - \text{Data Arrival}$$

Equation 9-5.

$$\text{Clock Arrival} = \text{Latch Edge} + \text{Shortest Clock Path to Destination}$$

Equation 9-6.

$$\text{Data Required} = \text{Clock Arrival} - \text{Micro } t_{su}$$

Equation 9-7.

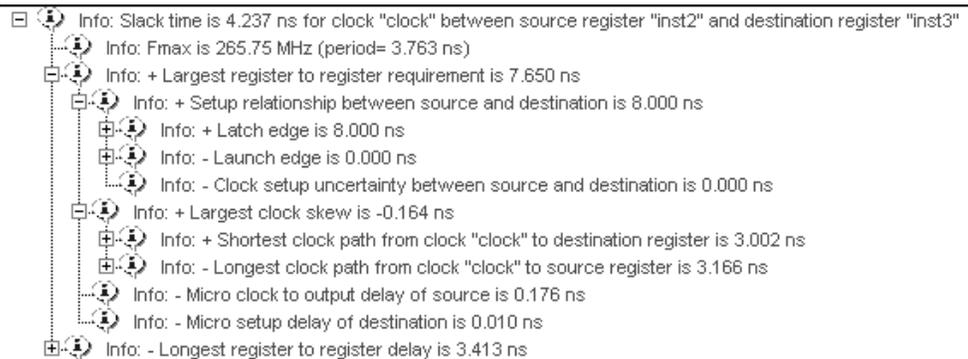
$$\text{Data Arrival} = \text{Launch Edge} + \text{Longest Clock Path to Source} + \text{Micro } t_{co} + \text{Longest Data Delay}$$



The longest data delay in the previous equation is equal to register-to-register data delay.

Figure 9-4 shows a clock setup check in the Quartus II software.

Figure 9-4. Clock Setup Check Reporting with the Classic Timing Analyzer



The results in Equation 9-8 are obtained by extracting the numbers from the Classic Timing Analyzer report and applying them to the clock setup slack equations from the Classic Timing Analyzer:

Equation 9-8.

$$\text{Setup Relationship between Source and Destination} = \text{Latch Edge} - \text{Launch Edge} - \text{Clock Setup Uncertainty}$$

$$8.0 - 0.0 - 0.0 = 8.0\text{ns}$$

$$\begin{aligned} \text{Clock Skew} &= \text{Shortest Clock Path to Destination} - \text{Longest Clock Path to Source} \\ 3.002 - 3.166 &= -0.164\text{ns} \end{aligned}$$

$$\begin{aligned} \text{Largest Register-to-Register Requirement} &= \\ \text{Setup Relationship between Source \& Destination} &+ \text{Largest Clock Skew} \\ -\text{Micro } t_{co} \text{ of Source Register} &- \text{Micro } t_{su} \text{ of Destination Register} \end{aligned}$$

$$8 + (-0.164) - 0.176 - 0.010 = 7.650\text{ns}$$

$$\text{Clock Setup Slack} = \text{Largest Register-to-Register Requirement} - \text{Longest Register-to-Register Delay}$$

$$7.650 - 3.413 = 4.237\text{ns}$$

For the same register-to-register path, the PrimeTime software generates a clock setup report as shown in [Example 9-3](#):

Example 9-3. Setup Path Report in PrimeTime

```

Startpoint: inst2~I.lereg
  (rising edge-triggered flip-flop clocked by clock)
Endpoint: inst3~I.lereg
  (rising edge-triggered flip-flop clocked by clock)
Path Group: clock
Path Type: max
Point          Incr      Path
-----
clock clock (rise edge)          0.000    0.000
clock network delay (propagated) 3.166    3.166
inst2~I.lereg.clk (stratix_lcell_register) 0.000    3.166r
inst2~I.lereg.regout (stratix_lcell_register) <- 0.176*   3.342r
inst2~I.regout (stratix_lcell) <- 0.000*   3.342r
inst3~I.dataac (stratix_lcell) <- 0.000*   3.342r
inst3~I.lereg.dataac (stratix_lcell_register) 3.413*   6.755r
data arrival time                    6.755
clock clock (rise edge)          8.000    8.000
clock network delay (propagated) 3.002    11.002
inst3~I.lereg.clk (stratix_lcell_register)      11.002r
library setup time                  -0.010*  10.992
data required time                   10.992
-----
data required time                    10.992
data arrival time                      -6.755
-----
slack (MET)                           4.237

```

Clock Hold Relationship and Slack

The Classic Timing Analyzer performs a hold time check along every register-to-register path in the design to ensure that no hold time violations have occurred. The hold time check verifies that data from the source register does not reach the destination until after the hold time of the destination register. The condition used for a hold check is shown in [Equation 9-9](#):

Equation 9-9.

$$\text{Data Arrival} - \text{Clock Arrival} \geq t_H$$

The Classic Timing Analyzer determines the clock hold slack with Equation 9-10, Equation 9-11, Equation 9-12, and Equation 9-13:

Equation 9-10.

$$\text{Clock Hold Slack} = \text{Shortest Register-to-Register Delay} - \text{Smallest Register-to-Register Requirement}$$

Equation 9-11.

$$\text{Smallest Register-to-Register Requirement} = \text{Hold Relationship between Source \& Destination} + \text{Smallest Clock Skew} - \text{Micro } t_{su} \text{ of Source} + \text{Micro } t_H \text{ of Destination}$$

Equation 9-12.

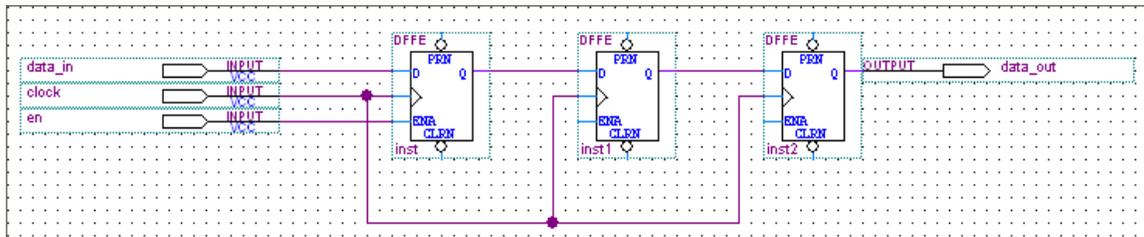
$$\text{Hold Relationship between Source \& Destination} = \text{Latch Edge} - \text{Launch Edge}$$

Equation 9-13.

$$\text{Smallest Clock Skew} = \text{Longest Clock Path from Clock to Destination Register} - \text{Shortest Clock Path from Clock to Source Register}$$

Figure 9-5 shows a simple three-register design.

Figure 9-5. Simple Three-Register Design



The Classic Timing Analyzer generates a report as shown in [Figure 9-6](#).

Figure 9-6. Timing Analyzer Report Generated from the Three-Register Design

	Minimum Slack	From	To	From Clock	To Clock	Required Hold Relationship	Required Shortest P2P Time	Actual Shortest P2P Time
1	3.149 ns	inst	inst2	clock	clock	0.000 ns	0.090 ns	3.239 ns
2	3.653 ns	inst2	inst3	clock	clock	0.000 ns	-0.240 ns	3.413 ns

The previous equations are similar to those found in the Quartus II software. [Equation 9-14](#) through [Equation 9-17](#) are the same equations that are used by the PrimeTime software, but they are rearranged.

Equation 9-14.

$$\text{Slack} = \text{Data Required} - \text{Data Arrival}$$

Equation 9-15.

$$\text{Clock Arrival} = \text{Latch Edge} + \text{Longest Clock Path to Destination}$$

Equation 9-16.

$$\text{Data Required} = \text{Clock Arrival} - \text{Micro } t_{\text{H}}$$

Equation 9-17.

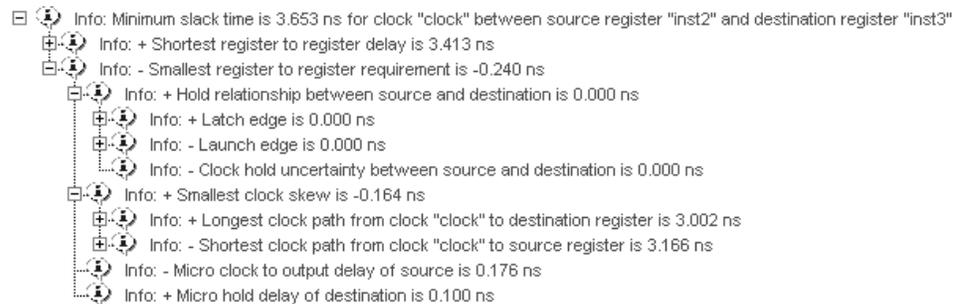
$$\text{Data Arrival} = \text{Launch Edge} + \text{Longest Clock Path to Source} + \text{Micro } t_{\text{co}} + \text{Shortest Data Delay}$$



The shortest register-to-register delay in the previous equation is equal to register-to-register data delay.

Figure 9-7 shows a clock setup check with the Classic Timing Analyzer.

Figure 9-7. Clock Hold Check Reporting with the Classic Timing Analyzer



The results in Equation 9-18 are obtained by extracting the numbers from the Timing Analysis report and applying the clock setup slack equations from the Classic Timing Analyzer.

Equation 9-18.

$$\text{Clock Hold Slack} = \text{Shortest Register-to-Register Delay} - \text{Smallest Register-to-Register Requirement}$$

$$3.413 - (-0.240) = 3.653\text{ns}$$

$$\text{Smallest Register-to-Register Requirement} = \text{Hold Relationship between Source \& Destination} + \text{Smallest Clock Skew} - \text{Micro } t_{co} \text{ of Source} + \text{Micro } t_H \text{ of Destination}$$

$$0 + (-0.164) - 0.176 + 0.100 = -0.240\text{ns}$$

$$\text{Hold Relationship between Source \& Destination} = \text{Latch} - \text{Launch}$$

$$0.0 - 0.0\text{ns}$$

$$\text{Smallest Clock Skew} = \text{Longest Clock Path from Clock to Destination Register} - \text{Shortest Clock Path from Clock to Source Register}$$

$$3.002 - 3.166 = -0.164\text{ns}$$

For the same register-to-register path, the PrimeTime software generates the report shown in [Example 9-4](#):

Example 9-4. Hold Path Report in PrimeTime

```

Startpoint: inst2~I.lereg
  (rising edge-triggered flip-flop clocked by clock)
Endpoint: inst3~I.lereg
  (rising edge-triggered flip-flop clocked by clock)
Path Group: clock
Path Type: min
Point          Incr      Path
-----
clock clock (rise edge)          0.000    0.000
clock network delay (propagated) 3.166    3.166
inst2~I.lereg.clk (stratix_lcell_register) 0.000    3.166r
inst2~I.lereg.regout (stratix_lcell_register) <- 0.176*   3.342r
inst2~I.regout (stratix_lcell)      0.000*   3.342r
inst3~I.dataac (stratix_lcell)      0.000*   3.342r
inst3~I.lereg.dataac (stratix_lcell_register) 3.413*   6.755r
data arrival time                    6.755

clock clock (rise edge)          0.000    0.000
clock network delay (propagated) 3.002    3.002
inst3~I.lereg.clk (stratix_lcell_register)      3.002r
library hold time                    0.100*   3.102
data required time                    3.102

-----
data required time                    3.102
data arrival time                    -6.755
-----
slack (MET)                          3.653

```

Both sets of hold slack equations can be used to determine the hold slack value of any path.

Input Delay and Output Delay Relationships and Slack

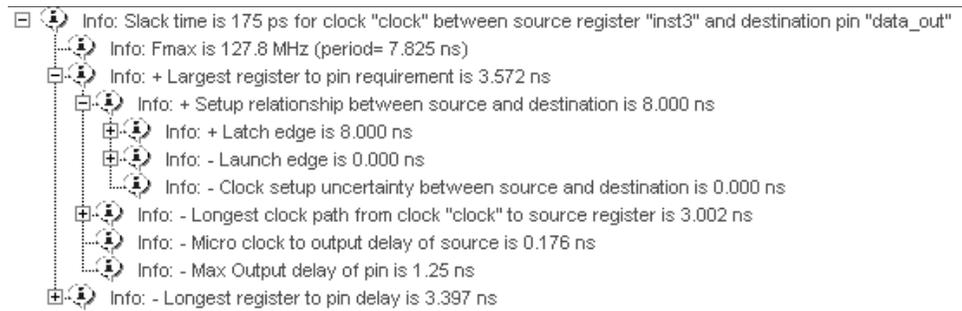
Input delay and output delay reports generated by the Classic Timing Analyzer are similar to the clock setup and clock hold relationship reports. [Figure 9-8](#) shows the input delay and output delay report for the design shown in [Figure 9-5](#) on page 9-13.

Figure 9-8. Input and Output Delay Reporting with the Classic Timing Analyzer

	Slack	Actual fmax (period)	From	To	From Clock	To Clock	Required Setup Relationship	Required Longest P2P Time	Actual Longest P2P Time
1	0.175 ns	127.80 MHz (period = 7.825 ns)	inst3	data_out	clock	clock	8.000 ns	3.572 ns	3.397 ns
2	5.380 ns	381.68 MHz (period = 2.620 ns)	data_in	inst	clock	clock	8.000 ns	8.490 ns	3.110 ns

Figure 9-9 shows the fully expanded view for the output delay path.

Figure 9-9. Output Delay Path Reporting with the Classic Timing Analyzer



For the same output delay path, the PrimeTime software generates a report similar to Example 9-5:

Example 9-5. Setup Path Report in PrimeTime

```
Startpoint: inst3~I.lereg
(rising edge-triggered flip-flop clocked by clock)
Endpoint: data_out
(output port clocked by clock)
Path Group: clock
Path Type: max
```

Point	Incr	Path
clock clock (rise edge)	0.000	0.000
clock network delay (propagated)	3.002	3.002
inst3~I.lereg.clk (stratix_lcell_register)	0.000	3.002r
inst3~I.lereg.regout (stratix_lcell_register)<-	0.176*	3.178r
inst3~I.regout (stratix_lcell)<-	0.000	3.178r
data_out~I.datain (stratix_io)<-	0.000	3.178r
data_out~I.out_mux3.A (mux21)<-	0.000	3.178r
data_out~I.out_mux3.MO (mux21)<-	0.000	3.178r
data_out~I.and2_22.IN1 (AND2)<-	0.000	3.178r
data_out~I.and2_22.Y (AND2)<-	0.000	3.178r
data_out~I.out_mux1.A (mux21)<-	0.000	3.178r
data_out~I.out_mux1.MO (mux21)<-	0.000	3.178r
data_out~I.inst1.datain (stratix_asynch_io)<-	0.902*	4.080r
data_out~I.inst1.padio (stratix_asynch_io)<-	2.495*	6.575r
data_out~I.padio (stratix_io)<-	0.000	6.575r
data_out (out)	0.000	6.575r
data arrival time		6.575
clock clock (rise edge)	8.000	8.000
clock network delay (propagated)	0.000	8.000
output external delay	1.250	6.750
data required time		6.750
data required time		6.750
data arrival time		6.575
slack (MET)		0.175

To generate a list of the 100 worst paths and place this data into a file called **file.timing**, type the following command at the `pt_shell` prompt:

```
report_timing -nworst 100 > file.timing ←
```

Timing paths in the PrimeTime software are listed in the order of most-negative-slack to most-positive-slack. The PrimeTime software does not categorize failing paths by default. Timing setup (tsu) and timing hold (th) times are not listed separately. In the PrimeTime software, each path is shown with a start and end point; for example, if it is a register-to-register or input-to-register type of path. If you only use the `report_timing` part of the command without adding a `-delay` option, only the setup-time-related timing paths are reported.

The following command is used to create a minimum timing report or a list of hold-time-related violations:

```
report_timing -delay_type min ←
```

Ensure that the correct `.sdo` file, either minimum or maximum delays, is loaded before running this command.

Static Timing Analyzer Differences

Under certain design conditions, several static timing analysis differences can exist between the Classic Timing Analyzer and the TimeQuest analyzer, and the PrimeTime software. The following sections explain the differences between the two static timing analysis engines and the PrimeTime software.

Classic Timing Analyzer and PrimeTime Software

The following section describes the differences between the Classic Timing Analyzer and the PrimeTime software.

Rise/Fall Support

The Classic Timing Analyzer does not support rise/fall analysis. However, rise/fall support is available in PrimeTime.

Minimum and Maximum Delays

The Classic Timing Analyzer calculates minimum and maximum delays for all device components with the exception of clock routing. PrimeTime does not model these delays. This can result in different slacks for a given path on average of 2 to 3%.

Recovery/Removal Analysis

The Classic Timing Analyzer performs a more pessimistic recovery/removal analysis for asynchronous paths than PrimeTime. This can result in different delays reported between the two tools.

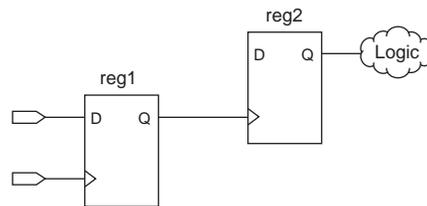
Encrypted Intellectual Property Blocks

The Quartus II software has the capability to decrypt all intellectual property (IP) blocks designed for Altera® devices that have been encrypted by their vendors. The decryption process allows the Quartus II software to perform a full compilation of the design that contains an encrypted IP block. This also allows the Classic Timing Analyzer to perform a complete static timing analysis on the design. However, licensed and encrypted IP blocks do not permit output netlists to be generated when using PrimeTime as the static timing analysis tool. (The EDA Netlist Writer does not generate `.vho` or `.vo` netlist files.)

Registered Clock Signals

Registered clock signals are clock signals that pass through a register before reaching the clock port of a sequential element. [Figure 9-10](#) shows an example of a registered clock signal.

Figure 9-10. Registered Clock Signal



If no clock setting is applied to the register on the clock path (shown as register `reg_1` in [Figure 9-10](#)), the Classic Timing Analyzer treats the register in the clock path as a buffer. The delay of the buffer is equal to the CELL delay of the register plus the t_{CO} of the register. The PrimeTime software does not treat the register as a buffer.

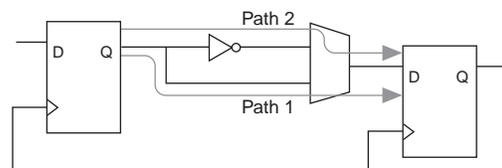


For more information about creating clock settings, refer to the [Quartus II Classic Timing Analyzer](#) chapter in volume 3 of the [Quartus II Handbook](#).

Multiple Source and Destination Register Pairs

In any design, multiple paths may exist from a source register to a destination register. Each path from the source register to the destination register may have a different delay value due to the different routes taken. For example, [Figure 9-11](#) shows a sample design that contains multiple path pairs between the source register and destination register.

Figure 9-11. Multiple Source and Destination Pairs



The Classic Timing Analyzer analyzes all source and destination pairs, but reports only the source and destination register pair with the worst slack. For example, if the Path 2 pair delay is greater than the Path 1 pair delay in [Figure 9-11](#), the Classic Timing Analyzer reports the slack value of the Path 2 pair and not the Path 1 pair. The PrimeTime software reports all possible source and destination register pairs.

Latches

By default, the Quartus II software implements all latches as combinational loops. The Classic Timing Analyzer can analyze such latches by treating them as registers with inverted clocks or analyze latches as a combinational loop modeled as a combinational delay.

 For more information about latch analysis, refer to the *Quartus II Classic Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

The PrimeTime software always analyzes these latches as combinational loops, as defined in the netlist file.

LVDS I/O

When analyzing the dedicated LVDS transceivers in your design, the Classic Timing Analyzer generates the Receiver Skew Margin (RSKM) report and a Channel-to-Channel Skew (TCCS) report. The PrimeTime software does not generate these reports.

Clock Latency

When a single clock signal feeds both the source and destination registers of a register-to-register path, and either an Early Clock Latency or a Late Clock Latency assignment has been applied to the clock signal, the Classic Timing Analyzer does not factor in the clock latency values when it calculates the clock skew between the two registers. The Classic Timing Analyzer factors in the clock latency values when the clock signal to the source and destination registers of a register-to-register path are different. The PrimeTime software applies the clock latency values when a single clock signal or different clock signals feeds the source and destination registers of a register-to-register path.

Input and Output Delay Assignments

When a purely combinational (non-registered) path exists between an input pin and output pin of the Altera FPGA and both pins have been constrained with an input delay and an output delay assignment applied, respectively, the Classic Timing Analyzer does not perform a clock setup or clock hold analysis. The PrimeTime software analyzes these paths.

Generated Clocks Derived from Generated Clocks

The Classic Timing Analyzer does not support a generated clock derived from a generated clock. This situation might occur if a generated clock feeds the input clock pin of a PLL. The output clock of the PLL is a generated clock.

TimeQuest Timing Analyzer and PrimeTime Software

The following sections describe the static timing analysis differences between the TimeQuest analyzer and the PrimeTime software.

Encrypted Intellectual Property Blocks

The Quartus II software has the capability to decrypt all IP blocks, designed for Altera devices that have been encrypted by their vendors. The decryption process allows the Quartus II software to perform a full compilation on the design containing an encrypted IP block. This also allows the TimeQuest analyzer to perform a complete static timing analysis on the design. However, licensed and encrypted IP blocks do not permit output netlists to be generated when using PrimeTime as the static timing analysis tool. (The EDA Netlist Writer does not generate `.who` or `.vo` netlist files.)

Latches

By default, the Quartus II software implements all latches as combinational loops. The TimeQuest analyzer can analyze such latches by treating them as registers with inverted clocks. The TimeQuest analyzer analyzes latches as a combinational loop modeled as a combinational delay.



For more information about latch analysis, refer to the *Quartus II TimeQuest Timing Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

The PrimeTime software always analyzes these latches as combinational loops, as defined in the netlist file.

LVDS I/O

When analyzing the dedicated LVDS transceivers in your design, the TimeQuest analyzer generates a Receiver Skew Margin (RSKM) report and a Channel-to-Channel Skew (TCCS) report. The PrimeTime software does not generate these reports.

The TimeQuest Timing Analyzer .sdc File and PrimeTime Compatibility

Because of differences between node naming conventions with the netlist generated by the EDA Netlist Writer and the internal netlist used by the Quartus II software, .sdc files generated for the Quartus II software or the TimeQuest analyzer are not compatible with the PrimeTime software.

Run the EDA Netlist Writer to generate a compatible .sdc file from the TimeQuest .sdc file for the PrimeTime software. After the files *<revision_name>.collections.sdc* and *<revision_name>.constraints.sdc* have been generated, both files can be read by the PrimeTime software for compatibility of constraints between the TimeQuest analyzer and the PrimeTime software.

Clock and Data Paths

If a timing path acts both as a clock path (a path that connects to a clock pin with a clock associated with it), and a data path (a path that feeds into the data-in port of a register), the TimeQuest analyzer reports the data paths, whereas PrimeTime does not.

Inverting and Non-Inverting Propagation

The TimeQuest analyzer always propagates non-inverting sense for clocks through non-unate paths in the clock network.

PrimeTime's default behavior is to propagate both inverting and non-inverting senses through a non-unate path in the clock network.

Multiple Rise/Fall Numbers For a Timing Arc

For a given timing path with a corresponding set of pins/ports that make up the path (including source and destination pair), if the individual components of that path have different rise/fall delays, there can potentially be many timing paths with different delays using the same set of pins. If this occurs, the TimeQuest analyzer reports only one timing path for the set of pins that make up the path.

Virtual Generated Clocks

PrimeTime does not support virtual generated clocks. To maintain compatibility between the TimeQuest analyzer and PrimeTime, all generated clocks should have an explicit target specified.

Generated Clocks Derived from Generated Clocks

The Classic Timing Analyzer does not support the creation of a generated clock derived from a generated clock. This situation might occur if a generated clock feeds the input clock pin of another generated clock. The output clock of the PLL is a generated clock.

Conclusion

The Quartus II software can export a netlist, constraints, and timing information for use with the PrimeTime software. The PrimeTime software can use data from either best-case or worst-case Quartus II timing models to measure timing. The PrimeTime software is controlled using a Tcl script generated by the Quartus II software that you can customize to direct the PrimeTime software to produce violation and slack reports.

Document Revision History

Table 9-4 shows the revision history for this chapter.

Table 9-4. Document Revision History

Date	Version	Changes Made
December 2010	10.0.1	Changed to new document template.
July 2010	10.0.0	■ Minor corrections throughout, and Quartus II interface changes.
November 2009	9.1.0	■ Updated “Setting the Quartus II Software to Generate the PrimeTime Software Files” figure for changes in the Quartus II software version 9.1
March 2009	9.0.0	■ This was chapter 10 in version 8.1. ■ Updated for the Quartus II software version 9.0 release.
November 2008	8.1.0	■ Changed to 8-1/2 x 11 page size. No change to content.
May 2008	8.0.0	■ Updated to Quartus II software version 8.0 and date. ■ Added hyperlinks to referenced Altera documentation throughout the chapter.

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

As FPGA designs grow larger and processes continue to shrink, power is an ever-increasing concern. When designing a PCB, the power consumed by a device must be accurately estimated to develop an appropriate power budget, and to design the power supplies, voltage regulators, heat sink, and cooling system.

The Quartus® II software allows you to estimate the power consumed by your current design during timing simulation. The power consumption of your design can be calculated using the Microsoft Excel-based power calculator, or the Simulation-Based Power Estimation features in the Quartus II software. This section explains how to use both.

This section includes the following chapter:

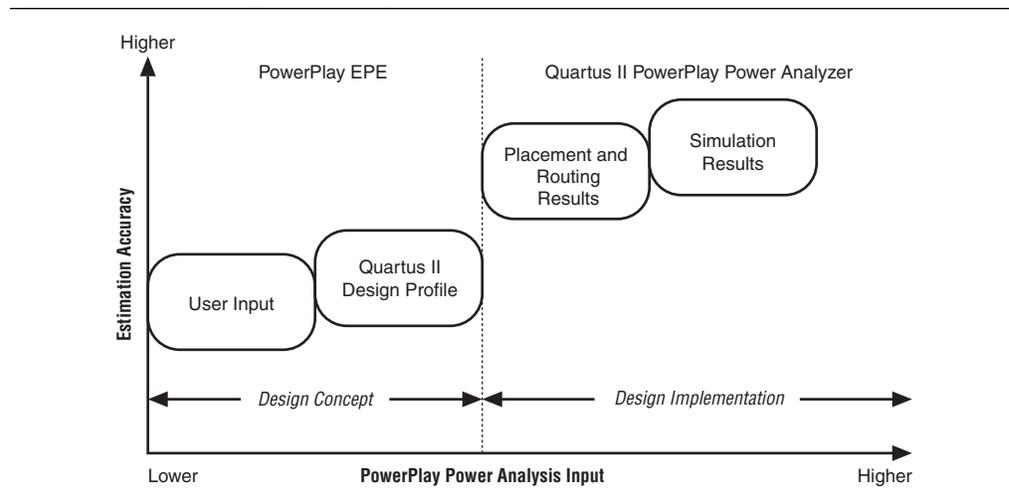
- **Chapter 10, PowerPlay Power Analysis**

This chapter describes the Altera® Quartus II PowerPlay power analysis tool and how to use the tools to accurately estimate device power consumption.

This chapter describes how to use the Altera® Quartus® II PowerPlay Power Analysis tools to accurately estimate device power consumption.

As designs grow larger and process technology continues to shrink, power becomes an increasingly important design consideration. When designing a PCB, the power consumed by a device must be accurately estimated to develop an appropriate power budget and to design the power supplies, voltage regulators, heat sink, and cooling system. As shown in Figure 10–1, the PowerPlay Power Analysis tools provide the ability to estimate power consumption from early design concept through design implementation.

Figure 10–1. PowerPlay Power Analysis



For more information about the PowerPlay suite of power analysis and optimizations tools, refer to *About Power Estimation and Analysis* in Quartus II Help. For more information about acquiring the PowerPlay EPE spreadsheet, refer to *PowerPlay Early Power Estimators (EPE) and Power Analyzer* on the Altera website.

This chapter discusses the following topics:

- “Types of Power Analyses” on page 10–2
- “Factors Affecting Power Consumption” on page 10–2
- “Creating PowerPlay EPE Spreadsheets” on page 10–5
- “PowerPlay Power Analyzer Flow” on page 10–7
- “Using Simulation Files in Modular Design Flows” on page 10–10
- “Using the PowerPlay Power Analyzer” on page 10–16
- “Conclusion” on page 10–24

© 2010 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, HARDCOPY, MAX, MEGACORE, NIOS, QUARTUS and STRATIX are Reg. U.S. Pat. & Tm. Off. and/or trademarks of Altera Corporation in the U.S. and other countries. All other trademarks and service marks are the property of their respective holders as described at www.altera.com/common/legal.html. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera’s standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

Types of Power Analyses

Understanding the uses of power analysis and the factors affecting power consumption helps you to use the PowerPlay Power Analyzer effectively. Power analysis meets two significant planning requirements:

- **Thermal planning**—The cooling solution must be sufficient to dissipate the heat generated by the device. The computed junction temperature must fall within normal device specifications.
- **Power supply planning**—Power supplies must provide adequate current to support device operation.

The two types of analyses are closely related because much of the power supplied to the device is dissipated as heat from the device; however, in some situations, the two types of analyses are not identical. For example, if you are using terminated I/O standards, some of the power drawn from the power supply of the device dissipates in termination resistors rather than in the device.

Power analysis also addresses the activity of your design over time as a factor that impacts the power consumption of the device. Static power is the power consumed regardless of design activity. Dynamic power is the additional power consumed due to signal activity or toggling.



For power supply planning, you can use the PowerPlay EPE at the early stages of your design cycle, or use the PowerPlay Power Analyzer reports when your design is complete to get an estimate of your design power requirement.

Factors Affecting Power Consumption

This section describes the factors affecting power consumption. Understanding these factors allows you to use the PowerPlay Power Analyzer and interpret its results effectively.

Device Selection

Different device families have different power characteristics. Many parameters affect the device family power consumption, including choice of process technology, supply voltage, electrical design, and device architecture. For example, the Cyclone II device family architecture consumes less static power than the high-performance and full-featured Stratix II device family.

Power consumption also varies in a single device family. A larger device consumes more static power than a smaller device in the same family because of its larger transistor count. Dynamic power can also increase with device size in devices that employ global routing architectures, for example, the MAX device family. Cyclone, MAX II, and Stratix devices do not exhibit significantly increased dynamic power as device size increases.

The choice of device package also affects the ability of the device to dissipate heat. This choice can impact your cooling solution choice required to meet junction temperature constraints.

Process variation can affect power consumption. Process variation primarily impacts static power because sub-threshold leakage current varies exponentially with changes in transistor threshold voltage. As a result, it is critical to consult device specifications for static power and not rely on empirical observation. Process variation has a weak effect on dynamic power.

Environmental Conditions

Operating temperature primarily affects device static power consumption. Higher junction temperatures result in higher static power consumption. The device thermal power and cooling solution that you use must result in the device junction temperature remaining within the maximum operating range for the device. The main environmental parameters affecting junction temperature are the cooling solution and ambient temperature.

Airflow

Airflow is a measure of how quickly heated air is removed from the vicinity of the device and replaced by air at ambient temperature. Airflow can either be specified as “still air” when no fan is used, or as the linear feet per minute rating of the fan used in the system. Higher airflow decreases thermal resistance.

Heat Sink and Thermal Compound

A heat sink allows more efficient heat transfer from the device to the surrounding area because of its large surface area exposed to the air. The thermal compound that interfaces the heat sink to the device also influences the rate of heat dissipation. The case-to-ambient thermal resistance (θ_{CA}) parameter describes the cooling capacity of the heat sink and thermal compound employed at a given airflow. Larger heat sinks and more effective thermal compounds reduce θ_{CA} .

Junction Temperature

The junction temperature of a device is equal to:

$$T_{Junction} = T_{Ambient} + P_{Thermal} \cdot \theta_{JA}$$

in which θ_{JA} is the total thermal resistance from the device transistors to the environment, having units of degrees Celsius per watt. The value θ_{JA} is equal to the sum of the junction-to-case (package) thermal resistance (θ_{JC}) and the case-to-ambient thermal resistance (θ_{CA}) of your cooling solution.

Board Thermal Model

The thermal resistance of the path through the board is referred to as the junction-to-board thermal resistance (θ_{JB}), having units of degrees Celsius per watt. It is used in conjunction with the board temperature, as well as the top-of-chip θ_{JA} and ambient temperatures, to compute junction temperature.

Device Resource Usage

The number and types of device resources used greatly affects power consumption.

Number, Type, and Loading of I/O Pins

Output pins drive off-chip components, resulting in high-load capacitance that leads to a high-dynamic power per transition. Terminated I/O standards require external resistors that generally draw constant (static) power from the output pin.

Number and Type of Logic Elements, Multiplier Elements, and RAM Blocks

A design with more logic elements (LEs), multiplier elements, and memory blocks tends to consume more power than a design with fewer circuit elements. The operating mode of each circuit element also affects its power consumption. For example, a DSP block performing 18×18 multiplications and a DSP block performing multiply-accumulate operations consume different amounts of dynamic power because of different amounts of internal capacitance being charged on each transition. The operating mode of a circuit element also affects static power.

Number and Type of Global Signals

Global signal networks span large portions of the device and have high capacitance, resulting in significant dynamic power consumption. The type of global signal is important as well. For example, Stratix II devices support several kinds of global clock networks that span either the entire device or a specific portion of the device (a regional clock network covers a quarter of the device). Clock networks that span smaller regions have lower capacitance and tend to consume less power. The location of the logic array blocks (LABs) driven by the clock network can also have an impact because the Quartus II software automatically disables unused branches of a clock.

Signal Activities

The final important factor in estimating power consumption is the behavior of each signal in your design. The two vital statistics are the toggle rate and the static probability.

The toggle rate of a signal is the average number of times that the signal changes value per unit of time. The units for toggle rate are transitions per second and a transition is a change from 1 to 0, or 0 to 1.

The static probability of a signal is the fraction of time that the signal is logic 1 during the period of device operation that is being analyzed. Static probability ranges from 0 (always at ground) to 1 (always at logic-high).

Dynamic power increases linearly with the toggle rate as the capacitive load is charged more frequently for logic and routing. The Quartus II software models full rail-to-rail switching. For high toggle rates, especially on circuit output I/O pins, the circuit can transition before fully charging the downstream capacitance. The result is a slightly conservative prediction of power by the PowerPlay Power Analyzer.

The static power consumed by both routing and logic can sometimes be affected by the static probabilities of their input signals. This effect is due to state-dependent leakage and has a larger effect on smaller process geometries. The Quartus II software models this effect on devices at 90 nm (or smaller) if it is important to the power estimate. The static power also varies with the static probability of a logic 1 or 0 on the I/O pin when output I/O standards drive termination resistors.

 To get accurate results from the power analysis, the signal activities for analysis must represent the actual operating behavior of your design. Inaccurate signal toggle rate data is the largest source of power estimation error.

Creating PowerPlay EPE Spreadsheets

You can use PowerPlay EPE spreadsheets to perform a preliminary thermal analysis and power consumption estimate for your design. You can either enter the data manually or use the tools in the Quartus II software to assist you with generating the device resources usage information for your design.

 For more information about generating a PowerPlay EPE File in the Quartus II software, refer to *Performing an Early Power Estimate Using the PowerPlay Early Power Estimator* in Quartus II Help.

Figure 10-2 shows an example of the contents of a PowerPlay EPE File generated for a design that targets a Stratix III device.

Figure 10-2. Example of a PowerPlay EPE File

	A	B	C	D	E	F	G	H
1	EARLY_POWER_ESTIMATOR_FILE_FORMAT_VERSION		6					
2	QUARTUS_II_VERSION	9.0 Build 235 06/17/2009 SP 2 SJ Full Version						
3	PROJECT	fractal						
4	REVISION	fractal_extra						
5	PROJECT_FILE	C:/Powe_Lab_and_Test_designs/fractal.qpf						
6	TIME	Fri Aug 14 11:45:17 2009						
7	TIME_SECONDS	1250275517						
8	FAMILY	Stratix III						
9	DEVICE	EP3SE260						
10	PACKAGE	FBGA						
11	PART	EP3SE260H780C4						
12	POWER_USE_DEVICE_CHARACTERISTICS	TYPICAL						
13	POWER_AUTO_COMPUTE_TJ	ON						
14	POWER_TJ_VALUE	25						
15	POWER_USE_CUSTOM_COOLING_SOLUTION	OFF						
16	MIN_JUNCTION_TEMPERATURE	0						
17	MAX_JUNCTION_TEMPERATURE	85						
18	POWER_PRESET_COOLING_SOLUTION	23 mm heat sink with 200 Lfpm airflow						
19	POWER_BOARD_THERMAL_MODEL	None (Conservative)						
20	POWER_USE_TA_VALUE	25						
21	POWER_BOARD_TEMPERATURE	-1						
22	POWER_OJC_VALUE	0.1						
23	POWER_OCS_VALUE	0.1						
24	POWER_OSA_VALUE	1.8						
25	POWER_OJB_VALUE	-1						
26	VCCIO	1A		2.5 1B		0 1C		2.5 2C
27	VCCPD	1A		2.5 1B		0 1C		2.5 2C
28	RAIL_VOLTAGES	VCC		1.1 VCCPT		2.5 VCCA_PLL		2.5 VC
29	HIGH_SPEED	NUM_HIGH_SPEED_M9K_block_TILES		35 NUM_M9K_block_TILES_USED		35		
30								
31								
32	BLOCK	M9K block	count	16 ram_mode		Simple Dual Port	ram_read_durinnew	rar
33	BLOCK	M9K block	count	3 ram_mode		Simple Dual Port	ram_read_durinnew	rar
34	BLOCK	M9K block	count	16 ram_mode		Simple Dual Port	ram_read_durinnew	rar
35	BLOCK	Combinational cell	count	28 avg_toggle_rate		181153.408	avg_toggle_rate	0 av
36	BLOCK	Combinational cell	count	2967 avg_toggle_rate		7292423.461	avg_toggle_rate	0.147695 av
37	BLOCK	Combinational cell	count	47 avg_toggle_rate		639906.5769	avg_toggle_rate	0.044742 av
38	BLOCK	Clock enable block	count	1 avg_toggle_rate		0	avg_toggle_rate	0 av
39	BLOCK	Clock enable block	count	1 avg_toggle_rate		0	avg_toggle_rate	0 av
40	BLOCK	Clock enable block	count	1 avg_toggle_rate		28500000	avg_toggle_rate	1.993007 av
41	BLOCK	Clock enable block	count	1 avg_toggle_rate		28500000	avg_toggle_rate	2 av
42	BLOCK	Clock enable block	count	1 avg_toggle_rate		98750000	avg_toggle_rate	2 av
43	BLOCK	Register cell	count	2567 avg_toggle_rate		3790981.691	avg_toggle_rate	0.076779 av
44	BLOCK	Register cell	count	71 avg_toggle_rate		249731.7324	avg_toggle_rate	0.017464 av
45	BLOCK	MLAB cell	count	1 mlab_width		8	mlab_depth	4 av
46	BLOCK	I/O pad	count	16 avg_toggle_rate		442715.8125	avg_toggle_rate	0 av
47	BLOCK	I/O pad	count	26 avg_toggle_rate		982211.5385	avg_toggle_rate	0 av
48	BLOCK	I/O pad	count	1 avg_toggle_rate		100850000	avg_toggle_rate	2.042532 av
49	BLOCK	I/O pad	count	4 avg_toggle_rate		12500	avg_toggle_rate	0.000253 av

The PowerPlay EPE spreadsheet includes the Import Data macro that parses the information in the PowerPlay EPE File and transfers it into the spreadsheet. If you do not want to use the macro, you can manually transfer the data into the PowerPlay EPE spreadsheet.

For example, after importing the PowerPlay EPE File information into the PowerPlay EPE spreadsheet, you can add additional device resource information at any time. If the existing Quartus II project represents only a portion of your full design, you must enter the additional device resources used in the final design manually.

PowerPlay EPE File Generator Compilation Report

After successfully generating the PowerPlay EPE File, you can locate a PowerPlay EPE File Generator report under the **Compilation Report** section. This report contains different sections, such as Summary, Settings, Generated Files, Confidence Metric Details, and Signal Activities. For more information about the PowerPlay EPE File Generator report, refer to “PowerPlay Power Analyzer Compilation Report” on page 10-20.

Table 10-1 lists the main differences between the PowerPlay EPE and the Quartus II PowerPlay Power Analyzer.

Table 10-1. Comparison of the PowerPlay EPE and Quartus II PowerPlay Power Analyzer

Characteristic	PowerPlay EPE	Quartus II PowerPlay Power Analyzer
Phase in the design cycle	Any time	Post-fit
Tool requirements	Spreadsheet program or the Quartus II software	The Quartus II software
Accuracy	Medium	Medium to very high
Data inputs	<ul style="list-style-type: none"> ■ Resource usage estimates ■ Clock requirements ■ Environmental conditions ■ Toggle rate 	<ul style="list-style-type: none"> ■ Post-fit design ■ Clock requirements ■ Signal activity defaults ■ Environmental conditions ■ Register transfer level (RTL) simulation results (optional) ■ Post-fit simulation results (optional) ■ Signal activities per node or entity (optional)
Data outputs (1)	<ul style="list-style-type: none"> ■ Total thermal power dissipation ■ Thermal static power ■ Thermal dynamic power ■ Off-chip power dissipation ■ Current drawn from voltage supplies 	<ul style="list-style-type: none"> ■ Total thermal power ■ Thermal static power ■ Thermal dynamic power ■ Thermal I/O power ■ Thermal power by design hierarchy ■ Thermal power by block type ■ Thermal power dissipation by clock domain ■ Off-chip (non-thermal) power dissipation ■ Device supply currents

Notes to Table 10-1:

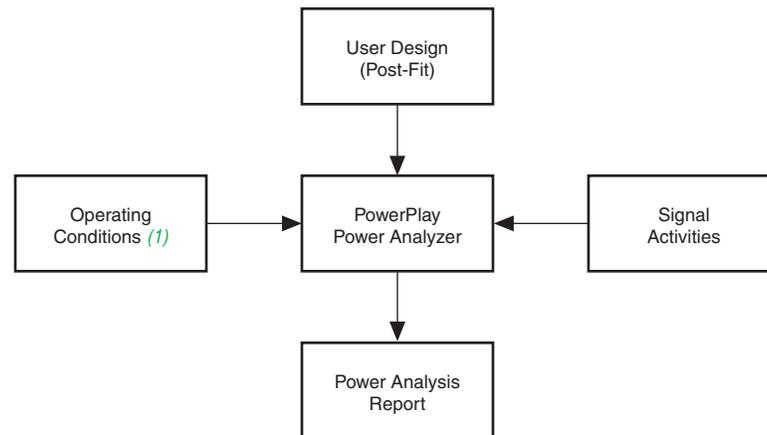
(1) PowerPlay EPE and PowerPlay Power Analyzer outputs vary by device family. For more information, refer to the [device-specific EPE User Guide](#) and [PowerPlay Power Analyzer Reports](#) in Quartus II Help.

The result of the PowerPlay Power Analyzer is only an estimation of power. Altera does not recommend using the result as a specification. The purpose of the estimation is to help you to establish a guide for the power budget of your design. Altera recommends measuring the actual power on the board. You must measure the total dynamic current of your design during device operation because the estimate is design dependent and depends on many variable factors, including input vector quantity, quality, and exact loading conditions of a PCB design. Static power consumption must not be based on empirical observation. The values reported by the PowerPlay Power Analyzer or data sheet must be used because the tested devices might not exhibit worst-case behavior.

PowerPlay Power Analyzer Flow

The PowerPlay Power Analyzer supports accurate power estimations by allowing you to specify all the important design factors affecting power consumption. [Figure 10-3](#) shows the high-level PowerPlay Power Analyzer flow.

Figure 10-3. PowerPlay Power Analyzer High-Level Flow



Note to Figure 10-3:

(1) Operating condition specifications are available only for some device families. For more information, refer to [Performing Power Analysis with the PowerPlay Power Analyzer](#) in Quartus II Help.

The PowerPlay Power Analyzer requires your design to be synthesized and fitted to the target device. You must specify the electrical standard used by each I/O cell and the capacitive load on each I/O standard in your design to obtain accurate I/O power estimates.

Operating Settings and Conditions

You can specify device power characteristics, operating voltage conditions, and operating temperature conditions for power analysis in the Quartus II software.

On the **Operating Settings and Conditions** page of the **Settings** dialog box, you can specify whether the device has typical power consumption characteristics or maximum power consumption characteristics.

- For more information, refer to *Operating Setting and Conditions Page (Settings Dialog Box)* in Quartus II Help.

On the **Voltage** page of the **Settings** dialog box, you can view the operating voltage conditions for each power rail in the device, and specify supply voltages for power rails with selectable supply voltages.

- For more information, refer to *Voltage Page (Settings Dialog Box)* in Quartus II Help.

On the **Temperature** page of the **Settings** dialog box, you can specify the thermal operating conditions of the device.

- For more information, refer to *Temperature Page (Settings Dialog Box)* in Quartus II Help.

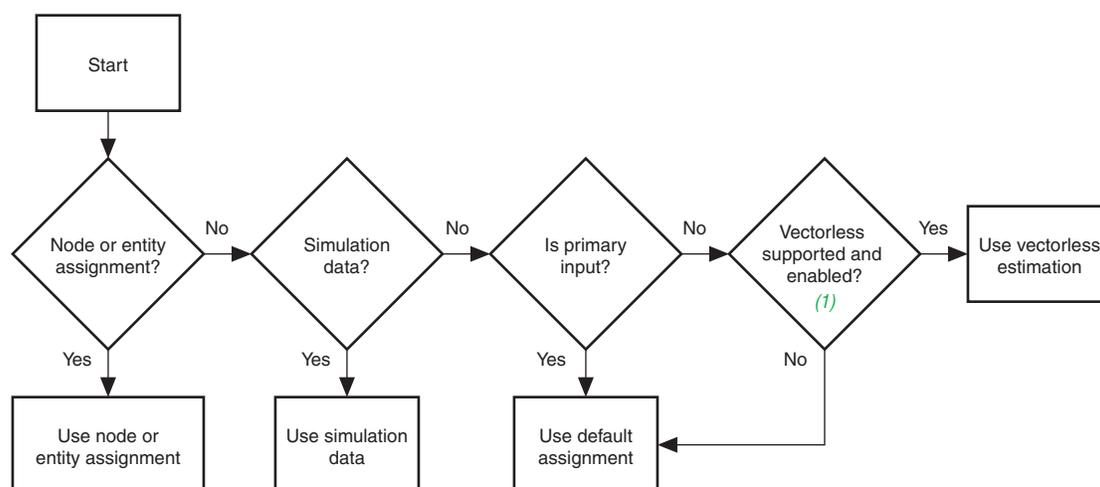
Signal Activities Data Sources

The PowerPlay Power Analyzer provides a flexible framework for specifying signal activities. It reflects the importance of using representative signal-activity data during power analysis. You can use the following sources to provide information about signal activity:

- Simulation results
- User-entered node, entity, and clock assignments
- User-entered default toggle rate assignment
- Vectorless estimation

The PowerPlay Power Analyzer allows you to mix and match the signal-activity data sources on a signal-by-signal basis. [Figure 10-4](#) shows the priority scheme. The following sections describe the data sources.

Figure 10-4. Signal-Activity Data Source Priority Scheme



Note to Figure 10-4:

- (1) Vectorless estimation is available only for some device families. For more information, refer to *Performing Power Analysis with the PowerPlay Power Analyzer*.

Simulation Results

The PowerPlay Power Analyzer directly reads the waveforms generated by a design simulation. The static probability and toggle rate for each signal are calculated from the simulation waveform. Power analysis is most accurate when you use representative input stimuli to generate simulations.

The PowerPlay Power Analyzer reads results generated by the following simulators:

- ModelSim®
- ModelSim-Altera
- QuestaSim
- Active-HDL
- NCSim
- VCS
- VCS MX
- Riviera-PRO

Signal activity and static probability information derive from a Verilog Value Change Dump File (.vcd). For more information, refer to [“Signal Activities” on page 10-4](#).

For third-party simulators, use the **Quartus II EDA Tool Settings for Simulation** to specify a **Generate Value Change Dump** file script. These scripts instruct the third-party simulators to generate a .vcd that encodes the simulated waveforms. The Quartus II PowerPlay Power Analyzer reads this file directly to derive the toggle rate and static probability data for each signal.

Third-party EDA simulators, other than those listed, can generate a .vcd that can then be used with the PowerPlay Power Analyzer. For those simulators, you must manually create a simulation script to generate the appropriate .vcd.



You can use a .vcd created for power analysis to optimize your design for power during fitting by utilizing the appropriate settings in the **PowerPlay power optimization** list, available in the **Fitter Settings** page of the **Settings** dialog box.



For more information about power optimization, refer to the [Power Optimization](#) chapter in volume 2 of the *Quartus II Handbook*.

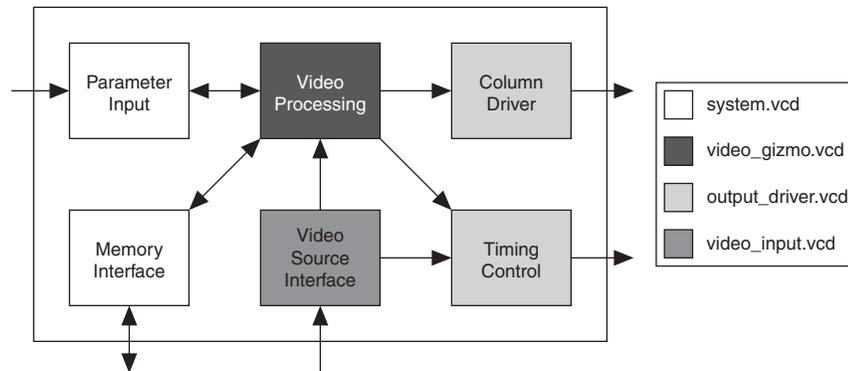


For more information about how to create a .vcd in other third-party EDA simulation tools, refer to [Section 1. Simulation](#) in volume 3 of the *Quartus II Handbook*.

Using Simulation Files in Modular Design Flows

A common design practice is to create modular or hierarchical designs in which you develop each design entity separately, and then instantiate it in a higher-level entity, forming a complete design. You can perform simulation on a complete design or on each modular design for verification. The PowerPlay Power Analyzer supports modular design flows when reading the signal activities generated from these simulation files. An example of a modular design flow is shown in [Figure 10-5](#).

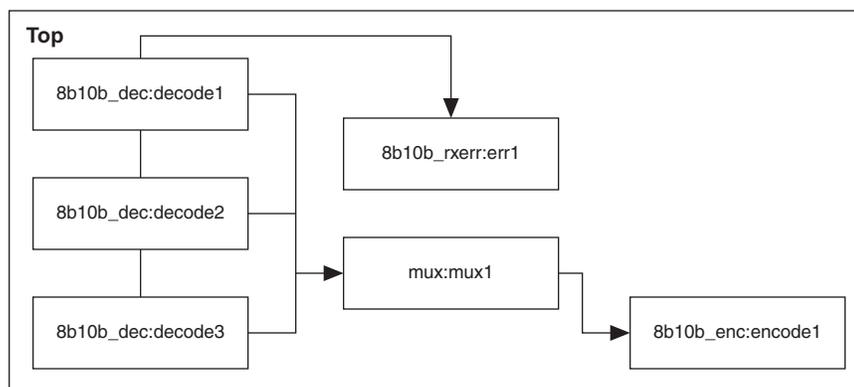
Figure 10-5. Modular Simulation Flow



When specifying a simulation file, an associated design entity name is given, such that the signal activities derived from the simulation file (**.vcd**) are imported into the PowerPlay Power Analyzer for that particular design entity. The PowerPlay Power Analyzer also supports the specification of multiple **.vcd** for power analysis, with each having an associated design entity name to allow the integration of partial design simulations into a complete design power analysis. When specifying multiple **.vcd** for your design, it is possible that more than one simulation file contains signal-activity information for the same signal. When you apply multiple **.vcd** to the same design entity, the signal activity used in the power analysis is the equal-weight arithmetic average of each **.vcd**. When you apply multiple simulation files to design entities at different levels in your design hierarchy, the signal activity in the power analysis derives from the simulation file that applies to the most specific design entity.

Figure 10-6 shows an example of a hierarchical design. The top-level module of your design, called **Top**, consists of three 8b/10b decoders, followed by a multiplexer. The output of the multiplexer is then encoded again before being the output from your design. There is also an error-handling module that handles any 8b/10b decoding errors. The top contains the top-level entity of your design and any logic not defined as part of another module. The design file for the top-level module might be just a wrapper for the hierarchical entities below it, or it might contain its own logic. The following usage scenarios show common ways that you can simulate your design and import .vcd into the PowerPlay Power Analyzer.

Figure 10-6. Example Hierarchical Design



Complete Design Simulation

You can simulate the entire design top, generating a .vcd from a third-party simulator. The .vcd can then be imported (specifying entity top) into the PowerPlay Power Analyzer. The resulting power analysis uses all the signal activities information from the generated .vcd, including those that apply to submodules, such as decode [1-3], err1, mux1, and encode1.

Modular Design Simulation

You can independently simulate submodules of the design top, and then import all the resulting .vcd into the PowerPlay Power Analyzer. For example, you can simulate the 8b10b_dec independent of the entire design, as well as multiplexer, 8b10b_rxerr, and 8b10b_enc. You can then import the .vcd generated from each simulation by specifying the appropriate instance name. For example, if the files produced by the simulations are 8b10b_dec.vcd, 8b10b_enc.vcd, 8b10b_rxerr.vcd, and mux.vcd, the import specifications in Table 10-2 are used.

Table 10-2. Import Specifications (Part 1 of 2)

File Name	Entity
8b10b_dec.vcd	Top 8b10b_dec:decode1
8b10b_dec.vcd	Top 8b10b_dec:decode2
8b10b_dec.vcd	Top 8b10b_dec:decode3
8b10b_rxerr.vcd	Top 8b10b_rxerr:err1

Table 10-2. Import Specifications (Part 2 of 2)

File Name	Entity
8b10b_enc.vcd	Top 8b10b_enc: encode1
mux.vcd	Top mux: mux1

The resulting power analysis applies the simulation vectors found in each file to the assigned entity. Simulation provides signal activities for the pins and for the outputs of functional blocks. If the inputs to an entity instance are input pins for the entire design, the simulation file associated with that instance does not provide signal activities for the inputs of that instance. For example, an input to an entity such as `mux1` has its signal activity specified at the output of one of the decode entities.

Multiple Simulations on the Same Entity

You can perform multiple simulations of an entire design or specific modules of a design. For example, in the process of verifying the design top, you can have three different simulation testbenches: one for normal operation and two for corner cases. Each of these simulations produces a separate `.vcd`. In this case, apply the different `.vcd` names to the same top-level entity, shown in [Table 10-3](#).

Table 10-3. Multiple Simulation File Names and Entities

File Name	Entity
normal.vcd	Top
corner1.vcd	Top
corner2.vcd	Top

The resulting power analysis uses an arithmetic average that the signal activities calculated from each simulation file to obtain the final signal activities used. If a signal `err_out` has a toggle rate of zero toggles per second in `normal.vcd`, 50 toggles per second in `corner1.vcd`, and 70 toggles per second in `corner2.vcd`, the final toggle rate in the power analysis is 40 toggles per second.

Overlapping Simulations

You can perform a simulation on the entire design top and more exhaustive simulations on a submodule, such as `8b10b_rxerr`. [Table 10-4](#) shows the import specification for overlapping simulations.

Table 10-4. Overlapping Simulation Import Specifications

File Name	Entity
full_design.vcd	Top
error_cases.vcd	Top 8b10b_rxerr: err1

In this case, signal activities from `error_cases.vcd` are used for all of the nodes in the generated `.vcd` and signal activities from `full_design.vcd` are used for only those nodes that do not overlap with nodes in `error_cases.vcd`. In general, the more specific hierarchy (the most bottom-level module) derives signal activities for overlapping nodes.

Partial Simulations

You can perform a simulation in which the entire simulation time is not applicable to signal-activity calculation. For example, run a simulation for 10,000 clock cycles and reset the chip for the first 2,000 clock cycles. If the signal-activity calculation is performed over all 10,000 cycles, the toggle rates are only 80% of their steady state value (because the chip is in reset for the first 20% of the simulation). In this case, you must specify the useful parts of the `.vcd` for power analysis. The **Limit VCD Period** option enables you to specify a start and end time to be used when performing signal-activity calculations.

Node Name Matching Considerations

Node name mismatches happen when you have `.vcd` applied to entities other than the top-level entity. In a modular design flow, the gate-level simulation files created in different Quartus II projects may not match their node names with the current Quartus II project.

For example, if you have a file named `8b10b_enc.vcd`, which was generated in a separate project called `8b10b_enc` and is simulating the 8b10b encoder, and you import that `.vcd` into another project called `Top`, you might encounter name mismatches when applying the `.vcd` to the `8b10b_enc` module in the `Top` project. This mismatch happens because all the combinational nodes in the `8b10b_enc.vcd` might be named differently in the `Top` project.

You can avoid name mismatching with only RTL simulation data, in which register names do not change, or with an incremental compilation flow that preserves node names in conjunction with a gate-level simulation.



To ensure the best accuracy, Altera recommends using an incremental compilation flow to preserve the node names of your design.



For more information about the incremental compilation flow, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.

Glitch Filtering

The PowerPlay Power Analyzer defines a glitch as two signal transitions so closely spaced in time that the pulse or glitch occurs faster than the logic and routing circuitry can respond. The output of a transport delay model simulator contains glitches for some signals. The logic and routing structures of the device form a low-pass filter that filters out glitches that are tens to hundreds of picoseconds long, depending on the device family.

Some third-party simulators use different models than the transport delay model as default model. Different models cause differences in signal activity and power estimation. The inertial delay model, which is the ModelSim default model, filters out more glitches than the transport delay model and usually yields a lower power estimate.

-  Altera recommends using the transport simulation model when using the Quartus II software glitch filtering support with third-party simulators. Simulation glitch filtering has little effect if you use the inertial simulation model.
-  For more information about how to set the simulation model type for your specific simulator, refer to Quartus II Help.

Glitch filtering in a simulator can also filter a glitch on one LE (or other circuit element) output from propagating to downstream circuit elements to ensure that the glitch does not affect simulated results. It prevents a glitch on one signal from producing non-physical glitches on all downstream logic, which can result in a signal toggle rate and a power estimate that are too high. Circuit elements in which every input transition produces an output transition, including multipliers and logic cells configured to implement XOR functions, are especially prone to glitches. Therefore, circuits with such functions can have power estimates that are too high when you do not use glitch filtering.

Altera recommends using the glitch filtering feature to obtain the most accurate power estimates. For .vcd, the PowerPlay Power Analyzer flows support two levels of glitch filtering, both of which are recommended for power estimation.

In the first level of glitch filtering, glitches are filtered during simulation. To enable this level of glitch filtering in the Quartus II software for supported third-party simulators, follow these steps:

1. On the Assignments menu, click **Settings**. The **Settings** dialog box appears.
2. In the **Category** list, select **Simulation** under **EDA Tool Settings**. The **Simulation** page appears.
3. Select the **Tool name** to use for the simulation.
4. Turn on **Enable glitch filtering**.

The second level of glitch filtering occurs while the PowerPlay Power Analyzer is reading the .vcd generated by a third-party simulator. To enable this level of glitch filtering, follow these steps:

On the Assignments menu, click **Settings**. The **Settings** dialog box appears.

1. In the **Category** list, select **PowerPlay Power Analyzer Settings**. The **PowerPlay Power Analyzer Settings** page appears.
2. Under **Input File(s)**, turn on **Perform glitch filtering on VCD files**.

The .vcd file reader performs complementary filtering to the filtering performed during simulation and is often not as effective. While the .vcd file reader can remove glitches on logic blocks, it has no way of determining how downstream logic and routing are affected by a given glitch, and may eliminate the impact of the glitch completely. Filtering the glitches during simulation avoids switching downstream routing and logic automatically.

-  When running simulation for design verification (rather than to produce input to the PowerPlay Power Analyzer), Altera recommends turning off the glitch filtering option to produce the most rigorous and conservative simulation from a functionality viewpoint. When performing simulation to produce input for the PowerPlay Power Analyzer, Altera recommends turning on the glitch filtering to produce the most accurate power estimates.

Node and Entity Assignments

You can assign specific toggle rates and static probabilities to individual nodes and entities in the design. These assignments have the highest priority, overriding data from all other signal-activity sources.

You must use the Assignment Editor or Tcl commands to create the **Power Toggle Rate** and **Power Static Probability** assignments. You can specify the power toggle rate as an absolute toggle rate in transitions using the **Power Toggle Rate** assignment or you can use the **Power Toggle Rate Percentage** assignment to specify a toggle rate relative to the clock domain of the assigned node for a more specific assignment made in terms of hierarchy level.

-  If the **Power Toggle Rate Percentage** assignment is used, and the given node does not have a clock domain, a warning is issued and the assignment is ignored.
-  For more information about how to use the Assignment Editor in the Quartus II software, refer to the *Assignment Editor* chapter in volume 2 of the *Quartus II Handbook*.

Assigning specific toggle rates and static probabilities to individual nodes and entities is appropriate for signals in which you have specific knowledge of the signal or entity being analyzed. For example, if you know that a 100 MHz data bus or memory output produces data that is essentially random (uncorrelated in time), you can directly enter a 0.5 static probability and a toggle rate of 50 million transitions per second.

Bidirectional I/O pins are treated specially. The combinational input port and the output pad for a given pin share the same name. However, those ports might not share the same signal activities. For the purpose of reading signal-activity assignments, the PowerPlay Power Analyzer creates a distinct name `<node_name~output>` when the bidirectional signal is configured as an output and `<node_name~result>` when the signal is configured as an input. For example, if a design has a bidirectional pin named MYPIN, assignments for the combinational input use the name MYPIN~result, and the assignments for the output pad use the name MYPIN~output.

-  When creating the logic assignment in the Assignment Editor, you will not find the MYPIN~result and MYPIN~output node names in the Node Finder. Therefore, to create the logic assignment, you must manually enter the two differentiating node names to create the specific assignment for the input and output port of the bidirectional pin.

Timing Assignments to Clock Nodes

For clock nodes, the PowerPlay Power Analyzer uses the timing requirements to derive the toggle rate when neither simulation data nor user-entered signal-activity data is available. f_{MAX} requirements specify full cycles per second, but each cycle represents a rising transition and a falling transition. For example, a clock f_{MAX} requirement of 100 MHz corresponds to 200 million transitions per second.

Default Toggle Rate Assignment

You can specify a default toggle rate for primary inputs and all other nodes in the design. The default toggle rate is used when no other method has specified the signal-activity data.

The toggle rate is specified in absolute terms (transitions per second) or as a fraction of the clock rate in effect for each particular node. The toggle rate for a given clock derives from the timing settings for the clock. For example, if a clock is specified with an f_{MAX} constraint of 100 MHz and a default relative toggle rate of 20%, nodes in this clock domain transition in 20% of the clock periods, or 20 million transitions occur per second. In some cases, the PowerPlay Power Analyzer cannot determine the clock domain for a given node because there is either no clock domain for the node or it is ambiguous. In these cases, the PowerPlay Power Analyzer substitutes and reports a toggle rate of zero.

Vectorless Estimation

For some device families, the PowerPlay Power Analyzer automatically derives estimates for signal activity on nodes with no simulation or user-entered signal-activity data. Vectorless estimation statistically estimates the signal activity of a node based on the signal activities of all nodes feeding that node, and on the actual logic function implemented by the node. Vectorless estimation cannot derive signal activities for primary inputs. Vectorless estimation is generally accurate for combinational nodes, but not for registered nodes. Therefore, simulation data for at least the registered nodes and I/O nodes is required for accuracy.

- ④ For more information, refer to *Performing Power Analysis with the PowerPlay Power Analyzer* in Quartus II Help.

The **PowerPlay Power Analyzer Settings** dialog box lets you disable vectorless estimation. When enabled, vectorless estimation takes priority over default toggle rates. Vectorless estimation does not override clock assignments.

Using the PowerPlay Power Analyzer

For all the flows that use the PowerPlay Power Analyzer, synthesize your design first and then fit it to the target device. You must either provide timing assignments for all the clocks in the design or use a simulation-based flow to generate activity data. The I/O standard used on each device input or output and the capacitive load on each output must be specified in the design.

- ④ For more information about using the PowerPlay Power Analyzer, refer to *Performing Power Analysis with the PowerPlay Power Analyzer* in Quartus II Help.

Common Analysis Flows

You can use the analysis flows in this section with the PowerPlay Power Analyzer. However, vectorless activity estimation is only available for some device families.

Signal Activities from Full Post-Fit Netlist (Timing) Simulation

This flow provides the most accuracy because all node activities reflect actual design behavior, provided that supplied input vectors are representative of typical design operation. Results are better if the simulation filters glitches. The disadvantage of this method is that the simulation time is long.

Signal Activities from Full Post-Fit Netlist (Zero Delay) Simulation

The zero delay simulation flow is used with designs for which signal activities from a full post-fit netlist (timing) simulation are not available. Zero delay simulation is as accurate as timing simulation in 95% of designs with no glitches.



If your design has glitches, power may be underestimated. Altera recommends using the signal activities from a full post-fit netlist (timing) simulation to achieve accurate power estimation of your design.

The following designs might exhibit glitches:

- Designs with many XOR gates (for example, an encryption core)
- Designs with arithmetic blocks without input and output registers (DSPs and carry chains)

For more information about creating zero delay simulation signal activities, refer to [“Generating a .vcd from Full Post-Fit Netlist \(Zero Delay\) Simulation”](#) on page 10-19.

Signal Activities from RTL (Functional) Simulation, Supplemented by Vectorless Estimation

In this flow, simulation provides toggle rates and static probabilities for all pins and registers in the design. Vectorless estimation fills in the values for all the combinational nodes between pins and registers, giving good results. This flow usually provides a compilation time benefit to the user in the third-party RTL simulator.



RTL simulation may not provide signal activities for all registers in the post-fitting netlist because some register names might be lost during synthesis. For example, synthesis might automatically transform state machines and counters, thus changing the names of registers in those structures.

Signal Activities from Vectorless Estimation and User-Supplied Input Pin Activities

This flow provides a low level of accuracy, because vectorless estimation for registers is not entirely accurate.

Signal Activities from User Defaults Only

This flow provides the lowest degree of accuracy.

Generating a .vcd

In previous versions of the Quartus II software, you could use either the Quartus II simulator or an EDA simulator to perform your simulation. The Quartus II software no longer supports a built-in simulator, and you must use EDA simulators to perform simulation. Use the .vcd as the input to the PowerPlay Power Analyzer to estimate power for your design.

 For more information about the supported third-party simulators, refer to [“Simulation Results” on page 10-9](#).

To create a .vcd for your design, follow these steps:

1. On the Assignments menu, click **Settings**. The **Settings** dialog box appears.
2. In the **Category** list, under **EDA Tool Settings**, click **Simulation**. The **Simulation** page appears.
3. In the **Tool name** list, select your preferred EDA simulator.
4. In the **Format for output netlist** list, select **Verilog HDL**, or **SystemVerilog HDL**, or **VHDL**.
5. Turn on **Generate Value Change Dump (VCD) file script**.

 This turns on the **Map illegal HDL characters** and **Enable glitch filtering** options. The **Map illegal HDL characters** option ensures that all signals have legal names and that signal toggle rates are available later in the PowerPlay Power Analyzer.

6. By turning on **Enable glitch filtering**, glitch filtering logic is the output when you generate an EDA netlist for simulation. This option is available regardless of whether or not you want to generate the .vcd scripts. For more information about glitch filtering, refer to [“Glitch Filtering” on page 10-13](#).

 When performing simulation using ModelSim, the **+nospecify** option for the vsim command disables the **specify path delays and timing checks** option in ModelSim. By enabling glitch filtering on the **Simulation** page, the simulation models include specified path delays. Thus, ModelSim might fail to simulate a design if glitch filtering is enabled, and the **+nospecify** option is specified. Altera recommends removing the **+nospecify** option from the ModelSim vsim command to ensure accurate simulation for power estimation.

7. Click **Script Settings**. The **Script Settings** dialog box appears.

Select which signals must be output to the .vcd. With **All signals** selected, the generated script instructs the third-party simulator to write all connected output signals to the .vcd. With **All signals except combinational lcell outputs** selected, the generated script tells the third-party simulator to write all connected output signals to the .vcd, except logic cell combinational outputs.

 The file can become extremely large if you write all output signals to the file because its size depends on the number of output signals being monitored and the number of transitions that occur.

8. Click **OK**.
9. Type a name for your testbench in the **Design instance name** box.
10. Compile your design with the Quartus II software and generate the necessary EDA netlist and script that instructs the third-party simulator to generate a **.vcd**.



For more information about the NativeLink feature, refer to *Section I. Simulation* in volume 3 of the *Quartus II Handbook*.

11. Perform a simulation with the third-party EDA simulation tool. Call the generated script in the simulation tool before running the simulation. The simulation tool generates the **.vcd** and places it in the project directory.

Generating a **.vcd** from ModelSim Software

To successfully produce a **.vcd** with the ModelSim software, follow these steps:

1. In the Quartus II software, on the Assignments menu, click **Settings**. The **Settings** dialog box appears.
2. In the **Category** list, under **EDA Tool Settings**, click **Simulation**. The **Simulation** page appears.
3. In the **Tool name** list, select your preferred EDA simulator.
4. In the **Format for output netlist** list, select **Verilog HDL**, or **SystemVerilog HDL**, or **VHDL**.
5. Turn on **Generate Value Change Dump (VCD) file script**.
6. To generate the **.vcd**, perform a full compilation.
7. In the ModelSim software, compile the files necessary for simulation.
8. Load your design by clicking **Start Simulation** on the Tools menu, or use the `vsim` command.
9. Use the **.vcd** script created in step 6 using the following command:

```
source <design>_dump_all_vcd_nodes.tcl
```
10. Run the simulation (for example, `run 2000ns` or `run -all`).
11. Quit the simulation using the `quit -sim` command, if required.
12. Exit the ModelSim software. If you do not exit the software, the ModelSim software might end the writing process of the **.vcd** improperly, resulting in a corrupt **.vcd**.

Generating a **.vcd** from Full Post-Fit Netlist (Zero Delay) Simulation

To successfully generate a **.vcd** from the full post-fit Netlist (zero delay) simulation, follow these steps:

1. Compile your design in the Quartus II software to generate the Netlist `<project_name>.vo`.
2. In `<project_name>.vo`, search for the include statement for `<project_name>.sdo`, comment the statement out, and save the file.

3. Generate a **.vcd** for power estimation by performing the steps in “[Generating a .vcd](#)” on page 10-18.



Altera recommends using the Standard Delay Format Output File (**.sdo**) for gate-level timing simulation. The **.sdo** contains the delay information of each architecture primitive and routing element specific to your design; however, you must exclude the **.sdo** for zero delay simulation.



For more information about how to create a **.vcd** in other third-party EDA simulation tools, refer to [Section I. Simulation](#) in volume 3 of the *Quartus II Handbook*.

Running the PowerPlay Power Analyzer Using the Quartus II GUI

To run the PowerPlay Power Analyzer using the Quartus II GUI, refer to [Performing Power Analysis with the PowerPlay Power Analyzer](#) in Quartus II Help.

PowerPlay Power Analyzer Compilation Report

The PowerPlay Power Analyzer section of the Compilation Report consists of the following sections.

Summary

This section of the report shows the estimated total thermal power consumption of your design. This includes dynamic, static, and I/O thermal power consumption. The I/O thermal power consumption is the total I/O power contributed by both the V_{CCIO} power supplies and some portion of the V_{CCINT} . The report also includes a confidence metric that reflects the overall quality of the data sources for the signal activities. For example, a **Low** power estimation confidence value reflects that you have provided insufficient toggle rate data, or most of the signal-activity information used for power estimation is from default or vectorless estimation settings. For more information about the input data, refer to the PowerPlay Power Analyzer Confidence Metric report.

Settings

This section of the report shows the PowerPlay Power Analyzer settings information of your design, including the default input toggle rates, operating conditions, and other relevant setting information.

Simulation Files Read

This section of the report lists the simulation output file (**.vcd**) used for power estimation. This section also includes the file ID, file type, entity, VCD start time, VCD end time, the unknown percentage, and the toggle percentage. The unknown percentage indicates the portion of the design module that is not exercised by the simulation vectors.

Operating Conditions Used

This section of the report shows device characteristics, voltages, temperature, and cooling solution, if any, that were used during the power estimation. This section also shows the entered junction temperature or auto-computed junction temperature that was used during the power analysis.

Thermal Power Dissipated by Block

This section of the report shows estimated thermal dynamic power and thermal static power consumption categorized by atoms. This information provides you with estimated power consumption for each atom in your design.

Thermal Power Dissipation by Block Type (Device Resource Type)

This section of the report shows the estimated thermal dynamic power and thermal static power consumption categorized by block types. This information is further categorized by estimated dynamic and static power that was used, as well as providing an average toggle rate by block type. Thermal power is the power dissipated as heat from the FPGA device.

Thermal Power Dissipation by Hierarchy

This section of the report shows estimated thermal dynamic power and thermal static power consumption categorized by design hierarchy. This information is further categorized by the dynamic and static power that was used by the blocks and routing in that hierarchy. This information is very useful when locating problem modules in your design.

Core Dynamic Thermal Power Dissipation by Clock Domain

This section of the report shows the estimated total core dynamic power dissipation by each clock domain, which provides designs with estimated power consumption for each clock domain in the design. If the clock frequency for a domain is unspecified by a constraint, the clock frequency is listed as “unspecified.” For all the combinational logic, the clock domain is listed as no clock with zero MHz.

Current Drawn from Voltage Supplies

This section of the report lists the current that was drawn from each voltage supply. The V_{CCIO} voltage supply is further categorized by I/O bank and by voltage. The minimum safe power supply size (current supply ability) is also listed for each supply voltage.

Transceiver-based devices have multiple voltage supplies, which are V_{CCH} , V_{CCT} , V_{CCR} , V_{CCA} , and V_{CCP} . The report also shows the static and dynamic current (in mA) drawn from each voltage supply. Total static and dynamic power consumed by the transceivers on all voltage supplies is listed under the “Thermal Power Dissipation by Block Type” report section, which contains a row that starts with “GXB Transceiver.”

The I/O thermal power dissipation, which is listed on the summary page, does not correlate directly to the power drawn from the V_{CCIO} voltage supply listed in this report. This is because the I/O thermal power dissipation value also includes portions of the V_{CCINT} power, such as the I/O element (IOE) registers, which are modeled as I/O power, but do not draw from the V_{CCIO} supply.

Confidence Metric Details

The confidence metric indicates the quality of the signal toggle rate data used to compute a power estimate. The confidence metric is low if the signal toggle rate data comes from sources that are considered poor predictors of real signal toggle rates in the device during an operation. Toggle rate data that comes from simulation, user-entered assignments on specific signals, or entities are considered reliable. Toggle rate data from default toggle rates (for example, 12.5% of the clock period) or vectorless estimation are considered relatively inaccurate. This section gives an overall confidence rating in the toggle rate data, from low to high. This section also summarizes how many pins, registers, and combinational nodes obtained their toggle rates from each of simulation, user entry, vectorless estimation, or default toggle rate estimations. This detailed information helps you understand how to increase the confidence metric, letting you determine your own confidence in the toggle rate data.

Signal Activities

This section lists toggle rates and static probabilities assumed by power analysis for all signals with fan-out and pins. The signal type is provided (pin, registered, or combinational), as well as the data source for the toggle rate and static probability. By default, all signal activities are reported, but can be turned off by turning off the **Write signal activities to report file** option on the **PowerPlay Power Analyzer Settings** page.



Altera recommends turning off the **Write signal activities to report file** option for a large design because of the large number of signals present. You can use the Assignment Editor to specify that activities for individual nodes or entities are reported by assigning an on value to those nodes for the **Power Report Signal Activities** assignment.

Messages

This section lists the messages generated by the Quartus II software during the analysis.

Specific Rules for Reporting

In a Stratix GX device, the XGM II state machine block is always used together with GXB transceivers, so its power is lumped into the power for the transceivers. Therefore, the power for the XGM II state machine block is reported as zero Watts.

Scripting Support

You can run procedures and create settings described in this chapter in a Tcl script. You can also run some procedures at a command prompt. For more information about scripting command options, refer to the Quartus II Command-Line and Tcl API Help browser. To run the Help browser, type the following command at the command prompt:

```
quartus_sh --qhelp←
```

-  For more information about Tcl scripting, refer to the *Tcl Scripting* chapter in volume 2 of the *Quartus II Handbook* and *API Functions for Tcl* in Quartus II Help. For more information about all settings and constraints in the Quartus II software, refer to the *Quartus II Settings File Reference Manual*. For more information about command-line scripting, refer to the *Command-Line Scripting* chapter in volume 2 of the *Quartus II Handbook*.

Running the PowerPlay Power Analyzer from the Command-Line

The separate executable used to run the PowerPlay Power Analyzer is `quartus_pow`. For a complete listing of all command-line options supported by `quartus_pow`, type the following command at a system command prompt:

```
quartus_pow --help or quartus_sh --qhelp ←
```

The following is an example of using the `quartus_pow` executable with project **sample.qpf**:

- To instruct the PowerPlay Power Analyzer to generate a PowerPlay EPE File, type the following command at a system command prompt:

```
quartus_pow sample --output_epe=sample.csv ←
```

- To instruct the PowerPlay Power Analyzer to generate a PowerPlay EPE File without performing the power estimate, type the following command at a system command prompt:

```
quartus_pow sample --output_epe=sample.csv --estimate_power=off ←
```

- To instruct the PowerPlay Power Analyzer to use a **.vcd** as input (**sample.vcd**), type the following command at a system command prompt:

```
quartus_pow sample --input_vcd=sample.vcd ←
```

- To instruct the PowerPlay Power Analyzer to use two **.vcd** files as input files (**sample1.vcd** and **sample2.vcd**), perform glitch filtering on the **.vcd** and use a default input I/O toggle rate of 10,000 transitions per second, type the following command at a system command prompt:

```
quartus_pow sample --input_vcd=sample1.vcd --input_vcd=sample2.vcd \  
--vcd_filter_glitches=on --\  
default_input_io_toggle_rate=10000transitions/s ←
```

- To instruct the PowerPlay Power Analyzer to not use any input file, a default input I/O toggle rate of 60%, no vectorless estimation, and a default toggle rate of 20% on all remaining signals, type the following command at a system command prompt:

```
quartus_pow sample --no_input_file --\  
default_input_io_toggle_rate=60% \  
--use_vectorless_estimation=off --default_toggle_rate=20% ←
```



There are no command-line options to specify the information found on the **PowerPlay Power Analyzer Settings Operating Conditions** page. You can use the Quartus II GUI to specify these options.

The `quartus_pow` executable creates a report file, `<revision name>.pow.rpt`. You can locate the report file in the main project directory. The report file contains the same information in “PowerPlay Power Analyzer Compilation Report” on page 10-20.

Conclusion

PowerPlay power analysis tools are designed for accurate estimation of power consumption from early design concept through design implementation. You can use the PowerPlay EPE to estimate power consumption during the design concept stage. You can use the PowerPlay Power Analyzer tool to refine power estimations during design implementation. The PowerPlay Power Analyzer produces detailed reports that you can use to optimize designs for lower power consumption and verify that the design is in your power budget.

Document Revision History

Table 10-5 shows the revision history for this chapter.

Table 10-5. Document Revision History

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Added links to Quartus II Help, removed redundant material. ■ Moved “Creating PowerPlay EPE Spreadsheets” to page 10-5. ■ Minor edits.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Removed references to the Quartus II Simulator. ■ Updated Table 10-1 on page 10-6, Table 10-2 on page 10-11, and Table 10-3 on page 10-12. ■ Updated Figure 10-3 on page 10-7, Figure 10-4 on page 10-8, and Figure 10-5 on page 10-10.
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Updated “Creating PowerPlay EPE Spreadsheets” on page 10-5 and “Simulation Results” on page 10-9. ■ Added “Signal Activities from Full Post-Fit Netlist (Zero Delay) Simulation” on page 10-17 and “Generating a .vcd from Full Post-Fit Netlist (Zero Delay) Simulation” on page 10-19. ■ Minor changes to “Generating a .vcd from ModelSim Software” on page 10-19. ■ Updated Figure 10-2 on page 10-5 and Figure 11-8 on page 11-24.
March 2009	9.0.0	<ul style="list-style-type: none"> ■ This chapter was chapter 11 in version 8.1. ■ Removed Figures 11-10, 11-11, 11-13, 11-14, and 11-17 from 8.1 version.
November 2008	8.1.0	<ul style="list-style-type: none"> ■ Updated for the Quartus II software version 8.1. ■ Replaced Figure 11-3. ■ Replaced Figure 11-14.
May 2008	8.0.0	<ul style="list-style-type: none"> ■ Updated Figure 11-5. ■ Updated “Types of Power Analyses” on page 11-5. ■ Updated “Operating Conditions” on page 11-9. ■ Updated “PowerPlay Power Analyzer Compilation Report” on page 11-31. ■ Updated “Current Drawn from Voltage Supplies” on page 11-32.

-  For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).
-  Take an [online survey](#) to provide feedback about this handbook chapter.

The Altera® Quartus® II design software provides a complete design debugging environment that easily adapts to your specific design requirements. This handbook is arranged in chapters, sections, and volumes that correspond to the major tools available for debugging your designs. For a general introduction to features and the standard design flow in the software, refer to the *Introduction to the Quartus II Software* manual.

This section is an introduction to System Debugging Tools and includes the following chapters:

■ [Chapter 11, System Debugging Tools Overview](#)

This chapter compares the various system debugging tools and explains when to use each of them.

■ [Chapter 12, Analyzing and Debugging Designs with the System Console](#)

This chapter describes the System Console Toolkit and compares the different capabilities within the toolkit.

[Chapter 13, Transceiver Link Debugging Using the System Console](#)

This chapter explains what functions are available within the Transceiver Toolkit and helps you decide which tool best meets your debugging needs.

■ [Chapter 14, Quick Design Debugging Using SignalProbe](#)

This chapter provides detailed instructions about how to use SignalProbe to quickly debug your design.

Use this chapter to verify your design more efficiently by routing internal signals to I/O pins quickly without affecting the design.

■ [Chapter 15, Design Debugging Using the SignalTap II Logic Analyzer](#)

This chapter describes how to debug your FPGA design during normal device operation without the need for external lab equipment. Use this chapter to learn how to examine the behavior of internal signals, without using extra I/O pins, while the design is running at full speed on an FPGA device.

■ [Chapter 16, In-System Debugging Using External Logic Analyzers](#)

This chapter explains how to use external logic analyzers to debug designs on Altera devices.

■ [Chapter 17, In-System Modification of Memory and Constants](#)

This chapter explains how to use the Quartus II In-System Memory Content Editor as part of your FPGA design and verification flow to easily view and debug your design in the hardware lab.

- [Chapter 18, Design Debugging Using In-System Sources and Probes](#)

This chapter provides detailed instructions about how to use the In-System Sources and Probes Editor and Tcl scripting in the Quartus®II software to debug your design.

The Altera® system debugging tools help you verify your FPGA designs. As your product requirements continue to increase in complexity, the time you spend on design verification continues to rise. This chapter provides a quick overview of the tools available in the system debugging suite and discusses the criteria for selecting the best tool for your design.

The Quartus® II software provides a portfolio of system design debugging tools for real-time verification of your design. Each tool in the system debugging portfolio uses a combination of available memory, logic, and routing resources to assist in the debugging process. The tools provide visibility by routing (or “tapping”) signals in your design to debugging logic. The debugging logic is then compiled with your design and downloaded into the FPGA or CPLD for analysis. Because different designs can have different constraints and requirements, such as the number of spare pins available or the amount of logic or memory resources remaining in the physical device, you can choose a tool from the available debugging tools that matches the specific requirements for your design.

System Debugging Tools

Table 11–1 summarizes the tools in the system debugging tool suite that are covered in this section.

Table 11–1. Available Tools in the In-System Verification Tools Suite (Part 1 of 2)

Tool	Description	Typical Usage
System Console	This is a Tcl console that communicates to hardware modules instantiated into your design. You can use it with the Transceiver Toolkit to monitor or debug your design.	You need to perform system-level debugging. For example, if you have an Avalon-MM slave or Avalon-ST interfaces, you can debug your design at a transaction level. The tool supports JTAG connectivity, but also supports PLI connectivity to a simulation model, as well as TCP/IP connectivity to the target FPGA you wish to debug.
Transceiver Toolkit	This toolkit allows you to calibrate PMA settings so that you can optimize settings for your board.	You need to debug or optimize transceiver link channels in your design.
SignalTap® II Logic Analyzer	This logic analyzer uses FPGA resources to sample tests nodes and outputs the information to the Quartus II software for display and analysis.	You have spare on-chip memory and you want functional verification of your design running in hardware.
SignalProbe	This tool incrementally routes internal signals to I/O pins while preserving results from your last place-and-routed design.	You have spare I/O pins and you would like to check the operation of a small set of control pins using either an external logic analyzer or an oscilloscope.

© 2010 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, HARDCOPY, MAX, MEGACORE, NIOS, QUARTUS and STRATIX are Reg. U.S. Pat. & Tm. Off. and/or trademarks of Altera Corporation in the U.S. and other countries. All other trademarks and service marks are the property of their respective holders as described at www.altera.com/common/legal.html. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera’s standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

Table 11-1. Available Tools in the In-System Verification Tools Suite (Part 2 of 2)

Tool	Description	Typical Usage
Logic Analyzer Interface (LAI)	This tool multiplexes a larger set of signals to a smaller number of spare I/O pins. LAI allows you to select which signals are switched onto the I/O pins over a JTAG connection.	You have limited on-chip memory, and have a large set of internal data buses that you would like to verify using an external logic analyzer. Logic analyzer vendors, such as Tektronics and Agilent, provide integration with the tool to improve the usability of the tool.
In-System Sources and Probes	This tool provides an easy way to drive and sample logic values to and from internal nodes using the JTAG interface.	You want to prototype a front panel with virtual buttons for your FPGA design.
In-System Memory Content Editor	This tool displays and allows you to edit on-chip memory.	You would like to view and edit the contents of on-chip memory that is not connected to a Nios II processor. You can also use the tool when you do not want to have a Nios II debug core in your system.
Virtual JTAG Interface	This megafunction allows you to communicate with the JTAG interface so that you can develop your own custom applications.	You have custom signals in your design that you want to be able to communicate with.

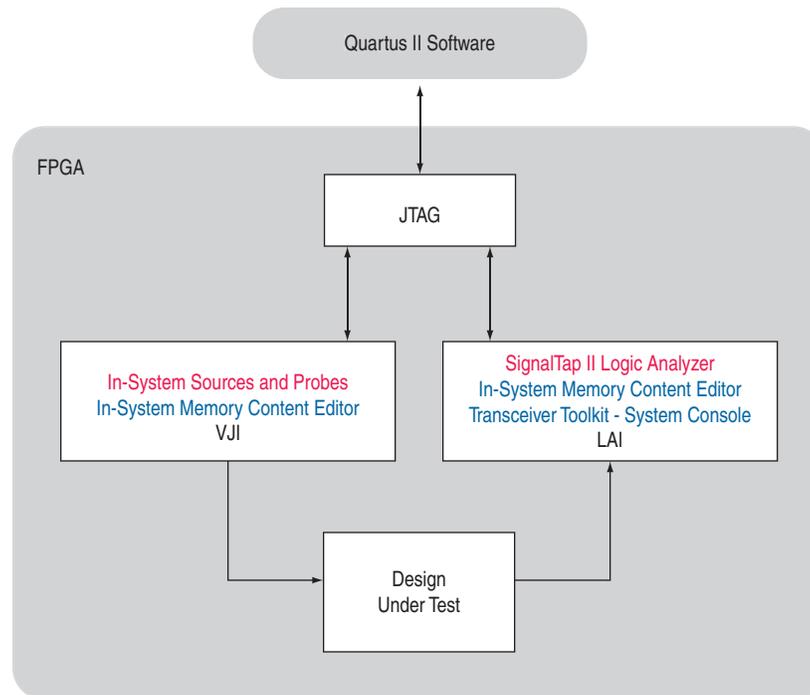
With the exception of SignalProbe, each of the on-chip debugging tools uses the JTAG port to control and read back data from debugging logic and signals under test. System Console uses JTAG and other interfaces as well. The JTAG resource is shared among all of the on-chip debugging tools.

For all system debugging tools except System Console, the Quartus II software compiles logic into your design automatically to distinguish between data and control information and each of the debugging logic blocks when the JTAG resource is required. This arbitration logic, also known as the System-Level Debugging (SLD) infrastructure, is shown in the design hierarchy of your compiled project as `sld_hub:sld_hub_inst`. The SLD logic allows you to instantiate multiple debugging blocks into your design and run them simultaneously. For System Console, you must explicitly insert IP cores into your design to enable debugging.

To maximize debugging closure, the Quartus II software allows you to use a combination of the debugging tools in tandem to fully exercise and analyze the logic under test. All of the tools described in [Table 11-1](#) have basic analysis features built in; that is, all of the tools enable you to read back information collected from the design nodes that are connected to the debugging logic. Out of the set of debugging tools, the SignalTap II Logic Analyzer, the LAI, and the SignalProbe feature are general purpose debugging tools optimized for probing signals in your RTL netlist. In-System Sources and Probes, the Virtual JTAG Interface, System Console, Transceiver Toolkit, and

In-System Memory Content Editor, in addition to being able to read back data from the debugging points, allow you to input values into your design during runtime. Taken together, the set of on-chip debugging tools form a debugging ecosystem. The set of tools can generate a stimulus to and solicit a response from the logic under test, providing a complete debugging solution (Figure 11-1).

Figure 11-1. Quartus II Debugging Ecosystem (Note 1)



Note to Figure 11-1:

- (1) The set of debugging tools offer end-to-end debugging coverage.

The tools in the toolchain offer different advantages and different trade-offs. To understand the selection criteria between the different tools, the following sections analyze the tools according to their typical applications.

The first section, “Analysis Tools for RTL Nodes”, compares the SignalTap II Logic Analyzer, SignalProbe, and the LAI. These three tools are logically grouped since they are intended for debugging nodes from your RTL netlist at system speed.

The second section, “Stimulus-Capable Tools” on page 11-8, compares System Console, In-System Memory Content Editor, Virtual JTAG Interface Megafunction, and In-System Sources and Probes. These tools are logically grouped since they offer the ability to both read and write transactions through the JTAG port.

Analysis Tools for RTL Nodes

The SignalTap II Logic Analyzer, the SignalProbe feature, and the LAI are designed specifically for probing and debugging RTL signals at system speed. They are general-purpose analysis tools that enable you to tap and analyze any routable node from the FPGA or CPLD. If you have spare logic and memory resources, the SignalTap II Logic Analyzer is useful for providing fast functional verification of your design running on actual hardware.

Conversely, if logic and memory resources are tight and you require the large sample depths associated with external logic analyzers, both the LAI and the SignalProbe feature make it easy to view internal design signals using external equipment.

The most important selection criteria for these three tools are the available resources remaining on your device after implementing your design and the number of spare pins available. You should evaluate your preferred debugging option early on in the design planning process to ensure that your board, your Quartus II project, and your design are all set up to support the appropriate options. Planning early can reduce time spent during debugging and eliminate the necessary late changes to accommodate your preferred debugging methodologies. The following two sections provide information to assist you in choosing the appropriate tool by comparing the tools according to their resource usage and their pin consumption.



The SignalTap II Logic Analyzer is not supported on CPLDs, because there are no memory resources available on these devices.

Resource Usage

Any debugging tool that requires the use of a JTAG connection requires the SLD infrastructure logic mentioned earlier, for communication with the JTAG interface and arbitration between any instantiated debugging modules. This overhead logic uses around 200 logic elements (LEs), a small fraction of the resources available in any of the supported devices. The overhead logic is shared between all available debugging modules in your design. Both the SignalTap II Logic Analyzer and the LAI use a JTAG connection.

SignalProbe requires very few on-chip resources. Because it requires no JTAG connection, SignalProbe uses no logic or memory resources. SignalProbe uses only routing resources to route an internal signal to a debugging test point.

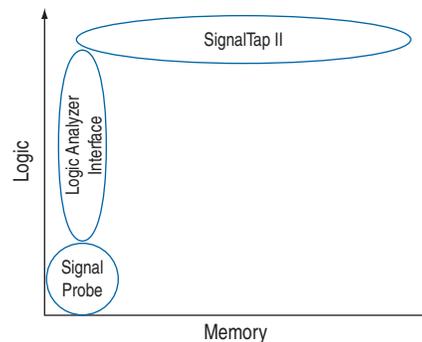
The LAI requires a small amount of logic to implement the multiplexing function between the signals under test, in addition to the SLD infrastructure logic. Because no data samples are stored on the chip, the LAI uses no memory resources.

The SignalTap II Logic Analyzer requires both logic and memory resources. The number of logic resources used depends on the number of signals tapped and the complexity of the trigger logic. However, the amount of logic resources that the SignalTap II Logic Analyzer uses is typically a small percentage of most designs. A baseline configuration consisting of the SLD arbitration logic and a single node with basic triggering logic contains approximately 300 to 400 Logic Elements (LEs). Each additional node you add to the baseline configuration adds about 11 LEs. Compared with logic resources, memory resources are a more important factor to consider for

your design. Memory usage can be significant and depends on how you configure your SignalTap II Logic Analyzer instance to capture data and the sample depth that your design requires for debugging. For the SignalTap II Logic Analyzer, there is the added benefit of requiring no external equipment, as all of the triggering logic and storage is on the chip.

Figure 11-2 shows a conceptual graph of the resource usage of the three analysis tools relative to each other.

Figure 11-2. Resource Usage per Debugging Tool (Note 1)



Note to Figure 11-2:

(1) Though resource usage is highly dependent on the design, this graph provides a rough guideline for tool selection.

The resource estimation feature for the SignalTap II Logic Analyzer and the LAI allows you to quickly judge if enough on-chip resources are available before compiling the tool with your design. Figure 11-3 shows the resource estimation feature for the SignalTap II Logic Analyzer and the LAI.

Figure 11-3. Resource Estimator

Instance Manager: Compile the project to continue					
Instance	Status	LEs: 652	Memory: 524288	M512/MLAB: 0/94	M4K/M9K: 128/60
auto_signaltap_0	Not running	652 cells	524288 bits	0 blocks	Can't Fit 128 blocks

Pin Usage

The ratio of the number of pins used to the number of signals tapped for the SignalProbe feature is one-to-one. Because this feature can consume free pins quickly, a typical application for this feature is routing control signals to spare debugging pins for debugging.

The ratio of the number of pins used to the number of signals tapped for the LAI is many-to-one. It can map up to 256 signals to each debugging pin, depending on available routing resources. The control of the active signals that are mapped to the spare I/O pins is performed via the JTAG port. The LAI is ideal for routing data buses to a set of test pins for analysis.

Other than the JTAG test pins, the SignalTap II Logic Analyzer uses no additional pins. All data is buffered using on-chip memory and communicated to the SignalTap II Logic Analyzer GUI via the JTAG test port.

Usability Enhancements

The SignalTap II Logic Analyzer, the SignalProbe feature, and the LAI tools can be added to your existing design with minimal effects. With the node finder, you can find signals to route to a debugging module without making any changes to your HDL files. SignalProbe inserts signals directly from your post-fit database. The SignalTap II Logic Analyzer and LAI support inserting signals from both pre-synthesis and post-fit netlists. All three tools allow you to find and configure your debugging setup quickly. In addition, the Quartus II incremental compilation feature and the Quartus II incremental routing feature allow for a fast turnaround time for your programming file, increasing productivity and enabling fast debugging closure.

Both LAI and the SignalTap II Logic Analyzer support incremental compilation. With incremental compilation, you can add a SignalTap II Logic Analyzer instance or an LAI instance incrementally into your placed-and-routed design. This has the benefit of both preserving your timing and area optimizations from your existing design, and decreasing the overall compilation time when any changes are necessary during the debugging process. With incremental compilation, you can save up to 70% compile time of a full compilation.

SignalProbe uses the incremental routing feature. The incremental routing feature runs only the Fitter stage of the compilation. This also leaves your compiled design untouched, except for the newly routed node or nodes. With SignalProbe, you can save as much as 90% compile time of a full compilation.

As another productivity enhancement, all tools in the on-chip debugging tool set support scripting via the `quartus_stp` Tcl package. For the SignalTap II Logic Analyzer and the LAI, scripting enables user-defined automation for data collection while debugging in the lab.

In addition, the JTAG server allows you to debug a design that is running on a device attached to a PC in a remote location. This allows you to set up your hardware in the lab environment, download any new `.sof` files, and perform any analysis from your desktop.

Table 11-2 compares common debugging features between these tools and provides suggestions about which is the best tool to use for a given feature.

Table 11-2. Suggested On-Chip Debugging Tools for Common Debugging Features (Part 1 of 2) (Note 1)

Feature	SignalProbe	Logic Analyzer Interface (LAI)	SignalTap II Logic Analyzer	Description
Large Sample Depth	N/A	✓	—	An external logic analyzer used with the LAI has a bigger buffer to store more captured data than the SignalTap II Logic Analyzer. No data is captured or stored with SignalProbe.
Ease in Debugging Timing Issue	✓	✓	—	External equipment, such as oscilloscopes and Mixed Signal Oscilloscopes (MSOs), can be used with either LAI or SignalProbe. When used with the LAI to provide you with access to timing mode, you can debug combined streams of data.

Table 11-2. Suggested On-Chip Debugging Tools for Common Debugging Features (Part 2 of 2) (Note 1)

Feature	SignalProbe	Logic Analyzer Interface (LAI)	SignalTap II Logic Analyzer	Description
Minimal Effect on Logic Design	✓	✓ (2)	✓ (2)	The LAI adds minimal logic to a design, requiring fewer device resources. The SignalTap II Logic Analyzer has little effect on the design, because it is set as a separate design partition. SignalProbe incrementally routes nodes to pins, not affecting the design at all.
Short Compile and Recompile Time	✓	✓ (2)	✓ (2)	SignalProbe attaches incrementally routed signals to previously reserved pins, requiring very little recompilation time to make changes to source signal selections. The SignalTap II Logic Analyzer and the LAI can take advantage of incremental compilation to refit their own design partitions to decrease recompilation time.
Triggering Capability	N/A	N/A	✓	The SignalTap II Logic Analyzer offers triggering capabilities that are comparable to commercial logic analyzers.
I/O Usage	—	—	✓	No additional output pins are required with the SignalTap II Logic Analyzer. Both the LAI and SignalProbe require I/O pin assignments.
Acquisition Speed	N/A	—	✓	The SignalTap II Logic Analyzer can acquire data at speeds of over 200 MHz. The same acquisition speeds are obtainable with an external logic analyzer used with the LAI, but may be limited by signal integrity issues.
No JTAG Connection Required	✓	—	—	An FPGA design with the SignalTap II Logic Analyzer or the LAI requires an active JTAG connection to a host running the Quartus II software. SignalProbe does not require a host for debugging purposes.
No External Equipment Required	—	—	✓	The SignalTap II Logic Analyzer logic is completely internal to the programmed FPGA device. No extra equipment is required other than a JTAG connection from a host running the Quartus II software or the stand-alone SignalTap II Logic Analyzer software. SignalProbe and the LAI require the use of external debugging equipment, such as multimeters, oscilloscopes, or logic analyzers.

Notes to Table 11-2:

- (1) ✓ indicates the recommended tools for the feature.
 — indicates that while the tool is available for that feature, that tool may not give the best results.
 N/A indicates that the feature is not applicable for the selected tool.
- (2) When used with incremental compilation.

Stimulus-Capable Tools

The In-System Memory Content Editor, the In-System Sources and Probes, and the Virtual JTAG interface each enable you to use the JTAG interface as a general-purpose communication port. Though all three tools can be used to achieve the same results, there are some considerations that make one tool easier to use in certain applications than others. In-System Sources and Probes is ideal for toggling control signals. The In-System Memory Content Editor is useful for inputting large sets of test data. Finally, the Virtual JTAG interface is well suited for more advanced users who want to develop their own customized JTAG solution.

System Console provides system-level debugging at a transaction level, such as with Avalon-MM slave or Avalon-ST interfaces. You can communicate to a chip through JTAG, PLI connectivity for simulation models, and TCP/IP protocols. System Console is a Tcl console that you use to communicate with hardware modules that you have instantiated into your design.

In-System Sources and Probes

In-System Sources and Probes is an easy way to access JTAG resources to both read and write to your design. You can start by instantiating a megafunction into your HDL code. The megafunction contains source ports and probe ports for driving values into and sampling values from the signals that are connected to the ports, respectively. Transaction details of the JTAG interface are abstracted away by the megafunction. During runtime, a GUI displays each source and probe port by instance and allows you to read from each probe port and drive to each source port. The GUI makes this tool ideal for toggling a set of control signals during the debugging process.

A good application of In-System Sources and Probes is to use the GUI as a replacement for the push buttons and LEDs used during the development phase of a project. Furthermore, In-System Sources and Probes supports a set of scripting commands for reading and writing using `quartus_stp`. When used with the Tk toolkit, you can build your own graphical interfaces. This feature is ideal for building a virtual front panel during the prototyping phase of the design.

In-System Memory Content Editor

The In-System Memory Content Editor allows you to quickly view and modify memory content either through a GUI interface or through Tcl scripting commands. The In-System Memory Content Editor works by turning single-port RAM blocks into dual-port RAM blocks. One port is connected to your clock domain and data signals, and the other port is connected to the JTAG clock and data signals for editing or viewing.

Because you can modify a large set of data easily, a useful application for the In-System Memory Content Editor is to generate test vectors for your design. For example, you can instantiate a free memory block, connect the output ports to the logic under test (using the same clock as your logic under test on the system side), and create the glue logic for the address generation and control of the memory. At runtime, you can modify the contents of the memory using either a script or the In-System Memory Content Editor GUI and perform a burst transaction of the data contents in the modified RAM block synchronous to the logic being tested.

Virtual JTAG Interface Megafunction

The Virtual JTAG Interface megafunction provides the finest level of granularity for manipulating the JTAG resource. This megafunction allows you to build your own JTAG scan chain by exposing all of the JTAG control signals and configuring your JTAG Instruction Registers (IRs) and JTAG Data Registers (DRs). During runtime, you control the IR/DR chain through a Tcl API, or with System Console. This feature is meant for users who have a thorough understanding of the JTAG interface and want precise control over the number and type of resources used.

System Console

System Console is a framework that you can launch from the Quartus II software to start services for performing various debugging tasks. System Console provides you with Tcl scripts and a GUI to access either the Qsys system integration tool or SOPC Builder modules to perform low-level hardware debugging of your design, as well as identify a module by its path, and open and close a connection to a Qsys or SOPC Builder module. You can access your design at a system level for purposes of loading, unloading, and transferring designs to multiple devices.

System Console also allows you to access commands that allow you to control how you generate test patterns, as well as verify the accuracy of data generated by test patterns. You can use JTAG debug commands in System Console to verify the functionality and signal integrity of your JTAG chain. You can test clock and reset signals.

You can use System Console to access programmable logic devices on your development board, as well as bring up a board and verify stages of setup. You can also access software running on a Nios II processor, as well as access modules that produce or consume a stream of bytes.

Transceiver Toolkit runs from the System Console framework, and allows you to run automatic tests of your transceiver links for debugging and optimizing your transceiver designs. You can use the Transceiver Toolkit GUI to set up channel links in your transceiver devices, and then automatically run EyeQ and Auto Sweep testing to view a graphical representation of your test data.

Conclusion

The Quartus II on-chip debugging tool suite allows you to reach debugging closure quickly by providing you with a set of powerful analysis tools and a set of tools that open up the JTAG port as a general purpose communication interface. The Quartus II software further broadens the scope of applications by giving you a comprehensive Tcl/Tk API. With the Tcl/Tk API, you can increase the level of automation for all of the analysis tools. You can also build virtual front panel applications quickly during the early prototyping phase.

In addition, all of the on-chip debugging tools have a tight integration with the rest of the productivity features within the Quartus II software. The incremental compilation and incremental routing features enable a fast turnaround time for programming file generation. The cross-probing feature allows you to find and identify nodes quickly. The SignalTap II Logic Analyzer, when used with the TimeQuest Timing Analyzer, is a best-in-class timing verification suite that allows fast functional and timing verification.

Document Revision History

Table 11-3 shows the revision history for this chapter.

Table 11-3. Document Revision History

Date	Version	Changes
December 2010	10.1.0	Maintenance release. Changed to new document template.
July 2010	10.0.0	Initial release

-  For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).
-  Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter provides a description of the System Console. The chapter lists Tcl commands and provides examples of how to use various applications in the System Console to analyze and debug your FPGA designs.

Introduction

The System Console performs low-level hardware debugging of either Qsys or SOPC Builder systems. You can use the System Console to debug systems that include IP cores instantiated in your Qsys or SOPC Builder system, as well as for initial bring-up of your printed circuit board, and for low-level testing. You access the System Console from the Tools menu of either Qsys or SOPC Builder. You can use the System Console in command-line mode or with the GUI.

- Command-line mode allows you to run Tcl scripts from the **Tcl Console** pane, located in the System Console window. You can also run Tcl scripts from the System Console GUI.
 - The System Console GUI allows you to view a main window with separate panes, including **System Explorer**, **Tcl Console**, **Messages**, and **Tools** panes.
- ❓ For more information about the System Console GUI, refer to *About System Console* in Quartus II Help.

This chapter contains the following sections:

- “System Console Overview”
- “Setting Up the System Console” on page 12–3
- “Interactive Help” on page 12–3
- “Using the System Console” on page 12–3
- “System Console Examples” on page 12–27
- “Device Support” on page 12–37

System Console Overview

The System Console allows you to use many types of services. When you interact with the Tcl console in the System Console, you have general commands related to finding and accessing instances of those services. Each service type has functions that are unique to that service.

Finding and Referring To Services

You can retrieve available service types with the `get_service_types` command.

The System Console uses a virtual file system to organize the available services, which is similar to the `/dev` location on Linux systems. Instances of services are referred to by their unique service path in the file system. You can retrieve service paths for a particular service with the command `get_service_paths <service-type>`.



Every System Console command requires a service path to identify the service instance you want to access.

Most System Console service instances are automatically discovered when you start the System Console. The System Console automatically scans for all JTAG-based service instances, and their services paths are immediately retrieved. Some other services, such as those connected by TCP/IP, are not automatically discovered. You can use the `add_service Tcl` command to inform the System Console about those services.

Accessing the Service Life Cycle

After you have a service path to a particular service instance, you can access the service for use.

The `open_service` command tells the System Console to start using a particular service instance. The `open_service` command works on every service type. The `open_service` command claims a service instance for exclusive use. The command does not tell the System Console which part of a service you are interested in. As such, service instances that you open are not safe for shared use among multiple users.

The `claim_service` command tells the System Console to start accessing a particular portion of a service instance. For example, if you use the master service to access memory, then use `claim_service` to tell the System Console that you only want to access the address space between `0x0` and `0x1000`. The System Console then allows other users to access other memory ranges and denies them access to your claimed memory range. The `claim_service` command returns a newly created service path that you can use to access your claimed resources.

Not all services support the `claim_service` command.

You can access a claim service after you open or claim it. When you finish accessing a service instance, use the `close_service` command to tell the System Console to free up resources.

Applying Services

The System Console provides extensive portfolios of services for various applications, such as real-time on-chip control and debugging, and system measurement. Examples of how to use these services are provided in this chapter. [Table 12-1](#) lists example applications included with the System Console and associated services.

The System Console functions by running Tcl commands that are described in [Table 12-3](#) through [Table 12-22](#).

Table 12–1. System Console Example Applications

Application	Services Used
Board Bring-Up	device, jtag_debug, sld
Processor Debug	processor, elf, bytestream, master
Active retrieval of dynamic information	master
Query static design information	marker, design
System Monitoring	monitor, master, dashboard
Transceiver Toolkit Direct Phy Control	transceiver_reconfig_analog, alt_xcvr_reconfig_dfe, alt_xcvr_reconfig_eye_viewer
Transceiver Toolkit System Level Control	transceiver_channel_rx, transceiver_channel_tx, transceiver_debug_link

Setting Up the System Console

You set up the System Console according to the hardware you have available in your system. You can access available debug IP on your system with the System Console. The debug IP allows you to access the running state of your system. The following sections discuss setting up and using debug IP in further detail.

“[System Console Examples](#)” on page 12–27 provides you with detailed examples of using the System Console in conjunction with SOPC Builder components.

-  Download the design files for the example designs from the [On-chip Debugging Design Examples](#) page on the Altera website.

These design examples demonstrate how to add debug IP blocks to your design and how to connect them before you can use the host application.

Interactive Help

Typing `help help` into the System Console lists all available commands. Typing `help <command name>` provides the syntax of individual commands. The System Console provides command completion if you type the beginning letters of a command and then press the Tab key.

-  The System Console interactive help commands only provide help for enabled services; consequently, typing `help help` does not display help for commands supplied by disabled plugins.

Using the System Console

The Quartus® II software expands the framework of the System Console by allowing you to start services for performing different types of tasks, as described in the following sections of this chapter. These sections provide Tcl scripting commands, arguments, and a brief description of the command functions.

 Many of the System Console commands do not work unless you connect to a system with a programming cable and with the proper debug IP.

Qsys and SOPC Builder Communications

You can use the System Console to help you debug Qsys or SOPC Builder systems. The System Console communicates with debug IP in your system design. You can instantiate debug IP cores using Qsys, SOPC Builder, or the MegaWizard Plug-In Manager.

 For more information about the Qsys system integration tool, refer to *System Design with Qsys* in volume 1 of the *Quartus II Handbook*.

Table 12-2 describes some of the IP cores that you can use to facilitate debugging your system with the System Console. When connected to the System Console, these components enable you to send commands and receive data.

Table 12-2. Qsys and SOPC Builder Components for Communication with the System Console (Note 1)

Component Name	Component Interface Types for Debugging
Nios® II processor with JTAG debug enabled	Components that include an Avalon® Memory-Mapped (Avalon-MM) slave interface. The JTAG debug module can also control the Nios II processor for debug functionality, including starting, stopping, and stepping the processor.
JTAG to Avalon master bridge	Components that include an Avalon-MM slave interface.
Avalon Streaming (Avalon-ST) JTAG Interface	Components that include an Avalon-ST interface.
JTAG UART	The JTAG UART is an Avalon-MM slave device that can be used in conjunction with the System Console to send and receive byte streams.
TCP/IP	For more information, refer to <i>AN624: Debugging with System Console over TCP/IP</i> .

Note to Table 12-2:

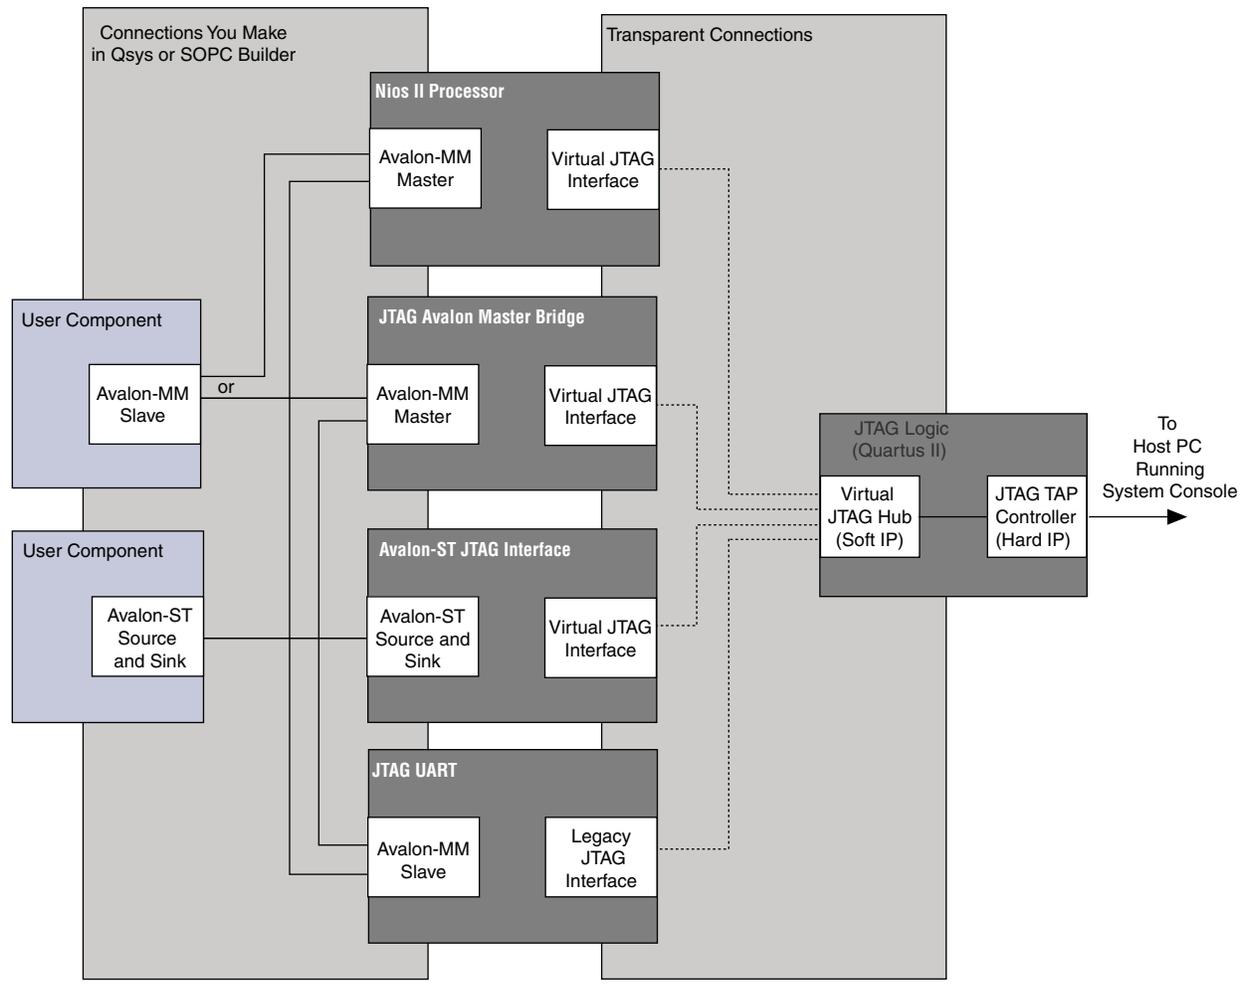
(1) The System Console can also send and receive byte streams from any system-level debugging (SLD) node whether it is instantiated in Qsys or SOPC Builder components provided by Altera, a custom component, or part of your Quartus II project; however, this approach requires detailed knowledge of the JTAG commands.

 For more information about Qsys and SOPC Builder components, refer to the following web pages and documents:

- [Nios II Processor](#) page of the Altera website
- *SPI Slave/JTAG to Avalon Master Bridge Cores* chapter in the *Embedded Peripherals IP User Guide*
- *Avalon Verification IP Suite User Guide*
- *Avalon-ST JTAG Interface Core* chapter in the *Embedded Peripherals IP User Guide*
- *Virtual JTAG (sld_virtual_jtag) Megafunction User Guide*
- *JTAG UART Core* chapter in the *Embedded Peripherals IP User Guide*
- *System Design with Qsys* in volume 1 of the *Quartus II Handbook*
- *About Qsys* in Quartus II Help

Figure 12-1 illustrates examples of interfaces of the components that the System Console can use.

Figure 12-1. Example Interfaces (Paths) the System Console Uses to Send Commands



Altera recommends that you include the following components in your system:

- On-chip memory
- JTAG UART
- System ID core

The System Console provides many different types of services. Different modules can provide the same type of service. For example, both the Nios II processor and the JTAG to Avalon Bridge master provide the master service; consequently, you can use the master commands to access both of these modules.



If your system includes a Nios II/f core with a data cache it may complicate the debugging process. If you suspect that writes to memory from the data cache at nondeterministic intervals are overwriting data written by the System Console, you can disable the cache of the Nios II/f core while debugging.

You can start the System Console from a Nios II command shell.

1. On the Windows Start menu, point to **All Programs**, then **Altera**, then **Nios II EDS <version>**, and then click **Nios II <version> Command Shell**.
2. To start the System Console, type the following command:

```
system-console ←
```

You can customize your System Console environment by adding commands to the configuration file called `system_console_rc.tcl`. This file can be located in either of the following places:

- `<quartus_install_dir>/sopc_builder/system_console_macros/system_console_rc.tcl`, known as the global configuration file, which affects all users of the system
- `<$HOME>/system_console/system_console_rc.tcl`, known as the user configuration file, which only affects the owner of that home directory

On startup, the System Console automatically runs any Tcl commands in these files. The commands in the global configuration file run first, followed by the commands in the user configuration file.

Console Commands

The console commands enable testing. You can use console commands to identify a module by its path, and to open and close a connection to it. The path that identifies a module is the first argument to most of the System Console commands. To exercise a module, follow these steps:

1. Identify a module by specifying the path to it, using the `get_service_paths` command.
2. Open a connection to the module using the `open_service` or `claim_service` command. (`claim_service` returns a new path for use).
3. Run Tcl and System Console commands to use a debug module that you insert to test another module of interest.
4. Close a connection to a module using the `close_service` command.

Table 12-3 describes the syntax of the console commands.

Table 12-3. Console Commands (Part 1 of 2)

Command	Arguments	Function
<code>get_service_types</code>	—	Returns a list of service types that the System Console manages. Examples of service types include master, bytestream, processor, sld, jtag_debug, device, and plugin.
<code>get_service_paths</code>	<code><service_type_name></code>	Returns a list of paths to nodes that implement the requested service type.
<code>open_service</code>	<code><service_type_name></code> <code><service_path></code>	Claims a service instance for exclusive use.

Table 12-3. Console Commands (Part 2 of 2)

Command	Arguments	Function
claim_service	<service-type> <service-path> <library-name> <claims>	Provides finer control of the portion of a service you want to use. Run <code>help claim_service</code> to get a list of <service-type>s. Then run <code>help claim_service <service-type></code> to get specific help on that service.
close_service	<service_type_name> <service_path>	Frees the System Console resources at the path you claimed.
is_service_open	<service_type_name> <service_path>	Returns 1 if the service type provided by the path is open, 0 if the service type is closed.
get_services_to_add	—	Returns a list of all services that are instantiable with the <code>add_service</code> command.
add_service	<service-type> <instance-name> <optional-parameters>	Adds a service of the specified service type with the given instance name. Run <code>get_services_to_add</code> to retrieve a list of instantiable services. This command returns the path where the service was added. Run <code>help add_service <service-type></code> to get specific help about that service-type, including any arguments that might be required for that service. Available services include: dashboard gdbserver nios2dpx pli_bytestream pli_master pli_packets tcp tcp_bytestream tcp_master transceiver_channel_rx transceiver_channel_tx transceiver_debug_link
get_version	—	Returns the current System Console version and build number.
add_help	<command> <help-text>	Adds help text for a given command. Use this when you write a Tcl script procedure (proc) and then want to provide help for others to use the script.
get_claimed_services	<library-name>	For the given library name, returns a list services claimed via <code>claim_service</code> with that library name. The returned list consists of pairs of paths and service types, each pair one claimed service.
refresh_connections	—	Scans for available hardware and updates the available service paths if there have been any changes.
send_message	<level> <message>	Sends a message of the given level to the message window. Available levels are info, warning, error, and debug.

Plugins

Plugins allow you to customize how you use the System Console services, and are enabled by default. [Table 12-4](#) lists Plugin commands.

Table 12-4. Plugin Commands

Command	Arguments	Function
Plugin Service Type Commands		
plugin_enable	<plugin-path>	Enables the plugin specified by the path. After a plugin is enabled, you can retrieve the <service-path> and <service_type_name> for additional services using the get_service_paths command.
plugin_disable	<plugin-path>	Disables the plugin specified by the path.
is_plugin_enabled	<plugin-path>	Returns a non-zero value when the plugin at the specified path is enabled.

Design Service Commands

Design Service commands allow you to use the System Console services to load and work with your design at a system level. [Table 12-5](#) lists Design Service commands.

Table 12-5. Design Service Commands

Command	Arguments	Function
design_load	<quartus-project-path> or <qpf-file-path>	Loads a model of a Quartus II design into the System Console. For example, if your .qpf file is in c:/projects/loopback , type the following command: design_load {c:\projects\loopback\}
design_instantiate	<design-path> <instance-name>	Allows you to apply the same design to multiple devices.
design_link	<design-instance-path> <device-service-path>	Links a Quartus II logical design with a physical device. For example, you can link a Quartus II design called 2c35_quartus_design to a 2c35 device. After you create this link, the System Console creates the appropriate correspondences between the logical and physical submodules of the Quartus II project. Example 12-4 on page 12-30 shows a transcript illustrating the design_load and design_link commands. Note that the System Console does not verify that the link is valid; if you create an incorrect link, the System Console does not report an error.

Data Pattern Generator Commands

Data Pattern Generator commands allow you control data patterns that you generate for testing, debugging, and optimizing your design. You must first insert debug IP to enable these commands. [Table 12-6](#) lists Data Pattern Generator commands.

Table 12-6. Data Pattern Generator Commands (Part 1 of 2)

Command	Arguments	Function
<code>data_pattern_generator_start</code>	<code><service-path></code>	Starts the data pattern generator.
<code>data_pattern_generator_stop</code>	<code><service-path></code>	Stops the data pattern generator.
<code>data_pattern_generator_is_generating</code>	<code><service-path></code>	Returns non-zero if the generator is running.
<code>data_pattern_generator_inject_error</code>	<code><service-path></code>	Injects a 1-bit error into the generator's output.
<code>data_pattern_generator_set_pattern</code>	<code><service-path></code> <code><pattern-name></code>	Sets the output pattern specified by the <code><pattern-name></code> . In all, 6 patterns are available, 4 are pseudo-random binary sequences (PRBS), 1 is high frequency and 1 is low frequency. The following pattern names are defined: <ul style="list-style-type: none"> ■ PRBS7 ■ PRBS15 ■ PRBS23 ■ PRBS31 ■ HF—outputs a high frequency, constant pattern of alternating 0s and 1s ■ LF—outputs a low frequency, constant pattern of 10b'1111100000 for 10-bit symbols and 8b'11110000 for 8-bit symbols
<code>data_pattern_generator_get_pattern</code>	<code><service-path></code>	Returns currently selected output pattern.
<code>data_pattern_generator_get_available_patterns</code>	<code><service-path></code>	Returns a list of available data patterns by name.
<code>data_pattern_generator_enable_preamble</code>	<code><service-path></code>	Enables the preamble mode at the beginning of generation.
<code>data_pattern_generator_disable_preamble</code>	<code><service-path></code>	Disables the preamble mode at the beginning of generation.
<code>data_pattern_generator_is_preamble_enabled</code>	<code><service-path></code>	Returns a non-zero value if preamble mode is enabled.
<code>data_pattern_generator_set_preamble_word</code>	<code><service-path></code> <code><preamble-word></code>	Sets the preamble word (could be 32-bit or 40-bit).
<code>data_pattern_generator_get_preamble_word</code>	<code><service-path></code>	Gets the preamble word.

Table 12-6. Data Pattern Generator Commands (Part 2 of 2)

Command	Arguments	Function
<code>data_pattern_generator_set_preamble_beats</code>	<code><service-path></code> <code><number-of-preamble-beats></code>	Sets the number of beats to send out the in the preamble word.
<code>data_pattern_generator_get_preamble_beats</code>	<code><service-path></code>	Returns the currently set number of beats to send out in the preamble word.

Data Pattern Checker Commands

Data Pattern Checker commands allow you verify the data patterns that you generate. You must first insert debug IP to enable these commands. [Table 12-7](#) lists Data Pattern Checker commands.

Table 12-7. Data Pattern Checker Commands

Command	Arguments	Function
<code>data_pattern_checker_start</code>	<code><service-path></code>	Starts the data pattern checker.
<code>data_pattern_checker_stop</code>	<code><service-path></code>	Stops the data pattern checker.
<code>data_pattern_checker_is_checking</code>	<code><service-path></code>	Returns a non-zero value if the checker is running.
<code>data_pattern_checker_is_locked</code>	<code><service-path></code>	Returns non-zero if the checker is locked onto the incoming data.
<code>data_pattern_checker_set_pattern</code>	<code><service-path></code> <code><pattern-name></code>	Sets the expected pattern to the one specified by the <code><pattern-name></code> .
<code>data_pattern_checker_get_pattern</code>	<code><service-path></code>	Returns the currently selected expected pattern by name.
<code>data_pattern_checker_get_available_patterns</code>	<code><service-path></code>	Returns a list of available data patterns by name.
<code>data_pattern_checker_get_data</code>	<code><service-path></code>	Returns the correct and incorrect counts of data received, providing the number of bits and the number of errors.
<code>data_pattern_checker_reset_counters</code>	<code><service-path></code>	Resets the bit and error counters inside the checker.

Programmable Logic Device (PLD) Commands

The PLD commands provide access to programmable logic devices on your board. Before you use these commands, identify the path to the programmable logic device on your board using the `get_service_paths` command described in [Table 12-3](#).

Table 12-8 describes the PLD commands.

Table 12-8. PLD Commands

Command	Arguments	Function
device_download_sof	<device_path> <sof_file>	Loads the specified SRAM object file (.sof) file to the device specified by the path.
device_load_jdi	<device_path> <jdi_file>	Renames the Tcl interface layer's nodes to the names specified in the JTAG debug interface (.jdi) file, making your design easier to understand.

Board Bring-Up Commands

The board bring-up commands allow you to test your system. These commands are presented in the order that you would use them during board bring-up, including the following setup work flow:

1. Verify JTAG connectivity.
2. Verify the clock and reset signals.
3. Verify memory and other peripheral interfaces.
4. Verify basic Nios II processor functionality.



The System Console is intended for debugging the basic hardware functionality of your Nios II processor, including its memories and pinout. If you are writing device drivers, you may want to use the System Console and the Nios II software build tools together to debug your code.



For more information about the hardware functioning correctly and software debugging, refer to *Nios II Software Build Tools Reference* in the *Nios II Software Developer's Handbook*.

JTAG Debug Commands

You can use JTAG debug commands to verify the functionality and signal integrity of your JTAG chain. Your JTAG chain must be functioning correctly to debug the rest of your system. To verify signal integrity of your JTAG chain, Altera recommends that you provide an extensive list of byte values. Table 12-9 lists these commands.

Table 12-9. JTAG Commands

Command	Arguments	Function
jtag_debug_loop	<path> <list_of_byte_values>	Loops the specified list of bytes through a loopback of tdi and tdo of a system-level debug (SLD) node. Returns the list of byte values in the order that they were received. Blocks until all bytes are received. Byte values are given with the 0x (hexadecimal) prefix and delineated by spaces.
jtag_debug_reset_system	<service-path>	Issues a reset request to the system.

Clock and Reset Signal Commands

The next stage of board bring-up tests the clock and reset signals. Table 12-10 lists the three commands to verify these signals. Use these commands to verify that your clock is toggling and that the reset signal has the expected value.

Table 12-10. Clock and Reset Commands

Command	Argument	Function
<code>jtag_debug_sample_clock</code>	<code><service-path></code>	Returns the value of the clock signal of the system clock that drives the module's system interface. The clock value is sampled asynchronously; consequently, you may need to sample the clock several times to guarantee that it is toggling.
<code>jtag_debug_sample_reset</code>	<code><service-path></code>	Returns the value of the reset signal of the system reset that drives the module's system interface.
<code>jtag_debug_sense_clock</code>	<code><service-path></code>	Returns the result of a sticky bit that monitors for system clock activity. If the clock has ever toggled, the bit is 1. Returns <code>true</code> if the bit has ever toggled and otherwise returns <code>false</code> . The sticky bit is reset to 0 on read.

Avalon-MM Commands

Avalon-MM commands allow you to access the memory-mapped or SLD modules that include Avalon-ST slave interfaces. You can read or write to the Avalon-MM interfaces using the master read and write commands, but the Avalon-MM slave must be connected to an Avalon MM master that you can control from the host.

PLI Master commands listed in Table 12-13 provide bytestream and master services over a PLI connection to a simulator. Tcl Master commands provide bytestream and master services over TCP IP.

Additionally, you can use the SLD commands to shift values into the instruction and data registers of SLD nodes and read the previous value. Table 12-11 lists these commands.

Table 12-11. Module Commands (Part 1 of 2) (Note 1)

Command	Arguments	Function
Avalon-MM Master Commands		
<code>master_write_memory</code>	<code><service-path></code> <code><address></code> <code><list_of_byte_values></code>	Writes the specified data to the specified path and address. Values are given in hexadecimal format with the 0x prefix and delineated by spaces.
<code>master_write_8</code>	<code><service-path></code> <code><address></code> <code><list_of_byte_values></code>	
<code>master_write_16</code>	<code><service-path></code> <code><address></code> <code><list_of_16_bit_words></code>	
<code>master_write_32</code>	<code><service-path></code> <code><address></code> <code><list_of_32_bit_words></code>	

Table 12-11. Module Commands (Part 2 of 2) (Note 1)

Command	Arguments	Function
master_read_memory	<service-path> <base_address> <size_in_bytes>	Returns a list of read values.
master_read_8	<service-path> <base_address> <size_in_bytes>	
master_read_16	<service-path> <base_address> <size_in_multiples_of_16_bits>	
master_read_32	<service-path> <base_address> <size_in_multiples_of_32_bits>	
SLD Commands		
sld_access_ir	<service-path> <value> <timeout> (in μ s)	Shifts the instruction value into the instruction register of the specified node. Returns the previous value of the instruction. If the timeout value is set to 0, the operation never times out. A suggested starting value for timeout is 1000 μ s.
sld_access_dr	<service-path> <size_in_bits> <timeout> (in μ s), <list_of_byte_values>	Shifts the byte values into the data register of the SLD node up to the size in bits specified. If the timeout value is set to 0, the operation never times out. Returns the previous contents of the data register. A suggested starting value for timeout is 1000 μ s.
sld_lock	sld_lock <service-path> <timeout> (in ms)	Locks the SLD chain to guarantee exclusive access. If the SLD chain is already locked, tries for <timeout> ms before returning -1, indicating an error. Returns 0 if successful.
sld_unlock	sld_unlock <service-path>	Unlocks the SLD chain.

Note to Table 12-11:

(1) Transfers performed in 16- and 32-bit sizes are packed in little endian format.

Processor Commands

These commands allow you to start, stop, and step through software running on a Nios II processor. The commands also allow you to read and write the registers of the processor. [Table 12-12](#) lists the commands.

Table 12-12. Processor Commands

Command	Arguments	Function
elf_download	<i><processor-service-path></i> <i><master-service-path></i> <i><elf-file-path></i>	Downloads the given Executable and Linking Format (.elf) File to memory using the specified master service. Sets the processor's program counter to the .elf entry point.
processor_in_debug_mode	<i><service-path></i>	Returns a non-zero value if the processor is in debug mode.
processor_reset	<i><service-path></i>	Resets the processor and places it in debug mode.
processor_run	<i><service-path></i>	Puts the processor into existing debug mode.
processor_stop	<i><service-path></i>	Puts the processor into debug mode.
processor_step	<i><service-path></i>	Executes one assembly instruction.
processor_get_register_names	<i><service-path></i>	Returns a list with the names of all of the processor's accessible registers.
processor_get_register	<i><service-path></i> <i><register_name></i>	Returns the value of the specified register.
processor_set_register	<i><service-path></i> <i><register_name></i>	Sets the value of the specified register.

Bytestream Commands

These commands provide access to modules that produce or consume a stream of bytes. You can use the bytestream service to communicate directly to IP that provides bytestream interfaces, such as the Altera JTAG UART. [Table 12-13](#) lists the commands.

Table 12-13. Bytestream Commands

Command	Arguments	Function
bytestream_send	<i><service-path></i> <i><list_of_byte_values></i>	Sends the list of byte values on the specified path. Values are in hexadecimal format and delineated by spaces.
bytestream_receive	<i><service-path></i> <i><number_of_bytes></i>	Limits the maximum number of bytes to return.

Transceiver Toolkit Commands

You can run the Transceiver Toolkit from the System Console window. Alternatively, you can open the Transceiver Toolkit from the Tools menu in the Quartus II software. You can debug, monitor, and optimize the transceiver channel links in your design with Tcl scripts, or you can launch the Transceiver Toolkit GUI from the System Console Tools menu. The GUI for the Transceiver Toolkit provides a graphical representation of automatic tests that you run. You can also view graphical control panels to change physical medium attachment (PMA) settings of channels, start and stop generators, and checkers.

 For further information about the Transceiver Toolkit, refer to the *Transceiver Link Debugging Using the System Console* chapter of the *Quartus II Handbook*.

Table 12-14 through Table 12-19 lists the Transceiver Toolkit commands..

Table 12-14. Transceiver Toolkit Channel_rx Commands (Part 1 of 2)

Command	Arguments	Function
transceiver_channel_rx_get_data	<service-path>	Returns a list of the current checker data. The results are in the order: number of bits, number of errors, bit error rate.
transceiver_channel_rx_get_dcgain	<service-path>	Gets the DC gain value on the receiver channel.
transceiver_channel_rx_get_dfe_tap_value	<service-path> <tap position>	Gets the current tap value of the specified channel at the specified tap position.
transceiver_channel_rx_get_eqctrl	<service-path>	Gets the equalization control value on the receiver channel.
transceiver_channel_rx_get_pattern	<service-path>	Returns the current data checker pattern by name.
transceiver_channel_rx_has_dfe	<service-path>	Gets whether this channel has the DFE feature available.
transceiver_channel_rx_is_checking	<service-path>	Returns non-zero if the checker is running.
transceiver_channel_rx_is_dfe_enabled	<service-path>	Gets whether the DFE feature is enabled on the specified channel.
transceiver_channel_rx_is_locked	<service-path>	Returns non-zero if the checker is locked onto the incoming data.
transceiver_channel_rx_reset_counters	<service-path>	Resets the bit and error counters inside the checker.
transceiver_channel_rx_set_dcgain	<service-path> <dc_gain setting>	Sets the DC gain value on the receiver channel.
transceiver_channel_rx_set_dfe_enabled	<service-path> <disable(0)/enable(1)>	Enables/disables the DFE feature on the specified channel.

Table 12-14. Transceiver Toolkit Channel_rx Commands (Part 2 of 2)

Command	Arguments	Function
transceiver_channel_rx_set_dfe_tap_value	<service-path> <tap position> <tap value>	Sets the current tap value of the specified channel at the specified tap position to the specified value.
transceiver_channel_rx_set_eqctrl	<service-path> <eqctrl setting>	Sets the equalization control value on the receiver channel.
transceiver_channel_rx_start_checking	<service-path>	Starts the checker.
transceiver_channel_rx_stop_checking	<service-path>	Stops the checker.
transceiver_channel_rx_get_eyeq_phase_step	<service-path>	Gets the current phase step of the specified channel.
transceiver_channel_rx_set_pattern	<service-path> <pattern-name>	Sets the expected pattern to the one specified by the pattern name.
transceiver_channel_rx_is_eyeq_enabled	<service-path>	Gets whether the EyeQ feature is enabled on the specified channel.
transceiver_channel_rx_set_eyeq_enabled	<service-path> <disable(0)/enable(1)>	Enables/disables the EyeQ feature on the specified channel.
transceiver_channel_rx_set_eyeq_phase_step	<service-path> <phase step>	Sets the phase step of the specified channel.
transceiver_channel_rx_set_word_aligner_enabled	<service-path> <disable(0)/enable(1)>	Enables/disables the word aligner of the specified channel.
transceiver_channel_rx_is_word_aligner_enabled	<service-path> <disable(0)/enable(1)>	Enables/disables the word aligner of the specified channel.

Table 12-15. Transceiver Toolkit Channel_tx Commands (Part 1 of 2)

Command	Arguments	Function
transceiver_channel_tx_disable_preamble	<service-path>	Disables the preamble mode at the beginning of generation.
transceiver_channel_tx_enable_preamble	<service-path>	Enables the preamble mode at the beginning of generation.
transceiver_channel_tx_get_number_of_preamble_beats	<service-path>	Returns the number of preamble beats.
transceiver_channel_tx_get_pattern	<service-path>	Returns the currently set pattern.
transceiver_channel_tx_get_preamble_word	<service-path>	Returns the currently set preamble word.
transceiver_channel_tx_get_preemph0t	<service-path>	Gets the pre-emphasis pre-tap value on the transmitter channel.
transceiver_channel_tx_get_preemph1t	<service-path>	Gets the pre-emphasis first post-tap value on the transmitter channel.

Table 12-15. Transceiver Toolkit Channel_tx Commands (Part 2 of 2)

Command	Arguments	Function
transceiver_channel_tx_get_preemph2t	<service-path>	Gets the pre-emphasis second post-tap value on the transmitter channel.
transceiver_channel_tx_get_vodctrl	<service-path>	Gets the VOD control value on the transmitter channel.
transceiver_channel_tx_inject_error	<service-path>	Injects a 1-bit error into the generator's output.
transceiver_channel_tx_is_generating	<service-path>	Returns non-zero if the generator is running.
transceiver_channel_tx_is_preamble_enabled	<service-path>	Returns non-zero if preamble mode is enabled.
transceiver_channel_tx_set_number_of_preamble_beats	<service-path> <number-of-preamble-beats>	Sets the number of beats to send out the preamble word.
transceiver_channel_tx_set_pattern	<service-path> <pattern-name>	Sets the output pattern to the one specified by the pattern name.
transceiver_channel_tx_set_preamble_word	<service-path> <preamble-word>	Sets the preamble word to be sent out.
transceiver_channel_tx_set_preemph0t	<service-path> <preemph0t value>	Sets the pre-emphasis pre-tap value on the transmitter channel.
transceiver_channel_tx_set_preemph1t	<service-path> <preemph1t value>	Sets the pre-emphasis first post-tap value on the transmitter channel.
transceiver_channel_tx_set_preemph2t	<service-path> <preemph2t value>	Sets the pre-emphasis second post-tap value on the transmitter channel.
transceiver_channel_tx_set_vodctrl	<service-path> <vodctrl value>	Sets the VOD control value on the transmitter channel.
transceiver_channel_tx_start_generation	<service-path>	Starts the generator.
transceiver_channel_tx_stop_generation	<service-path>	Stops the generator.

Table 12-16. Transceiver Toolkit Debug_Link Commands

Command	Arguments	Function
transceiver_debug_link_get_pattern	<service-path>	Gets the currently set pattern the link uses to run the test.
transceiver_debug_link_is_running	<service-path>	Returns non-zero if the test is running on the link.
transceiver_debug_link_set_pattern	<service-path> <data pattern>	Sets the pattern the link uses to run the test.
transceiver_debug_link_start_running	<service-path>	Starts running a test with the currently selected test pattern.
transceiver_debug_link_stop_running	<service-path>	Stops running the test.

Table 12-17. Transceiver Toolkit Reconfig_analog Commands (Part 1 of 2)

Command	Arguments	Function
transceiver_reconfig_analog_get_logical_channel_address	<service-path>	Gets the transceiver logical channel address currently set.
transceiver_reconfig_analog_get_rx_dcgain	<service-path>	Gets the DC gain value on the receiver channel specified by the current logical channel address.
transceiver_reconfig_analog_get_rx_eqctrl	<service-path>	Gets the equalization control value on the receiver channel specified by the current logical channel address.
transceiver_reconfig_analog_get_tx_preemph0t	<service-path>	Gets the pre-emphasis pre-tap value on the transmitter channel specified by the current logical channel address.
transceiver_reconfig_analog_get_tx_preemph1t	<service-path>	Gets the pre-emphasis first post-tap value on the transmitter channel specified by the current logical channel address.
transceiver_reconfig_analog_get_tx_preemph2t	<service-path>	Gets the pre-emphasis second post-tap value on the transmitter channel specified by the current logical channel address.
transceiver_reconfig_analog_get_tx_vodctrl	<service-path>	Gets the VOD control value on the transmitter channel specified by the current logical channel address.
transceiver_reconfig_analog_set_logical_channel_address	<service-path> <logical channel address>	Sets the transceiver logical channel address.
transceiver_reconfig_analog_set_rx_dcgain	<service-path> <dc_gain value>	Sets the transceiver logical channel address.
transceiver_reconfig_analog_set_rx_eqctrl	<service-path> <eqctrl value>	Sets the equalization control value on the receiver channel specified by the current logical channel address.
transceiver_reconfig_analog_set_tx_preemph0t	<service-path> <preemph0t value>	Sets the pre-emphasis pre-tap value on the transmitter channel specified by the current logical channel address.
transceiver_reconfig_analog_set_tx_preemph1t	<service-path> <preemph1t value>	Sets the pre-emphasis first post-tap value on the transmitter channel specified by the current logical channel address.

Table 12-17. Transceiver Toolkit Reconfig_analog Commands (Part 2 of 2)

Command	Arguments	Function
transceiver_reconfig_analog_set_tx_preemph2t	<service-path> <preemph2t value>	Sets the pre-emphasis second post-tap value on the transmitter channel specified by the current logical channel address.
transceiver_reconfig_analog_set_tx_vodctrl	<service-path> <vodctrl value>	Sets the VOD control value on the transmitter channel specified by the current logical channel address.

Table 12-18. Transceiver Toolkit DFE Feedback Equalization (DFE) Tcl Commands

Command	Arguments	Function
alt_xcvr_reconfig_dfe_get_logical_channel_address	<service-path>	Gets the logical channel address that other alt_xcvr_reconfig_dfe commands use to apply.
alt_xcvr_reconfig_dfe_is_enabled	<service-path>	Gets whether the DFE feature is enabled on the previously specified channel.
alt_xcvr_reconfig_dfe_set_enabled	<service-path> <disable(0)/enable(1)>	Enables/disables the DFE feature on the previously specified channel.
alt_xcvr_reconfig_dfe_set_logical_channel_address	<service-path> <tap position>	Gets the tap value of the previously specified channel at specified tap position.
alt_xcvr_reconfig_dfe_set_logical_channel_address	<service-path> <logical channel address>	Sets the logical channel address that other alt_xcvr_reconfig_eye_viewer commands use to apply.
alt_xcvr_reconfig_dfe_set_tap_value	<service-path> <tap position> <tap value>	Sets the tap value at the previously specified channel at specified tap position and value.

Table 12-19. Transceiver Toolkit Eye Monitor Tcl Commands (Part 1 of 2)

Command	Arguments	Function
alt_xcvr_custom_is_word_aligner_enabled	<service-path> <disable(0)/enable(1)>	Enables/disables the word aligner of the previously specified channel.
alt_xcvr_custom_set_word_aligner_enabled	<service-path> <disable(0)/enable(1)>	Enables/disables the word aligner of the previously specified channel.
alt_xcvr_reconfig_eye_viewer_get_logical_channel_address	<service-path>	Gets the logical channel address on which other alt_reconfig_eye_viewer commands will use to apply.

Table 12-19. Transceiver Toolkit Eye Monitor Tcl Commands (Part 2 of 2)

Command	Arguments	Function
alt_xcvr_reconfig_eye_viewer_get_phase_step	<service-path>	Gets the current phase step of the previously specified channel.
alt_xcvr_reconfig_eye_viewer_is_enabled	<service-path>	Gets whether the EyeQ feature is enabled on the previously specified channel.
alt_xcvr_reconfig_eye_viewer_set_enabled	<service-path> <disable(0)/enable(1)>	Enables/disables the EyeQ feature on the previously specified channel.
alt_xcvr_reconfig_eye_viewer_set_logical_channel_address	<service-path> <logical channel address>	Sets the logical channel address on which other alt_reconfig_eye_viewer commands will use to apply.
alt_xcvr_reconfig_eye_viewer_set_phase_step	<service-path> <phase step>	Sets the phase step of the previously specified channel.

In-System Sources and Probes Commands

You can use the In-System Sources and Probes commands to read source and probe data. [Table 12-20](#) lists the commands. You use these commands with the In-System Sources and Probes that you insert into your project from the Quartus II software main menu.

Table 12-20. In-System Sources and Probes Tcl Commands

Command	Arguments	Function
issp_get_instance_info	<service-path>	Returns a list of the configurations of the In-System Sources and Probes instance.
issp_read_probe_data	<service-path>	Retrieves the current value of the probes. A hex string is returned representing the probe port value.
issp_read_source_data	<service-path>	Retrieves the current value of the sources. A hex string is returned representing the source port value.
issp_write_source_data	<service-path> <source-value>	Sets values for the sources. The value can be either a hex string or a decimal value supported by System Console Tcl interpreter.

Monitor Commands

You can use the Monitor commands to efficiently read many Avalon-MM slave memory locations at a regular interval. For example, if you want to do 100 reads per second, every second, you get much better performance using the monitor service than if you call 100 separate `master_read_memory` commands every second. This is the primary difference between the monitor service and the master service. [Table 12-21](#) lists the commands.

Table 12-21. Monitoring Commands (Part 1 of 2)

Command	Arguments	Function
<code>monitor_add_range</code>	<code><service-path></code> <code><target-path></code> <code><address></code> <code><size></code>	Adds a contiguous memory addresses into the monitored memory list.
<code>monitor_set_callback</code>	<code><service-path></code> <code><Tcl-expression></code>	Defines a Tcl expression in a single string that will be evaluated after all the memories monitored by this service are read. Typically, this expression should be specified as a Tcl procedure call with necessary argument passed in.
<code>monitor_set_interval</code>	<code><service-path></code> <code><interval></code>	Specifies the frequency of the polling action by specifying the interval between two memory reads. The actual polling frequency varies depending on the system activity. The monitor service will try it keep it as close to this specification as possible.
<code>monitor_get_interval</code>	<code><service-path></code>	Returns the current interval set which specifies the frequency of the polling action.
<code>monitor_set_enabled</code>	<code><service-path></code> <code><enable(1)/disable(0)></code>	Enables/disables monitoring. Memory read starts after this is enabled, and Tcl callback is evaluated after data is read.
<code>monitor_read_data</code>	<code><service-path></code> <code><target-path></code> <code><address></code> <code><size></code>	Returns a list of 8-bit values read from the most recent values read from device. The memory range specified must be within the monitored memory range as defined by <code>monitor_add_range</code> .

Table 12–21. Monitoring Commands (Part 2 of 2)

Command	Arguments	Function
<code>monitor_read_all_data</code>	<code><service-path></code> <code><target-path></code> <code><address></code> <code><size></code>	Returns a list of 8-bit values read from all recent values read from device since last Tcl callback. The memory range specified must be within the monitored memory range as defined by <code>monitor_add_range</code> .
<code>monitor_get_read_interval</code>	<code><service-path></code> <code><target-path></code> <code><address></code> <code><size></code>	Returns the number of milliseconds between last two data reads returned by <code>monitor_read_data</code> .
<code>monitor_get_all_read_intervals</code>	<code><service-path></code> <code><target-path></code> <code><address></code> <code><size></code>	Returns a list of intervals in milliseconds between two reads within the data returned by <code>monitor_read_all</code> .
<code>monitor_get_missing_event_count</code>	<code><service-path></code>	Returns the number of callback events missed during the evaluation of last Tcl callback expression.

Dashboard Commands

The System Console dashboard allows you to create graphical tools that seamlessly integrate into the System Console. This section describes how to build your own dashboard with Tcl commands and the properties that you can assign to the widgets on your dashboard. The dashboard allows you to create tools that interact with live instances of an IP core on your device. [Table 12–22](#) lists the dashboard Tcl commands available from the System Console.

[Example 12–1](#) shows a Tcl command to create a dashboard. After you run the command, you get a path. You can then use the path on the commands listed in [Table 12–22](#).

Example 12–1. Example of Creating a Dashboard

```
add_service dashboard my_new_dashboard "This is a New Dashboard" "Tools/My New Dashboard"
```

Table 12–22. Dashboard Commands (Part 1 of 2)

Command	Arguments	Description
<code>dashboard_add</code>	<code><service-path></code> <code><string></code>	Allows you to add a specified widget to your GUI dashboard.
<code>dashboard_remove</code>	<code><service-path></code> <code><string></code>	Allows you to remove a specified widget from your GUI dashboard.
<code>dashboard_set_property</code>	<code><string></code>	Allows you to set the specified properties of the specified widget that has been added to your GUI dashboard.

Table 12–22. Dashboard Commands (Part 2 of 2)

Command	Arguments	Description
<code>dashboard_get_property</code>	<code><service-path></code> <code><string></code>	Allows you to determine the existing properties of a widget added to your GUI dashboard.
<code>dashboard_get_types</code>	—	Returns a list of all possible widgets that you can add to your GUI dashboard.
<code>dashboard_get_properties</code>	—	Returns a list of all possible properties of the specified widgets in your GUI dashboard.

Specifying Widgets

You can specify the widgets that you add to your dashboard. [Table 12–23](#) lists the widgets.

Table 12–23. Dashboard Widgets

Widget	Description
Group	Allows you to add a collection of widgets and control the general layout of the widgets.
Button	Allows you to add a button.
FileChooserButton	Allows you to define button actions.
Label	Allows you to add a text string.
Text	Allows you to specify text.
Table	Allows you to add a table.
LED	Allows you to add an LED with a label.
Dial	Allows you to add the shape of an analog dial.
TimeChart	Allows you to add a chart of historic values, with the X-axis of the chart representing time.
BarChart	Allows you to add a bar chart.
LineChart	Allows you to add a line chart.
PieChart	Allows you to add a pie chart.

[Example 12–2](#) is a Tcl script to instantiate a widget. In this example, the Tcl command adds a label to the dashboard. The first argument is the path to the dashboard. This path is returned by the `add_service` command. The next argument is the ID you assign to the widget. The ID must be unique within the dashboard. You use this ID later on to refer to the widget.

Following that argument is the type of widget you are adding, which in this example is a `label`. The type of labels in this example is `label`. The last argument to this command is the group where you want to put this widget. In this example, a special keyword `self` is used. `self` refers to the dashboard itself, the primary group. You can then add a group to `self`, which allows you to add other widgets to this group by using the ID of the new group, rather than using the ID of the `self` group.

Example 12–2. Example of Instantiating a Widget

```
dashboard_add $dash myLabel label self
```

Customizing Widgets

You can change widget properties at any time. The `dashboard_set_property` command allows you to interact with the widgets you instantiate. This functionality is most useful when you use you change part of the execution of a callback. [Example 12-3](#) shows how to change the text in a label.

In [Example 12-3](#), the first argument is the path to the dashboard. Next is the unique ID of the widget, which then allows you to target an arbitrary widget. Following that is the name of the property. Each type of widget has a defined set of properties, discussed later. You can change the properties. In this example, `myLabel` is of the type `label`, and the example shows how to set its `text` property. The last argument is the value that the property takes when the command is executed.

Example 12-3. Example of Customizing a Widget

```
dashboard_set_property $dash myLabel text "Hello World!"
```

Assigning Dashboard Widget Properties

In [Table 12-24](#) through [Table 12-36](#), the various properties are listed that you can apply to the widgets on your dashboard.

Table 12-24. Properties Common to All Widgets

Property	Description
<code>enabled</code>	Enables or disables the widget.
<code>expandableX</code>	Allows the widget to be resized horizontally if there's space available in the cell where it resides.
<code>expandableY</code>	Allows the widget to be resized vertically if there's space available in the cell where it resides.
<code>maxHeight</code>	If the widget's <code>expandableY</code> is set, this is the maximum height in pixels that the widget can take.
<code>minHeight</code>	If the widget's <code>expandableY</code> is set, this is the minimum height in pixels that the widget can take.
<code>maxWidth</code>	If the widget's <code>expandableX</code> is set, this is the maximum width in pixels that the widget can take.
<code>minWidth</code>	If the widget's <code>expandableX</code> is set, this is the minimum width in pixels that the widget can take.
<code>preferredHeight</code>	The height of the widget if <code>expandableY</code> is not set.
<code>preferredWidth</code>	The width of the widget if <code>expandableX</code> is not set.
<code>toolTip</code>	A tool tip string that appears once the mouse hovers above the widget.

Table 12-25. Button Properties

Property	Description
<code>onClick</code>	A Tcl command to run, usually a <code>proc</code> , every time the button is clicked.
<code>text</code>	The text on the button.

Table 12–26. FileChooserButton Properties

Property	Description
text	The text on the button.
onChoose	A Tcl command to run, usually a <code>proc</code> , every time the button is clicked.
title	The text of file chooser dialog box title.
chooserButtonText	The text of file chooser dialog box approval button. By default, it is "Open".
filter	The file filter based on extension. Only one extension is supported. By default, all file names are allowed. The filter is specified as [list filter_description file_extension], for example [list "Text Document (.txt)" ".txt"].
mode	Specifies what kind of files or directories can be selected. "files_only", by default. Possible options are "files_only" and "directories_only".
multiSelectionEnabled	Controls whether multiple files can be selected. False, by default.
paths	Returns a list of file paths selected in the file chooser dialog box. This is read only property. It is most useful when used within the <code>onClick</code> script or a procedure when the result is freshly updated after the dialog box closes.

Table 12–27. Dial Properties

Properties	Description
max	The maximum value that the dial can show.
min	The minimum value that the dial can show.
tickSize	The space between the different tick marks of the dial.
title	The title of the dial.
value	The value that the dial's needle should mark. It must be between min and max.

Table 12–28. Group Properties

Properties	Description
itemsPerRow	The number of widgets the group can position in one row, from left to right, before moving to the next row.
title	The title of the group. Groups with a title can have a border around them, and setting an empty title removes the border.

Table 12–29. Label Properties

Properties	Description
text	The text to show in the label.

Table 12–30. LED Properties

Properties	Description
color	The color of the LED. The options are: <code>red_off</code> , <code>red</code> , <code>yellow_off</code> , <code>yellow</code> , <code>green_off</code> , <code>green</code> , <code>blue_off</code> , <code>blue</code> , and <code>black</code> .
text	The text to show next to the LED.

Table 12-31. Text Properties

Properties	Description
editable	Controls whether the text box is editable.
htmlCapable	Controls whether the text box can format HTML.
text	Controls the text of the textbox.

Table 12-32. Time Chart Properties

Properties	Description
labelX	The label for the X axis.
labelY	The label for the Y axis.
latest	The latest value in the series.
maximumItemCount	The number of sample points to display in the historic record.
title	The title of the chart.

Table 12-33. Table Properties (Part 1 of 2)

Properties	Description
Table-wide Properties	
columnCount	The number of columns (Mandatory) (0, by default).
rowCount	The number of rows (Mandatory) (0, by default).
headerReorderingAllowed	Controls whether you can drag the columns (false, by default).
headerResizingAllowed	Controls whether you can resize all column widths. (false, by default) Note, each column can be individually configured to be resized by using property, columnWidthResizable.
rowSorterEnabled	Controls whether you can sort the cell values in a column (false, by default).
showGrid	Controls whether to draw both horizontal and vertical lines (true, by default).
showHorizontalLines	Controls whether to draw horizontal line (true, by default).
showVerticalLines	Controls whether to draw vertical line (true, by default).
rowIndex	Current row index. 0 based. This value affects some properties below (0, by default).
columnIndex	Current column index. Zero-based. This value affects all column specific properties below (0, by default).
cellText	Specifies the text to be filled in the cell specified the current rowIndex and columnIndex (Empty, by default).
selectedRows	Control or retrieve row selection.
Column-specific Properties	
columnHeader	The text to be filled in the column header.
columnHorizontalAlignment	The cell text alignment in the specified column. Supported types are "leading"(default), "left", "center", "right", "trailing".
columnRowSorterType	The type of sorting method used. This is applicable only if rowSorterEnabled is true. Each column has its own sorting type. Supported types are "string" (default), "int", and "float".

Table 12-33. Table Properties (Part 2 of 2)

Properties	Description
columnWidth	The number of pixels used for the column width.
columnWidthResizable	Controls whether the column width is resizable by you (false, by default).

Table 12-34. BarChart Properties

Properties	Description
title	Chart title.
labelX	X axis label text.
labelY	Y axis label text.
range	Y axis value range. By default, it is auto range. Range is specified in a Tcl list, for example [list lower_numerical_value upper_numerical_value].
itemValue	Item value. Value is specified in a Tcl list, for example [list bar_category_str numerical_value].

Table 12-35. Linechart Properties

Properties	Description
title	Chart title.
labelX	Axis X label text.
labelY	Axis Y label text.
range	Axis Y value range. By default, it is auto range. Range is specified in a Tcl list, for example [list lower_numerical_value upper_numerical_value].
itemValue	Item value. Value is specified in a Tcl list, for example [list bar_category_str numerical_value].

Table 12-36. PieChart Properties

Properties	Description
title	Chart title.
itemValue	Item value. Value is specified in a Tcl list, for example [list bar_category_str numerical_value].

System Console Examples

This section provides an example of how to load and link a design in the System Console, as well as three SOPC Builder system examples that show you how to use the System Console. The **System-Console.zip** file contains design files for the first two example systems. This zip file includes files for both the Nios II Development Kit Cyclone® II Edition and the Nios II Development Kit Stratix® II Edition.



Download the design files for the example designs from the [On-chip Debugging Design Examples](#) page on the Altera website.

The first example Tcl script creates a LED light show on your board. The SOPC Builder system for this example includes two modules: a JTAG to Avalon master bridge and a parallel I/O (PIO) core. The JTAG to Avalon master bridge provides a connection between your development board and either SOPC Builder system via JTAG. The serial peripheral interface (SPI) to Avalon master bridge provides connections via SPI. The PIO module provides a memory-mapped interface between an Avalon-MM slave port and general-purpose I/O ports.

 For more information about these components, refer to the *Embedded Peripherals IP User Guide*.

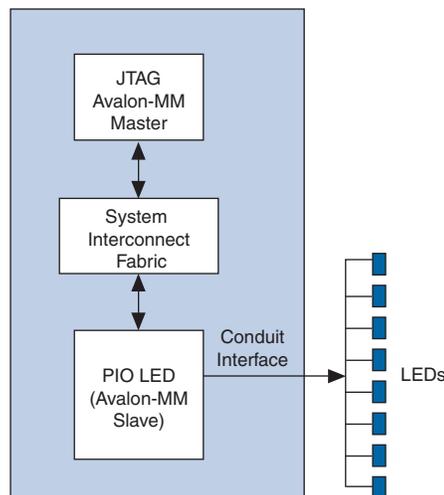
The first example Tcl script sends a series of `master_write_8` commands to the JTAG Avalon master bridge. The JTAG Avalon master sends these commands to the Avalon-MM slave port of the PIO module. The PIO I/O ports connect to FPGA pins that are, in turn, connected to the LEDs on your development board. The write commands to the PIO Avalon-MM slave port result in the light show.

 The instructions for these examples assume that you are familiar with the Quartus II software and either the SOPC Builder software.

LED Light Show Example

Figure 12-2 illustrates the SOPC Builder system for the first example.

Figure 12-2. SOPC Builder System for Light Show Example



To build this example system, perform the following steps:

1. On your host computer file system, locate the following directory: `<Nios II EDS install path>\examples\<verilog or vhdl>\<board version>\standard`. Each development board has a VHDL and Verilog HDL version of the design. You can use either of these design examples.

2. Copy the **standard** directory to a new location. By copying the design files, you avoid corrupting the original design and avoid issues with file permissions. This document refers to the newly created directory as the `c:\<projects>\standard` directory.

 For information on different board kits available from Altera, refer to the [All Development Kits](#) page on the Altera website.

3. Copy the **System_Console.zip** file to the `c:\<projects>\standard` directory and unzip it. Specific directories may be created for specific Altera development boards.
4. Choose **All Programs > Altera > Nios II EDS <version> Command Shell** (Windows Start menu) to run a Nios II command shell.
5. Change to the directory for your board.
6. To program your board with the **.sof** file, type the following command in the Nios II command shell:

```
nios2-configure-sof <sof_name>.sof ↵
```

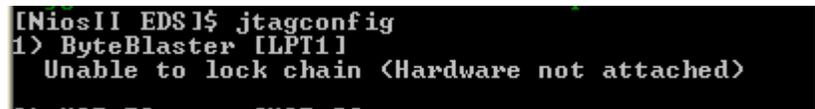
If your development board includes more than one JTAG cable, you must specify which cable you are communicating with as an argument to the `nios2-configure-sof <sof_name>.sof` command. To do so, type the following commands:

```
jtagconfig ↵
```

[Figure 12-3](#) gives sample output from the `jtagconfig` command. This output shows that the active JTAG cable is number 2. Substitute the number of your JTAG for the `<cable_number>` variable in the following command:

```
nios2-configure-sof -c <cable_number> <sof_name>.sof ↵
```

Figure 12-3. jtagconfig Output



```
[NiosII EDS]$ jtagconfig
1) ByteBlaster [LPT1]
Unable to lock chain (Hardware not attached)
```

7. You can then run the LED light show example by typing the following command:


```
system-console --script=led_lightshow.tcl ↵
```
8. You can see the LEDs performing a running light demonstration. Press **Ctrl+C** to stop the LED light show.
9. To see the commands that this script runs, open the **led_lightshow.tcl** file in your `\jtag_pio_<cii_or_sii>` directory.

Loading and Linking a Design

Example 12-4 shows how to load and link a Quartus II design set.

Example 12-4. Loading and Linking a Design

```
% get_service_paths device
{/connections/USB-Blaster [USB-0]/EP2C35}

% set device_path [lindex [get_service_paths device] 0]
/connections/USB-Blaster [USB-0]/EP2C35

% design_load /projects/9.1/standard
QuartusDesignFactory elaborating \projects\9.1\standard
QuartusDesignFactory found SOF File at NiosII_cycloneII_2c35_standard.sof
QuartusDesignFactory found JDI File at NiosII_cycloneII_2c35_standard.jdi
QuartusDesignFactory found SOPC Info File at
\projects\9.1\standard\NiosII_cycloneII_2c35_standard_sopc.sopcinfo

% set design_path [lindex [get_service_paths design] 0]
/designs/standard

% design_link $design_path $device_path
Created a link from /designs/standard to /connections/USB-Blaster [USB-0]/EP2C35.
Created a link from /designs/standard/NiosII
cycloneII_2c35_standard_sopc.sopcinfo/cpu.data_master to /connections/USB-Blaster
[USB-0]/EP2C35/cpu.
Created a link from
/designs/standard/NiosII_cycloneII_2c35_standard_sopc.sopcinfo/cpu.data_master/jtag_
uart.avalon_jtag_slave to /connections/USB-Blaster [USB-0]/EP2C35/jtag_uart
```

JTAG Examples

Two JTAG examples are described below. The first JTAG example gives you some practice working with the System Console as an interactive tool. The second example verifies that the clock is toggling.

Verify JTAG Chain

In this example, you verify the JTAG chain on your board. To run this example, perform the following steps:

1. On the Windows Start menu, point to **All Programs**, then point to **Altera**, and then click **Quartus II <version>** to run the Quartus II software. Open the Quartus II project file, **jtag_pio.qpf** or **jtag_pio_sii.qpf**.
2. On the Tools menu, click **SOPC Builder**.
3. On the SOPC Builder Tools menu, click **System Console**.
4. Set the path to the `jtag_debug` service by typing the following command:

```
set jd_path [lindex [get_service_paths jtag_debug] 0] ←
```

The `get_service_paths` command always returns a list, even if the list has a single item; consequently, you must index into the list using the `lindex` command. In this case, the variable `jd_path` is assigned the string that is the 0th element of the list.

5. Open the `jtag_debug` service by typing the following command:

```
open_service jtag_debug $jd_path ←
```

6. Set up a list of byte values to test the chain by typing the following command:

```
set values [list 0xaa 0x55 0xaa 0x55 0xaa 0x55 0xaa 0x55 0xaa 0x55  
0xaa 0x55 0xaa 0x55 0xaa 0x55 0xaa 0x55]↵
```

7. Loop the values by typing the following command:

```
jtag_debug_loop $jd_path $values ↵
```

If the `jtag_debug_loop` command is successful, you should see the values that you sent reflected in the System Console. [Figure 12-4](#) shows the transcript from this interactive session.

Figure 12-4. The `jtag_debug_loop` Command

```
⌘ set jd_path [lindex [get_service_paths jtag_debug] 0]  
/root/CONNECTIONS/USB-Blaster [USB-0]/EP2C35/[MFG:110 ID:132 INST:0 VER:1]  
  
⌘ open_service jtag_debug $jd_path  
  
⌘ set values [list 0xaa 0x55  
0xaa 0x55 0xaa 0x55]  
0xaa 0x55 0xaa 0x55
```

8. Close the `jtag_debug` service by typing the following command:

```
close_service jtag_debug $jd_path↵
```

This example provides the beginnings of a JTAG chain validation workflow. Depending on the number of FPGAs in your JTAG chain, you can expand upon this test by performing more operations, in which you can interleave access to JTAG chains with larger data sets, and potentially multiple devices.

Verify Clock

The command to verify that your clock is toggling samples the clock asynchronously. Consequently, you may need to use this command several times to determine if the clock is toggling. The `jtag_debug_sample_clock.tcl` script samples the clock 10 times. To run this script, type `source jtag_debug_sample_clock.tcl` at the System Console prompt. You should see 10 values for the JTAG clock printed to the System Console as [Figure 12-5](#) illustrates.

Figure 12-5. The `jtag_debug_sample_clock` Command

```

% source jtag_debug_sample_clock.tcl
Service Open Status is: 1

Multiple samples of clock status

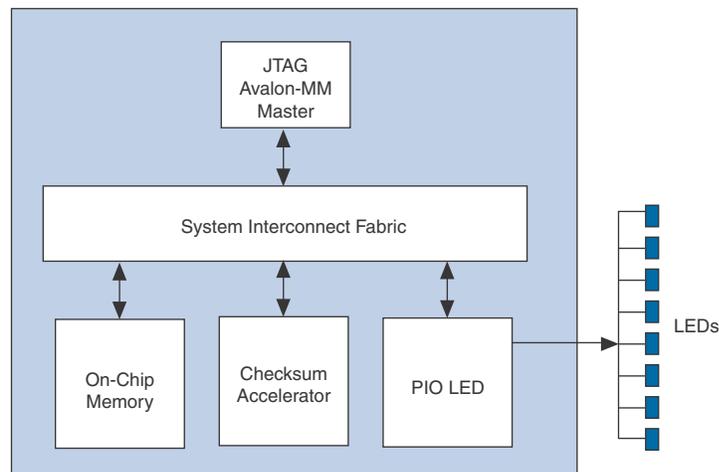
0
0
1
0
0
1
1
1
1
0
1

Closing jtag_debug service path ($jclk)
```

Checksum Example

In this example, you add an on-chip memory and hardware accelerator to the SOPC Builder system discussed in the previous example. The hardware accelerator calculates a checksum. Figure 12-6 illustrates this system.

Figure 12-6. SOPC Builder System for Checksum Accelerator Example



To build this example system, perform the following steps:

1. In the **System Contents** tab in SOPC Builder, double-click **On-Chip Memory (RAM or ROM)** in the **On-Chip** of the **Memories and Memory Controllers** folder to add this component to your system.
2. In the **On-Chip Memory (RAM or ROM)** wizard, for **Total memory size** type 128 to change the memory size to 128 bytes. Click **Finish** to accept the other default values.
3. To connect the on-chip memory to the master, click the open dot at the intersection of the onchip_mem s1 Avalon slave port and the JTAG to Avalon Master Bridge master port.
4. In the **System Contents** tab, double-click **Checksum Accelerator** in the **Custom Component** folder to add this component to your system.
5. To connect the checksum accelerator Slave port, click on the open dot at the intersection of the accelerator Slave and the master master port.
6. To connect the checksum accelerator Master port, click on the open dot at the intersection of the accelerator Master and the onchip_mem s1 port.

7. In the **Base** column, enter the base addresses for the slaves in your system.

- Onchip_mem s1 port—0x00000080
- Accelerator Slave port—0x00000020

Click on the **lock** icon next to each address to lock these values.

Figure 12-7 illustrates the completed system.

Figure 12-7. Checksum Accelerator Module Connections

Use	Con...	Module Name	Description	Clock
<input checked="" type="checkbox"/>		master	JTAG to Avalon Master Bridge	
		master	Avalon Memory Mapped Master	clk
<input checked="" type="checkbox"/>		led_pio	PIO (Parallel I/O)	
		s1	Avalon Memory Mapped Slave	clk
<input checked="" type="checkbox"/>		onchip_mem	On-Chip Memory (RAM or ROM)	
		s1	Avalon Memory Mapped Slave	clk
<input checked="" type="checkbox"/>		accelerator	Checksum Accelerator	
		Slave	Avalon Memory Mapped Slave	clk
		Master	Avalon Memory Mapped Master	

8. Save your system.
9. In the **System Contents** tab, click **Next**.
10. In the **System Generation** tab, click **Generate**.
11. On the Quartus II Processing menu, click **Start Compilation**.
12. When compilation completes, re-program your board by typing the following command in the Nios II command shell:


```
nios2-configure-sof jtag_pio.sof ↵
```
13. Type `system-console` ↵ in the Nios II command shell to start the System Console.

 If you reprogram your board, you must start a new System Console to receive the changes.

14. To run the checksum example, in the System Console, type:

```
source set_memory_and_run_checksum.tcl ↵
```

Figure 12-8 shows the output from a successful run.

Figure 12-8. System Console Output

```
System Console

% source set_memory_and_run_checksum.tcl

*****
Onchip RAM values out after filling with data.
*****
0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a
0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a
0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a
0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a 0x5a5a5a5a

*****
Starting Checksum operation.
*****
Writing to address and length registers.
Address register value = 0x00000080
Length register value = 0x00000020

Writing clear to status register.
Writing clear to control register.

Writing GO to control register.
Checksum DONE bit set.
Result register value (non-inverted) = 0xa5a5

Writing clear to status register.
Writing clear to control register.

Writing GO and INVERT to control register.
Checksum DONE bit set.
Result register value (inverted) = 0x5a5a

Checksum example finished.
```

You can change the value written into the RAM by changing the value given in the `fill_memory` routine in the `set_memory_and_run_checksum.tcl` file. Save the Tcl file after editing and rerun the command. (Because the system command uses `master_write_32`, if you use values that are less than 32 bits, they are filled with leading 0s.)

Nios II Processor Example

In this example, you program the Nios II processor on your board to run the count binary software example that is included in the Nios II installation. This is a simple program that uses an 8-bit variable to repeatedly count from 0x00 to 0xFF. The output of this variable is displayed on the LEDs on your board. After programming the Nios II processor, you use the System Console processor commands to start and stop the processor.

To run this example, perform the following steps:

1. Download the [Nios II Ethernet Standard Design Example](#) for your board from the Altera website.
2. Create a folder to extract the design. For this example, use `C:\Count_binary`.
3. Unzip the Nios II Ethernet Standard Design Example into `C:\Count_binary`.
4. In a Nios II command shell, change to the directory of your new project.
5. To program your board, type the following command in a Nios II command shell:


```
nios2-configure-sof niosii_ethernet_standard_<board_version>.sof ↵
```
6. Using Nios II Software Build Tools for Eclipse, create a new Nios II Application and BSP from Template using the **Count Binary** template and targeting the Nios II Ethernet Standard Design Example.
7. To build the executable and linkable format (ELF) file (`.elf`) for this application, right-click the **Count Binary** project and select **Build Project**.



For more information about creating Nios II applications, refer to the [Nios II Software Build Tools](#) chapter in the *Nios II Software Developer's Handbook*.

8. Download the `.elf` file to your board by right-clicking **Count Binary** project and selecting **Run As, Nios II Hardware**.

The LEDs on your board provide a new light show.

9. Start the System Console by typing `system-console` in your Nios II command shell.
10. Set the processor service path to the Nios II processor by typing the following command:

```
set niosii_proc [lindex [get_service_paths processor] 0] ↵
```

11. Open both services by typing the following commands:

```
open_service processor $niosii_proc ↵
```

12. Stop the processor by typing the following command:

```
processor_stop $niosii_proc ↵
```

The LEDs on your board freeze.

13. Start the processor by typing the following command:

```
processor_run $niosii_proc ↵
```

The LEDs on your board resume their previous activity.

14. Stop the processor by typing the following command:

```
processor_stop $niosii_proc ↵
```

15. Close the services by typing the following command:

```
close_service processor $niosii_proc ↵
```

The `processor_step`, `processor_set_register` and `processor_get_register` commands provide additional control over the Nios II processor.

Device Support

You can target all Altera device families with the System Console. Transceiver Toolkit commands, however, can only be targeted for Arria II GX and Stratix IV GX devices.

Conclusion

The System Console offers you a wide variety of options for communicating with modules in your design at a low level. You can use either Tcl scripting commands or the GUI to access and run services for setting up, running tests, optimizing design parameters, and debugging designs you have programmed into Altera supported device families without having to recompile you designs.

Document Revision History

Table 12-37 shows the revision history for this chapter.

Table 12-37. Document Revision History

Date	Version	Changes
December 2010	10.1.0	Maintenance release. This chapter adds new commands and references for Qsys.
July 2010	10.0.0	Initial release. Previously released as the System Console User Guide, which is being obsoleted. This new chapter adds new commands.

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter describes how to use the transceiver link debugging examples with the Transceiver Toolkit in the Quartus® II software. The Transceiver Toolkit in the Quartus II software allows you to quickly test the functionality of transceiver channels and helps you improve the signal integrity of transceiver links in your design.

You can use an example design available on the Altera® website if you want to immediately start using the Transceiver Toolkit, or you can create a custom design. In today's high-speed interfaces, stringent bit error rate (BER) requirements are not easy to meet and debug. You can use the Transceiver Toolkit in the Quartus II software to check and improve the signal integrity of transceiver links on your board before you complete the final design, saving you time and helping you find the best physical medium attachment (PMA) settings for your high-speed interfaces.

This chapter contains the following sections:

- “Transceiver Toolkit Overview”
- “Transceiver Link Debugging Design Examples” on page 13–2
- “Setting Up Tests for Link Debugging” on page 13–3
- “Tcl Script in System Console” on page 13–11
- “Usage Scenarios” on page 13–12

Transceiver Toolkit Overview

The underlying framework for the Transceiver Toolkit is the System Console. The System Console performs low-level hardware debugging of your design. The System Console provides read and write access to the IP cores instantiated in your design. Use the System Console for the initial bring-up of your PCB and low-level testing.

 For information about the System Console, refer to the *Analyzing and Debugging Designs with the System Console* chapter in volume 3 of the *Quartus II Handbook*.

The Transceiver Toolkit allows you to perform run-time tasks, including performing high-speed link tests for the transceivers in your devices. The Transceiver Toolkit allows you to test your high-speed interfaces in real-time. To launch the Transceiver Toolkit, in the main Quartus II window, on the Tools menu, click **Transceiver Toolkit**.

Auto Sweep

You can test the input sweep ranges for your transceiver PMA settings and run tests automatically with the auto sweep feature. You can store a history of the test runs and keep a record of the best PMA settings. You can then use these settings in your final design.

- ② For more information, refer to the *Transceiver Toolkit Auto Sweep Panel* in Quartus II Help.

EyeQ

You can determine signal integrity with the EyeQ feature. The EyeQ feature in the Transceiver Toolkit allows you to chart eye-width data and view the data in a bathtub curve. After you run the EyeQ feature you can view the data in the **Report** pane of the Transceiver Toolkit and export the data in Comma-Separated Value (.csv) format for further analysis.

- ② For more information about the EyeQ feature, refer to *Transceiver Toolkit EyeQ* in Quartus II Help.
-  For more information, refer to *AN 605: Using the On-Chip Signal Quality Monitoring Circuitry (EyeQ) Feature in Stratix IV Transceivers*.

Control Links

You can test the transmitter and receiver channel links in your design in manual mode with the channel control features. The channel control panels allow you to view and manually modify settings for transmitter and receiver channels while the channels are running.

- ② For more information about the Transceiver Toolkit, refer to *Working with the Transceiver Toolkit* in Quartus II Help.

Transceiver Link Debugging Design Examples

Altera provides design examples to assist you with setting up and using the Transceiver Toolkit. To learn more about the version of the Quartus II software used to create these design examples, the target device, and development board details, refer to the **readme.txt** of each example. Each example is verified and tested with the Quartus II software version referenced in the **readme.txt**. However, you may be able to use these examples with a later version of the Quartus II software.

If you are recompiling the design examples for a different board, refer to “*Changing Pin Assignments*” on page 13-7 to determine which pin assignments you must edit.

-  Download the Transceiver Toolkit design examples from the *On-Chip Debugging Design Examples* page of the Altera website.

Setting Up Tests for Link Debugging

Testing signal integrity for high-speed transceiver links involves using data patterns, such as pseudo-random binary sequences (PRBS). Although the sequences appear to be random, they have specific properties that you can use to measure the quality of a link. In the example designs available on the Altera website, data patterns are generated by a pattern generator and are then transmitted by the transmitter. The transceiver on the far end can then be looped back so that the same data is then received by the receiver portion of the transceiver. The data obtained is then checked by a data checker to verify any bit errors.

Figure 13–1 and Figure 13–2 show examples of the test setup for the transceiver link debugging tool. The figures show a setup that is similar to the design examples that you can download from the On-Chip Debugging Design Example page of the Altera website. You can also have the transmitter on one FPGA and the receiver on a different FPGA.

Figure 13–1. Transceiver Link Debugging Tool Test Setup

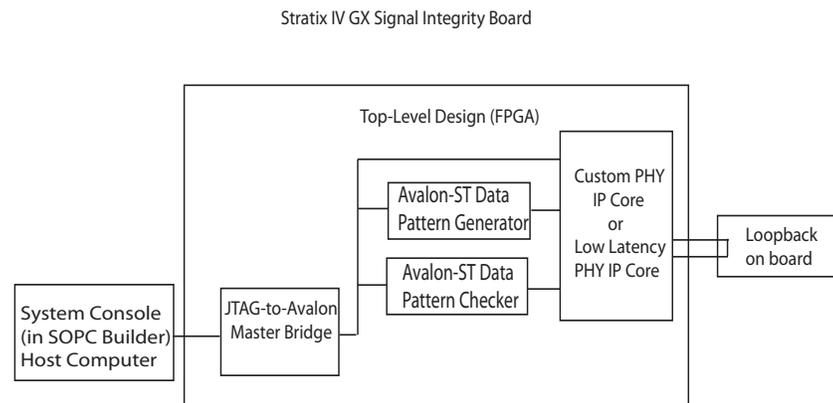
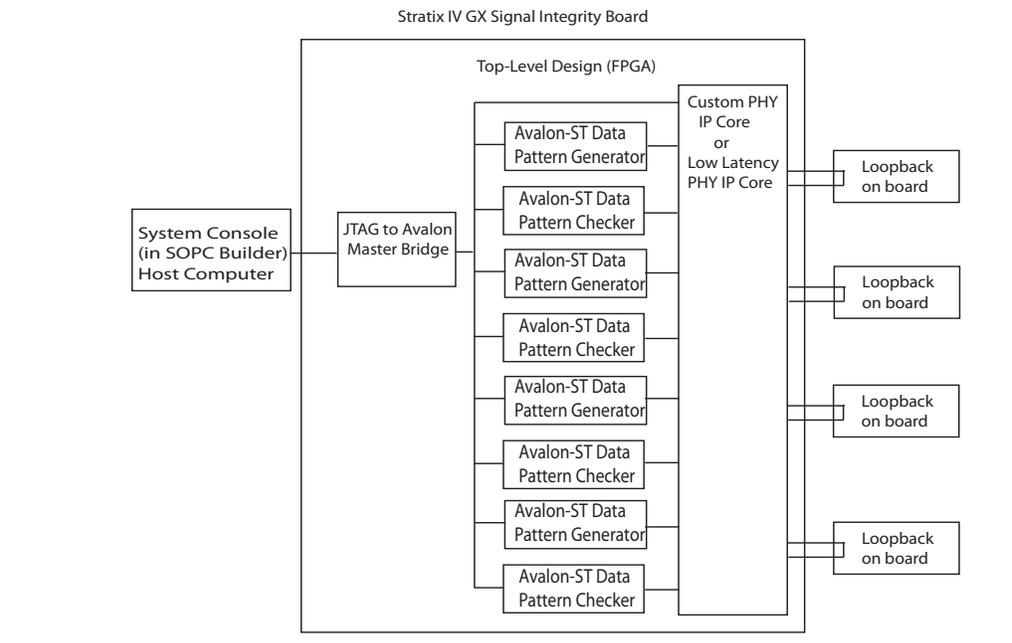


Figure 13-2 shows a similar test setup for the second design example described in this section, except that there are four sets of transceivers and receivers rather than one.

Figure 13-2. Transceiver Link Debugging Tool Test Setup (Four Channels)



The design examples include the Qsys system integration tool or SOPC Builder system, which contains the following components:

- Custom PHY IP Core or Low Latency PHY IP Core
- Avalon-ST Data Pattern Generator
- Avalon-ST Data Pattern Checker
- JTAG-to-Avalon Master Bridge
- The Avalon[®] Streaming interface (Avalon-ST)

The Avalon Streaming interface (Avalon-ST) accommodates the development of high-bandwidth, low-latency components for Qsys and SOPC Builder systems. The Avalon-ST interface protocol provides component designers with a framework to create interfaces that support the unidirectional flow of data, including multiplexed streams, packets, and DSP data.

Custom PHY IP Core

You can use the Custom PHY IP core to test all possible parallel data widths of the transceivers in these design examples. You can configure the Custom PHY IP core as 8, 10, 16, 20, 32 or 40-bit. The sweep tools disable word alignment during sweep, which is enabled to simplify timing closure. You can also use the Data Format Adapter IP component as required. You can have one or multiple channels in your design.

You use SOPC Builder to define and generate the Custom PHY IP core. The Custom PHY IP core in the design examples that you can download from the On-Chip Debugging Design Example page of the Altera website are generated for Stratix IV GX devices.

To use the Custom PHY IP core with the Transceiver Toolkit, perform the following steps:

1. Set the following parameters to meet your project requirements:
 - Number of lanes
 - Bonded group size
 - Serialization factor
 - Data rate
 - Input clock frequency
2. Turn on **Avalon data interfaces**.
3. Disable **8B/10B**
4. Set **Word alignment mode** to manual
5. Disable **rate match FIFO**
6. Disable **byte ordering block**

 For more information about the protocol settings used in the Custom PHY IP core, refer to the “Custom PHY IP User Core” section of the *Altera Transceiver PHY IP Core User Guide*.

Low Latency PHY IP Core

Use Low Latency PHY IP Core as follows:

- To get more than 8.5 gbps in GT devices.
- To use PMA direct mode, such as when using six channels in one quad.

To meet your project requirements, use the same set of parameters that you would use with the Custom PHY IP core.

The phase compensation FIFO mode must be set to embedded, above certain data rates. The Transceiver Toolkit provides a warning when you exceed the data rate. To be in PMA direct mode, you must set the phase compensation FIFO mode to none, which means that only certain data rates support the PMA direct mode.

 For more information about the protocol settings used in the Low Latency PHY IP core, refer to the “Low Latency PHY IP User Core” section of the *Altera Transceiver PHY IP Core User Guide*.

Avalon-ST Data Pattern Generator

This component produces data patterns in the test flow. You can use a variety of popular patterns to test transceiver signal integrity (SI). The data pattern generator component is provided as an Altera IP core component. You can use any pattern, but you must have a checker to verify that you receive that pattern properly.

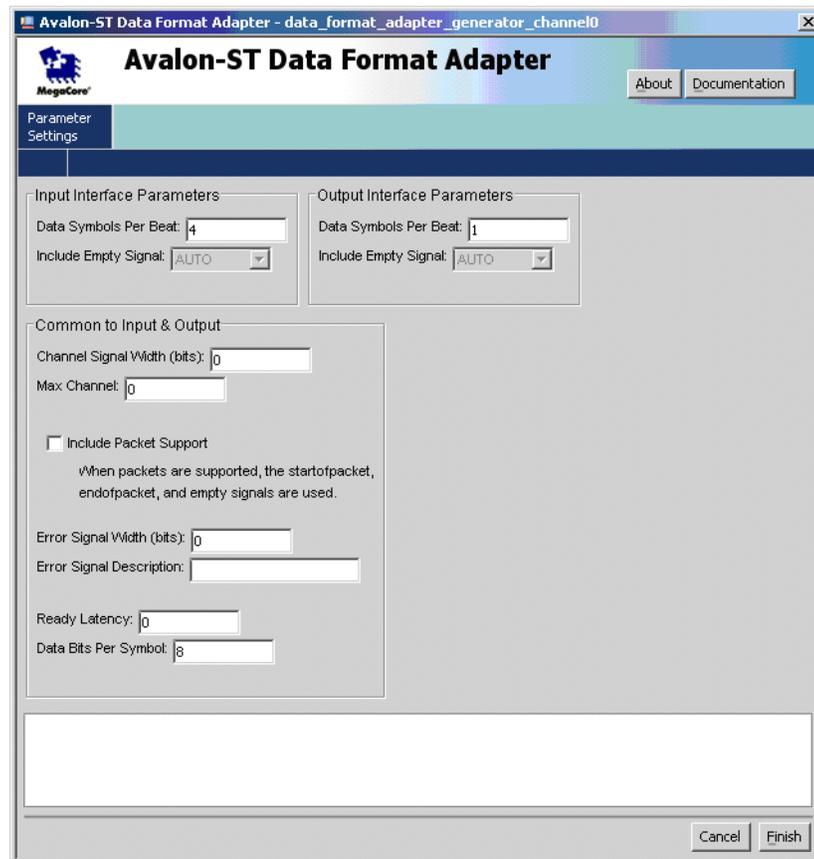
When you use the Avalon[®]-ST Data Pattern Generator, the width may be different in the pattern that you use, and you may need to use a data format adaptor. The Avalon-ST Data Pattern Generator component is available in the Qsys or SOPC Builder library tree. The adaptor, however, does not form a new data generation component package, but serves as a data formatter.

The Avalon-ST Data Pattern Generator generates industry-standard data patterns. Data patterns are generated on a 32-bit or 40-bit wide Avalon streaming source port.

For more information, refer to the *System Interconnect Fabric for Streaming Interfaces* chapter in the *SOPC Builder User Guide*.

Figure 13-3 shows the wizard page that you use to set parameters for the Avalon-ST Data Format Adapter.

Figure 13-3. Avalon-ST Data Format Adapter



Data Checker

The Avalon-ST Data Pattern Checker components are provided as Qsys or SOPC components. These components include Avalon-ST Data Pattern Checker, Avalon-ST Data Format Adapter, and Avalon-ST Timing Adapter. When the data checker runs, it attempts to find the selected data pattern in the incoming data by comparing the incoming pattern with the expected pattern.

The Avalon-ST Data Pattern Checker is available under **Debug and Performance** in the Qsys or SOPC Builder component library tree. The checker is the core logic for data pattern checking. Data patterns are accepted on 32-bit or 40-bit wide Avalon streaming sink ports.

 For more information, refer to *Avalon Streaming Data Pattern Generator and Checker Cores* in the *Embedded Peripherals User Guide*.

Use the design examples to quickly test the functionality of the receiver and transmitter channels in your design without creating any custom designs with data generators and checkers. When working with these examples, assume that your design is not complete, and that you want to check the transceiver link signal integrity (SI). You can quickly change the transceiver settings in the design examples to see how they affect transceiver link performance. You can also use the design examples to isolate and verify the transceiver links without having to debug other logic in your design.

Compiling Design Examples

Once you have downloaded the design examples, open the Quartus II software version 10.0 or later and unarchive the project in the example. If you have access to the same development board with the same device as mentioned in the **readme.txt** file of the example, you can directly program your device with the provided programming file in that example. If you want to recompile the design, you must regenerate the system in SOPC builder and recompile the design in the Quartus II software to get a new programming file.

If you have the same board as mentioned in **readme.txt** file, but a different device on your board, you must choose the appropriate device and recompile the design. For example, some early development boards are shipped with the engineering sample devices.

If you have a different board, you must edit the necessary pin assignments and recompile the design examples provided here.

Changing Pin Assignments

Table 13-1 shows the pin-assignment edits for the Stratix IV Transceiver Signal Integrity Development Kit (DK-SI-4SGX230N). You must make these assignments before you recompile your design.

Table 13-1. Stratix IV GX Top-Level Pin Assignments (DK-SI-4SGX230N)

Top-Level Signal Name	I/O Standard	Pin Number on DK-SI-4SGX230N Board
REFCLK_GXB2_156M25 (input)	2.5 V LVTTTL/LVCMOS	PIN_G38
S4GX_50M_CLK4P (input)	2.5 V LVTTTL/LVCMOS	PIN_AR22
GXB1_RX1 (input)	1.4-V PCML	PIN_R38
GXB1_TX1 (output)	1.4-V PCML	PIN_P36

Table 13-2 shows the pin-assignment edits for the Stratix IV GX development kit (DK-DEV-4SGX230N). You must make these assignments before you recompile your design.

Table 13-2. Stratix IV GX Top-Level Pin Assignments (DK-DEV-4SGX230N)

Top-Level Signal Name	I/O Standard	Pin Number on DK-DEV-4SGX230N Board
REFCLK_GXB2_156M25 (input)	LVDS	PIN_AA2
S4GX_50M_CLK4P (input)	2.5 V LVTTTL/LVCMOS	PIN_AC34
GXB1_RX1 (input)	1.4-V PCML	PIN_AU2
GXB1_TX1 (output)	1.4-V PCML	PIN_AT4

Similarly, you can change the pin assignment for your own board and recompile the design examples.

Transceiver Toolkit Link Test Setup

To set up testing with the Transceiver Toolkit, perform the following steps:

1. Program the device either with the programming file provided with the **.zip** file or with the new programming file after you recompile your project.
2. Open the Transceiver Toolkit from the Tools menu in the Quartus II software.
3. Ensure that the **.sof** of the device you use is programmed and that board settings, such as the jumper settings, are correct.

Loading the Project in System Console

From the Transceiver Toolkit, load the Quartus project file (***.qpf**) by going to the File menu and selecting **Load Project**. Whether you have recompiled the design or not, the unzipped contents of the examples contain this file. After the project is loaded, you can review the design information under **System Explorer** in the System Console.

Linking the Hardware Resource

Ensure the device board is turned ON and is available before you start the System Console or the Transceiver Toolkit. The connections are detected at startup of the tool. Right click on a **Design Instance** under the **System Explorer** and link that instance to the device you want. Each design example contains one transceiver Custom PHY IP core instance, the first one configured as one channel and the second one configured with four channels in bonded mode. Right-click on the design instance and click on the device to link that design instance with your device. If you are designing your own design example with more than one design instance and your board has more than one device, then you have the option of linking a certain device to a certain design instance. This is useful when you want to perform a link test between a transmitter and receiver on two separate devices.

To set up the transmitter channels, receiver channels, and transceiver links, go to the Tools menu and select **Transceiver Toolkit**, or from the **Welcome to the Transceiver Toolkit** tab, click on the **Transceiver Toolkit**. You can then create the channels, as described in the following section.

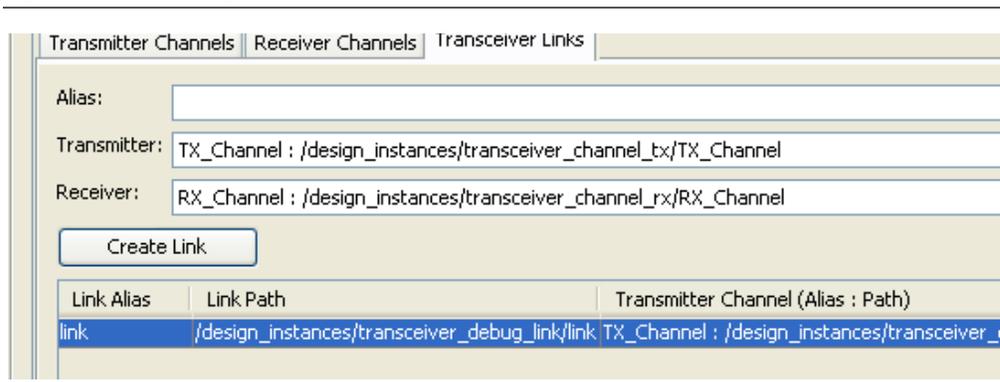
Creating the Channels

After you open the Transceiver Toolkit, there are three tabs in the System Console—**Transmitter Channels**, **Receiver Channels**, and **Transceiver Links** tabs.

The **Transmitter Channel** and **Receiver Channel** tabs are automatically populated with the existing transmitter and receiver channels in the design example, respectively. However, you can create your own additional transmitter and receiver channels in a situation where the expected configuration was not automatically populated. In the **Link Channel** tab you can create a link between the transmitter and receiver channels.

Links are automatically created when a receiver channel and transmitter channel share a transceiver channel. However, if you are not actually looping data back, but using it to transmit or receive to another transceiver channel, you must define and create a new link. For example, in [Figure 13-4](#), a link is created in the **Link Channel** tab between a transmitter and receiver channels of the same device.

Figure 13-4. Creating a Link Channel in Transceiver Toolkit



You can also perform a physical link test without loopback by connecting one device transmitter channel to another device receiver channel. In this channel you would define that connection into a link, and tests would run off that link. For example, as shown in [Figure 13-1](#), use the transmitter and receiver channels of the same device and loop them back on the far side of the board trace to check the signal integrity of your high-speed interface on the board trace. You can select the link that you created in the Transceiver Toolkit and use the different buttons to start and control link tests.

 You can also communicate with other devices that have the capability to generate and verify test patterns that Altera supports.

Running the Link Tests

Use the Link Channel tab options to control how you want to test the link. For example, use the **Auto Sweep** feature to sweep various transceiver settings parameters through a range of values to find the results that give the best BER value. You can also open Transmitter, Receiver and Link Control panels to manually control the PMA settings and run individual tests. You can change the various controls in the panels to suit your requirements.

To perform an auto sweep link test, perform the following steps:

1. Select the link to test and click **Auto Sweep** to open the auto sweep panel for the selected link.
2. Set the test pattern to test.
3. Select either **Smart AutoSweep** or **Full AutoSweep**. **Full AutoSweep** runs a test of every combination of settings that falls within the bounds that you set.



Smart AutoSweep minimizes the number of tests run to reach a good setting, which saves you time. **Smart AutoSweep** does not test every possible setting, so may not achieve the very best setting possible; however, it can quickly find an acceptable setting.

4. Set up run lengths for each test iteration. The run length limits you set are ignored when you run a smart sweep.
5. Set up the PMA sweep limits within the range you want to test.
6. Press **Start** and let the sweep run until complete.

After an Auto Sweep test has finished at least one iteration, you can create a report by clicking **Create Reports**, which opens the **Reports** tab. The report shows data from all tests that you have completed. You can sort by columns and filter data by regular expression in the **Reports** tab. You can also export the reports to a **.csv** file by right clicking on a report.

If you found a setting with the **Smart Sweep** mode, use the reported best case PMA settings, and apply a +/- 1 setting to them, returning the test to **Full AutoSweep** mode. This helps you determine if the settings you chose are the best.

You can then choose your own runtime conditions, reset the AutoSweep feature by pressing **Reset**, and re-running the tests to generate BER data based on your own run length conditions. This procedure allows you to obtain a good setting for PMA faster. Review the **BER** column at the end of the full sweep to determine the best case.



When setting smart sweep ranges, try to include the 0 setting as often as possible, as that setting often provides the best results.

Viewing Results in the EyeQ Feature

Advanced FPGA devices such as Stratix IV have built-in EyeQ circuitry, which is used with the EyeQ feature in Transceiver Toolkit. The EyeQ feature allows you to estimate the horizontal eye opening at the receiver of the transceiver. With this feature, you can tune the PMA settings of your transceiver, which results in the best eye and BER at high data rates.



For more information about the EyeQ feature, refer to *AN 605: Using the On-Chip Signal Quality Monitoring Circuitry (EyeQ) Feature in Stratix IV Transceivers*.

To use the EyeQ feature in the Transceiver Toolkit to view the results of the link tests, perform the following steps:

1. Select a **Transceiver Link** or **Receiver Channel** in the main Transceiver Toolkit panel that you want to run EyeQ against.

2. Click **Control Link** or **Control Channel** to open control panels. Use these control panels to set the PMA settings and test pattern that you want the EyeQ feature to run against. You can check the report panel or **Best Case** column from an Autosweep run for the settings with the best BER value and enter those PMA values through the Transmitter and Receiver control panels.
3. Click **EyeQ** to open the EyeQ feature.
4. Review the test stop conditions and set them to your preference.
5. Click **Run**. The EyeQ feature gathers current settings of the channel and uses those settings to start a phase sweep. Thirty-two iterations of the phase sweep will run.

As the run progresses, you can view the status, which is displayed on the top section of the EyeQ feature. You can read current progress and the chart is updated with values each time a run completes. When the run is completed, you should see a bathtub curve.

When at least one run has completed, you can click **Create Report** to view past data runs in a report panel.

6. When the run finishes, if the bathtub curve looks like a hill instead, click the center eye button, which then reorganizes the data into a bathtub curve.
7. If the data does not display on the chart, you can right-click and choose **Auto-Range**, which should make the data appear.

If you want to change PMA settings and re-run the EyeQ feature, make sure you first stop and reset the EyeQ feature. If you do not reset, the EyeQ feature continues testing based on the original PMA settings of the current test that it had begun, and overwrites any setting you may have changed through the control panel.

After you stop and reset the EyeQ feature, change settings in the link or receiver channel control panel. Then click **Start** on the EyeQ feature to start a new set of tests.

When you can see that you are running good PMA settings, the bathtub curve is wide, with sharp slopes near the edges. The curve may be up to 30 units wide. If the bathtub is narrow, maybe as small as two units wide, then the signal quality may be poor. The wider the bathtub curve, the wider the eye you have. Conversely, the smaller the bathtub curve, the smaller the eye.

- ❓ For more information about how to use the EyeQ feature, refer to *Working with the Transceiver Toolkit* in Quartus II Help.

Tcl Script in System Console

System Console is the framework in the Quartus II software that supports the Transceiver Toolkit. System Console provides a number of Tcl commands that you can use to build a custom test routine script to test the transceiver link using the data generator and checker. System Console allows you to tune PMA parameter, such as those for changing DC gain. To get help on the Tcl commands available, type `help` in the Tcl console in System Console. To run a transceiver link test flow, perform the following. You can perform all the tasks with Tcl commands in System Console.

1. Load the Quartus II project.

2. Find and link the design and service path.
3. Find and open links to transmitter channels and receiver channels.
4. Set up PRBS patterns to run on the link.
5. Set up PMA settings on transmitter and receiver channels.
6. Use PIO logic to generate a rising edge to enable the word aligner.
7. Start the link test.
8. Poll the receiver for BER data.
9. When the test is finished, stop the link test.
10. Close the link.



For more information about Tcl commands, refer to the *Analyzing and Debugging Designs with the System Console* chapter in volume 3 of the *Quartus II Handbook*.

The Tcl scripts provided with the design examples help you understand Tcl commands associated with each task.

Running Tcl Scripts

The tasks that you perform with the Transceiver Toolkit GUI for setting up your test environment you can also save as Tcl scripts. For example, you can save the steps you perform for loading your project in the Transceiver Toolkit by clicking **Create Tcl setup script** on the File menu.

You can save your steps at various stages when you add projects, create design instances, link designs to devices, and create transceiver links., so that you do not have to redo all the steps in the GUI to reach the initial setup you created to get your specific configuration ready for debugging. All the saved scripts are shown under the **Scripts** in the **System Explorer**. You can execute the script by right-clicking on **Scripts** under the **System Explorer** in the Transceiver Toolkit, and then clicking **Run**, or you can double-click on the script. You can also execute the script from the **System Console command line**.

Usage Scenarios

You can use the Transceiver Toolkit in many scenarios. For example you can use the Transceiver Toolkit if you are debugging one device on one board or more than one device on a single or multiple boards. Usually for a device you have a single Quartus II design or project, but you can have one design targeted for two or more similar devices on the same or different boards.

Possible scenarios for how you can use the Transceiver Toolkit in those situations follows. The scenarios assume that you have programmed the relevant **.sof** to the correct device that you are testing.

- [“Linking One Design to One Device Connected By One USB Blaster Cable” on page 13-13](#)
- [“Linking Two Designs to Two Separate Devices on Same Board \(JTAG Chained\), Connected By One USB Blaster Cable” on page 13-13](#)

- “Linking Two Designs to Two Separate Devices on Separate Boards, Connected to Separate USB Blaster Cables” on page 13-13
 - “Linking Same Design on Two Separate Devices” on page 13-14
 - “Linking Unrelated Designs” on page 13-14
 - “Saving Your Setup As Tcl Scripts” on page 13-14
 - “Verifying Channels Are Correct When Creating Link” on page 13-14
 - “Using the Recommended DFE Flow” on page 13-15
 - “Running Simultaneous Tests” on page 13-15
- ❓ For further information on how to use the Transceiver Toolkit GUI to perform the following scenarios, refer to *Working with the Transceiver Toolkit* in Quartus II Help.

Linking One Design to One Device Connected By One USB Blaster Cable

1. In the Transceiver Toolkit, open the Quartus II project file (.qpf).
2. Link the design instance to the device through the JTAG hardware to which you connect the device.
3. Create the link between channels on the device to test.

Linking Two Designs to Two Separate Devices on Same Board (JTAG Chained), Connected By One USB Blaster Cable

1. In the Transceiver Toolkit, open the Quartus II project file (.qpf) for the first device on the JTAG chain.
2. Link the design instance to the first device on the chain through the JTAG hardware to which you connected the device.
3. Open the project for the second device.
4. Link the second design instance to the second device you use on the JTAG chain.
5. Create a link between the channels on the devices you want to test.

Linking Two Designs to Two Separate Devices on Separate Boards, Connected to Separate USB Blaster Cables

1. In the Transceiver Toolkit, open the Quartus II project file (.qpf) for the device on the first USB blaster.
2. Link the design instance to the device connected to the first USB blaster.
3. Open the project for the device on the second USB blaster.
4. Link the design instance to the device you connected to the second USB blaster.
5. Create a link between the channels on the devices you want to test.

Linking Same Design on Two Separate Devices

1. In the Transceiver Toolkit, open the Quartus II project file (.qpf) you are using on both devices.
2. Right-click on the design to instantiate a second instance.
3. Link the first design instance to the first device, wherever the device is located. Follow the same linking method that you used on previous steps.
4. Link the second design instance to the second device, wherever the device is located. Follow the same linking method that you used on previous steps.
5. Create a link between the channels on the devices you want to test.

Linking Unrelated Designs

Use a combination of the above steps to load multiple Quartus II projects and make links between different systems. You can perform tests on completely separate systems that are not related to one another. All tests run through the same tool instance.



Do not attempt to start multiple instances of the Transceiver Toolkit. You can only control multiple devices and run multiple tests simultaneously through the same instance of the Transceiver Toolkit.

Saving Your Setup As Tcl Scripts

After you open projects and define links for the system so that the entire physical system is correctly described, use the command **Save Tcl Script** to create a setup script.

Close and reopen the Transceiver Toolkit.

Open the scripts folder in **System Explorer** and double-click the script to reload the system. You can also right-click and choose **Run Script**, or use the menu command **Load Script** to run the appropriate script.

Verifying Channels Are Correct When Creating Link

After you load your design and link your hardware, you should verify that the channels you have created are correct and looped back properly on the hardware. You should be able to send the data patterns and receive them correctly.

Perform the following steps before you perform Autosweep or EyeQ tests to verify your link and correct channel, which may save time in the work flow.

1. Assuming that you have completed the system setup, choose the transmitter channel, and click **Control Transmitter Channel**.
2. Set the test pattern to **prbs 7**.
3. Start the pattern generator, press **Start**.
4. Navigate to the control panel, choose the receiver channel, and click **Control Receiver Channel**.
5. Set the test pattern to **prbs 7**.

6. Press **Start**.
7. Verify channels that the channels are correct, based on the following conditions:
 - a. If the **Data Locked** check box goes high, the receiver is receiving data. To verify that the data is coming from the expected transmitter, you can navigate to the transmitter and do either of the following:
 - Click **Stop** on the transmitter and see if **Data Locked** on the receiver turns off.
 - If the receiver shows 0 error bits, click **Inject Error** on the transmitter and see if that error shows up on the receiver.
 - b. If the **Data Locked** check box does not go high, do either of the following:
 - The data quality is too poor to lock. You can manually adjust the PMA settings to see if you can get a lock. If not, use the Autosweep tool if you are certain the channel is correct.
 - The receiver and transmitter are not connected together. You either picked the wrong pair, or you have not made the physical connection between the pair.

After you have verified that the transmitter and receiver are talking to each other, create a link in the link tab with these two transceivers so that you can perform Autosweep and EyeQ tests with this pair.

Using the Recommended DFE Flow

To use the DFE flow recommended by Altera, perform the following steps:

1. Use the Autosweep flow to find optimal PMA settings while leaving the DFE setting **OFF**.
2. Take the best PMA setting achieved, if BER = 0. Then you do not have to do anything if you use this setting.
3. If BER > 0, then use this PMA setting and set minimum and maximum values in the AutoSweep tool to match this setting. Set the **DFE MAX** range to limits for each of the three DFE settings.
4. Run the AutoSweep tool to determine which DFE setting results in the best BER. Use these settings in conjunction with the PMA settings for the best results.

Running Simultaneous Tests

To run link tests simultaneously in one instance of the Transceiver Toolkit, perform the following steps:

1. Set up your system correctly with one of the previous set-up scenarios.
2. In the control panel for the link you work on, run either the Autosweep or EyeQ tool.
3. After you start the test, return to the Transceiver Toolkit control panel.
4. Select the control panel tab.

5. Open the Tools menu and click **Transceiver Toolkit**, which returns you to the control panel.
6. Repeat step 2 until all tests are run.

Conclusion

You gain productivity when optimizing high-speed transceiver links in your board designs with the Transceiver Toolkit and design examples provided from Altera. You can easily set up automatic testing of your transceiver channels so that you can monitor, debug, and optimize transceiver link channels in your board design quicker than before. You then know the optimal PMA settings to use for each channel in your final FPGA design. You can use standard design examples that you download from Altera's website, and then customize the examples to match your own requirements.

Document Revision History

Table 13-3 lists the revision history for this handbook chapter.

Table 13-3. Document Revision History

Date	Version	Changes
December 2010	10.1.0	Changed to new document template. Added new 10.1 release features.
August 2010	10.0.1	Corrected links
July 2010	10.0.0	Initial release

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter provides detailed instructions about how to use SignalProbe to quickly debug your design. The SignalProbe incremental routing feature helps reduce the hardware verification process and time-to-market for system-on-a-programmable-chip (SOPC) designs.

Easy access to internal device signals is important in the design or debugging process. The SignalProbe feature makes design verification more efficient by routing internal signals to I/O pins quickly without affecting the design. When you start with a fully routed design, you can select and route signals for debugging to either previously reserved or currently unused I/O pins.

The SignalProbe feature is supported with the Arria® GX, Stratix® series, Cyclone® series, and MAX® II, device families.

-  The Quartus® II software provides a portfolio of on-chip debugging solutions. For an overview and comparison of all of the tools available in the Quartus II software on-chip debugging tool suite, refer to *Section IV. System Debugging Tools* in volume 3 of the *Quartus II Handbook*.

Debugging Using the SignalProbe Feature

The SignalProbe feature allows you to reserve available pins and route internal signals to those reserved pins, while preserving the behavior of your design. SignalProbe is an effective debugging tool that provides visibility into your FPGA.

You can reserve pins for SignalProbe and assign I/O standards before or after a full compilation. Each SignalProbe-source to SignalProbe-pin connection is implemented as an engineering change order (ECO) change that is applied to your netlist after a full compilation.

To route the internal signals to the device's reserved pins for SignalProbe, perform the following tasks:

1. [Reserve the SignalProbe Pins](#), described on [page 14–2](#).
2. [Perform a Full Compilation](#), described on [page 14–2](#).
3. [Assign a SignalProbe Source](#), described on [page 14–2](#).
4. [Add Registers to the Pipeline Path to SignalProbe Pin](#), described on [page 14–3](#).
5. [Perform a SignalProbe Compilation](#), described on [page 14–3](#).
6. [Analyze the Results of the SignalProbe Compilation](#), described on [page 14–4](#).

Reserve the SignalProbe Pins

SignalProbe pins can only be reserved after compiling your design. You can also reserve any unused I/Os of the device for SignalProbe pins after compilation. Assigning sources is a simple process after reserving SignalProbe pins. The sources for SignalProbe pins are the internal nodes and registers in the post-compilation netlist that you want to probe.

 Although you can reserve SignalProbe pins using many features within the Quartus II software, including the Pin Planner and the Tcl interface, you should use the **SignalProbe Pins** dialog box to create and edit your SignalProbe pins.

 For more information, refer to *About SignalProbe* in Quartus II Help.

Perform a Full Compilation

You must complete a full compilation to generate an internal netlist containing a list of internal nodes to probe to a SignalProbe output pin.

To perform a full compilation, on the Processing menu, click **Start Compilation**.

Assign a SignalProbe Source

A SignalProbe source can be any combinational node, register, or pin in your post-compilation netlist. To find a SignalProbe source, in the Node Finder, use the SignalProbe filter to remove all sources that cannot be probed. You might not be able to find a particular internal node because the node can be optimized away during synthesis, or the node cannot be routed to the SignalProbe pin. For example, nodes and registers within Gigabit transceivers in Stratix IV devices cannot be probed because there are no physical routes available to the pins.

 To probe virtual I/O pins generated in low-level partitions in an incremental compilation flow, select the source of the logic that feeds the virtual pin as your SignalProbe source pin.

 For more information, refer to *SignalProbe Pins Dialog Box* and *Add SignalProbe Pins Dialog Box* in Quartus II Help.

Because SignalProbe pins are implemented and routed as ECOs, turning the **SignalProbe enable** option on or off is the same as selecting **Apply Selected Change** or **Restore Selected Change** in the Change Manager window. (If the Change Manager window is not visible at the bottom of your screen, on the View menu, point to **Utility Windows** and click **Change Manager**.)

 For more information about the Change Manager for the Chip Planner and Resource Property Editor, refer to the *Engineering Change Management with the Chip Planner* chapter in volume 2 of the *Quartus II Handbook*.

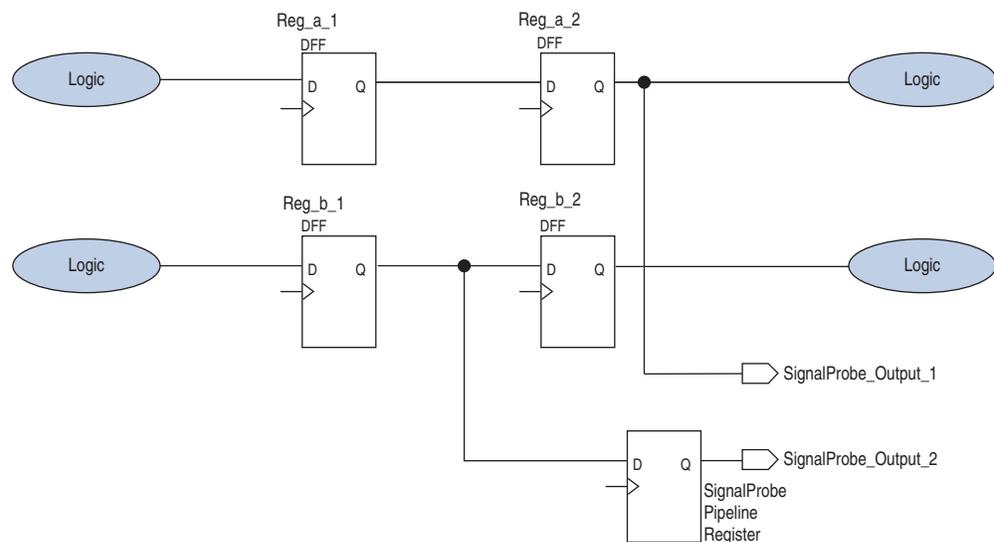
Add Registers to the Pipeline Path to SignalProbe Pin

You can specify the number of registers placed between a SignalProbe source and a SignalProbe pin to synchronize the data with a clock and to control the latency of the SignalProbe outputs. The SignalProbe feature automatically inserts the number of registers specified into the SignalProbe path.

Figure 14-1 shows a single register between the SignalProbe source Reg_b_1 and SignalProbe SignalProbe_Output_2 output pin added to synchronize the data between the two SignalProbe output pins.

 When you add a register to a SignalProbe pin, the SignalProbe compilation attempts to place the register to best fit timing requirements. You can place SignalProbe registers either near the SignalProbe source to meet f_{MAX} requirements, or near the I/O to meet t_{CO} requirements.

Figure 14-1. Synchronizing SignalProbe Outputs with a SignalProbe Register



 To pipeline an existing SignalProbe connection, refer to *Add SignalProbe Pins Dialog Box* in Quartus II Help.

In addition to clock input for pipeline registers, you can also specify a reset signal pin for pipeline registers. To specify a reset pin for pipeline registers, use the Tcl command `make_sp`, as described in “Scripting Support” on page 14-6.

Perform a SignalProbe Compilation

Perform a SignalProbe compilation to route your SignalProbe pins. A SignalProbe compilation saves and checks all netlist changes without recompiling the other parts of the design and completes compilation in a fraction of the time of a full compilation. The design’s current placement and routing are preserved.

To perform a SignalProbe compilation, on the Processing menu, point to **Start** and click **Start SignalProbe Compilation**.

Analyze the Results of the SignalProbe Compilation

After a SignalProbe compilation, the results are available in the compilation report file. Each SignalProbe pin is displayed in the **SignalProbe Fitting Result** page in the **Fitter** section of the Compilation Report. To view the status of each SignalProbe pin in the **SignalProbe Pins** dialog box, on the Tools menu, click **SignalProbe Pins**.

The status of each SignalProbe pin appears in the Change Manager window (Figure 14-2). (If the Change Manager window is not visible at the bottom of your GUI, from the View menu, point to **Utility Windows** and click **Change Manager**.)

Figure 14-2. Change Manager Window with SignalProbe Pins



Index	Node Name	Change Type	Old Value	Target Value	Current Value	Disk Value
1	signalprobe_1	SignalProbe	Disconnected	[/filter/state_m_inst1/filter.idle	[/filter/state_m_inst1/filter.idle	[/filter/state_m_inst1/filter.idle
2	signalprobe_2	SignalProbe	Disconnected	[/filter/state_m_inst1/filter.tap1	[/filter/state_m_inst1/filter.tap1	[/filter/state_m_inst1/filter.tap1
3	signalprobe_3	SignalProbe	Disconnected	[/filter/state_m_inst1/filter.tap2	[/filter/state_m_inst1/filter.tap2	[/filter/state_m_inst1/filter.tap2
4	signalprobe_4	SignalProbe	Disconnected	[/filter/state_m_inst1/filter.tap3	[/filter/state_m_inst1/filter.tap3	[/filter/state_m_inst1/filter.tap3
5	signalprobe_5	SignalProbe	Disconnected	[/filter/state_m_inst1/filter.tap4	[/filter/state_m_inst1/filter.tap4	[/filter/state_m_inst1/filter.tap4

For more information about how to use the Change Manager, refer to the *Engineering Change Management with the Chip Planner* chapter in volume 2 of the *Quartus II Handbook*.

To view the timing results of each successfully routed SignalProbe pin, on the Processing menu, point to **Start** and click **Start Timing Analysis**.

Performing a SignalProbe Compilation

After a full compilation, you can start a SignalProbe compilation either manually or automatically. A SignalProbe compilation performs the following functions:

- Validates SignalProbe pins
- Validates your specified SignalProbe sources
- If applicable, adds registers into SignalProbe paths
- Attempts to route from SignalProbe sources through registers to SignalProbe pins

To run the SignalProbe compilation automatically after a full compilation, on the Tools menu, click **SignalProbe Pins**. In the **SignalProbe Pins** dialog box, click **Start Check & Save All Netlist Changes**.

To run a SignalProbe compilation manually after a full compilation, on the Processing menu, point to **Start** and click **Start SignalProbe Compilation**.

You must run the Fitter before a SignalProbe compilation. The Fitter generates a list of all internal nodes that can be used as SignalProbe sources.

Turn the **SignalProbe enable** option on or off in the **SignalProbe Pins** dialog box to enable or disable each SignalProbe pin.

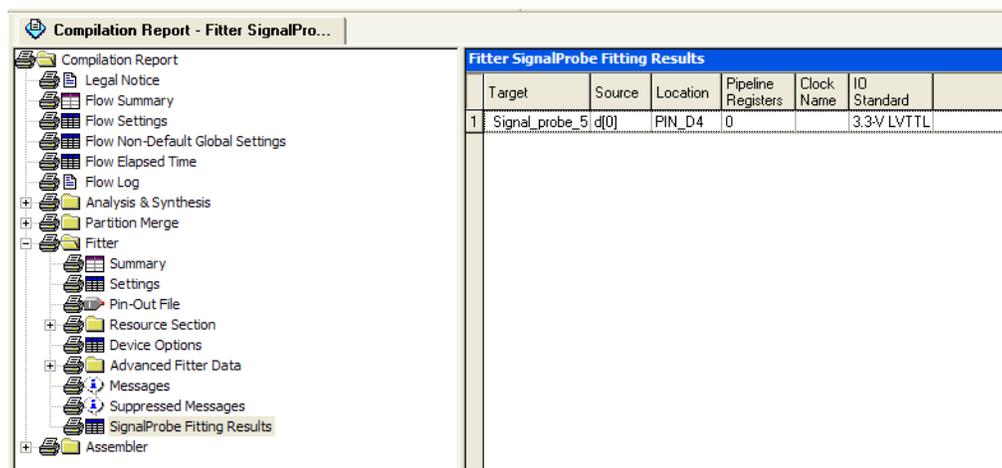
Understanding the Results of a SignalProbe Compilation

After a SignalProbe compilation, the results appear in two sections of the compilation report file. The fitting results and status (Table 14-1) of each SignalProbe pin is displayed in the **SignalProbe Fitting Result** screen in the Fitter section of the Compilation Report (Figure 14-3).

Table 14-1. Status Values

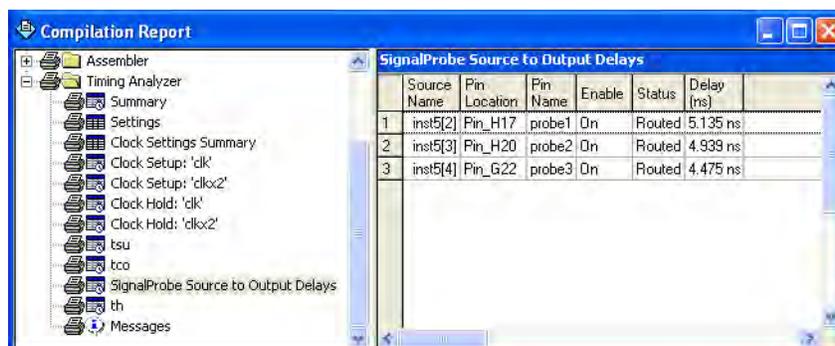
Status	Description
Routed	Connected and routed successfully
Not Routed	Not enabled
Failed to Route	Failed routing during last SignalProbe compilation
Need to Compile	Assignment changed since last SignalProbe compilation

Figure 14-3. SignalProbe Fitting Results Page in the Compilation Report Window



The timing results of each successfully routed SignalProbe pin is displayed in the **SignalProbe source to output delays** screen in the Timing Analysis section of the Compilation Report (Figure 14-4).

Figure 14-4. SignalProbe Source to Output Delays Page in the Compilation Report Window



 After a SignalProbe compilation, the processing screen of the Messages window also provides the results of each SignalProbe pin and displays slack information for each successfully routed SignalProbe pin.

Analyzing SignalProbe Routing Failures

The SignalProbe can begin compilation; however, one of the following reasons can prevent complete compilation:

- **Route unavailable**—the SignalProbe compilation failed to find a route from the SignalProbe source to the SignalProbe pin because of routing congestion
- **Invalid or nonexistent SignalProbe source**—you entered a SignalProbe source that does not exist or is invalid
- **Unusable output pin**—the output pin selected is found to be unusable

Routing failures can occur if the SignalProbe pin's I/O standard conflicts with other I/O standards in the same I/O bank.

If routing congestion prevents a successful SignalProbe compilation, you can allow the compiler to modify routing to the specified SignalProbe source. On the Tools menu, click **SignalProbe Pins** and turn on **Modify latest fitting results during SignalProbe compilation**. This setting allows the Fitter to modify existing routing channels used by your design.

 Turning on **Modify latest fitting results during SignalProbe compilation** can change the performance of your design.

Scripting Support

Running procedures and make settings using a Tcl script are described in this chapter. You can also run some procedures at a command prompt. For detailed information about scripting command options, refer to the Quartus II command-line and Tcl API Help browser. To run the Help browser, type the following command at the command prompt:

```
quartus_sh --qhelp ↵
```

 The Tcl commands in this section are part of the `::quartus::chip_planner` Quartus II Tcl API. Source or include the `::quartus::chip_planner` Tcl package in your scripts to make these commands available.

 For more information about Tcl scripting, refer to the *Tcl Scripting* chapter in volume 2 of the *Quartus II Handbook*. For more information about all settings and constraints in the Quartus II software, refer to the *Quartus II Settings File Reference Manual*. For more information about command-line scripting, refer to the *Command-Line Scripting* chapter in volume 2 of the *Quartus II Handbook*.

Make a SignalProbe Pin

To reserve a SignalProbe pin, type the following command:

```
make_sp [-h | -help] [-long_help] [-clk <clk>] [-io_std <io_std>] \  
-loc <loc> -pin_name <pin name> [-regs <regs>] [-reset <reset>] \  
-src_name <source name>
```

Delete a SignalProbe Pin

To delete a SignalProbe pin, use the following Tcl command:

```
delete_sp [-h | -help] [-long_help] -pin_name <pin name>
```

Enable a SignalProbe Pin

To enable a SignalProbe pin, use the following Tcl command:

```
enable_sp [-h | -help] [-long_help] -pin_name <pin name>
```

Disable a SignalProbe Pin

To disable a SignalProbe pin, use the following Tcl command:

```
disable_sp [-h | -help] [-long_help] -pin_name <pin name>
```

Perform a SignalProbe Compilation

To perform a SignalProbe compilation, type the following command:

```
quartus_sh --flow signalprobe <project name>
```

Script Example

[Example 14-1](#) shows a script that creates a SignalProbe pin called `sp1` and connects the `sp1` pin to source node `reg1` in a project that was already compiled.

Example 14-1. Creating a SignalProbe Pin Called `sp1`

```
package require ::quartus::chip_planner
project_open project
read_netlist
make_sp -pin_name sp1 -src_name reg1
check_netlist_and_save
project_close
```

Reserving SignalProbe Pins

To reserve a SignalProbe pin, add the commands shown in [Example 14-2](#) to the Quartus II Settings File `.qsf` for your project.

Example 14-2. Reserving a SignalProbe Pin

```
set_location_assignment <location> -to <SignalProbe pin name>
set_instance_assignment -name RESERVE_PIN \
"AS SIGNALPROBE OUTPUT" -to <SignalProbe pin name>
```

Valid locations are pin location names, such as `Pin_A3`.

For more information about reserving SignalProbe pins, refer to [“Reserve the SignalProbe Pins” on page 14-2](#).

Common Problems When Reserving a SignalProbe Pin

If you cannot reserve a SignalProbe pin in the Quartus II software, it is likely that one of the following is true:

- You have selected multiple pins.
- A compile is running in the background. Wait until the compilation is complete before reserving the pin.
- You have the Quartus II Web Edition software, in which the SignalProbe feature is not enabled by default. You must turn on TalkBack to enable the SignalProbe feature in the Quartus II Web Edition software.
- You have not set the pin reserve type to **As Signal Probe Output**. To reserve a pin, on the Assignments menu, in the **Assign Pins** dialog box, select **As SignalProbe Output**.
- The pin is reserved from a previous compilation. During a compilation, the Quartus II software reserves each pin on the targeted device. If you end the Quartus II process during a compilation, for example, with the **Windows Task Manager End Process** command or the UNIX `kill` command, perform a full recompilation before reserving pins as SignalProbe outputs.
- The pin does not support the SignalProbe feature. Select another pin.
- The current family does not support the SignalProbe feature.

Adding SignalProbe Sources

Use the following Tcl commands to add SignalProbe sources.

To assign the node name to a SignalProbe pin, use the following Tcl command:

```
set_instance_assignment -name SIGNALPROBE_SOURCE <node name> -to \
<SignalProbe pin name>
```

The next command turns on SignalProbe routing. To turn off individual SignalProbe pins, specify OFF instead of ON with the following command:

```
set_instance_assignment -name SIGNALPROBE_ENABLE ON -to \
<SignalProbe pin name>
```

- ② For more information about adding SignalProbe sources, refer to [SignalProbe Pins Dialog Box](#) and [Add SignalProbe Pins Dialog Box](#) in Quartus II Help.

Assigning I/O Standards

To assign an I/O standard to a pin, use the following Tcl command:

```
set_instance_assignment -name IO_STANDARD <I/O standard> -to \
<SignalProbe pin name>
```

- ② For a list of valid I/O standards, refer to the I/O Standards general description in the Quartus II Help.

Adding Registers for Pipelining

To add registers for pipelining, use the following Tcl command:

```
set_instance_assignment -name SIGNALPROBE_CLOCK <clock name> -to \
<SignalProbe pin name>
```

```
set_instance_assignment \  
-name SIGNALPROBE_NUM_REGISTERS <number of registers> -to \  
<SignalProbe pin name>
```

Run SignalProbe Automatically

To run SignalProbe automatically after a full compile, type the following Tcl command:

```
set_global_assignment -name SIGNALPROBE_DURING_NORMAL_COMPILATION ON
```

For more information about running SignalProbe automatically, refer to [“Performing a SignalProbe Compilation” on page 14-4](#).

Run SignalProbe Manually

To run SignalProbe as part of a scripted flow using Tcl, use the following in your script:

```
execute_flow -signalprobe
```

To perform a Signal Probe compilation interactively at a command prompt, type the following command:

```
quartus_sh_fit --flow signalprobe <project name>
```

For more information about running SignalProbe manually, refer to [“Performing a SignalProbe Compilation” on page 14-4](#).

Enable or Disable All SignalProbe Routing

Use the Tcl command in [Example 14-3](#) to turn on or turn off SignalProbe routing. When using this command, to turn SignalProbe routing on, specify ON. To turn SignalProbe routing off, specify OFF.

Example 14-3. Turning SignalProbe On or Off with Tcl Commands

```
set spe [get_all_assignments -name SIGNALPROBE_ENABLE] \  
foreach_in_collection asgn $spe {  
    set signalprobe_pin_name [lindex $asgn 2]  
    set_instance_assignment -name SIGNALPROBE_ENABLE -to \  
$signalprobe_pin_name <ON|OFF> }  
}
```

For more information about enabling or disabling SignalProbe routing, refer to [page 14-4](#).

Allow SignalProbe to Modify Fitting Results

To turn on **Modify latest fitting results**, type the following Tcl command:

```
set_global_assignment -name SIGNALPROBE_ALLOW_OVERUSE ON
```

For more information, refer to [“Analyzing SignalProbe Routing Failures” on page 14-6](#).

Conclusion

Using the SignalProbe feature can significantly reduce the time required compared to a full recompilation. Use the SignalProbe feature for quick access to internal design signals to perform system-level debugging.

Document Revision History

Table 14-2 shows the revision history for this chapter.

Table 14-2. Document Revision History

Date	Version	Changes
December 2010	10.0.1	Changed to new document template.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Revised for new UI. ■ Removed section SignalProbe ECO flows ■ Removed support for SignalProbe pin preservation when recompiling with incremental compilation turned on. ■ Removed outdated FAQ section. ■ Added links to Quartus II Help for procedural content.
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Removed all references and procedures for APEX devices. ■ Style changes.
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Removed the “Generate the Programming File” section ■ Removed unnecessary screenshots ■ Minor editorial updates
November 2008	8.1.0	<ul style="list-style-type: none"> ■ Modified description for preserving SignalProbe connections when using Incremental Compilation ■ Added plausible scenarios where SignalProbe connections are not reserved in the design
May 2008	8.0.0	<ul style="list-style-type: none"> ■ Added “Arria GX” to the list of supported devices ■ Removed the “On-Chip Debugging Tool Comparison” and replaced with a reference to the Section V Overview on page 13-1 ■ Added hyperlinks to referenced documents throughout the chapter ■ Minor editorial updates

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

To help with the process of design debugging, Altera provides a solution that allows you to examine the behavior of internal signals, without using extra I/O pins, while the design is running at full speed on an FPGA device.

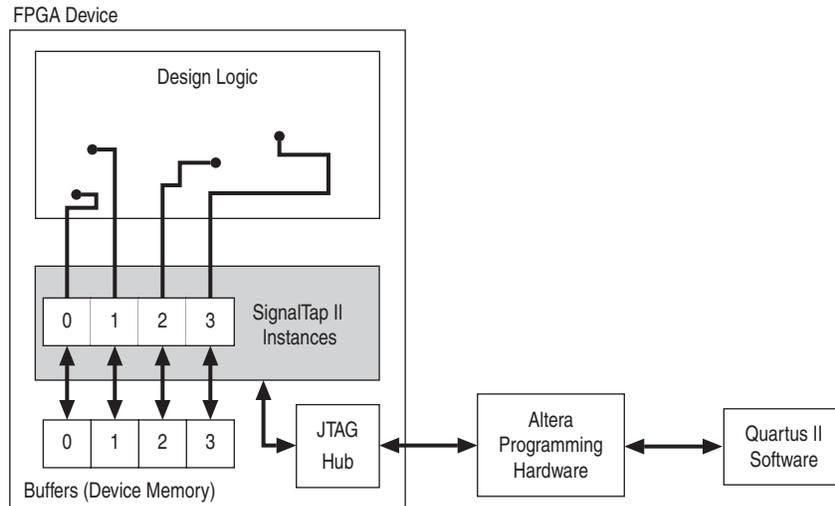
The SignalTap® II Logic Analyzer is scalable, easy to use, and is available as a stand-alone package or included with the Quartus® II software subscription. This logic analyzer helps debug an FPGA design by probing the state of the internal signals in the design without the use of external equipment. Defining custom trigger-condition logic provides greater accuracy and improves the ability to isolate problems. The SignalTap II Logic Analyzer does not require external probes or changes to the design files to capture the state of the internal nodes or I/O pins in the design. All captured signal data is conveniently stored in device memory until you are ready to read and analyze the data.

The topics in this chapter include:

- “Design Flow Using the SignalTap II Logic Analyzer” on page 15–4
- “SignalTap II Logic Analyzer Task Flow” on page 15–5
- “Configure the SignalTap II Logic Analyzer” on page 15–8
- “Define Triggers” on page 15–25
- “Compile the Design” on page 15–44
- “Program the Target Device or Devices” on page 15–49
- “Run the SignalTap II Logic Analyzer” on page 15–50
- “View, Analyze, and Use Captured Data” on page 15–55
- “Other Features” on page 15–61
- “Design Example: Using SignalTap II Logic Analyzers in SOPC Builder Systems” on page 15–66
- “Custom Triggering Flow Application Examples” on page 15–67
- “SignalTap II Scripting Support” on page 15–69

The SignalTap II Logic Analyzer is a next-generation, system-level debugging tool that captures and displays real-time signal behavior in a system-on-a-programmable-chip (SOPC) or any FPGA design. The SignalTap II Logic Analyzer supports the highest number of channels, largest sample depth, and fastest clock speeds of any logic analyzer in the programmable logic market. [Figure 15-1](#) shows a block diagram of the components that make up the SignalTap II Logic Analyzer.

Figure 15-1. SignalTap II Logic Analyzer Block Diagram *(Note 1)*



Note to Figure 15-1:

- (1) This diagram assumes that you compiled the SignalTap II Logic Analyzer with the design as a separate design partition using the Quartus II incremental compilation feature. This is the default setting for new projects in the Quartus II software. If incremental compilation is disabled or not used, the SignalTap II logic is integrated with the design. For information about the use of incremental compilation with SignalTap II, refer to [“Faster Compilations with Quartus II Incremental Compilation”](#) on page 15-45.

This chapter is intended for any designer who wants to debug an FPGA design during normal device operation without the need for external lab equipment. Because the SignalTap II Logic Analyzer is similar to traditional external logic analyzers, familiarity with external logic analyzer operations is helpful but not necessary. To take advantage of faster compile times when making changes to the SignalTap II Logic Analyzer, knowledge of the Quartus II incremental compilation feature is helpful.

 For information about using the Quartus II incremental compilation feature, refer to the [Quartus II Incremental Compilation for Hierarchical and Team-Based Design](#) chapter in volume 1 of the *Quartus II Handbook*.

Hardware and Software Requirements

You need the following components to perform logic analysis with the SignalTap II Logic Analyzer:

- Quartus II design software
- or
- Quartus II Web Edition (with the TalkBack feature enabled)
- or
- SignalTap II Logic Analyzer standalone software

- Download/upload cable
- Altera® development kit or your design board with JTAG connection to device under test



The Quartus II software Web Edition does not support the SignalTap II Logic Analyzer with the incremental compilation feature.

The device's memory blocks store captured data and transfers the data to the Quartus II software waveform display with a JTAG communication cable, such as EthernetBlaster or USB-Blaster™. Table 15-1 summarizes features and benefits of the SignalTap II Logic Analyzer.

Table 15-1. SignalTap II Logic Analyzer Features and Benefits

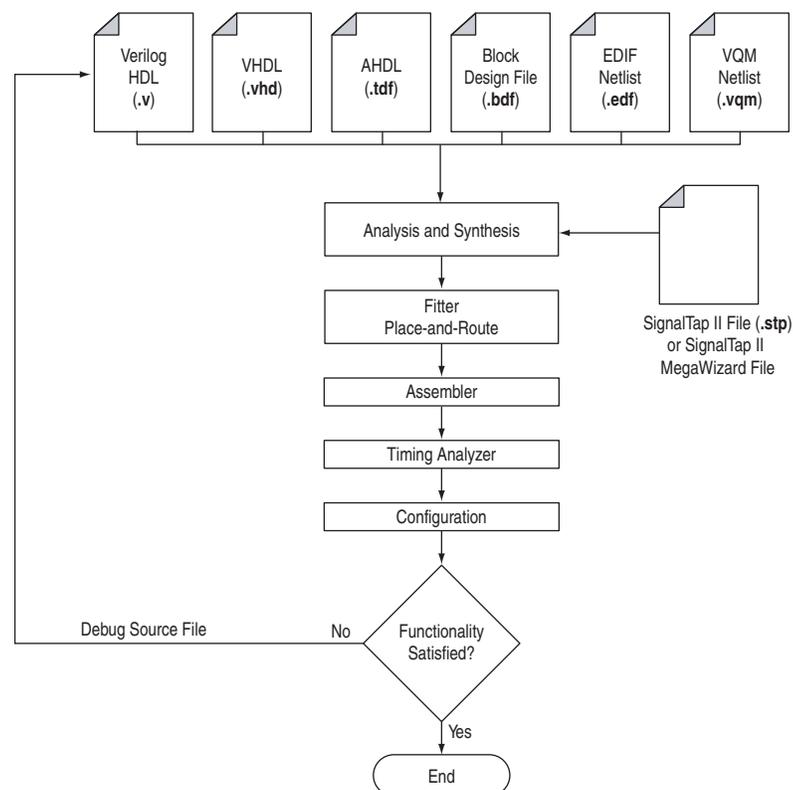
Feature	Benefit
Multiple logic analyzers in a single device	Captures data from multiple clock domains in a design at the same time.
Multiple logic analyzers in multiple devices in a single JTAG chain	Simultaneously captures data from multiple devices in a JTAG chain.
Plug-In Support	Easily specifies nodes, triggers, and signal mnemonics for IP, such as the Nios® II processor.
Up to 10 basic or advanced trigger conditions for each analyzer instance	Enables sending more complex data capture commands to the logic analyzer, providing greater accuracy and problem isolation.
Power-Up Trigger	Captures signal data for triggers that occur after device programming, but before manually starting the logic analyzer.
State-based Triggering Flow	Enables you to organize your triggering conditions to precisely define what your logic analyzer captures.
Incremental compilation	Modifies the SignalTap II Logic Analyzer monitored signals and triggers without performing a full compilation, saving time.
Flexible buffer acquisition modes	The buffer acquisition control allows you to precisely control the data that is written into the acquisition buffer. Both segmented buffers and non-segmented buffers with storage qualification allow you to discard data samples that are not relevant to the debugging of your design.
MATLAB integration with included MEX function	Collects the SignalTap II Logic Analyzer captured data into a MATLAB integer matrix.
Up to 2,048 channels per logic analyzer instance	Samples many signals and wide bus structures.
Up to 128K samples in each device	Captures a large sample set for each channel.
Fast clock frequencies	Synchronous sampling of data nodes using the same clock tree driving the logic under test.
Resource usage estimator	Provides estimate of logic and memory device resources used by SignalTap II Logic Analyzer configurations.
No additional cost	The SignalTap II Logic Analyzer is included with a Quartus II subscription and with the Quartus II Web Edition (with TalkBack enabled).
Compatibility with other on-chip debugging utilities	You can use the SignalTap II Logic Analyzer in tandem with any JTAG based on-chip debugging tool, such as an In-System Memory Content editor allowing you to change signal values in real-time while you are running an analysis with the SignalTap II Logic Analyzer.

- The Quartus II software offers a portfolio of on-chip debugging solutions. For an overview and comparison of all tools available in the In-System Verification Tool set, refer to *Section V. In-System Design Debugging*.

Design Flow Using the SignalTap II Logic Analyzer

Figure 15-2 shows a typical overall FPGA design flow for using the SignalTap II Logic Analyzer in your design. A SignalTap II file (.stp) is added to and enabled in your project, or a SignalTap II HDL function, created with the MegaWizard™ Plug-In Manager, is instantiated in your design. The figure shows the flow of operations from initially adding the SignalTap II Logic Analyzer to your design to final device configuration, testing, and debugging.

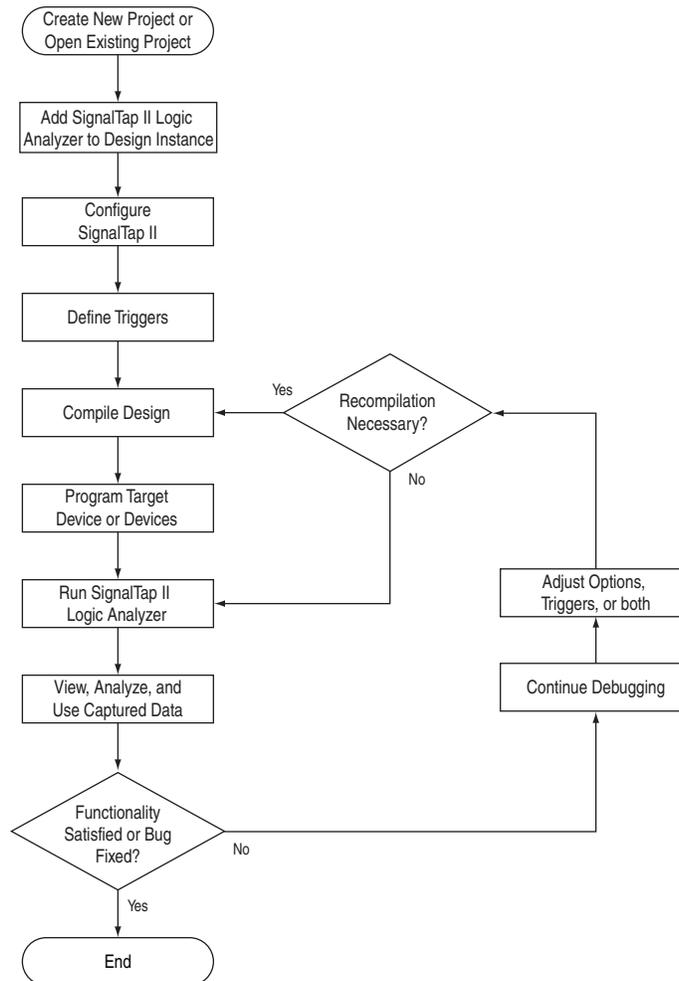
Figure 15-2. SignalTap II FPGA Design and Debugging Flow



SignalTap II Logic Analyzer Task Flow

To use the SignalTap II Logic Analyzer to debug your design, you perform a number of tasks to add, configure, and run the logic analyzer. Figure 15-3 shows a typical flow of the tasks you complete to debug your design. Refer to the appropriate section of this chapter for more information about each of these tasks.

Figure 15-3. SignalTap II Logic Analyzer Task Flow



Add the SignalTap II Logic Analyzer to Your Design

Create an `.stp` or create a parameterized HDL instance representation of the logic analyzer using the MegaWizard Plug-In Manager. If you want to monitor multiple clock domains simultaneously, add additional instances of the logic analyzer to your design, limited only by the available resources in your device.

- ④ For information about creating an `.stp`, refer to *Setting Up the SignalTap II Logic Analyzer* in Quartus II Help.

Configure the SignalTap II Logic Analyzer

After you add the SignalTap II Logic Analyzer to your design, configure it to monitor the signals you want. You can manually add signals or use a plug-in, such as the Nios II processor plug-in, to quickly add entire sets of associated signals for a particular intellectual property (IP). You can also specify settings for the data capture buffer, such as its size, the method in which data is captured and stored, and the device memory type to use for the buffer in devices that support memory type selection.

- ① For information about configuring the SignalTap II Logic Analyzer, refer to *Setting Up the SignalTap II Logic Analyzer* in Quartus II Help.

Define Trigger Conditions

The SignalTap II Logic Analyzer captures data continuously while it is running. To capture and store specific signal data, set up triggers that tell the logic analyzer under what conditions to stop capturing data. The SignalTap II Logic Analyzer lets you define trigger conditions that range from very simple, such as the rising edge of a single signal, to very complex, involving groups of signals, extra logic, and multiple conditions. Power-Up Triggers allow you to capture data from trigger events occurring immediately after the device enters user-mode after configuration.

- ① For information about defining trigger conditions, refer to *Setting Up the SignalTap II Logic Analyzer* in Quartus II Help.

Compile the Design

With the `.stp` configured and trigger conditions defined, compile your project as usual to include the logic analyzer in your design. Because you may need to change monitored signal nodes or adjust trigger settings frequently during debugging, Altera recommends that you use the incremental compilation feature built into the SignalTap II Logic Analyzer, along with Quartus II incremental compilation, to reduce recompile times.

- ① For information about compiling your design, refer to *Compiling a Design that Contains a SignalTap II Logic Analyzer* in Quartus II Help.

Program the Target Device or Devices

When you are debugging a design with the SignalTap II Logic Analyzer, you can program a target device directly from the `.stp` without using the Quartus II Programmer. You can also program multiple devices with different designs and simultaneously debug them.

-  The SignalTap II Logic Analyzer supports all current Altera FPGA device families including Arria, Cyclone, HardCopy, and Stratix devices.

- ① For instructions on programming devices in the Quartus II software, refer to *Running the SignalTap II Logic Analyzer* in Quartus II Help.

Run the SignalTap II Logic Analyzer

In normal device operation, you control the logic analyzer through the JTAG connection, specifying when to start looking for trigger conditions to begin capturing data. With Runtime or Power-Up Triggers, read and transfer the captured data from the on-chip buffer to the `.stp` for analysis.

- For information about analyzing results from the SignalTap II Logic Analyzer, refer to *Analyzing Data in the SignalTap II Logic Analyzer* in Quartus II Help.

View, Analyze, and Use Captured Data

After you have captured data and read it into the `.stp`, that data is available for analysis and debugging. Set up mnemonic tables, either manually or with a plug-in, to simplify reading and interpreting the captured signal data. To speed up debugging, use the **Locate** feature in the **SignalTap II node** list to find the locations of problem nodes in other tools in the Quartus II software. Save the captured data for later analysis, or convert the data to other formats for sharing and further study.

- For information about analyzing results from the SignalTap II Logic Analyzer, refer to *Analyzing Data in the SignalTap II Logic Analyzer* in Quartus II Help.

Embedding Multiple Analyzers in One FPGA

The SignalTap II Logic Analyzer Editor includes support for adding multiple logic analyzers by creating instances in the `.stp`. You can create a unique logic analyzer for each clock domain in the design.

- For information about creating instances, refer to *Running the SignalTap II Logic Analyzer* in Quartus II Help.

Monitoring FPGA Resources Used by the SignalTap II Logic Analyzer

The SignalTap II Logic Analyzer has a built-in resource estimator that calculates the logic resources and amount of memory that each logic analyzer instance uses. Furthermore, because the most demanding on-chip resource for the logic analyzer is memory usage, the resource estimator reports the ratio of total RAM usage in your design to the total amount of RAM available, given the results of the last compilation. The resource estimator provides a warning if a potential for a “no-fit” occurs.

You can see resource usage of each logic analyzer instance and total resources used in the columns of the **Instance Manager** pane of the SignalTap II Logic Analyzer Editor. Use this feature when you know that your design is running low on resources.

The logic element value reported in the resource usage estimator may vary by as much as 10% from the actual resource usage.

Table 15-2 shows the SignalTap II Logic Analyzer M4K memory block resource usage for the listed devices per signal width and sample depth.

Table 15-2. SignalTap II Logic Analyzer M4K Block Utilization (Note 1)

Signals (Width)	Samples (Depth)			
	256	512	2,048	8,192
8	< 1	1	4	16
16	1	2	8	32
32	2	4	16	64
64	4	8	32	128
256	16	32	128	512

Note to Table 15-2:

- (1) When you configure a SignalTap II Logic Analyzer, the Instance Manager reports an estimate of the memory bits and logic elements required to implement the given configuration.

Using the MegaWizard Plug-In Manager to Create Your Logic Analyzer

You can create a SignalTap II Logic Analyzer instance by using the MegaWizard Plug-In Manager. The MegaWizard Plug-In Manager generates an HDL file that you instantiate in your design.

 The State-based trigger flow, the state machine debugging feature, and the storage qualification feature are not supported when using the MegaWizard Plug-In Manager to create the logic analyzer. These features are described in the following sections:

- “Adding Finite State Machine State Encoding Registers” on page 15-13
- “Using the Storage Qualifier Feature” on page 15-17
- “State-Based Triggering” on page 15-29

-  For information about creating a SignalTap II instance with the MegaWizard Plug-In Manager, refer to *Setting Up the SignalTap II Logic Analyzer* in Quartus II Help.

Configure the SignalTap II Logic Analyzer

The `.stp` provides many options for configuring instances of the logic analyzer. Some of the settings are similar to those found on traditional external logic analyzers. Other settings are unique to the SignalTap II Logic Analyzer because of the requirements for configuring a logic analyzer. All settings allow you to configure the logic analyzer the way you want to help debug your design.

 Some settings can only be adjusted when you are viewing Run-Time Trigger conditions instead of Power-Up Trigger conditions. To learn about Power-Up Triggers and viewing different trigger conditions, refer to “Creating a Power-Up Trigger” on page 15-40.

Assigning an Acquisition Clock

Assign a clock signal to control the acquisition of data by the SignalTap II Logic Analyzer. The logic analyzer samples data on every positive (rising) edge of the acquisition clock. The logic analyzer does not support sampling on the negative (falling) edge of the acquisition clock. You can use any signal in your design as the acquisition clock. However, for best results, Altera recommends that you use a global, non-gated clock synchronous to the signals under test for data acquisition. Using a gated clock as your acquisition clock can result in unexpected data that does not accurately reflect the behavior of your design. The Quartus II static timing analysis tools show the maximum acquisition clock frequency at which you can run your design. Refer to the Timing Analysis section of the Compilation Report to find the maximum frequency of the logic analyzer clock.

 For information about assigning an acquisition clock, refer to *Working with Nodes in the SignalTap II Logic Analyzer* in Quartus II Help.

 Altera recommends that you exercise caution when using a recovered clock from a transceiver as an acquisition clock for the SignalTap II Logic Analyzer. Incorrect or unexpected behavior has been noted, particularly when a recovered clock from a transceiver is used as an acquisition clock with the power-up trigger feature.

If you do not assign an acquisition clock in the SignalTap II Logic Analyzer Editor, the Quartus II software automatically creates a clock pin called `auto_stp_external_clk`.

You must make a pin assignment to this pin independently from the design. Ensure that a clock signal in your design drives the acquisition clock.

 For information about assigning signals to pins, refer to the *I/O Management* chapter in volume 2 of the *Quartus II Handbook*.

Adding Signals to the SignalTap II File

While configuring the logic analyzer, add signals to the node list in the `.stp` to select which signals in your design you want to monitor. Selected signals are also used to define triggers. You can assign the following two types of signals to your `.stp` file:

- **Pre-synthesis**—This signal exists after design elaboration, but before any synthesis optimizations are done. This set of signals should reflect your Register Transfer Level (RTL) signals.
- **Post-fitting**—This signal exists after physical synthesis optimizations and place-and-route.

 If you are not using incremental compilation, add only pre-synthesis signals to your `.stp`. Using pre-synthesis is particularly useful if you want to add a new node after you have made design changes. Source file changes appear in the Node Finder after an Analysis and Elaboration has been performed. On the Processing Menu, point to **Start** and click **Start Analysis & Elaboration**.

 For more information about incremental compilation, refer to *About Incremental Compilation* in Quartus II Help.

The Quartus II software does not limit the number of signals available for monitoring in the SignalTap II window waveform display. However, the number of channels available is directly proportional to the number of logic elements (LEs) or adaptive logic modules (ALMs) in the device. Therefore, there is a physical restriction on the number of channels that are available for monitoring. Signals shown in blue text are post-fit node names. Signals shown in black text are pre-synthesis node names.

After successful Analysis and Elaboration, the signals shown in red text are invalid signals. Unless you are certain that these signals are valid, remove them from the `.stp` for correct operation. The SignalTap II Status Indicator also indicates if an invalid node name exists in the `.stp`.

As a general guideline, signals can be tapped if a routing resource (row or column interconnects) exists to route the connection to the SignalTap II instance. For example, signals that exist in the I/O element (IOE) cannot be directly tapped because there are no direct routing resources from the signal in an IOE to a core logic element. For input pins, you can tap the signal that is driving a logic array block (LAB) from an IOE, or, for output pins, you can tap the signal from the LAB that is driving an IOE.

When adding pre-synthesis signals, all connections made to the SignalTap II Logic Analyzer are made prior to synthesis. Logic and routing resources are allocated during recompilation to make the connection as if a change in your design files had been made. As such, pre-synthesis signal names for signals driving to and from IOEs coincide with the signal names assigned to the pin.

In the case of post-fit signals, connections that you make to the SignalTap II Logic Analyzer are the signal names from the actual atoms in your post-fit netlist. A connection can only be made if the signals are part of the existing post-fit netlist and existing routing resources are available from the signal of interest to the SignalTap II Logic Analyzer. In the case of post-fit output signals, tap the `COMBOUT` or `REGOUT` signal that drives the IOE block. For post-fit input signals, signals driving into the core logic coincide with the signal name assigned to the pin.



If you are tapping the signal from the atom that is driving an IOE, be aware that the signal may be inverted due to NOT-gate push back. You can check this by locating the signal in either the Resource Property Editor or the Technology Map Viewer. The Technology Map viewer and the Resource Property Editor are also helpful in finding post-fit node names.



For information about cross-probing to source design files and other Quartus II windows, refer to the *Analyzing Designs with Quartus II Netlist Viewers* chapter in volume 1 of the *Quartus II Handbook*.

For more information about the use of incremental compilation with the SignalTap II Logic Analyzer, refer to “[Faster Compilations with Quartus II Incremental Compilation](#)” on page 15-45.

Signal Preservation

Many of the RTL signals are optimized during the process of synthesis and place-and-route. RTL signal names frequently may not appear in the post-fit netlist after optimizations. For example, the compilation process can add tildes (“~”) to nets that are fanning out from a node, making it difficult to decipher which signal nets they actually represent. This can cause a problem when you use the incremental

compilation flow with the SignalTap II Logic Analyzer. Because only post-fitting signals can be added to the SignalTap II Logic Analyzer in partitions of type **post-fit**, RTL signals that you want to monitor may not be available, preventing their usage. To avoid this issue, use synthesis attributes to preserve signals during synthesis and place-and-route. When the Quartus II software encounters these synthesis attributes, it does not perform any optimization on the specified signals, forcing them to continue to exist in the post-fit netlist. However, if you do this, you could see an increase in resource utilization or a decrease in timing performance. The two attributes you can use are:

- **keep**—Ensures that combinational signals are not removed
- **preserve**—Ensures that registers are not removed

 For more information about using these attributes, refer to the *Quartus II Integrated Synthesis* chapter in volume 1 of the *Quartus II Handbook*.

If you are debugging an IP core, such as the Nios II CPU or other encrypted IP, you might need to preserve nodes from the core to make them available for debugging with the SignalTap II Logic Analyzer. This is often necessary when a plug-in is used to add a group of signals for a particular IP.

If you use incremental compilation flow with the SignalTap II Logic Analyzer, pre-synthesis nodes may not be connected to the SignalTap II Logic Analyzer if the affected partition is of the post-fit type. A critical warning is issued for all pre-synthesis node names that are not found in the post-fit netlist.

 For more information about node preservation or how to avoiding these warnings, refer to *Working with Nodes in the SignalTap II Logic Analyzer* in Quartus II Help.

Assigning Data Signals Using the Technology Map Viewer

You can easily add post-fit signal names that you find in the Technology map viewer. To do so, launch the Technology map viewer (post-fitting) after compiling your design. When you find the desired node, copy the node to either the active **.stp** for your design or a new **.stp**.

Node List Signal Use Options

When a signal is added to the node list, you can select options that specify how the signal is used with the logic analyzer. You can turn off the ability of a signal to trigger the analyzer by disabling the **Trigger Enable** option for that signal in the node list in the **.stp**. This option is useful when you want to see only the captured data for a signal and you are not using that signal as part of a trigger.

You can turn off the ability to view data for a signal by disabling the **Data Enable** column. This option is useful when you want to trigger on a signal, but have no interest in viewing data for that signal.

For information about using signals in the node list to create SignalTap II trigger conditions, refer to “*Define Triggers*” on page 15-25.

Untappable Signals

Not all of the post-fitting signals in your design are available in the **SignalTap II : post-fitting filter** in the **Node Finder** dialog box. The following signal types cannot be tapped:

- **Post-fit output pins**—You cannot tap a post-fit output pin directly. To make an output signal visible, tap the register or buffer that drives the output pin. This includes pins defined as bidirectional.
- **Signals that are part of a carry chain**—You cannot tap the carry out (`cout0` or `cout1`) signal of a logic element. Due to architectural restrictions, the carry out signal can only feed the carry in of another LE.
- **JTAG Signals**—You cannot tap the JTAG control (`TCK`, `TDI`, `TDO`, and `TMS`) signals.
- **ALTGXB megafunction**—You cannot directly tap any ports of an ALTGXB instantiation.
- **LVDS**—You cannot tap the data output from a serializer/deserializer (SERDES) block.
- **DQ, DQS Signals**—You cannot directly tap the `DQ` or `DQS` signals in a DDR/DDR2 design.

Adding Signals with a Plug-In

Instead of adding individual or grouped signals through the **Node Finder**, you can add groups of relevant signals of a particular type of IP with a plug-in. The SignalTap II Logic Analyzer comes with one plug-in already installed for the Nios II processor. Besides easy signal addition, plug-ins also provide features such as pre-designed mnemonic tables, useful for trigger creation and data viewing, as well as the ability to disassemble code in captured data.

The Nios II plug-in, for example, creates one mnemonic table in the **Setup** tab and two tables in the **Data** tab:

- **Nios II Instruction (Setup tab)**—Capture all the required signals for triggering on a selected instruction address.
- **Nios II Instance Address (Data tab)**—Display address of executed instructions in hexadecimal format or as a programming symbol name if defined in an optional Executable and Linking Format (`.elf`) file.
- **Nios II Disassembly (Data tab)**—Displays disassembled code from the corresponding address.

For information about the other features plug-ins provided, refer to [“Define Triggers” on page 15-25](#) and [“View, Analyze, and Use Captured Data” on page 15-55](#).

To add signals to the `.stp` using a plug-in, perform the following steps after running Analysis and Elaboration on your design:

1. Right-click in the node list. On the Add Nodes with Plug-In submenu, choose the plug-in you want to use, such as the included plug-in named **Nios II**.



If the IP for the selected plug-in does not exist in your design, a message informs you that you cannot use the selected plug-in.

2. The **Select Hierarchy Level** dialog box appears showing the IP hierarchy of your design. Select the IP that contains the signals you want to monitor with the plug-in and click **OK**.
3. If all the signals in the plug-in are available, a dialog box might appear, depending on the plug-in selected, where you can specify options for the plug-in. With the Nios II plug-in, you can optionally select an `.elf` containing program symbols from your Nios II Integrated Development Environment (IDE) software design. Specify options for the selected plug-in as desired and click **OK**.

 To make sure all the required signals are available, in the Quartus II **Analysis & Synthesis** settings, turn on **Create debugging nodes for IP cores**.

All the signals included in the plug-in are added to the node list.

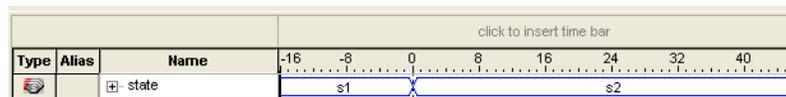
Adding Finite State Machine State Encoding Registers

Finding the signals to debug Finite State Machines (FSM) can be challenging. Finding nodes from the post-fit netlist may be impossible, as FSM encoding signals may be changed or optimized away during synthesis and place-and-route. If you can find all of the relevant nodes in the post-fit netlist or you used the nodes from the pre-synthesis netlist, an additional step is required to find and map FSM signal values to the state names that you specified in your HDL.

The SignalTap II Logic Analyzer GUI can detect FSMs in your compiled design. The SignalTap II Logic Analyzer configuration automatically tracks the FSM state signals as well as state encoding through the compilation process. Shortcut menu commands from the SignalTap II Logic Analyzer GUI allow you to add all of the FSM state signals to your logic analyzer with a single command. For each FSM added to your SignalTap II configuration, the FSM debugging feature adds a mnemonic table to map the signal values to the state enumeration that you provided in your source code. The mnemonic tables enable you to visualize state machine transitions in the waveform viewer. The FSM debugging feature supports adding FSM signals from both the pre-synthesis and post-fit netlists.

Figure 15-4 shows the waveform viewer with decoded signal values from a state machine added with the FSM debugging feature.

Figure 15-4. Decoded FSM Mnemonics



-  For coding guidelines for specifying FSM in Verilog and VHDL, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*.
-  For information about adding FSM signals to the configuration file, refer to *Setting Up the SignalTap II Logic Analyzer* in Quartus II Help.

Modifying and Restoring Mnemonic Tables for State Machines

When you add FSM state signals via the FSM debugging feature, the SignalTap II Logic Analyzer GUI creates a mnemonic table using the format `<StateSignalName>_table`, where **StateSignalName** is the name of the state signals that you have declared in your RTL. You can edit any mnemonic table using the **Mnemonic Table Setup** dialog box.

If you want to restore a mnemonic table that was modified, right-click anywhere in the node list window and select **Recreate State Machine Mnemonics**. By default, restoring a mnemonic table overwrites the existing mnemonic table that you modified. To restore a FSM mnemonic table to a new record, turn off **Overwrite existing mnemonic table** in the **Recreate State Machine Mnemonics** dialog box.



If you have added or deleted a signal from the FSM state signal group from within the setup tab, delete the modified register group and add the FSM signals back again.

For more information about using Mnemonics, refer to [“Creating Mnemonics for Bit Patterns” on page 15-59](#).

Additional Considerations

The SignalTap II configuration GUI recognizes state machines from your design only if you use Quartus II Integrated Synthesis (QIS). The state machine debugging feature is not able to track the FSM signals or state encoding if you use other EDA synthesis tools.

If you add post-fit FSM signals, the SignalTap II Logic Analyzer FSM debug feature may not track all optimization changes that are a part of the compilation process. If the following two specific optimizations are enabled, the SignalTap II FSM debug feature may not list mnemonic tables for state machines in the design:

- If you have physical synthesis turned on, state registers may be resource balanced (register retiming) to improve f_{MAX} . The FSM debug feature does not list post-fit FSM state registers if register retiming occurs.
- The FSM debugging feature does not list state signals that have been packed into RAM and DSP blocks during QIS or Fitter optimizations.

You can still use the FSM debugging feature to add pre-synthesis state signals.

Specifying the Sample Depth

The sample depth specifies the number of samples that are captured and stored for each signal in the captured data buffer. To specify the sample depth, select the desired number of samples to store in the **Sample Depth** list. The sample depth ranges from 0 to 128K.

If device memory resources are limited, you may not be able to successfully compile your design with the sample buffer size you have selected. Try reducing the sample depth to reduce resource usage.

Capturing Data to a Specific RAM Type

When you use the SignalTap II Logic Analyzer with some devices, you have the option to select the RAM type where acquisition data is stored. Once SignalTap II Logic Analyzer is allocated to a particular RAM block, the entire RAM block becomes a dedicated resource for the logic analyzer. RAM selection allows you to preserve a specific memory block for your design and allocate another portion of memory for SignalTap II Logic Analyzer data acquisition. For example, if your design has an application that requires a large block of memory resources, such a large instruction or data cache, you would choose to use MLAB, M512, or M4k blocks for data acquisition and leave the M9k blocks for the rest of your design.

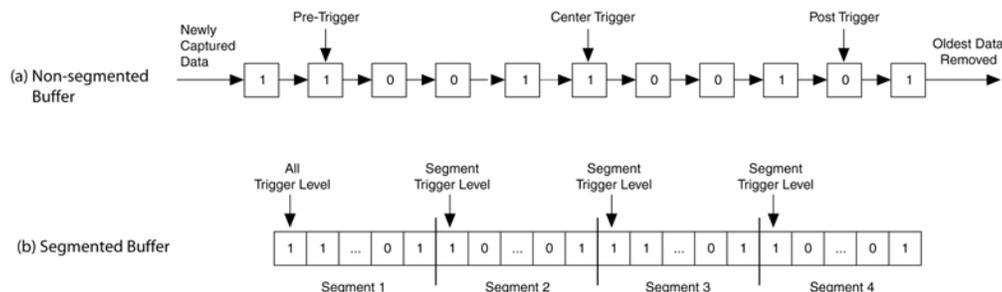
To select the RAM type to use for the SignalTap II Logic Analyzer buffer, select it from the RAM type list. Use this feature when the acquired data (as reported by the SignalTap II resource estimator) is not larger than the available memory of the memory type that you have selected in the FPGA.

Choosing the Buffer Acquisition Mode

The Buffer Acquisition Type Selection feature in the SignalTap II Logic Analyzer lets you choose how the captured data buffer is organized and can potentially reduce the amount of memory that is required for SignalTap II data acquisition. There are two types of acquisition buffer within the SignalTap II Logic Analyzer—a non-segmented buffer and a segmented buffer. With a non-segmented buffer, the SignalTap II Logic Analyzer treats entire memory space as a single FIFO, continuously filling the buffer until the logic analyzer reaches a defined set of trigger conditions. With a segmented buffer, the memory space is split into a number of separate buffers. Each buffer acts as a separate FIFO with its own set of trigger conditions. Only a single buffer is active during an acquisition. The SignalTap II Logic Analyzer advances to the next segment after the trigger condition or conditions for the active segment has been reached.

When using a non-segmented buffer, you can use the storage qualification feature to determine which samples are written into the acquisition buffer. Both the segmented buffers and the non-segmented buffer with the storage qualification feature help you maximize the use of the available memory space. Figure 15-5 illustrates the differences between the two buffer types.

Figure 15-5. Buffer Type Comparison in the SignalTap II Logic Analyzer (Note 1), (Note 2)



Notes to Figure 15-5:

- (1) Both non-segmented and segmented buffers can use a predefined trigger (Pre-Trigger, Center Trigger, Post-Trigger) position or define a custom trigger position using the **State-Based Triggering** tab. Refer to “[Specifying the Trigger Position](#)” on page 15-39 for more details.
- (2) Each segment is treated like a FIFO, and behaves as the non-segmented buffer shown in (a).

For more information about the storage qualification feature, refer to “Using the Storage Qualifier Feature” on page 15-17.

Non-Segmented Buffer

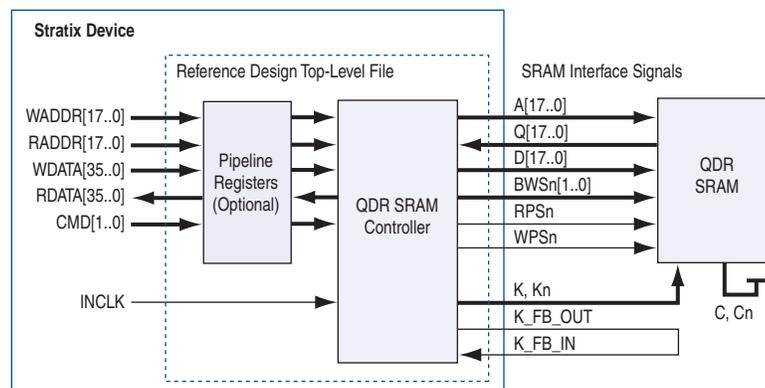
The non-segmented buffer (also known as a circular buffer) shown in Figure 15-5 (a) is the default buffer type used by the SignalTap II Logic Analyzer. While the logic analyzer is running, data is stored in the buffer until it fills up, at which point new data replaces the oldest data. This continues until a specified trigger event, consisting of a set of trigger conditions, occurs. When the trigger event happens, the logic analyzer continues to capture data after the trigger event until the buffer is full, based on the trigger position setting in the **Signal Configuration** pane in the `.stp`. To capture the majority of the data before the trigger occurs, select **Post trigger position** from the list. To capture the majority of the data after the trigger, select **Pre-trigger position**. To center the trigger position in the data, select **Center trigger position**. Alternatively, use the custom State-based triggering flow to define a custom trigger position within the capture buffer.

For more information, refer to “Specifying the Trigger Position” on page 15-39.

Segmented Buffer

A segmented buffer allows you to debug systems that contain relatively infrequent recurring events. The acquisition memory is split into evenly sized segments, with a set of trigger conditions defined for each segment. Each segment acts as a non-segmented buffer. Figure 15-6 shows an example of a segmented buffer system.

Figure 15-6. Example System that Generates Recurring Events



The SignalTap II Logic Analyzer verifies the functionality of the design shown in Figure 15-6 to ensure that the correct data is written to the SRAM controller. Buffer acquisition in the SignalTap II Logic Analyzer allows you to monitor the RDATA port when `H'0F0F0F0F` is sent into the RADDR port. You can monitor multiple read transactions from the SRAM device without running the SignalTap II Logic Analyzer again. The buffer acquisition feature allows you to segment the memory so you can capture the same event multiple times without wasting allocated memory. The number of cycles that are captured depends on the number of segments specified under the **Data** settings.

To enable and configure buffer acquisition, select **Segmented** in the SignalTap II Logic Analyzer Editor and select the number of segments to use. In the example in [Figure 15-6](#), selecting sixty-four 64-sample segments allows you to capture 64 read cycles when the RADDR signal is H'0F0F0F0F.

- ① For more information about buffer acquisition mode, refer to [Configuring the Trigger Flow in the SignalTap II Logic Analyzer](#) in the Quartus II Help.

Using the Storage Qualifier Feature

Both non-segmented and segmented buffers described in the previous section offer a snapshot in time of the data stream being analyzed. The default behavior for writing into acquisition memory with the SignalTap II Logic Analyzer is to sample data on every clock cycle. With a non-segmented buffer, there is one data window that represents a comprehensive snapshot of the datastream. Similarly, segmented buffers use several smaller sampling windows spread out over more time, with each sampling window representing a contiguous data set.

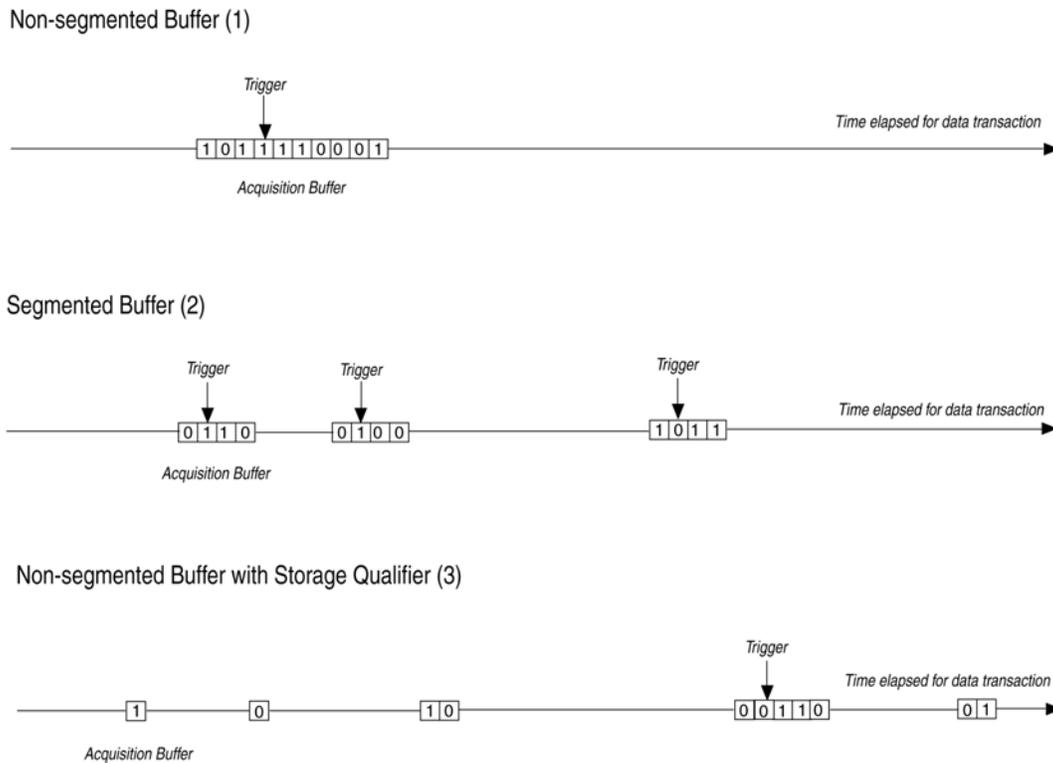
With carefully chosen trigger conditions and a generous sample depth for the acquisition buffer, analysis using segmented and non-segmented buffers captures a majority of functional errors in a chosen signal set. However, each data window can have a considerable amount of redundancy associated with it; for example, a capture of a data stream containing long periods of idle signals between data bursts. With default behavior using the SignalTap II Logic Analyzer, you cannot discard the redundant sample bits.

The Storage Qualification feature allows you to filter out individual samples not relevant to debugging the design. With this feature, a condition acts as a write enable to the buffer during each clock cycle of data acquisition. Through fine tuning the data that is actually stored in acquisition memory, the Storage Qualification feature allows for a more efficient use of acquisition memory in the specified number of samples over a longer period of analysis.

Use of the Storage Qualification feature is similar to an acquisition using a segmented buffer, in that you can create a discontinuity in the capture buffer. Because you can create a discontinuity between any two samples in the buffer, the Storage Qualification feature is equivalent to being able to create a customized segmented buffer in which the number and size of segment boundaries are adjustable. [Figure 15-7](#) illustrates three ways the SignalTap II Logic Analyzer writes into acquisition memory.

 You can only use the Storage Qualification feature with a non-segmented buffer. The MegaWizard Plug-In Manager instantiated flow only supports the Input Port mode for the Storage Qualification feature.

Figure 15-7. Data Acquisition Using Different Modes of Controlling the Acquisition Buffer



Notes to Figure 15-7:

- (1) Non-segmented Buffers capture a fixed sample window of contiguous data.
- (2) Segmented buffers divide the buffer into fixed sized segments, with each segment having an equal sample depth.
- (3) Storage Qualification allows you to define a custom sampling window for each segment you create with a qualifying condition. Storage qualification potentially allows for a larger time scale of coverage.

There are six types available under the Storage Qualification feature:

- Continuous
- Input port
- Transitional
- Conditional
- Start/Stop
- State-based

Continuous (the default mode selected) turns the Storage Qualification feature off.

Each selected storage qualifier type is active when an acquisition starts. Upon the start of an acquisition, the SignalTap II Logic Analyzer examines each clock cycle and writes the data into the acquisition buffer based upon storage qualifier type and condition. The acquisition stops when a defined set of trigger conditions occur.

 Trigger conditions are evaluated independently of storage qualifier conditions. The SignalTap II Logic Analyzer evaluates the data stream for trigger conditions on every clock cycle after the acquisition begins.

Trigger conditions are defined in “Define Trigger Conditions” on page 15-6.

The storage qualifier operates independently of the trigger conditions.

The following subsections describe each storage qualification mode from the acquisition buffer.

Input Port Mode

When using the Input port mode, the SignalTap II Logic Analyzer takes any signal from your design as an input. When the design is running, if the signal is high on the clock edge, the SignalTap II Logic Analyzer stores the data in the buffer. If the signal is low on the clock edge, the data sample is ignored. A pin is created and connected to this input port by default if no internal node is specified.

If you are using an .stp to create a SignalTap II Logic Analyzer instance, specify the storage qualifier signal using the input port field located on the **Setup** tab. You must specify this port for your project to compile.

If you use the MegaWizard Plug-In Manager flow, the storage qualification input port, if specified, appears in the MegaWizard-generated instantiation template. You can then connect this port to a signal in your RTL.

Figure 15-8 shows a data pattern captured with a segmented buffer. Figure 15-9 shows a capture of the same data pattern with the storage qualification feature enabled.

Figure 15-8. Data Acquisition of a Recurring Data Pattern in Continuous Capture Mode (to illustrate Input port mode)

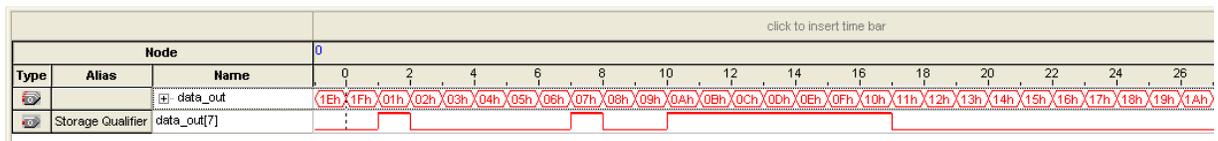
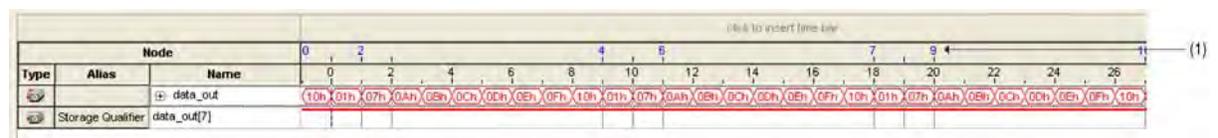


Figure 15-9. Data Acquisition of a Recurring Data Pattern Using an Input Signal as a Storage Qualifier



(1) Markers display samples when the logic analyzer paused a write into acquisition memory. These markers are enabled with the option “Record data discontinuities.”

Transitional Mode

In Transitional mode, you choose a set of signals for inspection using the node list check boxes in the **Storage Qualifier** column. During acquisition, if any of the signals marked for inspection have changed since the previous clock cycle, new data is written to the acquisition buffer. If none of the signals marked have changed since the previous clock cycle, no data is stored. [Figure 15-10](#) shows the transitional storage qualifier setup. [Figure 15-11](#) and [Figure 15-12](#) show captures of a data pattern in continuous capture mode and a data pattern using the Transitional mode for storage qualification.

Figure 15-10. Transitional Storage Qualifier Setup

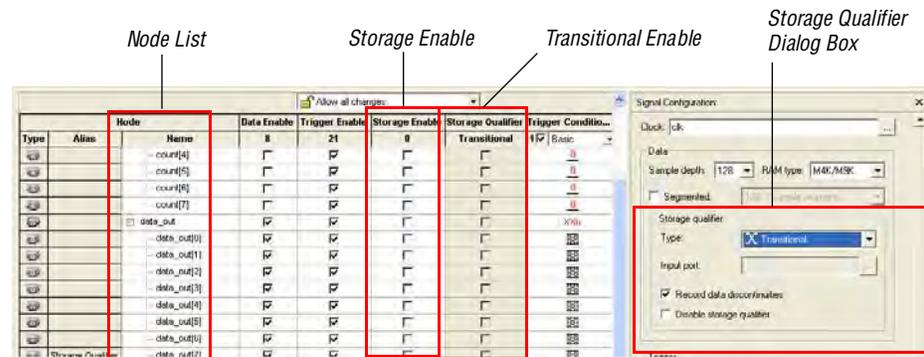


Figure 15-11. Data Acquisition of a Recurring Data Pattern in Continuous Capture Mode (to illustrate Transitional mode)

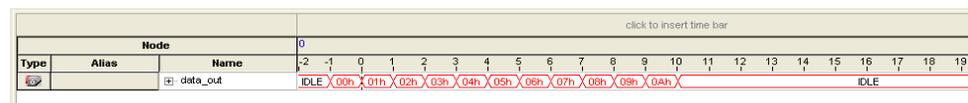
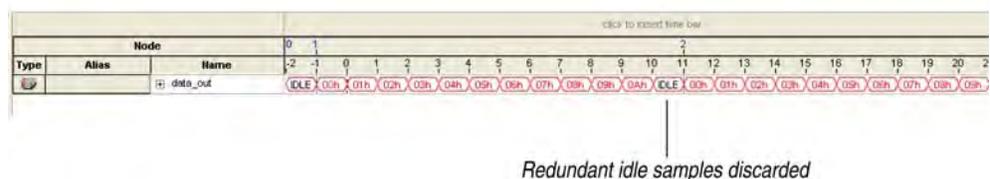


Figure 15-12. Data Acquisition of Recurring Data Pattern Using a Transitional Mode as a Storage Qualifier



Conditional Mode

In Conditional mode, the SignalTap II Logic Analyzer evaluates a combinational function of storage qualifier enabled signals within the node list to determine whether a sample is stored. The SignalTap II Logic Analyzer writes into the buffer during the clock cycles in which the condition you specify evaluates TRUE.

You can select either **Basic** or **Advanced** storage qualifier conditions. A **Basic** storage qualifier condition matches each signal to one of the following:

- Don't Care
- Low

- High
- Falling Edge
- Either Edge

If you specify a Basic Storage qualifier condition for more than one signal, the SignalTap II Logic Analyzer evaluates the logical AND of the conditions.

Any other combinational or relational operators that you may want to specify with the enabled signal set for storage qualification can be done with an advanced storage condition. Figure 15-13 details the conditional storage qualifier setup in the .stp.

You can specify up storage qualification conditions similar to the manner in which trigger conditions are specified. For details about basic and advanced trigger conditions, refer to the sections “Creating Basic Trigger Conditions” on page 15-25 and “Creating Advanced Trigger Conditions” on page 15-25. Figure 15-14 and Figure 15-15 show a data capture with continuous sampling, and the same data pattern using the conditional mode for analysis, respectively.

Figure 15-13. Conditional Storage Qualifier Setup

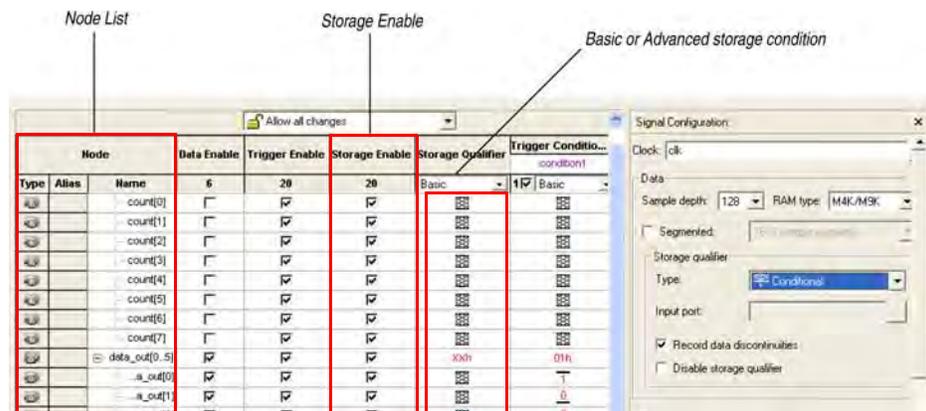
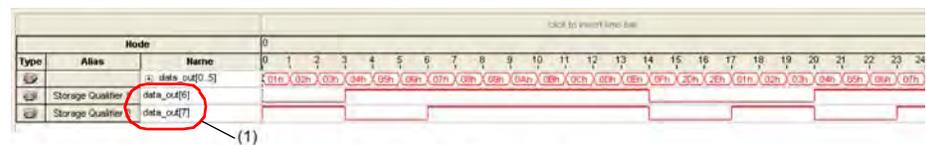
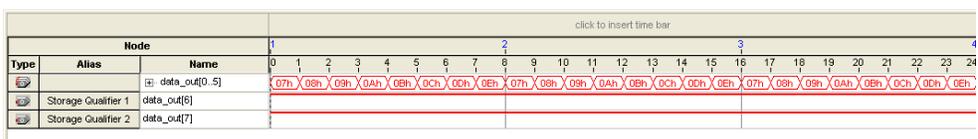


Figure 15-14. Data Acquisition of a Recurring Data Pattern in Continuous Capture Mode (to illustrate Conditional capture)



(1) Storage Qualifier condition is set up to evaluate data_out[6] AND data_out[7].

Figure 15-15. Data Acquisition of a Recurring Data Pattern in Conditional Capture Mode



Start/Stop Mode

The Start/Stop mode is similar to the Conditional mode for storage qualification. However, in this mode there are two sets of conditions, one for start and one for stop. If the start condition evaluates to TRUE, data begins is stored in the buffer every clock cycle until the stop condition evaluates to TRUE, which then pauses the data capture. Additional start signals received after the data capture has started are ignored. If both start and stop evaluate to TRUE at the same time, a single cycle is captured.

 You can force a trigger by pressing the **Stop** button if the buffer fails to fill to completion due to a stop condition.

Figure 15-16 shows the Start/Stop mode storage qualifier setup. Figure 15-17 and Figure 15-18 show captures data pattern in continuous capture mode and a data pattern in using the Start/Stop mode for storage qualification.

Figure 15-16. Start/Stop Mode Storage Qualifier Setup

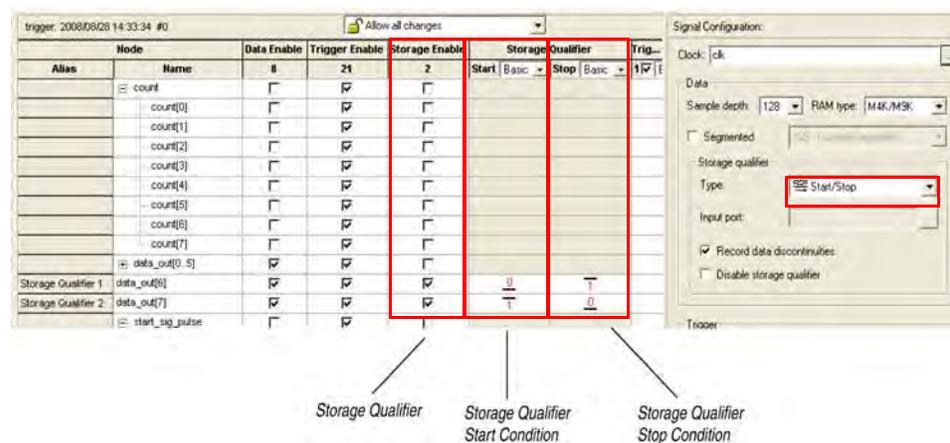


Figure 15-17. Data Acquisition of a Recurring Data Pattern in Continuous Mode (to illustrate Start/Stop mode)

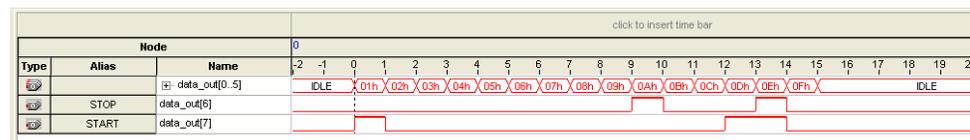
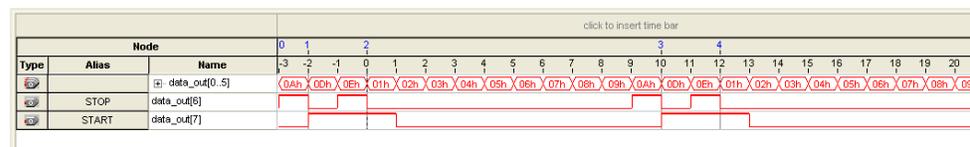


Figure 15-18. Data Acquisition of a Recurring Data Pattern with Start/Stop Storage Qualifier Enabled



State-Based

The State-based storage qualification mode is used with the State-based triggering flow. The state based triggering flow evaluates an if-else based language to define how data is written into the buffer. With the State-based trigger flow, you have command over boolean and relational operators to guide the execution flow for the target acquisition buffer. When the storage qualifier feature is enabled for the State-based flow, two additional commands are available, the `start_store` and `stop_store` commands. These commands operate similarly to the Start/Stop capture conditions described in the previous section. Upon the start of acquisition, data is not written into the buffer until a `start_store` action is performed. The `stop_store` command pauses the acquisition. If both `start_store` and `stop_store` actions are performed within the same clock cycle, a single sample is stored into the acquisition buffer.

For more information about the State-based flow and storage qualification using the State-based trigger flow, refer to the section [“State-Based Triggering” on page 15–29](#).

Showing Data Discontinuities

When you turn on **Record data discontinuities**, the SignalTap II Logic Analyzer marks the samples during which the acquisition paused from a storage qualifier. This marker is displayed in the waveform viewer after acquisition completes.

Disable Storage Qualifier

You can turn off the storage qualifier quickly with the **Disable Storage Qualifier** option, and perform a continuous capture. This option is run-time reconfigurable; that is, the setting can be changed without recompiling the project. Changing storage qualifier mode from the **Type** field requires a recompilation of the project.



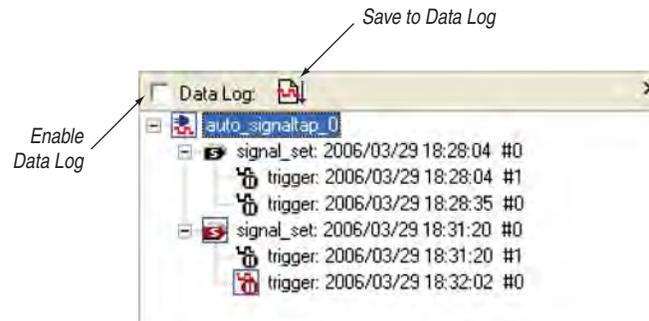
For a detailed explanation of Runtime Reconfigurable options available with the SignalTap II Logic Analyzer, and storage qualifier application examples using runtime reconfigurable options, refer to [“Runtime Reconfigurable Options” on page 15–52](#).

Managing Multiple SignalTap II Files and Configurations

You may have more than one `.stp` in one design. Each file potentially has a different group of monitored signals. These signal groups make it possible to debug different blocks in your design. In turn, each group of signals can also be used to define different sets of trigger conditions. Along with each `.stp`, there is also an associated programming file (SRAM Object File [`.sof`]). The settings in a selected SignalTap II file must match the SignalTap II logic design in the associated `.sof` for the logic analyzer to run properly when the device is programmed. Use the Data Log feature and the SOF Manager to manage all of the `.stp` files and their associated settings and programming files.

The Data Log allows you to store multiple SignalTap II configurations within a single **.stp**. Figure 15-19 shows two signal set configurations with multiple trigger conditions in one **.stp**. To toggle between the active configurations, double-click on an entry in the Data Log. As you toggle between the different configurations, the signal list and trigger conditions change in the **Setup** tab of the **.stp**. The active configuration displayed in the **.stp** is indicated by the blue square around the signal specified in the Data Log. To store a configuration in the Data Log, on the Edit menu, click **Save to Data Log** or click **Save to Data Log** at the top of the Data Log.

Figure 15-19. Data Log



The SOF Manager allows you to embed multiple **SOFs** into one **.stp**. Embedding an **SOF** in an **.stp** lets you move the **.stp** to a different location, either on the same computer or across a network, without the need to include the associated **.sof** as a separate. To embed a new **SOF** in the **.stp**, right-click in the SOF Manager, and click **Attach SOF File** (Figure 15-20).

Figure 15-20. SOF Manager



As you switch between configurations in the Data Log, you can extract the **SOF** that is compatible with that particular configuration. You can use the programmer in the SignalTap II Logic Analyzer to download the new **SOF** to the FPGA, ensuring that the configuration of your **.stp** always matches the design programmed into the target device.

Define Triggers

When you start the SignalTap II Logic Analyzer, it samples activity continuously from the monitored signals. The SignalTap II Logic Analyzer “triggers”—that is, the logic analyzer stops and displays the data—when a condition or set of conditions that you specified has been reached. This section describes the various types of trigger conditions that you can specify using the SignalTap II Logic Analyzer on the **Signal Configuration** pane.

Creating Basic Trigger Conditions

The simplest kind of trigger condition is a basic trigger. Select this from the list at the top of the **Trigger Conditions** column in the node list in the SignalTap II Logic Analyzer Editor. If you select the **Basic** trigger type, you must specify the trigger pattern for each signal you have added in the **.stp**. To specify the trigger pattern, right-click in the **Trigger Conditions** column and click the desired pattern. Set the trigger pattern to any of the following conditions:

- Don't Care
- Low
- High
- Falling Edge
- Rising Edge
- Either Edge

For buses, type a pattern in binary, or right-click and select **Insert Value** to enter the pattern in other number formats. Note that you can enter X to specify a set of “don't care” values in either your hexadecimal or your binary string. For signals added to the **.stp** that have an associated mnemonic table, you can right-click and select an entry from the table to specify pre-defined conditions for the trigger.

For more information about creating and using mnemonic tables, refer to [“View, Analyze, and Use Captured Data” on page 15-55](#), and to the Quartus II Help.

For signals added with certain plug-ins, you can create basic triggers easily using predefined mnemonic table entries. For example, with the Nios II plug-in, if you have specified an **.elf** from your Nios II IDE design, you can type the name of a function from your Nios II code. The logic analyzer triggers when the Nios II instruction address matches the address of the specified code function name.

Data capture stops and the data is stored in the buffer when the logical AND of all the signals for a given trigger condition evaluates to TRUE.

Creating Advanced Trigger Conditions

With the SignalTap II Logic Analyzer's basic triggering capabilities, you can build more complex triggers with extra logic that enables you to capture data when a combination of conditions exist. If you select the **Advanced** trigger type at the top of the **Trigger Conditions** column in the node list of the SignalTap II Logic Analyzer Editor, a new tab named **Advanced Trigger** appears where you can build a complex

trigger expression using a simple GUI. Drag-and-drop operators into the Advanced Trigger Configuration Editor window to build the complex trigger condition in an expression tree. To configure the operators' settings, double-click or right-click the operators that you have placed and select **Properties**. Table 15-3 lists the operators you can use.

Table 15-3. Advanced Triggering Operators (Note 1)

Name of Operator	Type
Less Than	Comparison
Less Than or Equal To	Comparison
Equality	Comparison
Inequality	Comparison
Greater Than	Comparison
Greater Than or Equal To	Comparison
Logical NOT	Logical
Logical AND	Logical
Logical OR	Logical
Logical XOR	Logical
Reduction AND	Reduction
Reduction OR	Reduction
Reduction XOR	Reduction
Left Shift	Shift
Right Shift	Shift
Bitwise Complement	Bitwise
Bitwise AND	Bitwise
Bitwise OR	Bitwise
Bitwise XOR	Bitwise
Edge and Level Detector	Signal Detection

Note to Table 15-3:

- (1) For more information about each of these operators, refer to the Quartus II Help.

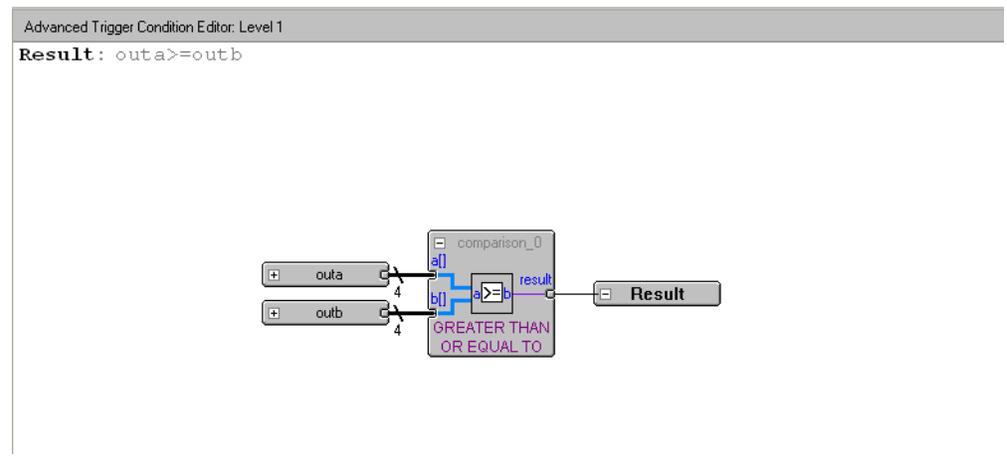
Adding many objects to the Advanced Trigger Condition Editor can make the work space cluttered and difficult to read. To keep objects organized while you build your advanced trigger condition, use the shortcut menu and select **Arrange All Objects**. You can also use the **Zoom-Out** command to fit more objects into the Advanced Trigger Condition Editor window.

Examples of Advanced Triggering Expressions

The following examples show how to use Advanced Triggering:

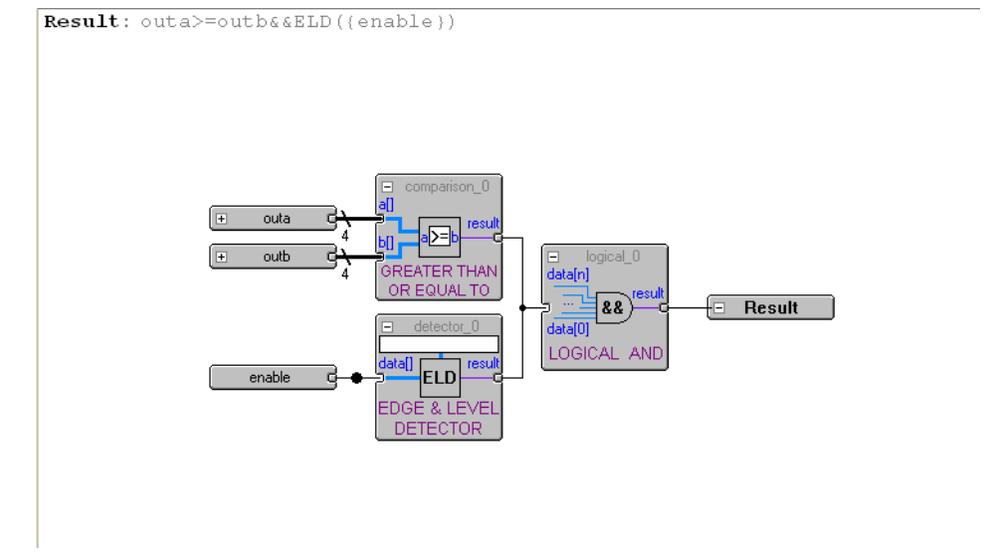
- Trigger when bus outa is greater than or equal to outb (Figure 15-21).

Figure 15-21. Bus outa is Greater Than or Equal to Bus outb



- Trigger when bus outa is greater than or equal to bus outb, and when the enable signal has a rising edge (Figure 15-22).

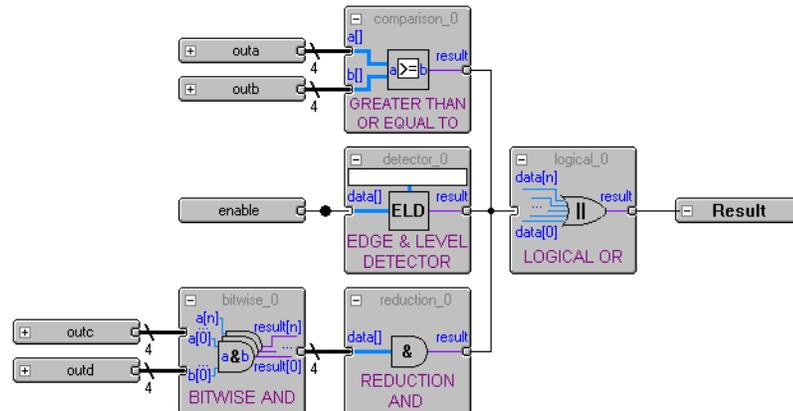
Figure 15-22. Enable Signal has a Rising Edge



- Trigger when bus outa is greater than or equal to bus outb, or when the enable signal has a rising edge. Or, when a bitwise AND operation has been performed between bus outc and bus outd, and all bits of the result of that operation are equal to 1 (Figure 15-23).

Figure 15-23. Bitwise AND Operation

```
Result: outa>=outb||ELD({enable})||(&outc&outd)
```



Trigger Condition Flow Control

The SignalTap II Logic Analyzer offers multiple triggering conditions to give you precise control of the method in which data is captured into the acquisition buffers. Trigger Condition Flow allows you to define the relationship between a set of triggering conditions. The SignalTap II Logic Analyzer **Signal Configuration** pane offers two flow control mechanisms for organizing trigger conditions:

- **Sequential Triggering**—The default triggering flow. Sequential triggering allows you to define up to 10 triggering levels that must be satisfied before the acquisition buffer finishes capturing.
- **State-Based Triggering**—Allows you the greatest control over your acquisition buffer. Custom-based triggering allows you to organize trigger conditions into states based on a conditional flow that you define.

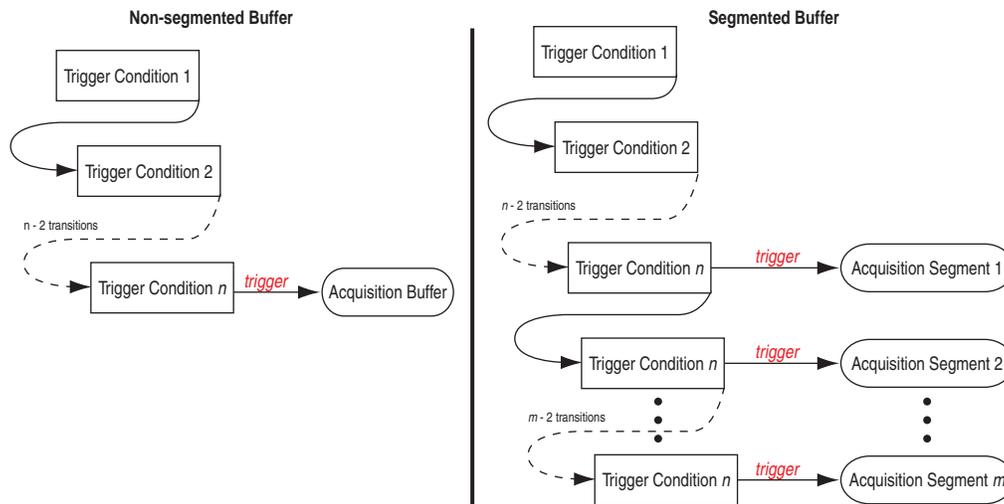
You can use sequential or state based triggering with either a segmented or a non-segmented buffer.

Sequential Triggering

Sequential triggering flow allows you to cascade up to 10 levels of triggering conditions. The SignalTap II Logic Analyzer sequentially evaluates each of the triggering conditions. When the last triggering condition evaluates to TRUE, the SignalTap II Logic Analyzer triggers the acquisition buffer. For segmented buffers, every acquisition segment after the first segment triggers on the last triggering condition that you have specified. Use the Simple Sequential Triggering feature with basic triggers, advanced triggers, or a mix of both. Figure 15-24 illustrates the simple sequential triggering flow for non-segmented and segmented buffers.

 The external trigger is considered as trigger level 0. The external trigger must be evaluated before the main trigger levels are evaluated.

Figure 15–24. Sequential Triggering Flow (Note 1), (2)



Notes to Figure 15–24:

- (1) The acquisition buffer stops capture when all n triggering levels are satisfied, where $n \leq 10$.
- (2) An external trigger input, if defined, is evaluated before all other defined trigger conditions are evaluated. For more information about external triggers, refer to “Using External Triggers” on page 15–42.

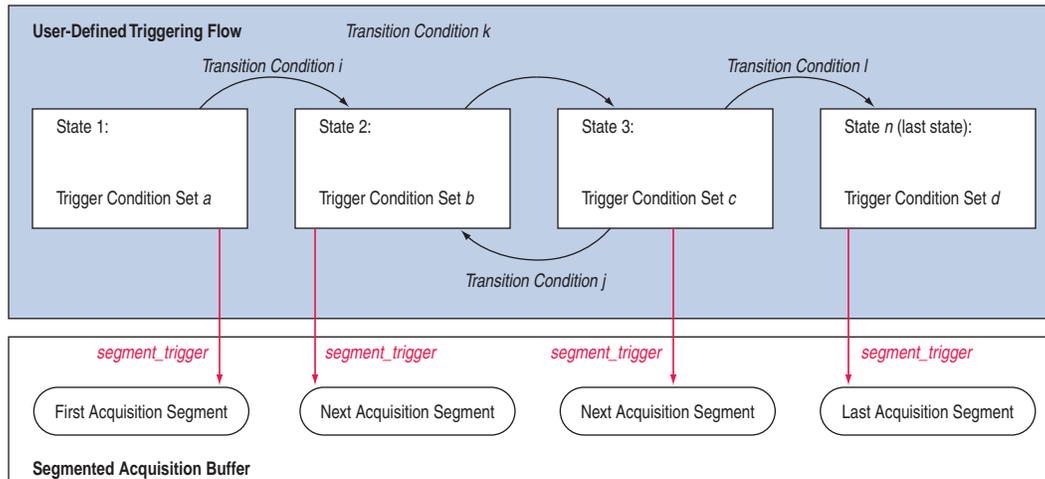
To configure the SignalTap II Logic Analyzer for Sequential triggering, in the SignalTap II editor on the **Trigger flow control** list, select **Sequential**. Select the desired number of trigger conditions from the **Trigger Conditions** list. After you select the desired number of trigger conditions, configure each trigger condition in the node list. To disable any trigger condition, turn on the trigger condition at the top of the column in the node list.

State-Based Triggering

Custom State-based triggering provides the most control over triggering condition arrangement. The State-Based Triggering flow allows you to describe the relationship between triggering conditions precisely, using an intuitive GUI and the SignalTap II Trigger Flow Description Language, a simple description language based upon conditional expressions. Tooltips within the custom triggering flow GUI allow you to describe your desired flow quickly. The custom State-based triggering flow allows for more efficient use of the space available in the acquisition buffer because only specific samples of interest are captured.

Figure 15-25 illustrates the custom State-based triggering flow. Events that trigger the acquisition buffer are organized by a state diagram that you define. All actions performed by the acquisition buffer are captured by the states and all transition conditions between the states are defined by the conditional expressions that you specify within each state.

Figure 15-25. State-Based Triggering Flow (Note 1), (2)



Notes to Figure 15-25:

- (1) You are allowed up to 20 different states.
- (2) An external trigger input, if defined, is evaluated before any conditions in the custom State-based triggering flow are evaluated. For more information, refer to "Using External Triggers" on page 15-42.

Each state allows you to define a set of conditional expressions. Each conditional expression is a Boolean expression dependent on a combination of triggering conditions (configured within the **Setup** tab), counters, and status flags. Counters and status flags are resources provided by the SignalTap II Logic Analyzer custom-based triggering flow.

Within each conditional expression you define a set of actions. Actions include triggering the acquisition buffer to stop capture, a modification to either a counter or status flag, or a state transition.

Trigger actions can apply to either a single segment of a segmented acquisition buffer or to the entire non-segmented acquisition buffer. Each trigger action provides you with an optional count that specifies the number of samples captured before stopping acquisition of the current segment. The count argument allows you to control the amount of data captured precisely before and after triggering event.

Resource manipulation actions allow you to increment and decrement counters or set and clear status flags. The counter and status flag resources are used as optional inputs in conditional expressions. Counters and status flags are useful for counting the number of occurrences of particular events and for aiding in triggering flow control.

This SignalTap II custom State-based triggering flow allows you to capture a sequence of events that may not necessarily be contiguous in time; for example, capturing a communication transaction between two devices that includes a handshaking protocol containing a sequence of acknowledgements.

The **State-Based Trigger Flow** tab is the control interface for the custom state-based triggering flow. To enable this tab, select **State-based** on the **Trigger Flow Control** list. (Note that when **Trigger Flow Control** is specified as **Sequential**, the **State-Based Trigger Flow** tab is hidden.)

The **State-Based Trigger Flow** tab is partitioned into the following three panes:

- [State Diagram Pane](#)
- [Resources Pane](#)
- [State Machine Pane](#)

State Diagram Pane

The **State Diagram** pane provides a graphical overview of the triggering flow that you define. It shows the number of states available and the state transitions between the states. You can adjust the number of available states by using the menu above the graphical overview.

State Machine Pane

The **State Machine** pane contains the text entry boxes where you can define the triggering flow and actions associated with each state. You can define the triggering flow using the SignalTap II Trigger Flow Description Language, a simple language based on “if-else” conditional statements. Tooltips appear when you move the mouse over the cursor, to guide command entry into the state boxes. The GUI provides a syntax check on your flow description in real-time and highlights any errors in the text flow.



For a full description of the SignalTap II Trigger Flow Description Language, refer to “[SignalTap II Trigger Flow Description Language](#)” on page 15–32.



You can also refer to [SignalTap II Trigger Flow Description Language](#) in Quartus II Help.

The State Machine description text boxes default to show one text box per state. You can also have the entire flow description shown in a single text field. This option can be useful when copying and pasting a flow description from a template or an external text editor. To toggle between one window per state, or all states in one window, select the appropriate option under **State Display mode**.

Resources Pane

The **Resources** pane allows you to declare Status Flags and Counters for use in the conditional expressions in the Custom Triggering Flow. Actions to decrement and increment counters or to set and clear status flags are performed within the triggering flow that you define.

You can specify up to 20 counters and 20 status flags. Counter and status flags values may be initialized by right-clicking the status flag or counter name after selecting a number of them from the respective pull-down list, and selecting **Set Initial Value**. To specify a counter width, right-click the counter name and select **Set Width**. Counters and flag values are updated dynamically after acquisition has started to assist in debugging your trigger flow specification.

The **configurable at runtime** options in the **Resources** pane allows you to configure the custom-flow control options that can be changed at runtime without requiring a recompilation. Table 15-4 contains a description of options for the State-based trigger flow that can be reconfigured at runtime.



For a broader discussion about all options that can be changed without incurring a recompile refer to [“Runtime Reconfigurable Options”](#) on page 15-52.

Table 15-4. Runtime Reconfigurable Settings, State-Based Triggering Flow

Setting	Description
Destination of goto action	Allows you to modify the destination of the state transition at runtime.
Comparison values	Allows you to modify comparison values in Boolean expressions at runtime. In addition, you can modify the <code>segment_trigger</code> and trigger action post-fill count argument at runtime.
Comparison operators	Allows you to modify the operators in Boolean expressions at runtime.
Logical operators	Allows you to modify the logical operators in Boolean expressions at runtime.

You can restrict changes to your SignalTap configuration to include only the options that do not require a recompilation by using the menu above the trigger list in the **Setup** tab. **Allow trigger condition changes only** restricts changes to only the configuration settings that have the **configurable at runtime** specified. With this option enabled, to modify Trigger Flow conditions in the **Custom Trigger Flow** tab, click the desired parameter in the text box and select a new parameter from the menu that appears.



The runtime configurable settings for the **Custom Trigger Flow** tab are on by default. You may get some performance advantages by disabling some of the runtime configurable options. For details about the effects of turning off the runtime modifiable options, refer to [“Performance and Resource Considerations”](#) on page 15-48.

SignalTap II Trigger Flow Description Language

The Trigger Flow Description Language is based on a list of conditional expressions per state to define a set of actions. Each line in [Example 15-1](#) shows a language format. Keywords are shown in bold. Non-terminals are delimited by “<>” and are further explained in the following sections. Optional arguments are delimited by “[]” ([Example 15-1](#)).

 Examples of Triggering Flow descriptions for common scenarios using the SignalTap II Custom Triggering Flow are provided in “Custom Triggering Flow Application Examples” on page 15-67.

Example 15-1. Trigger Flow Description Language Format (Note 1)

```
state <State_label>:  
<action_list>  
  
if( <Boolean_expression> )  
<action_list>  
[else if ( <boolean_expression> )  
<action_list>] (1)  
[else  
<action_list>]
```

Note to Example 15-1:

(1) Multiple `else if` conditions are allowed.

The priority for evaluation of conditional statements is assigned from top to bottom. The `<boolean_expression>` in an `if` statement can contain a single event, or it can contain multiple event conditions. The `action_list` within an `if` or an `else if` clause must be delimited by the begin and end tokens when the action list contains multiple statements. When the boolean expression is evaluated `TRUE`, the logic analyzer analyzes all of the commands in the action list concurrently. The possible actions include:

- Triggering the acquisition buffer
- Manipulating a counter or status flag resource
- Defining a state transition

State Labels

State labels are identifiers that can be used in the action `goto`.

`state <state_label>`: begins the description of the actions evaluated when this state is reached.

The description of a state ends with the beginning of another state or the end of the whole trigger flow description.

Boolean_expression

`Boolean_expression` is a collection of logical operators, relational operators, and their operands that evaluate into a Boolean result. Depending on the operator, the operand can be a reference to a trigger condition, a counter and a register, or a numeric value. Within an expression, parentheses can be used to group a set of operands.

Logical operators accept any boolean expression as an operand. The supported logical operators are shown in [Table 15-5](#).

Table 15-5. Logical Operators

Operator	Description	Syntax
!	NOT operator	! expr1
&&	AND operator	expr1 && expr2
	OR operator	expr1 expr2

Relational operators are performed on counters or status flags. The comparison value, the right operator, must be a numerical value. The supported relational operators are shown in [Table 15-6](#).

Table 15-6. Relational Operators

Operator	Description	Syntax <i>(Note 1) (2)</i>
>	Greater than	<identifier> > <numerical_value>
>=	Greater than or Equal to	<identifier> >= <numerical_value>
==	Equals	<identifier> == <numerical_value>
!=	Does not equal	<identifier> != <numerical_value>
<=	Less than or equal to	<identifier> <= <numerical_value>
<	Less than	<identifier> < <numerical_value>

Notes to Table 15-6:

- (1) <identifier> indicates a counter or status flag.
- (2) <numerical_value> indicates an integer.

Action_list

Action_list is a list of actions that can be performed when a state is reached and a condition is also satisfied. If more than one action is specified, they must be enclosed by `begin` and `end`. The actions can be categorized as resource manipulation actions, buffer control actions, and state transition actions. Each action is terminated by a semicolon (`;`).

Resource Manipulation Action

The resources used in the trigger flow description can be either counters or status flags. [Table 15-7](#) shows the description and syntax of each action.

Table 15-7. Resource Manipulation Action

Action	Description	Syntax
increment	Increments a counter resource by 1	increment <counter_identifier>;
decrement	Decrements a counter resource by 1	decrement <counter_identifier>;
reset	Resets counter resource to initial value	reset <counter_identifier>;
set	Sets a status Flag to 1	set <register_flag_identifier>;
clear	Sets a status Flag to 0	clear <register_flag_identifier>;

Buffer Control Action

Buffer control actions specify an action to control the acquisition buffer. [Table 15-8](#) shows the description and syntax of each action.

Table 15-8. Buffer Control Action

Action	Description	Syntax
trigger	Stops the acquisition for the current buffer and ends analysis. This command is required in every flow definition.	trigger <post-fill_count>;
segment_trigger	Ends the acquisition of the current segment. The SignalTap II Logic Analyzer starts acquiring from the next segment on evaluating this command. If all segments are filled, the oldest segment is overwritten with the latest sample. The acquisition stops when a trigger action is evaluated. This action cannot be used in non-segmented acquisition mode.	segment_trigger <post-fill_count>;
start_store	Asserts the write_enable to the SignalTap II acquisition buffer. This command is active only when the State-based storage qualifier mode is enabled.	start_store
stop_store	De-asserts the write_enable signal to the SignalTap II acquisition buffer. This command is active only when the State-based storage qualifier mode is enabled.	stop_store

Both `trigger` and `segment_trigger` actions accept an optional post-fill count argument. If provided, the current acquisition acquires the number of samples provided by post-fill count and then stops acquisition. If no post-count value is specified, the trigger position for the affected buffer defaults to the trigger position specified in the **Setup** tab.



In the case of `segment_trigger`, acquisition of the current buffer stops immediately if a subsequent triggering action is issued in the next state, regardless of whether or not the post-fill count has been satisfied for the current buffer. The remaining unfilled post-count acquisitions in the current buffer are discarded and displayed as grayed-out samples in the data window.

State Transition Action

The State Transition action specifies the next state in the custom state control flow. It is specified by the `goto` command. The syntax is as follows:

```
goto <state_label>;
```

Using the State-Based Storage Qualifier Feature

When you select State-based for the storage qualifier type, the `start_store` and `stop_store` actions are enabled in the State-based trigger flow. These commands, when used in conjunction with the expressions of the State-based trigger flow, give you maximum flexibility to control data written into the acquisition buffer.

 The `start_store` and `stop_store` commands can only be applied to a non-segmented buffer.

The `start_store` and `stop_store` commands function similar to the `start` and `stop` conditions when using the **start/stop** storage qualifier mode conditions. If storage qualification is enabled, the `start_store` command must be issued for SignalTap II to write data into the acquisition buffer. No data is acquired until the `start_store` command is performed. Also, a `trigger` command must be included as part of the trigger flow description. The `trigger` command is necessary to complete the acquisition and display the results on the waveform display.

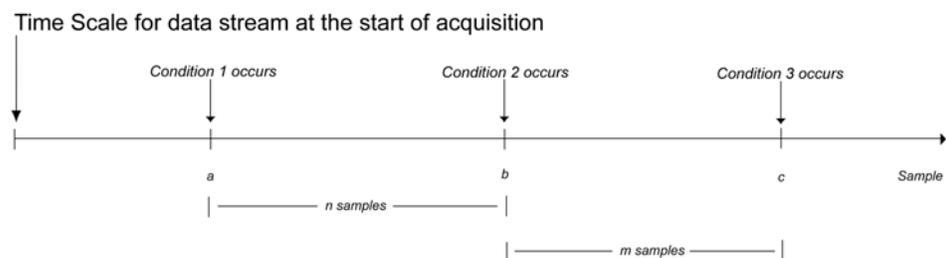
The following examples illustrate the behavior of the State-based trigger flow with the storage qualification commands.

Figure 15-26 shows a hypothetical scenario with three trigger conditions that happen at different times after you click **Start Analysis**. The trigger flow description in Example 15-2, when applied to the scenario shown in Figure 15-26, illustrates the functionality of the storage qualification feature for the state-based trigger flow.

Example 15-2. Trigger Flow Description 1

```
State 1: ST1:
if ( condition1 )
    start_store;
else if ( condition2 )
    trigger value;
else if ( condition3 )
    stop_store;
```

Figure 15-26. Capture Scenario for Storage Qualification with the State-Based Trigger Flow



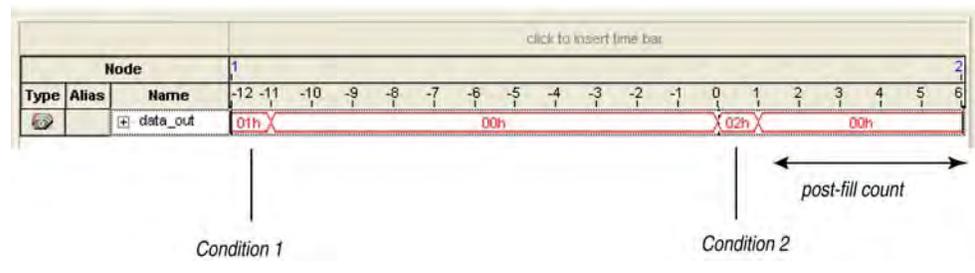
In this example, the SignalTap II Logic Analyzer does not write into the acquisition buffer until sample *a*, when Condition 1 occurs. Once sample *b* is reached, the `trigger value` command is evaluated. The logic analyzer continues to write into the buffer to finish the acquisition. The trigger flow specifies a `stop_store` command at sample *c*, *m* samples after the trigger point occurs.

The logic analyzer finishes the acquisition and displays the contents of the waveform if it can successfully finish the post-fill acquisition samples before Condition 3 occurs. In this specific case, the capture ends if the post-fill count value is less than *m*.

If the post-fill count value specified in Trigger Flow description 1 is greater than m samples, the buffer pauses acquisition indefinitely, provided there is no recurrence of Condition 1 to trigger the logic analyzer to start capturing data again. The SignalTap II Logic Analyzer continues to evaluate the **stop_store** and **start_store** commands even after the **trigger** command is evaluated. If the acquisition has paused, you can click **Stop Analysis** to manually stop and force the acquisition to trigger. You can use counter values, flags, and the State diagram to help you perform the trigger flow. The counter values, flags, and the current state are updated in real-time during a data acquisition.

Figure 15-27 and Figure 15-28 show a real data acquisition of the scenario. Figure 15-27 illustrates a scenario where the data capture finishes successfully. It uses a buffer with a sample depth of 64, $m = n = 10$, and the post-fill count value = 5. Figure 15-28 illustrates a scenario where the logic analyzer pauses indefinitely even after a trigger condition occurs due to a **stop_store** condition. This scenario uses a sample depth of 64, with $m = n = 10$ and post-fill count = 15.

Figure 15-27. Storage Qualification with Post-Fill Count Value Less than m (Acquisition Successfully Completes)



The combination of using counters, Boolean and relational operators in conjunction with the `start_store` and `stop_store` commands can give a clock-cycle level of resolution to controlling the samples that are written into the acquisition buffer. [Example 15-3](#) shows a trigger flow description that skips three clock cycles of samples after hitting condition 1. [Figure 15-30](#) shows the data transaction on a continuous capture and [Figure 15-32](#) shows the data capture with the Trigger flow description in [Example 15-3](#) applied.

Example 15-3. Trigger Flow Description 2

```

State 1: ST1
start_store
if ( condition1 )
begin
    stop_store;
    goto ST2;
end

State 2: ST2
if ( c1 < 3 )
    increment c1; //skip three clock cycles; c1 initialized to 0

else if ( c1 == 3 )
begin
    start_store; //start_store necessary to enable writing to finish
                //acquisition
    trigger;
end
    
```

Figure 15-30. Continuous Capture of Data Transaction for Example 2

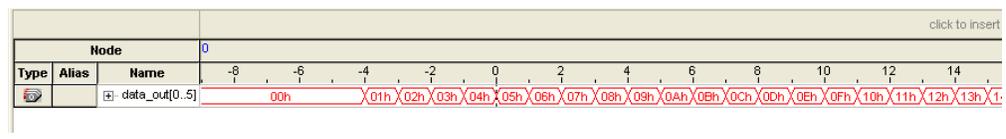
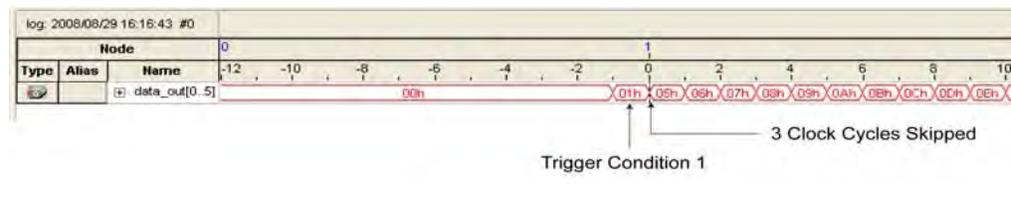


Figure 15-31. Capture of Data Transaction with Trigger Flow Description Applied



Specifying the Trigger Position

The SignalTap II Logic Analyzer allows you to specify the amount of data that is acquired before and after a trigger event. You can specify the trigger position independently between a Runtime and Power-Up Trigger. Select the desired ratio of pre-trigger data to post-trigger data by choosing one of the following ratios:

- **Pre**—Saves signal activity that occurred after the trigger (12% pre-trigger, 88% post-trigger).
- **Center**—Saves 50% pre-trigger and 50% post-trigger data.
- **Post**—Saves signal activity that occurred before the trigger (88% pre-trigger, 12% post-trigger).

These pre-defined ratios apply to both non-segmented buffers and segmented buffers.

If you use the custom-state based triggering flow, you can specify a custom trigger position. The `segment_trigger` and `trigger` actions accept a post-fill count argument. The post-fill count specifies the number of samples to capture before stopping data acquisition for the non-segmented buffer or a data segment when using the `trigger` and `segment_trigger` commands, respectively. When the captured data is displayed in the SignalTap II data window, the trigger position appears as the number of post-count samples from the end of the acquisition segment or buffer. Refer to [Equation 15-1](#):

Equation 15-1.

$$\text{Sample Number of Trigger Position} = (N - \text{Post-Fill Count})$$

In this case, N is the sample depth of either the acquisition segment or non-segmented buffer.

For segmented buffers, the acquisition segments that have a post-count argument define use of the post-count setting. Segments that do not have a post-count setting default to the trigger position ratios defined in the **Setup** tab.

For more details about the custom State-based triggering flow, refer to [“State-Based Triggering” on page 15-29](#).

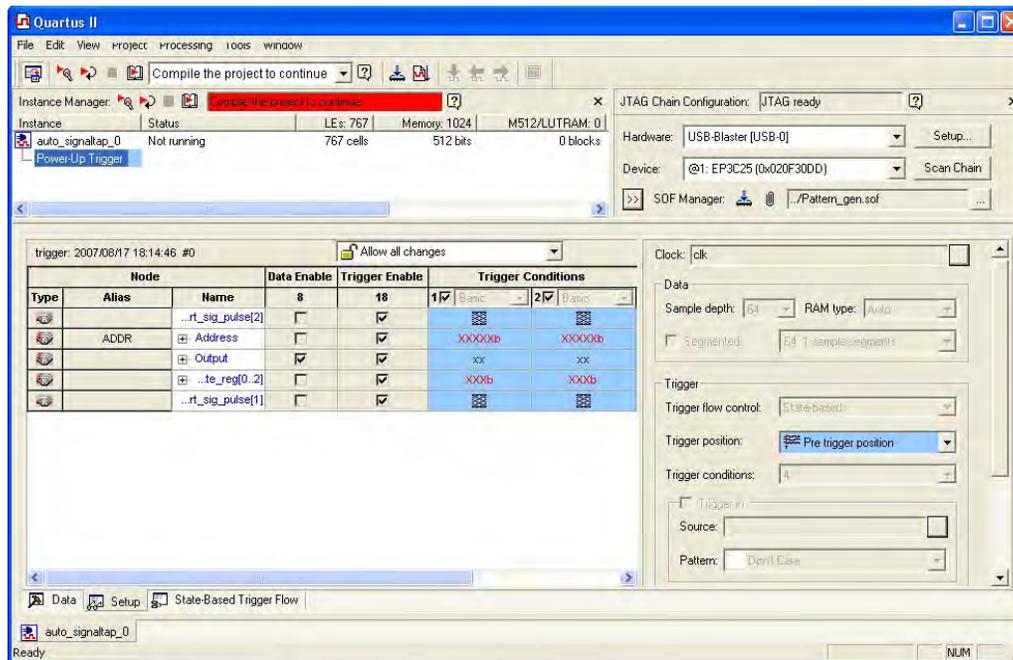
Creating a Power-Up Trigger

Typically, the SignalTap II Logic Analyzer is used to trigger on events that occur during normal device operation. You start an analysis manually once the target device is fully powered on and the device’s JTAG connection is available. However, there may be cases when you would like to capture trigger events that occur during device initialization, immediately after the FPGA is powered on or reset. With the SignalTap II Power-Up Trigger feature, you arm the SignalTap II Logic Analyzer and capture data immediately after device programming.

Enabling a Power-Up Trigger

You can add a different Power-Up Trigger to each logic analyzer instance in the **SignalTap II Instance Manager** pane. To enable the Power-Up Trigger for a logic analyzer instance, right-click the instance and click **Enable Power-Up Trigger**, or select the instance, and on the **Edit** menu, click **Enable Power-Up Trigger**. To disable a Power-Up Trigger, click **Disable Power-Up Trigger** in the same locations. Power-Up Trigger is shown as a child instance below the name of the selected instance with the default trigger conditions specified in the node list. [Figure 15-32](#) shows the SignalTap II Logic Analyzer Editor when Power-Up Trigger is enabled.

Figure 15-32. SignalTap II Logic Analyzer Editor with Power-Up Trigger Enabled



Managing and Configuring Power-Up and Runtime Trigger Conditions

When the Power-Up Trigger is enabled for a logic analyzer instance, you can create basic and advanced trigger conditions for the trigger as you do with a Run-Time Trigger. Power-Up Trigger conditions that you can adjust are color coded light blue, while Run-Time Trigger conditions you cannot adjust remain white. Since each instance now has two sets of trigger conditions—the Power-Up Trigger and the Run-Time Trigger—you can differentiate between the two with color coding. To switch between the trigger conditions of the Power-Up Trigger and the Run-Time Trigger, double-click the instance name or the Power-Up Trigger name in the **Instance Manager**.

You cannot make changes to Power-Up Trigger conditions that would normally require a full recompile with Runtime Trigger conditions, such as adding signals, deleting signals, or changing between basic and advanced triggers. To apply these changes to the Power-Up Trigger conditions, first make the changes using the Runtime Trigger conditions.

-  Any change made to the Power-Up Trigger conditions requires that you recompile the SignalTap II Logic Analyzer instance, even if a similar change to the Runtime Trigger conditions does not require a recompilation.

While creating or making changes to the trigger conditions for the Run-Time Trigger or the Power-Up Trigger, you may want to copy these conditions to the other trigger. This enables you to look for the same trigger during both power-up and runtime. To do this, right-click the instance name or the Power-Up Trigger name in the **Instance Manager** and click **Duplicate Trigger**, or select the instance name or the Power-Up Trigger name and on the **Edit** menu, click **Duplicate Trigger**.

You can also use In-System Sources and Probes in conjunction with the SignalTap II Logic Analyzer to force trigger conditions. The In-System Sources and Probes feature allows you to drive and sample values on to selected nets over the JTAG chain. For more information, refer to the *Design Debugging Using In-System Sources and Probes* chapter in volume 3 of the *Quartus II Handbook*.

Using External Triggers

You can create a trigger input that allows you to trigger the SignalTap II Logic Analyzer from an external source. The external trigger input behaves like trigger condition 1, is evaluated, and must be `TRUE` before any other configured trigger conditions are evaluated. The logic analyzer supplies a signal to trigger external devices or other SignalTap II Logic Analyzer instances. These features allow you to synchronize external logic analysis equipment with the internal logic analyzer. Power-Up Triggers can use the external triggers feature, but they must use the same source or target signal as their associated Run-Time Trigger.

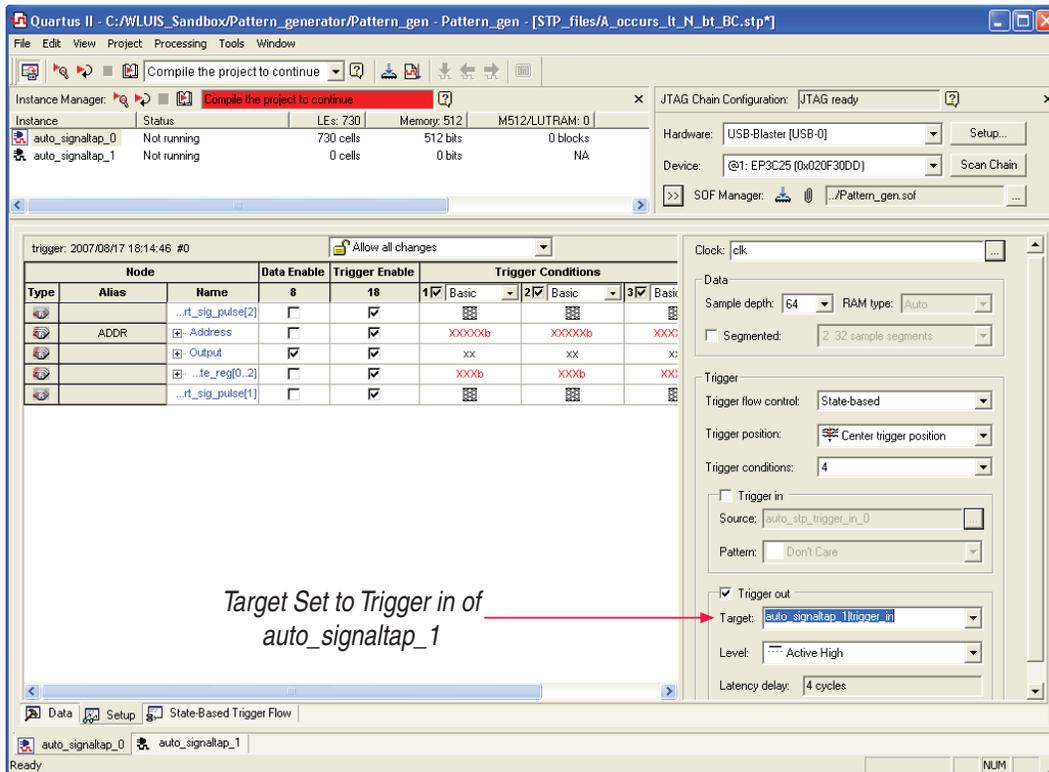
-  For more information about setting up external triggers, refer to *Signal Configuration Pane* in Quartus II Help.

Using the Trigger Out of One Analyzer as the Trigger In of Another Analyzer

An advanced feature of the SignalTap II Logic Analyzer is the ability to use the **Trigger out** of one analyzer as the **Trigger in** to another analyzer. This feature allows you to synchronize and debug events that occur across multiple clock domains.

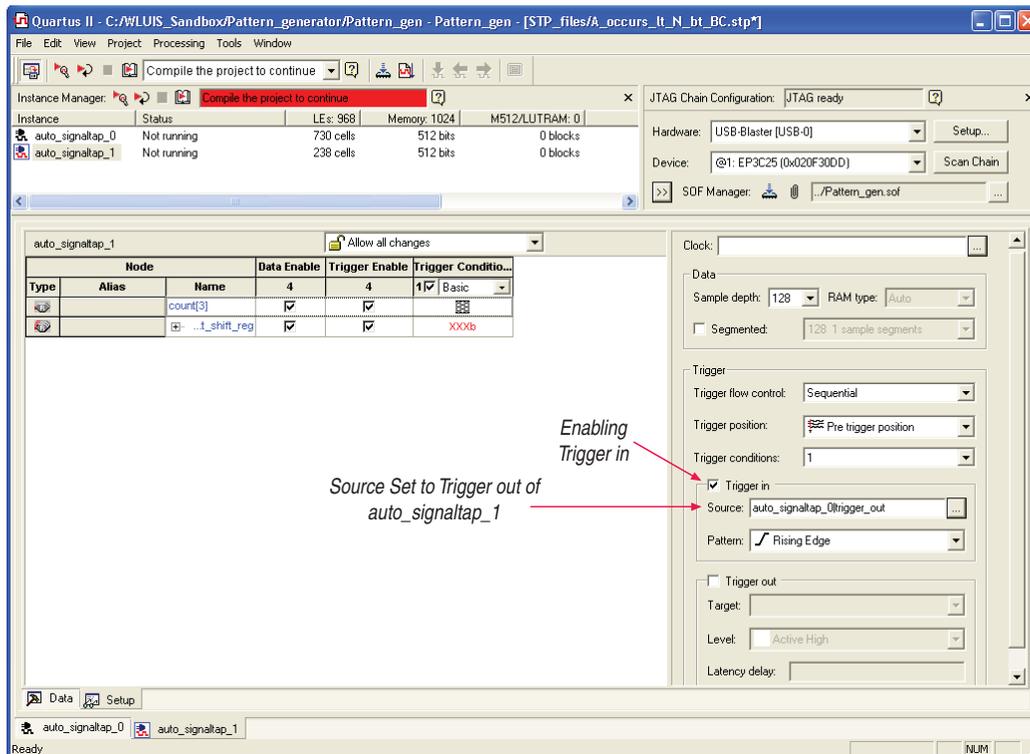
To perform this operation, first turn on **Trigger out** for the source logic analyzer instance. On the **Target** list of the **Trigger out** trigger, select the targeted logic analyzer instance. For example, if the instance named `auto_signaltap_0` should trigger `auto_signaltap_1`, select `auto_signaltap_1|trigger_in` from the list (Figure 15-33).

Figure 15-33. Configuring the Trigger Out Signal



- Turning on **Trigger out** automatically enables the **Trigger in** of the targeted logic analyzer instance and fills in the **Source** field of the **Trigger in** trigger with the **Trigger out** signal from the source logic analyzer instance. In this example, `auto_sigtap_0` is targeting `auto_sigtap_1`. The Trigger In Source field of `auto_sigtap_1` is automatically filled in with `auto_sigtap_0 | trigger_out` (Figure 15-34).

Figure 15-34. Configuring the Trigger In Signal



Compile the Design

When you add an `.stp` to your project, the SignalTap II Logic Analyzer becomes part of your design. You must compile your project to incorporate the SignalTap II logic and enable the JTAG connection you use to control the logic analyzer. When you are debugging with a traditional external logic analyzer, you must often make changes to the signals monitored as well as the trigger conditions. Because these adjustments require that you recompile your design when using the SignalTap II Logic Analyzer, use the SignalTap II Logic Analyzer feature along with incremental compilation in the Quartus II software to reduce recompilation time.

- For more information on reducing your recompilation burden with incremental compilation, refer to *Using the Incremental Compilation Design Flow* in Quartus II Help.

Faster Compilations with Quartus II Incremental Compilation

When you compile your design with an **.stp**, the `sld_signaltap` and `sld_hub` entities are automatically added to the compilation hierarchy. These two entities are the main components of the SignalTap II Logic Analyzer, providing the trigger logic and JTAG interface required for operation.

Incremental compilation enables you to preserve the synthesis and fitting results of your original design and add the SignalTap II Logic Analyzer to your design without recompiling your original source code. Incremental compilation is also useful when you want to modify the configuration of the **.stp**. For example, you can modify the buffer sample depth or memory type without performing a full compilation after the change is made. Only the SignalTap II Logic Analyzer, configured as its own design partition, must be recompiled to reflect the changes.

To use incremental compilation, first enable **Full Incremental Compilation** for your design if it is not already enabled, assign design partitions if necessary, and set the design partitions to the correct preservation levels. Incremental compilation is the default setting for new projects in the Quartus II software, so you can establish design partitions immediately in a new project. However, it is not necessary to create any design partitions to use the SignalTap II incremental compilation feature. When your design is set up to use full incremental compilation, the SignalTap II Logic Analyzer acts as its own separate design partition. You can begin taking advantage of incremental compilation by using the **SignalTap II: post-fitting filter** in the Node Finder to add signals for logic analysis.

Enabling Incremental Compilation for Your Design

Your project is fully compiled the first time, establishing the design partitions you have created. When enabled for your design, the SignalTap II Logic Analyzer is always a separate partition. After the first compilation, you can use the SignalTap II Logic Analyzer to analyze signals from the post-fit netlist. If your partitions are designed correctly, subsequent compilations due to SignalTap II Logic Analyzer settings take less time.

The netlist type for the top-level partition defaults to **source**. To take advantage of incremental compilation, specify the Netlist types for the partitions you wish to tap as **Post-fit**.



For more information about configuring and performing incremental compilation, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.

Using Incremental Compilation with the SignalTap II Logic Analyzer

The SignalTap II Logic Analyzer is automatically configured to work with the incremental compilation flow. For all signals that you want to connect to the SignalTap II Logic Analyzer from the post-fit netlist, set the netlist type of the partition containing the desired signals to **Post-Fit** or **Post-Fit (Strict)** with a Fitter Preservation Level of **Placement and Routing** using the Design Partitions window.

Use the **SignalTap II: post-fitting filter** in the **Node Finder** to add the signals of interest to your SignalTap II configuration file. If you want to add signals from the pre-synthesis netlist, set the netlist type to **Source File** and use the **SignalTap II: pre-synthesis filter** in the **Node Finder**. Do not use the netlist type **Post-Synthesis** with the SignalTap II Logic Analyzer.



Be sure to conform to the following guidelines when using post-fit and pre-synthesis nodes:

- Read all incremental compilation guidelines to ensure the proper partition of a project.
- To speed compile time, use only post-fit nodes for partitions specified as to preservation-level post-fit.
- Do not mix pre-synthesis and post-fit nodes in any partition. If you must tap pre-synthesis nodes for a particular partition, make all tapped nodes in that partition pre-synthesis nodes and change the netlist type to **source** in the design partitions window.

Node names may be different between a pre-synthesis netlist and a post-fit netlist. In general, registers and user input signals share common names between the two netlists. During compilation, certain optimizations change the names of combinational signals in your RTL. If the type of node name chosen does not match the netlist type, the compiler may not be able to find the signal to connect to your SignalTap II Logic Analyzer instance for analysis. The compiler issues a critical warning to alert you of this scenario. The signal that is not connected is tied to ground in the **SignalTap II data** tab.

If you do use incremental compile flow with the SignalTap II Logic Analyzer and source file changes are necessary, be aware that you may have to remove compiler-generated post-fit net names. Source code changes force the affected partition to go through resynthesis. During synthesis, the compiler cannot find compiler-generated net names from a previous compilation.

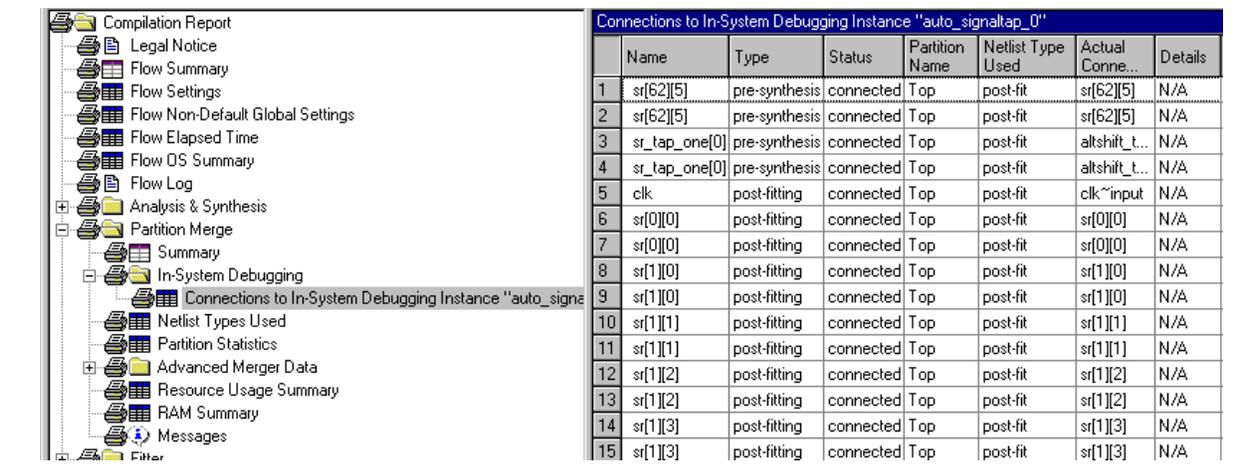


Altera recommends using only registered and user-input signals as debugging taps in your **.stp** whenever possible.

Both registered and user-supplied input signals share common node names in the pre-synthesis and post-fit netlist. As a result, using only registered and user-supplied input signals in your **.stp** limits the changes you need to make to your SignalTap II Logic Analyzer configuration.

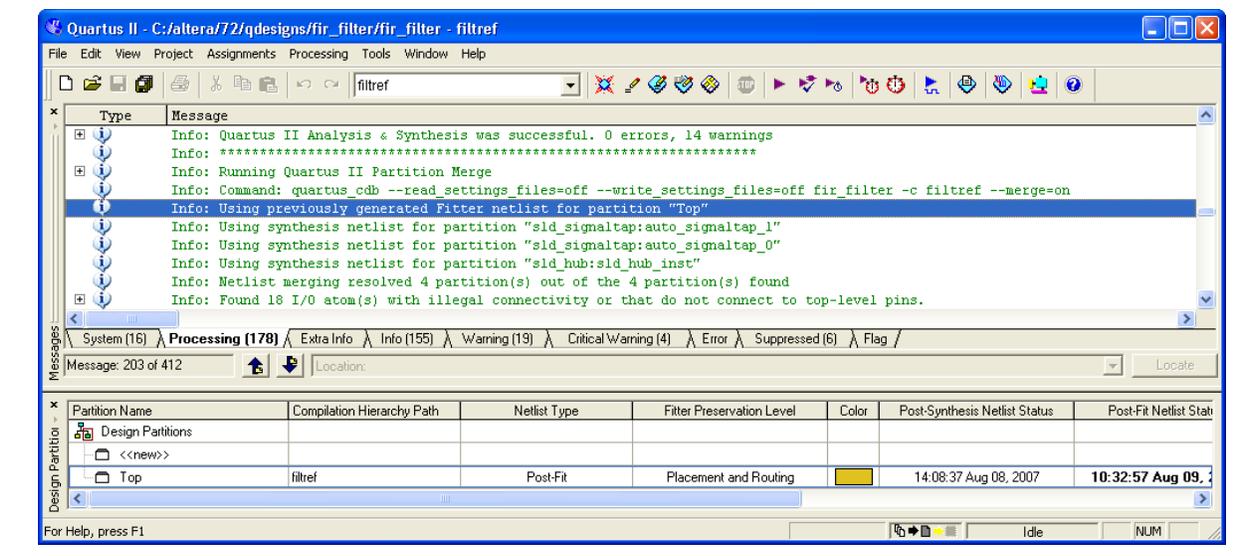
You can check the nodes that are connected to each SignalTap II instance using the In-System Debugging compilation reports. These reports list each node name you selected to connect to a SignalTap II instance, the netlist type used for the particular connection, and the actual node name used after compilation. If incremental compile is turned off, the In-System Debugging reports are located in the Analysis & Synthesis folder. If incremental compile is turned on, this report is located in the Partition Merge folder. Figure 15-35 shows an example of an In-System Debugging compilation report for a design using incremental compilation.

Figure 15-35. Compilation Report Showing Connectivity to SignalTap II Instance



To verify that your original design was not modified, examine the messages in the Partition Merge section of the Compilation Report. Figure 15-36 shows an example of the messages displayed.

Figure 15-36. Compilation Report Messages



Unless you make changes to your design partitions that require recompilation, only the SignalTap II design partition is recompiled. If you make subsequent changes to only the `.stp`, only the SignalTap II design partition must be recompiled, reducing your recompilation time.

Preventing Changes Requiring Recompilation

You can configure the `.stp` to prevent changes that normally require recompilation. To do this, select a lock mode from above the node list in the **Setup** tab. To lock your configuration, choose to allow only trigger condition changes, regardless of whether you use incremental compilation.

- ❓ For more information about the use of lock modes, refer to *Setup Tab (SignalTap II Logic Analyzer)* in Quartus II Help.

Timing Preservation with the SignalTap II Logic Analyzer

In addition to verifying functionality, timing closure is one of the most crucial processes in successfully completing a design. When you compile a project with a SignalTap II Logic Analyzer without the use of incremental compilation, you add IP to your existing design. Therefore, you can affect the existing placement, routing, and timing of your design. To minimize the effect that the SignalTap II Logic Analyzer has on your design, Altera recommends that you use incremental compilation for your project. Incremental compilation is the default setting in new designs and can be easily enabled and configured in existing designs. With the SignalTap II Logic Analyzer instance in its own design partition, it has little to no affect on your design.

In addition to using the incremental compilation flow for your design, you can use the following techniques to help maintain timing:

- Avoid adding critical path signals to your `.stp`.
- Minimize the number of combinational signals you add to your `.stp` and add registers whenever possible.
- Specify an f_{MAX} constraint for each clock in your design.

-  For an example of timing preservation with the SignalTap II Logic Analyzer, refer to the *Area and Timing Optimization* chapter in volume 2 of the *Quartus II Handbook*.

Performance and Resource Considerations

There is a necessary trade-off between the runtime flexibility of the SignalTap II Logic Analyzer, the timing performance of the SignalTap II Logic Analyzer, and resource usage. The SignalTap II Logic Analyzer allows you to select the runtime configurable parameters to balance the need for runtime flexibility, speed, and area. The default values have been chosen to provide maximum flexibility so you can complete debugging as quickly as possible; however, you can adjust these settings to determine whether there is a more optimal configuration for your design.

The following tips provide extra timing slack if you have determined that the SignalTap II logic is in your critical path, or to alleviate the resource requirements that the SignalTap II Logic Analyzer consumes if your design is resource-constrained.

If SignalTap II logic is part of your critical path, follow these tips to speed up the performance of the SignalTap II Logic Analyzer:

- **Disable runtime configurable options**—Certain resources are allocated to accommodate for runtime flexibility. If you use either advanced triggers or State-based triggering flow, disable runtime configurable parameters for a boost in f_{MAX} of the SignalTap II logic. If you are using State-based triggering flow, try disabling the **Goto state destination** option and performing a recompilation before disabling the other runtime configurable options. The **Goto state destination** option has the greatest impact on f_{MAX} , as compared to the other runtime configurable options.
- **Minimize the number of signals that have Trigger Enable selected**—All signals that you add to the .stp have **Trigger Enable** turned on. Turn off **Trigger Enable** for signals that you do not plan to use as triggers.
- **Turn on Physical Synthesis for register retiming**—If you have a large number of triggering signals enabled (greater than the number of inputs that would fit in a LAB) that fan-in logic to a gate-based triggering condition, such as a basic trigger condition or a logical reduction operator in the advanced trigger tab, turn on **Perform register retiming**. This can help balance combinational logic across LABs.

If your design is resource constrained, follow these tips to reduce the amount of logic or memory used by the SignalTap II Logic Analyzer:

- **Disable runtime configurable options**—Disabling runtime configurability for advanced trigger conditions or runtime configurable options in the State-based triggering flow results in using fewer LEs.
- **Minimize the number of segments in the acquisition buffer**—You can reduce the number of logic resources used for the SignalTap II Logic Analyzer by limiting the number of segments in your sampling buffer to only those required.
- **Disable the Data Enable for signals that are used for triggering only**—By default, both the **data enable** and **trigger enable** options are selected for all signals. Turning off the **data enable** option for signals used as trigger inputs only saves on memory resources used by the SignalTap II Logic Analyzer.

Because performance results are design-dependent, try these options in different combinations until you achieve the desired balance between functionality, performance, and utilization.

 For more information about area and timing optimization, refer the *Area and Timing Optimization* chapter in volume 2 of the *Quartus II Handbook*.

Program the Target Device or Devices

After you compile your project, including the SignalTap II Logic Analyzer, configure the FPGA target device. When you are using the SignalTap II Logic Analyzer for debugging, configure the device from the .stp instead of the Quartus II Programmer. Because you configure from the .stp, you can open more than one .stp and program multiple devices to debug multiple designs simultaneously.

The settings in an **.stp** must be compatible with the programming **.sof** used to program the device. An **.stp** is considered compatible with an **.sof** when the settings for the logic analyzer, such as the size of the capture buffer and the signals selected for monitoring or triggering, match the way the target device is programmed. If the files are not compatible, you can still program the device, but you cannot run or control the logic analyzer from the SignalTap II Logic Analyzer Editor.

 When the SignalTap II Logic Analyzer detects incompatibility after analysis is started, a system error message is generated containing two CRC values, the expected value and the value retrieved from the **.stp** instance on the device. The CRC values are calculated based on all SignalTap II settings that affect the compilation.

To ensure programming compatibility, make sure to program your device with the latest **.sof** created from the most recent compilation. Checking whether or not a particular **SOF** is compatible with the current SignalTap II configuration is achieved quickly by attaching the **SOF** to the SOF manager. For more details about using the SOF manager, refer to [“Managing Multiple SignalTap II Files and Configurations”](#) on page 15-23.

Before starting a debugging session, do not make any changes to the **.stp** settings that would require recompiling the project. You can check the SignalTap II status display at the top of the **Instance Manager** pane to verify whether a change you made requires recompiling the project, producing a new **.sof**. This gives you the opportunity to undo the change, so that you do not need to recompile your project. To prevent any such changes, select **Allow trigger condition changes only** to lock the **.stp**.

Although the Quartus II project is not required when using an **.stp**, it is recommended. The project database contains information about the integrity of the current SignalTap II Logic Analyzer session. Without the project database, there is no way to verify that the current **.stp** matches the **.sof** that is downloaded to the device. If you have an **.stp** that does not match the **.sof**, incorrect data is captured in the SignalTap II Logic Analyzer.

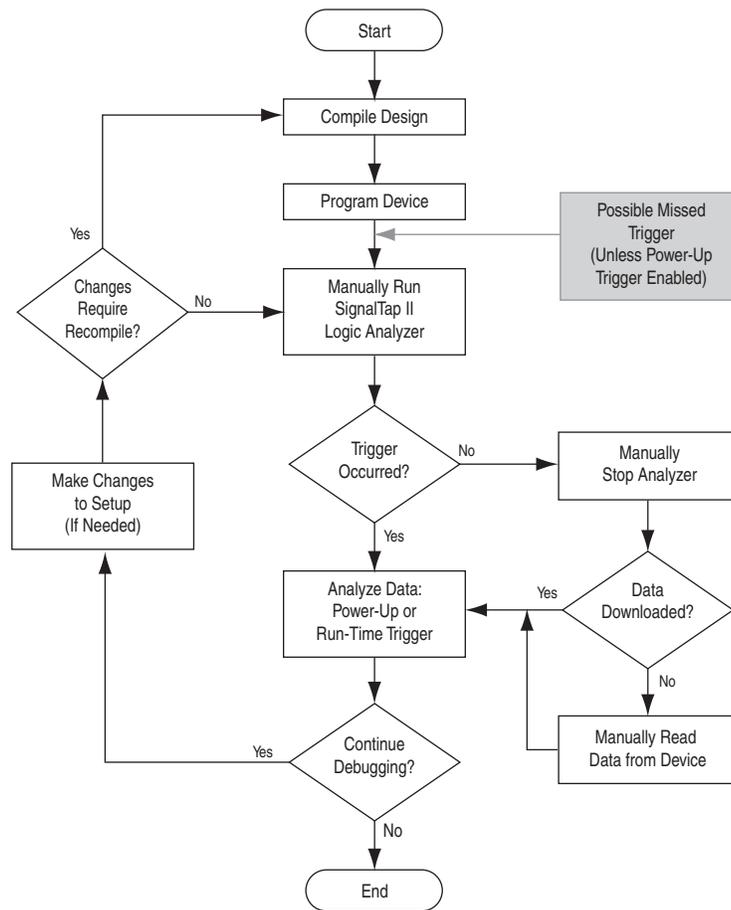
 For instructions on programming devices in the Quartus II software, refer to [Running the SignalTap II Logic Analyzer](#) in Quartus II Help.

Run the SignalTap II Logic Analyzer

After the device is configured with your design that includes the SignalTap II Logic Analyzer, perform debugging operations in a manner similar to when you use an external logic analyzer. You initialize the logic analyzer by starting an analysis. When your trigger event occurs, the captured data is stored in the memory buffer on the device and then transferred to the **.stp** with the JTAG connection.

You can also perform the equivalent of a force trigger instruction that lets you view the captured data currently in the buffer without a trigger event occurring. [Figure 15-37](#) illustrates a flow that shows how you operate the SignalTap II Logic Analyzer. The flowchart indicates where Power-Up and Runtime Trigger events occur and when captured data from these events is available for analysis.

Figure 15-37. Power-Up and Runtime Trigger Events Flowchart



❓ For information on running the analyzer from the **Instance Manager** pane, refer to *Running the SignalTap II Logic Analyzer* in Quartus II Help.

🔧 You can also use In-System Sources and Probes in conjunction with the SignalTap II Logic Analyzer to force trigger conditions. The In-System Sources and Probes feature allows you to drive and sample values on to selected signals over the JTAG chain. For more information, refer to the *Design Debugging Using In-System Sources and Probes* chapter in volume 3 of the *Quartus II Handbook*.

Runtime Reconfigurable Options

Certain settings in the `.stp` are changeable without recompiling your design when you use Runtime Trigger mode. Runtime Reconfigurable features are described in [Table 15-9](#).

Table 15-9. Runtime Reconfigurable Features

Runtime Reconfigurable Setting	Description
Basic Trigger Conditions and Basic Storage Qualifier Conditions	All signals that have the Trigger condition turned on can be changed to any basic trigger condition value without recompiling.
Advanced Trigger Conditions and Advanced Storage Qualifier Conditions	Many operators include runtime configurable settings. For example, all comparison operators are runtime-configurable. Configurable settings are shown with a white background in the block representation. This runtime reconfigurable option is turned on in the Object Properties dialog box.
Switching between a storage-qualified and a continuous acquisition	Within any storage-qualified mode, you can switch to continuous capture mode without recompiling the design. To enable this feature, turn on disable storage qualifier .
State-based trigger flow parameters	Table 15-4 lists Reconfigurable State-based trigger flow options.

Runtime Reconfigurable options can potentially save time during the debugging cycle by allowing you to cover a wider possible scenario of events without the need to recompile the design. You may experience a slight impact to the performance and logic utilization of the SignalTap II IP core. You can turn off Runtime re-configurability for Advanced Trigger Conditions and the State-based trigger flow parameters, boosting performance and decreasing area utilization.

You can configure the `.stp` to prevent changes that normally require recompilation. To do this, in the **Setup** tab, select **Allow Trigger Condition changes only** above the node list.

Example 15-4 illustrates a potential use case for Runtime Reconfigurable features. This example provides a storage qualified enabled State-based trigger flow description and shows how you can modify the size of a capture window at runtime without a recompile. This example gives you equivalent functionality to a segmented buffer with a single trigger condition where the segment sizes are runtime reconfigurable.

Example 15-4. Trigger Flow Description Providing Runtime Reconfigurable “Segments”

```
state ST1:
if ( condition1 && (c1 <= m) ) // each "segment" triggers on condition
//1
begin // m = number of total "segments"
start_store;
increment c1;
goto ST2:
End

else (c1 > m ) //This else condition handles the last
//segment.
begin
start_store
Trigger (n-1)
end

state ST2:
if ( c2 >= n) //n = number of samples to capture in each
//segment.
begin
reset c2;
stop_store;
goto ST1;
end

else (c2 < n)
begin
increment c2;
goto ST2;
end
```

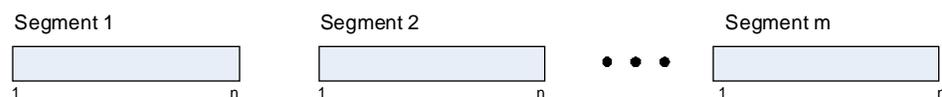
Note to Example 15-4:

(1) $m \times n$ must equal the sample depth to efficiently use the space in the sample buffer.

Figure 15-38 shows a segmented buffer described by the trigger flow in Example 15-4.

During runtime, the values m and n are runtime reconfigurable. By changing the m and n values in the preceding trigger flow description, you can dynamically adjust the segment boundaries without incurring a recompile.

Figure 15-38. Segmented Buffer Created with Storage Qualifier and State-Based Trigger (Note 1)



Note to Figure 15-38:

(1) Total sample depth is fixed, where $m \times n$ must equal sample depth.

You can add states into the trigger flow description and selectively mask out specific states and enable other ones at runtime with status flags.

[Example 15-5](#) shows a modified description of [Example 15-4](#) with an additional state inserted. You use this extra state to specify a different trigger condition that does not use the storage qualifier feature. You insert status flags into the conditional statements to control the execution of the trigger flow.

Example 15-5. Modified Trigger Flow Description of Example 16-4 with Status Flags to Selectively Enable States

```
state ST1 :
if (condition2 && f1) //additional state added for a non-segmented
                    //acquisition Set f1 to enable state
begin
    start_store;
    trigger
end

else if (! f1)
    goto ST2;

state ST2:
if ( (condition1 && (c1 <= m) && f2) // f2 status flag used to mask state. Set f2
                    //to enable.
begin
    start_store;
    increment c1;
    goto ST3:
end

else (c1 > m )
    start_store
    Trigger (n-1)
end

state ST3:
if ( c2 >= n)
begin
    reset c2;
    stop_store;
    goto ST1;
end

else (c2 < n)
begin
    increment c2;
    goto ST2;
end
```

SignalTap II Status Messages

Table 15-10 describes the text messages that might appear in the SignalTap II Status Indicator in the **Instance Manager** pane before, during, and after a data acquisition. Use these messages to monitor the state of the logic analyzer or what operation it is performing.

Table 15-10. Text Messages in the SignalTap II Status Indicator

Message	Message Description
Not running	The SignalTap II Logic Analyzer is not running. There is no connection to a device or the device is not configured.
(Power-Up Trigger) Waiting for clock (1)	The SignalTap II Logic Analyzer is performing a Runtime or Power-Up Trigger acquisition and is waiting for the clock signal to transition.
Acquiring (Power-Up) pre-trigger data (1)	The trigger condition has not been evaluated yet. A full buffer of data is collected if using the non-segmented buffer acquisition mode and storage qualifier type is continuous.
Trigger In conditions met	Trigger In condition has occurred. The SignalTap II Logic Analyzer is waiting for the condition of the first trigger condition to occur. This can appear if Trigger In is specified.
Waiting for (Power-up) trigger (1)	The SignalTap II Logic Analyzer is now waiting for the trigger event to occur.
Trigger level <x> met	The condition of trigger condition x has occurred. The SignalTap II Logic Analyzer is waiting for the condition specified in condition $x + 1$ to occur.
Acquiring (power-up) post-trigger data (1)	The entire trigger event has occurred. The SignalTap II Logic Analyzer is acquiring the post-trigger data. The amount of post-trigger data collected is you define between 12%, 50%, and 88% when the non-segmented buffer acquisition mode is selected.
Offload acquired (Power-Up) data (1)	Data is being transmitted to the Quartus II software through the JTAG chain.
Ready to acquire	The SignalTap II Logic Analyzer is waiting for you to initialize the analyzer.

Note to Table 15-10:

(1) This message can appear for both Runtime and Power-Up Trigger events. When referring to a Power-Up Trigger, the text in parentheses is added.

 In segmented acquisition mode, pre-trigger and post-trigger do not apply.

View, Analyze, and Use Captured Data

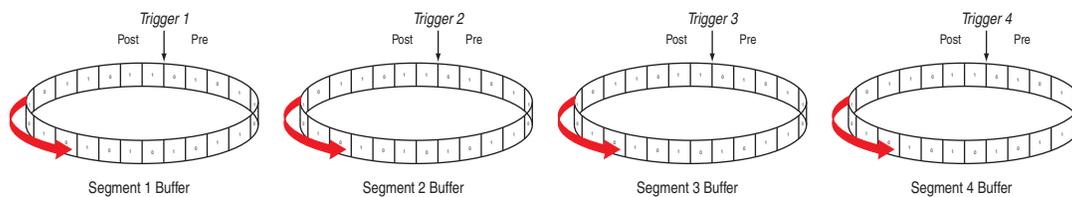
Once a trigger event has occurred or you capture data manually, you can use the SignalTap II interface to examine the data, and use your findings to help debug your design.

 For information about what you can do with captured data, refer to *Analyzing Data in the SignalTap II Logic Analyzer* in Quartus II Help.

Capturing Data Using Segmented Buffers

Segmented Acquisition buffers allow you to perform multiple captures with a separate trigger condition for each acquisition segment. This feature allows you to capture a recurring event or sequence of events that span over a long period time efficiently. Each acquisition segment acts as a non-segmented buffer, continuously capturing data when it is activated. When you run an analysis with the **segmented buffer** option enabled, the SignalTap II Logic Analyzer performs back-to-back data captures for each acquisition segment within your data buffer. The trigger flow, or the type and order in which the trigger conditions evaluate for each buffer, is defined by either the Sequential trigger flow control or the Custom State-based trigger flow control. [Figure 15-39](#) shows a segmented acquisition buffer with four segments represented as four separate non-segmented buffers.

Figure 15-39. Segmented Acquisition Buffer



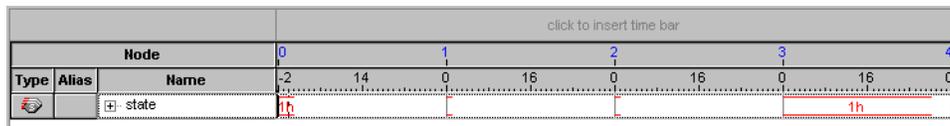
The SignalTap II Logic Analyzer finishes an acquisition with a segment, and advances to the next segment to start a new acquisition. Depending on when a trigger condition occurs, it may affect the way the data capture appears in the waveform viewer.

[Figure 15-39](#) illustrates the method in which data is captured. The Trigger markers in [Figure 15-39](#)—Trigger 1, Trigger 2, Trigger 3 and Trigger 4—refer to the evaluation of the `segment_trigger` and `trigger` commands in the Custom State-based trigger flow. If you use a sequential flow, the Trigger markers refer to trigger conditions specified within the **Setup** tab.

If the Segment 1 Buffer is the active segment and Trigger 1 occurs, the SignalTap II Logic Analyzer starts evaluating Trigger 2 immediately. Data Acquisition for Segment 2 buffer starts when either Segment Buffer 1 finishes its post-fill count, or when Trigger 2 evaluates as `TRUE`, whichever condition occurs first. Thus, trigger conditions associated with the next buffer in the data capture sequence can preempt the post-fill count of the current active buffer. This allows the SignalTap II Logic Analyzer to accurately capture all of the trigger conditions that have occurred. Samples that have not been used appear as a blank space in the waveform viewer.

Figure 15-40 shows an example of a capture using sequential flow control with the trigger condition for each segment specified as **Don't Care**. Each segment before the last captures only one sample, because the next trigger condition immediately preempts capture of the current buffer. The trigger position for all segments is specified as pre-trigger (10% of the data is before the trigger condition and 90% of the data is after the trigger position). Because the last segment starts immediately with the trigger condition, the segment contains only post-trigger data. The three empty samples in the last segment are left over from the pre-trigger samples that the SignalTap II Logic Analyzer allocated to the buffer.

Figure 15-40. Segmented Capture with Preemption of Acquisition Segments (Note 1)



Note to Figure 15-40:

- (1) A segmented acquisition buffer using the sequential trigger flow with a trigger condition specified as Don't Care. All segments, with the exception of the last segment, capture only one sample because the next trigger condition preempts the current buffer from filling to completion.

For the sequential trigger flow, the **Trigger Position** option applies to every segment in the buffer. For maximum flexibility on how the trigger position is defined, use the custom state-based trigger flow. By adjusting the trigger position specific to your debugging requirements, you can help maximize the use of the allocated buffer space.

Differences in Pre-fill Write Behavior Between Different Acquisition Modes

The SignalTap II Logic Analyzer uses one of the following three modes when writing into the acquisition memory:

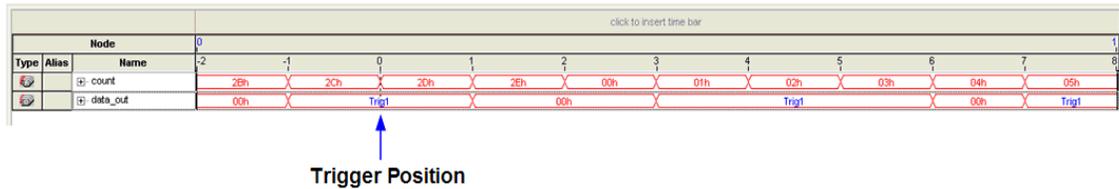
- Non-segmented buffer
- Non-segmented buffer with a storage qualifier
- Segmented buffer

There are subtle differences in the amount of data captured immediately after running the SignalTap II Logic Analyzer and before any trigger conditions occur. A non-segmented buffer, running in continuous mode, completely fills the buffer with sampled data before evaluating any trigger conditions. Thus, a non-segmented capture without any storage qualification enabled always shows a waveform with a full buffer's worth of data captured.

Filling the buffer provides you with as much data as possible within the capture window. The buffer gets pre-filled with data samples prior to evaluating the trigger condition. As such, SignalTap requires that the buffer be filled at least once before any data can be retrieved through the JTAG connection and prevents the buffer from being dumped during the first acquisition prior to a trigger condition when you perform a **Stop Analysis**.

Notice in Figure 15-41 that Trig1 occurs several times in the data buffer before the SignalTap II Logic Analyzer actually triggers. A full buffer's worth of data is captured before the logic analyzer evaluates any trigger conditions. After the trigger condition occurs, the logic analyzer continues acquisition until it captures eight additional samples (12% of the buffer, as defined by the "post-trigger" position).

Figure 15-42. SignalTap II Logic Analyzer Conditional Data Capture (Note 1)



Note to Figure 15-42:

- (1) Conditional capture, storage always enabled, post-fill count.
- (2) SignalTap II Logic Analyzer capture of a recurring pattern using a non-segmented buffer in conditional mode. The logic analyzer is configured with a basic trigger condition (shown in the figure as "Trig1"), with a sample depth of 64 bits. The "Trigger in" condition is specified as "Don't care", which means that every sample is captured.

Notice in Figure 15-42 that the logic analyzer triggers immediately. As in Figure 15-41, the logic analyzer completes the acquisition with eight samples, or 12% of 64, the sample capacity of the acquisition buffer.

Creating Mnemonics for Bit Patterns

The mnemonic table feature allows you to assign a meaningful name to a set of bit patterns, such as a bus. To create a mnemonic table, right-click in the **Setup** or **Data** tab of an **.stp** and click **Mnemonic Table Setup**. You create a mnemonic table by entering sets of bit patterns and specifying a label to represent each pattern. Once you have created a mnemonic table, assign the table to a group of signals. To assign a mnemonic table, right-click on the group, click **Bus Display Format** and select the desired mnemonic table.

You use the labels you create in a table in different ways on the **Setup** and **Data** tabs. On the **Setup** tab, you can create basic triggers with meaningful names by right-clicking an entry in the **Trigger Conditions** column and selecting a label from the table you assigned to the signal group. On the **Data** tab, if any captured data matches a bit pattern contained in an assigned mnemonic table, the signal group data is replaced with the appropriate label, making it easy to see when expected data patterns occur.

Automatic Mnemonics with a Plug-In

When you use a plug-in to add signals to an **.stp**, mnemonic tables for the added signals are automatically created and assigned to the signals defined in the plug-in. To enable these mnemonic tables manually, right-click on the name of the signal or signal group. On the **Bus Display Format** shortcut menu, then click the name of the mnemonic table that matches the plug-in.

As an example, the Nios II plug-in helps you to monitor your design's signal activity as code is executed. If you set up the logic analyzer to trigger on a function name in your Nios II code based on data from an `.elf`, you can see the function name in the **Instance Address** signal group at the trigger sample, along with the corresponding disassembled code in the **Disassembly** signal group, as shown in Figure 15-43. Captured data samples around the trigger are referenced as offset addresses from the trigger function name.

Figure 15-43. Data Tab when the Nios II Plug-In is Used

Type	Alias	Name	Value
PCNios II Inst Address	alt_main+0x8
DISNios II Disassembly	mov fp, sp
			alt_main+0xc
			movi r2, 2

Locating a Node in the Design

When you find the source of an error in your design using the SignalTap II Logic Analyzer, you can use the node locate feature to locate that signal in many of the tools found in the Quartus II software, as well as in your design files. This lets you find the source of the problem quickly so you can modify your design to correct the flaw. To locate a signal from the SignalTap II Logic Analyzer in one of the Quartus II software tools or your design files, right-click on the signal in the `.stp`, and click **Locate in <tool name>**.

You can locate a signal from the node list with the following tools:

- Assignment Editor
- Pin Planner
- Timing Closure Floorplan
- Chip Planner
- Resource Property Editor
- Technology Map Viewer
- RTL Viewer
- Design File



For more information about using these tools, refer to each of the corresponding chapters in the *Quartus II Handbook*.

Saving Captured Data

The data log shows the history of captured data and the triggers used to capture the data. The SignalTap II Logic Analyzer acquires data, stores it in a log, and displays it as waveforms. When the logic analyzer is in auto-run mode and a trigger event occurs more than once, captured data for each time the trigger occurred is stored as a separate entry in the data log. This allows you to review the captured data for each trigger event. The default name for a log is based on the time when the data was acquired. Altera recommends that you rename the data log with a more meaningful name.

The logs are organized in a hierarchical manner; similar logs of captured data are grouped together in trigger sets. To open the **Data Log** pane, on the View menu, select **Data Log**. To turn on data logging, turn on **Enable data log** in the **Data Log** (Figure 15-19). To recall and activate a data log for a given trigger set, double-click the name of the data log in the list.

You can use the Data Log feature for organizing different sets of trigger conditions and different sets of signal configurations. For more information, refer to “[Managing Multiple SignalTap II Files and Configurations](#)” on page 15-23.

Exporting Captured Data to Other File Formats

You can export captured data to the following file formats, for use with other EDA simulation tools:

- Comma Separated Values File (.csv)
- Table File (.tbl)
- Value Change Dump File (.vcd)
- Vector Waveform File (.vwf)
- Graphics format files (.jpg, .bmp)

To export the SignalTap II Logic Analyzer’s captured data, on the File menu, click **Export** and specify the **File Name**, **Export Format**, and **Clock Period**.

Creating a SignalTap II List File

Captured data can also be viewed in an **.stp** list file. An **.stp** list file is a text file that lists all the data captured by the logic analyzer for a trigger event. Each row of the list file corresponds to one captured sample in the buffer. Columns correspond to the value of each of the captured signals or signal groups for that sample. If a mnemonic table was created for the captured data, the numerical values in the list are replaced with a matching entry from the table. This is especially useful with the use of a plug-in that includes instruction code disassembly. You can immediately see the order in which the instruction code was executed during the same time period of the trigger event. To create an **.stp** list file in the Quartus II software, on the File menu, select **Create/Update** and click **Create SignalTap II List File**.

Other Features

The SignalTap II Logic Analyzer has other features that do not necessarily belong to a particular task in the task flow.

Using the SignalTap II MATLAB MEX Function to Capture Data

If you use MATLAB for DSP design, you can call the MATLAB MEX function `alt_signaltap_run`, built into the Quartus II software, to acquire data from the SignalTap II Logic Analyzer directly into a matrix in the MATLAB environment. If you use the MATLAB MEX function in a loop, you can perform as many acquisitions in the same amount of time as you can when using SignalTap II in the Quartus II software environment.



The SignalTap II MATLAB MEX function is available only in the Windows version of the Quartus II software. It is compatible with MATLAB Release 14 Original Release Version 7 and later.

To set up the Quartus II software and the MATLAB environment to perform SignalTap II acquisitions, perform the following steps:

1. In the Quartus II software, create an **.stp** file.
2. In the node list in the **Data** tab of the SignalTap II Logic Analyzer Editor, organize the signals and groups of signals into the order in which you want them to appear in the MATLAB matrix. Each column of the imported matrix represents a single SignalTap II acquisition sample, while each row represents a signal or group of signals in the order they are organized in the **Data** tab.



Signal groups acquired from the SignalTap II Logic Analyzer and transferred into the MATLAB MEX function are limited to a width of 32 signals. If you want to use the MATLAB MEX function with a bus or signal group that contains more than 32 signals, split the group into smaller groups that do not exceed the 32-signal limit.

3. Save the **.stp** and compile your design. Program your device and run the SignalTap II Logic Analyzer to ensure your trigger conditions and signal acquisition work correctly.
4. In the MATLAB environment, add the Quartus II binary directory to your path with the following command:

```
addpath <Quartus install directory>\win ←
```

You can view the help file for the MEX function by entering the following command in MATLAB without any operators:

```
alt_signaltap_run ←
```

Use the MATLAB MEX function to open the JTAG connection to the device and run the SignalTap II Logic Analyzer to acquire data. When you finish acquiring data, close the JTAG connection.

To open the JTAG connection and begin acquiring captured data directly into a MATLAB matrix called `stp`, use the following command:

```
stp = alt_signaltap_run \
(' <stp filename>' [, ('signed' | 'unsigned') [, '<instance names>' [, \
' <signalset name>' [, '<trigger name>' ]]]]); ←
```

When capturing data you must assign a filename, for example, *<stp filename>* as a requirement of the MATLAB MEX function. Other MATLAB MEX function options are described in [Table 15-11](#).

Table 15-11. SignalTap II MATLAB MEX Function Options

Option	Usage	Description
signed unsigned	'signed' 'unsigned'	The signed option turns signal group data into 32-bit two's-complement signed integers. The MSB of the group as defined in the SignalTap II Data tab is the sign bit. The unsigned option keeps the data as an unsigned integer. The default is signed .
<i><instance name></i>	'auto_signaltap_0'	Specify a SignalTap II instance if more than one instance is defined. The default is the first instance in the .stp , <i>auto_signaltap_0</i> .
<i><signal set name></i> <i><trigger name></i>	'my_signalset' 'my_trigger'	Specify the signal set and trigger from the SignalTap II data log if multiple configurations are present in the .stp . The default is the active signal set and trigger in the file.

You can enable or disable verbose mode to see the status of the logic analyzer while it is acquiring data. To enable or disable verbose mode, use the following commands:

```
alt_signaltap_run('VERBOSE_ON'); ←  
alt_signaltap_run('VERBOSE_OFF'); ←
```

When you finish acquiring data, close the JTAG connection with the following command:

```
alt_signaltap_run('END_CONNECTION'); ←
```



For more information about the use of MATLAB MEX functions in MATLAB, refer to the MATLAB Help.

Using SignalTap II in a Lab Environment

You can install a stand-alone version of the SignalTap II Logic Analyzer. This version is particularly useful in a lab environment in which you do not have a workstation that meets the requirements for a complete Quartus II installation, or if you do not have a license for a full installation of the Quartus II software. The standalone version of the SignalTap II Logic Analyzer is included with the Quartus II stand-alone Programmer and is available from the Downloads page of the Altera website (www.altera.com).

Remote Debugging Using the SignalTap II Logic Analyzer

You can use the SignalTap II Logic Analyzer to debug a design that is running on a device attached to a PC in a remote location.

To perform a remote debugging session, you must have the following setup:

- The Quartus II software installed on the local PC
- Stand-alone SignalTap II Logic Analyzer or the full version of the Quartus II software installed on the remote PC
- Programming hardware connected to the device on the PCB at the remote location
- TCP/IP protocol connection

Equipment Setup

On the PC in the remote location, install the standalone version of the SignalTap II Logic Analyzer or the full version of the Quartus II software. This remote computer must have Altera programming hardware connected, such as the EthernetBlaster or USB-Blaster.

On the local PC, install the full version of the Quartus II software. This local PC must be connected to the remote PC across a LAN with the TCP/IP protocol.

- ❓ For information about enabling remote access to a JTAG server, refer to *Using the JTAG Server* in Quartus II Help.

Using the SignalTap II Logic Analyzer in Devices with Configuration Bitstream Security

Certain device families support bitstream decryption during configuration using an on-device AES decryption engine. You can still use the SignalTap II Logic Analyzer to analyze functional data within the FPGA. However, note that JTAG configuration is not possible after the security key has been programmed into the device.

Altera recommends that you use an unencrypted bitstream during the prototype and debugging phases of the design. Using an unencrypted bitstream allows you to generate new programming files and reconfigure the device over the JTAG connection during the debugging cycle.

If you must use the SignalTap II Logic Analyzer with an encrypted bitstream, first configure the device with an encrypted configuration file using Passive Serial (PS), Fast Passive Parallel (FPP), or Active Serial (AS) configuration modes. The design must contain at least one instance of the SignalTap II Logic Analyzer. After the FPGA is configured with a SignalTap II Logic Analyzer instance in the design, when you open the SignalTap II Logic Analyzer in the Quartus II software, you then scan the chain and are ready to acquire data with the JTAG connection.

Backward Compatibility with Previous Versions of Quartus II Software

You can open an **.stp** created in a previous version in a current version of the Quartus II software. However, opening an **.stp** modifies it so that it cannot be opened in a previous version of the Quartus II software.

If you have a Quartus II project file from a previous version of the software, you may have to update the **.stp** configuration file to recompile the project. You can update the configuration file by opening the SignalTap II Logic Analyzer. If you need to update your configuration, a prompt appears asking if you would like to update the **.stp** to match the current version of the Quartus II software.

SignalTap II Command-Line Options

To compile your design with the SignalTap II Logic Analyzer using the command prompt, use the `quartus_stp` command. Table 15-12 shows the options that help you use the `quartus_stp` executable.

Table 15-12. SignalTap II Command-Line Options

Option	Usage	Description
<code>stp_file</code>	<code>quartus_stp</code> <code>--stp_file <stp_filename></code>	Assigns the specified <code>.stp</code> to the <code>USE_SIGNALTAP_FILE</code> in the <code>.qsf</code> .
<code>enable</code>	<code>quartus_stp --enable</code>	Creates assignments to the specified <code>.stp</code> in the <code>.qsf</code> and changes <code>ENABLE_SIGNALTAP</code> to <code>ON</code> . The SignalTap II Logic Analyzer is included in your design the next time the project is compiled. If no <code>.stp</code> is specified in the <code>.qsf</code> , the <code>--stp_file</code> option must be used. If the <code>--enable</code> option is omitted, the current value of <code>ENABLE_SIGNALTAP</code> in the <code>.qsf</code> is used.
<code>disable</code>	<code>quartus_stp --disable</code>	Removes the <code>.stp</code> reference from the <code>.qsf</code> and changes <code>ENABLE_SIGNALTAP</code> to <code>OFF</code> . The SignalTap II Logic Analyzer is removed from the design database the next time you compile your design. If the <code>--disable</code> option is omitted, the current value of <code>ENABLE_SIGNALTAP</code> in the <code>.qsf</code> is used.
<code>create_signaltap_hdl_file</code>	<code>quartus_stp</code> <code>--create_signaltap_hdl_file</code>	Creates an <code>.stp</code> representing the SignalTap II instance in the design generated by the SignalTap II Logic Analyzer megafunction created with the MegaWizard Plug-In Manager. The file is based on the last compilation. You must use the <code>--stp_file</code> option to create an <code>.stp</code> properly. Analogous to the Create SignalTap II File from Design Instance(s) command in the Quartus II software.

Example 15-6 illustrates how to compile a design with the SignalTap II Logic Analyzer at the command line.

Example 15-6.

```
quartus_stp filtref --stp_file stp1.stp --enable ←
quartus_map filtref --source=filtref.bdf --family=CYCLONE ←
quartus_fit filtref --part=EP1C12Q240C6 --fmax=80MHz --tsu=8ns ←
quartus_asm filtref ←
```

The `quartus_stp --stp_file stp1.stp --enable` command creates the QSF variable and instructs the Quartus II software to compile the `stp1.stp` file with your design. The `--enable` option must be applied for the SignalTap II Logic Analyzer to compile properly into your design.

Example 15-7 shows how to create a new **.stp** after building the SignalTap II Logic Analyzer instance with the MegaWizard Plug-In Manager.

Example 15-7.

```
quartus_stp filtref --create_signaltap_hdl_file --stp_file stp1.stp ←
```



For information about the other command line executables and options, refer to the *Command-Line Scripting* chapter in volume 2 of the *Quartus II Handbook*.

SignalTap II Tcl Commands

The **quartus_stp** executable supports a Tcl interface that allows you to capture data without running the Quartus II GUI. You cannot execute SignalTap II Tcl commands from within the Tcl console in the Quartus II software. They must be executed from the command line with the **quartus_stp** executable. To execute a Tcl file that has SignalTap II Logic Analyzer Tcl commands, use the following command:

```
quartus_stp -t <Tcl file> ←
```



For information about Tcl commands that you can use with the SignalTap II Logic Analyzer Tcl package, refer to `::quartus::stp` in Quartus II Help.

Example 15-8 is an excerpt from a script you can use to continuously capture data. Once the trigger condition is met, the data is captured and stored in the data log.

Example 15-8.

```
#opens signaltap session
open_session -name stp1.stp
#start acquisition of instance auto_signaltap_0 and
#auto_signaltap_1 at the same time
#calling run_multiple_end will start all instances
#run after run_multiple_start call
run_multiple_start
run -instance auto_signaltap_0 -signal_set signal_set_1 -trigger /
trigger_1 -data_log log_1 -timeout 5
run -instance auto_signaltap_1 -signal_set signal_set_1 -trigger /
trigger_1 -data_log log_1 -timeout 5
run_multiple_end
#close signaltap session
close_session
```

When the script is completed, open the **.stp** that you used to capture data to examine the contents of the Data Log.

Design Example: Using SignalTap II Logic Analyzers in SOPC Builder Systems

The system in this example contains many components, including a Nios processor, a direct memory access (DMA) controller, on-chip memory, and an interface to external SDRAM memory. In this example, the Nios processor executes a simple C program from on-chip memory and waits for you to press a button. After you press a button, the processor initiates a DMA transfer, which you analyze using the SignalTap II Logic Analyzer.

-  For more information about this example and using the SignalTap II Logic Analyzer with SOPC builder systems, refer to [AN 323: Using SignalTap II Logic Analyzers in SOPC Builder Systems](#) and [AN 446: Debugging Nios II Systems with the SignalTap II Logic Analyzer](#).

Custom Triggering Flow Application Examples

The custom triggering flow in the SignalTap II Logic Analyzer is most useful for organizing a number of triggering conditions and for precise control over the acquisition buffer. This section provides two application examples for defining a custom triggering flow within the SignalTap II Logic Analyzer. Both examples can be easily copied and pasted directly into the state machine description box by using the state display mode **All states in one window**.

-  For additional triggering flow design examples for on-chip debugging, refer to the [On-chip Debugging Design Examples](#) page on the Altera website.

Design Example 1: Specifying a Custom Trigger Position

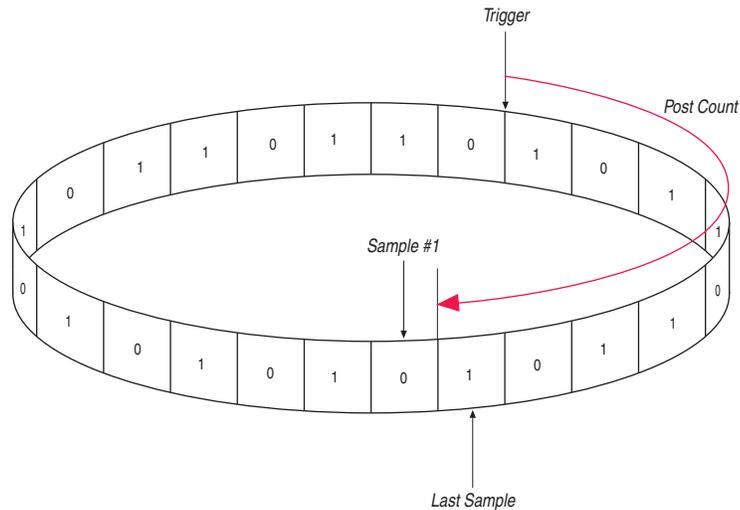
Actions to the acquisition buffer can accept an optional post-count argument. This post-count argument enables you to define a custom triggering position for each segment in the acquisition buffer. [Example 15-9](#) shows an example that applies a trigger position to all segments in the acquisition buffer. The example describes a triggering flow for an acquisition buffer split into four segments. If each acquisition segment is 64 samples in depth, the trigger position for each buffer will be at sample #34. The acquisition stops after all four segments are filled once.

Example 15-9.

```
if (c1 == 3 && condition1)
    trigger 30;
else if ( condition1 )
begin
    segment_trigger 30;
    increment c1;
end
```

Each segment acts as a non-segmented buffer that continuously updates the memory contents with the signal values. The last acquisition before stopping the buffer is displayed on the **Data** tab as the last sample number in the affected segment. The trigger position in the affected segment is then defined by $N - \text{post count fill}$, where N is the number of samples per segment. [Figure 15-44](#) illustrates the triggering position.

Figure 15-44. Specifying a Custom Trigger Position



Design Example 2: Trigger When triggercond1 Occurs Ten Times between triggercond2 and triggercond3

The custom trigger flow description is often useful to count a sequence of events before triggering the acquisition buffer. [Example 15-10](#) shows such a sample flow. This example uses three basic triggering conditions configured in the SignalTap II Setup tab.

This example triggers the acquisition buffer when condition1 occurs after condition3 and occurs ten times prior to condition3. If condition3 occurs prior to ten repetitions of condition1, the state machine transitions to a permanent wait state.

Example 15-10.

```
state ST1:

if ( condition2 )
begin
    reset c1;
    goto ST2;
end

State ST2 :
if ( condition1 )
    increment c1;

else if (condition3 && c1 < 10)
    goto ST3;

else if ( condition3 && c1 >= 10)
    trigger;

ST3:
goto ST3;
```

SignalTap II Scripting Support

You can run procedures and make settings described in this chapter in a Tcl script. You can also run some procedures at a command prompt. For detailed information about scripting command options, refer to the Quartus II Command-Line and Tcl API Help browser. To run the Help browser, type the following at the command prompt:

```
quartus_sh --qhelp ←
```

 For more information about Tcl scripting, refer to the *Tcl Scripting* chapter in volume 2 of the *Quartus II Handbook*

 You can also refer to *About Quartus II Tcl Scripting* in Quartus II Help.

Conclusion

As the FPGA industry continues to make technological advancements, outdated methodologies are replaced with new technologies that maximize productivity. The SignalTap II Logic Analyzer gives you the same benefits as a traditional logic analyzer, without the many shortcomings of a piece of dedicated test equipment. The SignalTap II Logic Analyzer provides many new and innovative features that allow you to capture and analyze internal signals in your FPGA, allowing you to quickly debug your design.

Document Revision History

Table 15-13 shows the revision history for this chapter.

Table 15-13. Document Revision History

Date	Version	Changes Made
December 2010	10.0.1	Changed to new document template.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Add new acquisition buffer content to the “View, Analyze, and Use Captured Data” section. ■ Added script sample for generating hexadecimal CRC values in programmed devices. ■ Created cross references to Quartus II Help for duplicated procedural content.
November 2009	9.1.0	No change to content.
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Updated Table 15-1 ■ Updated “Using Incremental Compilation with the SignalTap II Logic Analyzer” on page 15-45 ■ Added new Figure 15-35 ■ Made minor editorial updates
November 2008	8.1.0	Updated for the Quartus II software version 8.1 release: <ul style="list-style-type: none"> ■ Added new section “Using the Storage Qualifier Feature” on page 14-25 ■ Added description of <code>start_store</code> and <code>stop_store</code> commands in section “Trigger Condition Flow Control” on page 14-36 ■ Added new section “Runtime Reconfigurable Options” on page 14-63
May 2008	8.0.0	Updated for the Quartus II software version 8.0: <ul style="list-style-type: none"> ■ Added “Debugging Finite State machines” on page 14-24 ■ Documented various GUI usability enhancements, including improvements to the resource estimator, the bus find feature, and the dynamic display updates to the counter and flag resources in the State-based trigger flow control tab ■ Added “Capturing Data Using Segmented Buffers” on page 14-16 ■ Added hyperlinks to referenced documents throughout the chapter ■ Minor editorial updates

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

The Quartus II Logic Analyzer Interface (LAI) allows you to use an external logic analyzer and a minimal number of Altera-supported device I/O pins to examine the behavior of internal signals while your design is running at full speed on your Altera®-supported device.

The LAI connects a large set of internal device signals to a small number of output pins. You can connect these output pins to an external logic analyzer for debugging purposes. In the Quartus II LAI, the internal signals are grouped together, distributed to a user-configurable multiplexer, and then output to available I/O pins on your Altera-supported device. Instead of having a one-to-one relationship between internal signals and output pins, the Quartus II LAI enables you to map many internal signals to a smaller number of output pins. The exact number of internal signals that you can map to an output pin varies based on the multiplexer settings in the Quartus II LAI.

This chapter details the following topics:

- “Choosing a Logic Analyzer”
- “Debugging Your Design Using the LAI” on page 16–4
- “Working with LAI Files” on page 16–4
- “Controlling the Active Bank During Runtime” on page 16–7
- “Using the LAI with Incremental Compilation” on page 16–7



The term “logic analyzer” when used in this chapter includes both logic analyzers and oscilloscopes equipped with digital channels, commonly referred to as mixed signal analyzers or MSOs.

- ❓ Refer to *Devices and Adapters* in Quartus II Help for a list of Altera-supported devices.

Choosing a Logic Analyzer

The Quartus II software offers the following two general purpose on-chip debugging tools for debugging a large set of RTL signals from your design:

- The SignalTap® II Logic Analyzer
- An external logic analyzer, which connects to internal signals in your Altera-supported device by using the Quartus II LAI



Table 16–1 describes the advantages of each debugging tool.

Table 16–1. Comparing the SignalTap II Logic Analyzer with the Logic Analyzer Interface

Feature and Description	Logic Analyzer Interface	SignalTap II Logic Analyzer
Sample Depth You have access to a wider sample depth with an external logic analyzer. In the SignalTap II Logic Analyzer, the maximum sample depth is set to 128 Kb, which is a device constraint. However, with an external logic analyzer, there are no device constraints, providing you a wider sample depth.	✓	—
Debugging Timing Issues Using an external logic analyzer provides you with access to a “timing” mode, which enables you to debug combined streams of data.	✓	—
Performance You frequently have limited routing resources available to place and route when you use the SignalTap II Logic Analyzer with your design. An external logic analyzer adds minimal logic, which removes resource limits on place-and-route.	✓	—
Triggering Capability The SignalTap II Logic Analyzer offers triggering capabilities that are comparable to external logic analyzers.	✓	✓
Use of Output Pins Using the SignalTap II Logic Analyzer, no additional output pins are required. Using an external logic analyzer requires the use of additional output pins.	—	✓
Acquisition Speed With the SignalTap II Logic Analyzer, you can acquire data at speeds of over 200 MHz. You can achieve the same acquisition speeds with an external logic analyzer; however, you must consider signal integrity issues.	—	✓

 The Quartus II software offers a portfolio of on-chip debugging tools. For an overview and comparison of all tools available in the Quartus II software on-chip debugging tool suite, refer to *Section V. In-System Debugging* in volume 3 of the *Quartus II Handbook*.

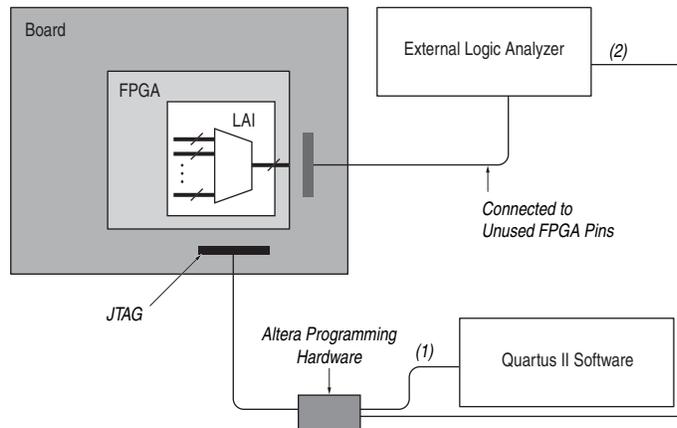
Required Components

You must have the following components to perform analysis using the Quartus II LAI:

- The Quartus II software starting with version 5.1 and later
- The device under test
- An external logic analyzer
- An Altera communications cable
- A cable to connect the Altera-supported device to the external logic analyzer

Figure 16-1 shows the LAI and the hardware setup.

Figure 16-1. LAI and Hardware Setup



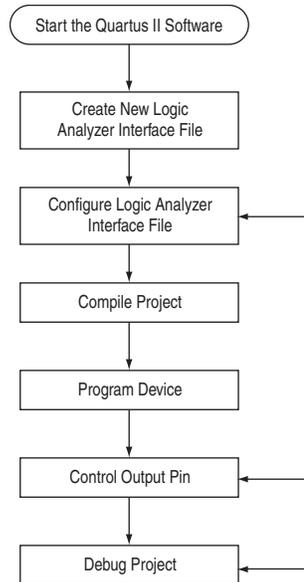
Notes to Figure 16-1:

- (1) Configuration and control of the LAI using a computer loaded with the Quartus II software via the JTAG port.
- (2) Configuration and control of the LAI using a third-party vendor logic analyzer via the JTAG port. Support varies by vendor.

Debugging Your Design Using the LAI

Figure 16-2 shows the steps you must follow to debug your design with the Quartus II LAI.

Figure 16-2. LAI and Hardware Setup



Notes to Figure 16-1:

- (1) Configuration and control of the LAI using a computer loaded with the Quartus II software via the JTAG port.
- (2) Configuration and control of the LAI using a third-party vendor logic analyzer via the JTAG port. Support varies by vendor.

Working with LAI Files

The **.lai** file stores the configuration of an LAI instance. The **.lai** file opens in the LAI editor. The editor allows you to group multiple internal signals to a set of external pins. The configuration parameters are described in the following sections.

- ① To create a new **.lai** file or open an existing **.lai** file, refer to *Setting Up the Logic Analyzer Interface* in Quartus II Help.

Configuring the File Core Parameters

After you create the .lai file, you must configure the .lai file core parameters by clicking on the **Setup View** list, and then selecting **Core Parameters**. Table 16-2 lists the .lai file core parameters.

Table 16-2. LAI File Core Parameters

Parameter	Description
Pin Count	The Pin Count parameter signifies the number of pins you want dedicated to your LAI. The pins must be connected to a debug header on your board. Within the Altera-supported device, each pin is mapped to a user-configurable number of internal signals. The Pin Count parameter can range from 1 to 255 pins.
Bank Count	The Bank Count parameter signifies the number of internal signals that you want to map to each pin. For example, a Bank Count of 8 implies that you will connect eight internal signals to each pin. The Bank Count parameter can range from 1 to 255 banks.
Output/Capture Mode	The Output/Capture Mode parameter signifies the type of acquisition you perform. There are two options that you can select: Combinational/Timing —This acquisition uses your external logic analyzer's internal clock to determine when to sample data. Because Combinational/Timing acquisition samples data asynchronously to your Altera-supported device, you must determine the sample frequency you should use to debug and verify your system. This mode is effective if you want to measure timing information, such as channel-to-channel skew. For more information about the sampling frequency and the speeds at which it can run, refer to the data sheet for your external logic analyzer. Registered/State —This acquisition uses a signal from your system under test to determine when to sample. Because Registered/State acquisition samples data synchronously with your Altera-supported device, it provides you with a functional view of your Altera-supported device while it is running. This mode is effective when you verify the functionality of your design.
Clock	The Clock parameter is available only when Output/Capture Mode is set to Registered State. You must specify the sample clock in the Core Parameters view. The sample clock can be any signal in your design. However, for best results, Altera recommends that you use a clock with an operating frequency fast enough to sample the data you would like to acquire.
Power-Up State	The Power-Up State parameter specifies the power-up state of the pins you have designated for use with the LAI. You have the option of selecting tri-stated for all pins, or selecting a particular bank that you have enabled.

Mapping the LAI File Pins to Available I/O Pins

To configure the .lai file I/O pin parameters, select **Pins** in the **Setup View** list. To assign pin locations for the LAI, double-click the **Location** column next to the reserved pins in the **Name** column, and the Pin Planner opens.



For information about how to use the Pin Planner, refer to the *Pin Planner* section in the *I/O Management* chapter in volume 2 of the *Quartus II Handbook*.

Mapping Internal Signals to the LAI Banks

After you have specified the number of banks to use in the **Core Parameters** settings page, you must assign internal signals for each bank in the LAI. Click the **Setup View** arrow and select **Bank n** or **All Banks**.

To view all of your bank connections, click **Setup View** and select **All Banks**.

Using the Node Finder

Before making bank assignments, on the View menu, point to **Utility Windows** and click **Node Finder**. Find the signals that you want to acquire, then drag and drop the signals from the **Node Finder** dialog box into the bank **Setup View**. When adding signals, use **SignalTap II: pre-synthesis** for non-incrementally routed instances and **SignalTap II: post-fitting** for incrementally routed instances.

As you continue to make assignments in the bank **Setup View**, the schematic of your LAI in the **Logical View** of your .lai file begins to reflect your assignments. Continue making assignments for each bank in the **Setup View** until you have added all of the internal signals for which you wish to acquire data.

Compiling Your Quartus II Project

When you save your .lai file, a dialog box prompts you to enable the LAI instance for the active project. Alternatively, you can specify the .lai file your project uses in the **Global Project Settings** dialog box.

After you specify the name of your .lai file, you must compile your project. To compile your project, on the Processing menu, click **Start Compilation**.

To ensure that the LAI is properly compiled with your project, expand the entity hierarchy in the Project Navigator. (To display the Project Navigator, on the View menu, point to **Utility Windows** and click **Project Navigator**.) If the LAI is compiled with your design, the sld_hub and sld_multitap entities are shown in the project navigator (Figure 16-3).

Figure 16-3. Project Navigator

Entity	Logic Cells	LC Registers
Stratix: EP1S10B672C7		
test	136 (1)	81
sld_multitap:auto_lai_0	35 (11)	15
sld_hub:sld_hub_inst	100 (25)	65

Programming Your Altera-Supported Device Using the LAI

After compilation completes, you must configure your Altera-supported device before using the LAI.

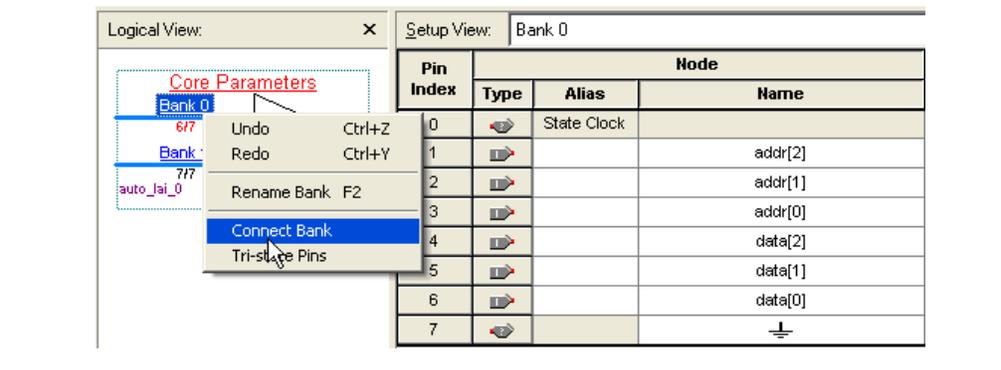
You can use the LAI with multiple devices in your JTAG chain. Your JTAG chain can also consist of devices that do not support the LAI or non-Altera, JTAG-compliant devices. To use the LAI in more than one Altera-supported device, create an .lai file and configure an .lai file for each Altera-supported device that you want to analyze.

- ❓ To configure a device or a set of devices for use with LAI, refer to *Enabling the Logic Analyzer Interface* in Quartus II Help.

Controlling the Active Bank During Runtime

When you have programmed your Altera-supported device, you can control which bank you map to the reserved .jai file output pins. To control which bank you map, in the schematic in the logical view, right-click the bank and click **Connect Bank** (Figure 16-4).

Figure 16-4. Configuring Banks



Acquiring Data on Your Logic Analyzer

To acquire data on your logic analyzer, you must establish a connection between your device and the external logic analyzer.

- For more information about this process and for guidelines about how to establish connections between debugging headers and logic analyzers, refer to the documentation for your logic analyzer.

Using the LAI with Incremental Compilation

The Incremental Compilation feature in the Quartus II software allows you to preserve the synthesis and fitting results of your design. This is an effective feature for reducing compilation times if you only modify a portion of a design or you wish to preserve the optimization results from a previous compilation.

The Incremental Compilation feature is well suited for use with LAI since LAI comprises a small portion of most designs. Because LAI consists of only a small portion of your design, incremental compilation helps to minimize your compilation time. Incremental compilation works best when you are only changing a small portion of your design. Incremental compilation yields an accurate representation of your design behavior when changing the .jai file through multiple compilations.

- For further details on how to use Incremental Compilation with the LAI, refer to *Enabling the Logic Analyzer Interface* in Quartus II Help.

Conclusion

As the device industry continues to make technological advancements, outdated debugging methodologies must be replaced with new technologies that maximize productivity. The LAI feature enables you to connect many internal signals within your Altera-supported device to an external logic analyzer with the use of a small number of I/O pins. This new technology in the Quartus II software enables you to use feature-rich external logic analyzers to debug your Altera-supported device design, ultimately enabling you to deliver your product in the shortest amount of time.

Document Revision History

Table 16-3 shows the revision history for this chapter.

Table 16-3. Document Revision History

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Minor editorial updates ■ Changed to new document template
August 2010	10.0.1	Corrected links
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Created links to the Quartus II Help ■ Editorial updates ■ Removed Referenced Documents section
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Removed references to APEX devices ■ Editorial updates
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Minor editorial updates ■ Removed Figures 15-4, 15-5, and 15-11 from 8.1 version
November 2008	8.1.0	Changed to 8-1/2 x 11 page size. No change to content
May 2008	8.0.0	<ul style="list-style-type: none"> ■ Updated device support list on page 15-3 ■ Added links to referenced documents throughout the chapter ■ Added "Referenced Documents" ■ Added reference to <i>Section V. In-System Debugging</i> ■ Minor editorial updates

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter explains how to use the Quartus®II In-System Memory Content Editor as part of your FPGA design and verification flow.

The In-System Memory Content Editor allows you to view and update memories and constants with the JTAG port connection.

The In-System Memory Content Editor allows access to dense and complex FPGA designs. When you program devices, you have read and write access to the memories and constants through the JTAG interface. You can then identify, test, and resolve issues with your design by testing changes to memory contents in the FPGA while your design is running.

Overview

This chapter contains the following sections:

- “Updating Memory and Constants in Your Design” on page 17–2
- “Updating Memory and Constants in Your Design” on page 17–2
- “Creating In-System Modifiable Memories and Constants” on page 17–2
- “Running the In-System Memory Content Editor” on page 17–2

When you use the In-System Memory Content Editor in conjunction with the SignalTap II Logic Analyzer, you can more easily view and debug your design in the hardware lab.

 For more information about the SignalTap II Logic Analyzer, refer to the *Design Debugging Using the SignalTap II Logic Analyzer* chapter in volume 3 of the *Quartus II Handbook*.

The ability to read data from memories and constants allows you to quickly identify the source of problems. The write capability allows you to bypass functional issues by writing expected data. For example, if a parity bit in your memory is incorrect, you can use the In-System Memory Content Editor to write the correct parity bit values into your RAM, allowing your system to continue functioning. You can also intentionally write incorrect parity bit values into your RAM to check the error handling functionality of your design.

 The Quartus II software offers a variety of on-chip debugging tools. For an overview and comparison of all tools available in the Quartus II software on-chip debugging tool suite, refer to *Section IV. System Debugging Tools* in volume 3 of the *Quartus II Handbook*.

 For a list of the types of memories and constants currently supported by the Quartus II software, refer to *Megafunctions/LPM* in Quartus II Help.

Updating Memory and Constants in Your Design

To use the In-System Updating of Memory and Constants feature, perform the following steps:

1. Identify the memories and constants that you want to access.
2. Edit the memories and constants to be run-time modifiable.
3. Perform a full compilation.
4. Program your device.
5. Launch the In-System Memory Content Editor.

Creating In-System Modifiable Memories and Constants

When you specify that a memory or constant is run-time modifiable, the Quartus II software changes the default implementation. A single-port RAM is converted to a dual-port RAM, and a constant is implemented in registers instead of look-up tables (LUTs). These changes enable run-time modification without changing the functionality of your design.

- ❓ To enable your memory or constant to be modifiable, refer to *Setting up the In-System Memory Content Editor* in Quartus II Help.

If you instantiate a memory or constant megafunction directly with ports and parameters in VHDL or Verilog HDL, add or modify the `lpm_hint` parameter as follows:

In VHDL code, add the following:

```
lpm_hint => "ENABLE_RUNTIME_MOD = YES,  
           INSTANCE_NAME = <instantiation name>";
```

In Verilog HDL code, add the following:

```
defparam <megafunction instance name>.lpm_hint =  
"ENABLE_RUNTIME_MOD = YES,  
INSTANCE_NAME = <instantiation name>";
```

Running the In-System Memory Content Editor

The In-System Memory Content Editor has three separate panes: the **Instance Manager**, the **JTAG Chain Configuration**, and the **Hex Editor**.

The **Instance Manager** pane displays all available run-time modifiable memories and constants in your FPGA device. The **JTAG Chain Configuration** pane allows you to program your FPGA and select the Altera® device in the chain to update.

Using the In-System Memory Content Editor does not require that you open a project. The In-System Memory Content Editor retrieves all instances of run-time configurable memories and constants by scanning the JTAG chain and sending a query to the specific device selected in the **JTAG Chain Configuration** pane.

If you have more than one device with in-system configurable memories or constants in a JTAG chain, you can launch multiple In-System Memory Content Editors within the Quartus II software to access the memories and constants in each of the devices. Each In-System Memory Content Editor can access the in-system memories and constants in a single device.

Instance Manager

When you scan the JTAG chain to update the **Instance Manager** pane, you can view a list of all run-time modifiable memories and constants in the design. The **Instance Manager** pane displays the Index, Instance, Status, Width, Depth, Type, and Mode of each element in the list.

- ② You can read and write to in-system memory with the **Instance Manager** pane. For more information refer to *Instance Manager Pane* in Quartus II Help.

 In addition to the buttons available in the **Instance Manager** pane, you can read and write data by selecting commands from the Processing menu, or the right-click menu in the **Instance Manager** pane or **Hex Editor** pane.

The status of each instance is also displayed beside each entry in the **Instance Manager** pane. The status indicates if the instance is **Not running**, **Offloading data**, or **Updating data**. The health monitor provides information about the status of the editor.

The Quartus II software assigns a different index number to each in-system memory and constant to distinguish between multiple instances of the same memory or constant function. View the **In-System Memory Content Editor Settings** section of the Compilation Report to match an index number with the corresponding instance ID.

Editing Data Displayed in the Hex Editor Pane

You can edit data read from your in-system memories and constants displayed in the **Hex Editor** pane by typing values directly into the editor or by importing memory files.

- ② For more information, refer to *Working with In-System Memory Content Editor Data* in Quartus II Help.

Importing and Exporting Memory Files

The In-System Memory Content Editor allows you to import and export data values for memories that have the In-System Updating feature enabled. Importing from a data file enables you to quickly load an entire memory image. Exporting to a data file enables you to save the contents of the memory for future use.

- ② For more information, refer to *Working with In-System Memory Content Editor Data* in Quartus II Help.

Scripting Support

The In-System Memory Content Editor supports reading and writing of memory contents via a Tcl script or Tcl commands entered at a command prompt. For detailed information about scripting command options, refer to the Quartus II command-line and Tcl API Help browser.

To run the Help browser, type the following command at the command prompt:

```
quartus_sh --qhelp ←
```



For more information about Tcl scripting, refer to the *Tcl Scripting* chapter in volume 2 of the *Quartus II Handbook* and *API Functions for Tcl* in Quartus II Help. For more information about command-line scripting, refer to the *Command-Line Scripting* chapter in volume 2 of the *Quartus II Handbook*.

The commonly used commands for the In-System Memory Content Editor are as follows:

- Reading from memory:

```
read_content_from_memory  
[-content_in_hex]  
-instance_index <instance index>  
-start_address <starting address>  
-word_count <word count>
```

- Writing to memory:

```
write_content_to_memory
```

- Save memory contents to file:

```
save_content_from_memory_to_file
```

- Update memory contents from File:

```
update_content_to_memory_from_file
```

⓪ For descriptions of the command options and scripting examples, refer to the Tcl API Help Browser and the *API Functions for Tcl* in Quartus II Help.

Programming the Device with the In-System Memory Content Editor

If you make changes to your design, you can program the device from within the In-System Memory Content Editor.

⓪ To program the device, refer to *Setting up the In-System Memory Content Editor* in Quartus II Help.

Example: Using the In-System Memory Content Editor with the SignalTap II Logic Analyzer

The following scenario describes how you can use the In-System Updating of Memory and Constants feature with the SignalTap II Logic Analyzer to efficiently debug your design in-system. You can use the In-System Memory Content Editor and the SignalTap II Logic Analyzer simultaneously with the JTAG interface.

Scenario: After completing your FPGA design, you find that the characteristics of your FIR filter design are not as expected.

1. To locate the source of the problem, change all your FIR filter coefficients to be in-system modifiable and instantiate the SignalTap II Logic Analyzer.
2. Using the SignalTap II Logic Analyzer to tap and trigger on internal design nodes, you find the FIR filter to be functioning outside of the expected cutoff frequency.
3. Using the **In-System Memory Content Editor**, you check the correctness of the FIR filter coefficients. Upon reading each coefficient, you discover that one of the coefficients is incorrect.
4. Because your coefficients are in-system modifiable, you update the coefficients with the correct data with the **In-System Memory Content Editor**.

In this scenario, you can quickly locate the source of the problem using both the In-System Memory Content Editor and the SignalTap II Logic Analyzer. You can also verify the functionality of your device by changing the coefficient values before modifying the design source files.

You can also modify the coefficients with the In-System Memory Content Editor to vary the characteristics of the FIR filter, for example, filter attenuation, transition bandwidth, cut-off frequency, and windowing function.

Conclusion

The In-System Updating of Memory and Constants feature provides access to a device for efficient debugging in a hardware lab. You can use the In-System Memory and Content Editor with the SignalTap II Logic Analyzer to maximize the visibility into an Altera FPGA. By maximizing visibility and access to internal logic of the device, you can identify and resolve problems with your design more easily.

Document Revision History

Table 17-1 shows the revision history of this chapter.

Table 17-1. Document Revision History

Date	Version	Changes
December 2010	10.0.2	■ Changed to new document template. No change to content
August 2010	10.0.1	■ Corrected links
July 2010	10.0.0	■ Inserted links to Quartus II Help ■ Removed Reference Documents section
November 2009	9.1.0	■ Delete references to APEX devices ■ Style changes
March 2009	9.0.0	■ No change to content

Table 17-1. Document Revision History

November 2008	8.1.0	<ul style="list-style-type: none"> ■ Changed to 8-1/2 x 11 page size. No change to content
May 2008	8.0.0	<ul style="list-style-type: none"> ■ Added reference to Section V. In-System Debugging in volume 3 of the Quartus II Handbook on page 16-1 ■ Removed references to the Mercury device, as it is now considered to be a “Mature” device ■ Added links to referenced documents throughout document ■ Minor editorial updates

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter provides detailed instructions about how to use the In-System Sources and Probes Editor and Tcl scripting in the Quartus® II software to debug your design.

Traditional debugging techniques often involve using an external pattern generator to exercise the logic and a logic analyzer to study the output waveforms during run time. The SignalTap® II Logic Analyzer and SignalProbe allow you to read or “tap” internal logic signals during run time as a way to debug your logic design. You can make the debugging cycle more efficient when you can drive any internal signal manually within your design, which allows you to perform the following actions:

- Force the occurrence of trigger conditions set up in the SignalTap II Logic Analyzer
- Create simple test vectors to exercise your design without using external test equipment
- Dynamically control run time control signals with the JTAG chain

The In-System Sources and Probes Editor in the Quartus II software extends the portfolio of verification tools, and allows you to easily control any internal signal and provides you with a completely dynamic debugging environment. Coupled with either the SignalTap II Logic Analyzer or SignalProbe, the In-System Sources and Probes Editor gives you a powerful debugging environment in which to generate stimuli and solicit responses from your logic design.

-  The Virtual JTAG Megafunction and the In-System Memory Content Editor also give you the capability to drive virtual inputs into your design. The Quartus II software offers a variety of on-chip debugging tools. For an overview and comparison of all the tools available in the Quartus II software on-chip debugging tool suite, refer to *Section IV. System Debugging Tools* in volume 3 of the *Quartus II Handbook*.

Overview

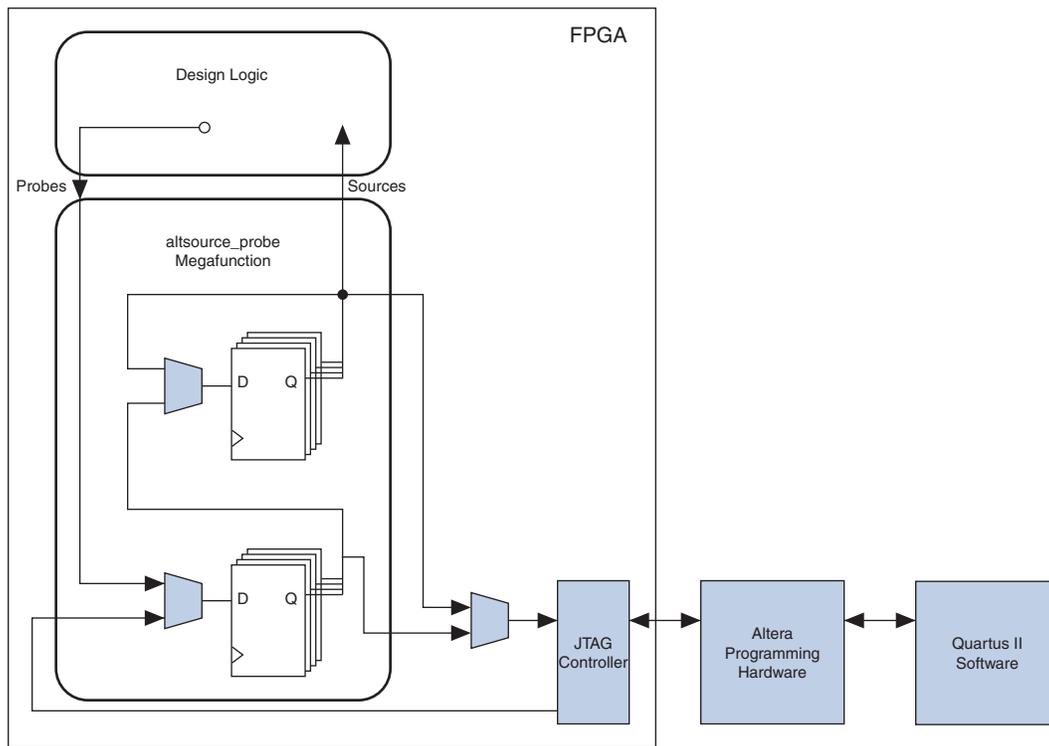
This chapter includes the following topics:

- “Design Flow Using the In-System Sources and Probes Editor” on page 18–4
- “Running the In-System Sources and Probes Editor” on page 18–7
- “Tcl interface for the In-System Sources and Probes Editor” on page 18–9
- “Design Example: Dynamic PLL Reconfiguration” on page 18–13

The In-System Sources and Probes Editor consists of the ALTSOURCE_PROBE megafunction and an interface to control the ALTSOURCE_PROBE megafunction instances during run time. Each ALTSOURCE_PROBE megafunction instance provides you with source output ports and probe input ports, where source ports drive selected signals and probe ports sample selected signals. When you compile

your design, the ALTSOURCE_PROBE megafunction sets up a register chain to either drive or sample the selected nodes in your logic design. During run time, the In-System Sources and Probes Editor uses a JTAG connection to shift data to and from the ALTSOURCE_PROBE megafunction instances. Figure 18-1 shows a block diagram of the components that make up the In-System Sources and Probes Editor.

Figure 18-1. In-System Sources and Probes Editor Block Diagram



The ALTSOURCE_PROBE megafunction hides the detailed transactions between the JTAG controller and the registers instrumented in your design to give you a basic building block for stimulating and probing your design. Additionally, the In-System Sources and Probes Editor provides single-cycle samples and single-cycle writes to selected logic nodes. You can use this feature to input simple virtual stimuli and to capture the current value on instrumented nodes. Because the In-System Sources and Probes Editor gives you access to logic nodes in your design, you can toggle the inputs of low-level components during the debugging process. If used in conjunction with the SignalTap II Logic Analyzer, you can force trigger conditions to help isolate your problem and shorten your debugging process.

The In-System Sources and Probes Editor allows you to easily implement control signals in your design as virtual stimuli. This feature can be especially helpful for prototyping your design, such as in the following operations:

- Creating virtual push buttons
- Creating a virtual front panel to interface with your design
- Emulating external sensor data
- Monitoring and changing run time constants on the fly

The In-System Sources and Probes Editor supports Tcl commands that interface with all your ALTSOURCE_PROBE megafunction instances to increase the level of automation.

Hardware and Software Requirements

The following components are required to use the In-System Sources and Probes Editor:

- Quartus II software

or

- Quartus II Web Edition (with the TalkBack feature turned on)
- Download Cable (USB-Blaster™ download cable or ByteBlaster™ cable)
- Altera® development kit or user design board with a JTAG connection to device under test

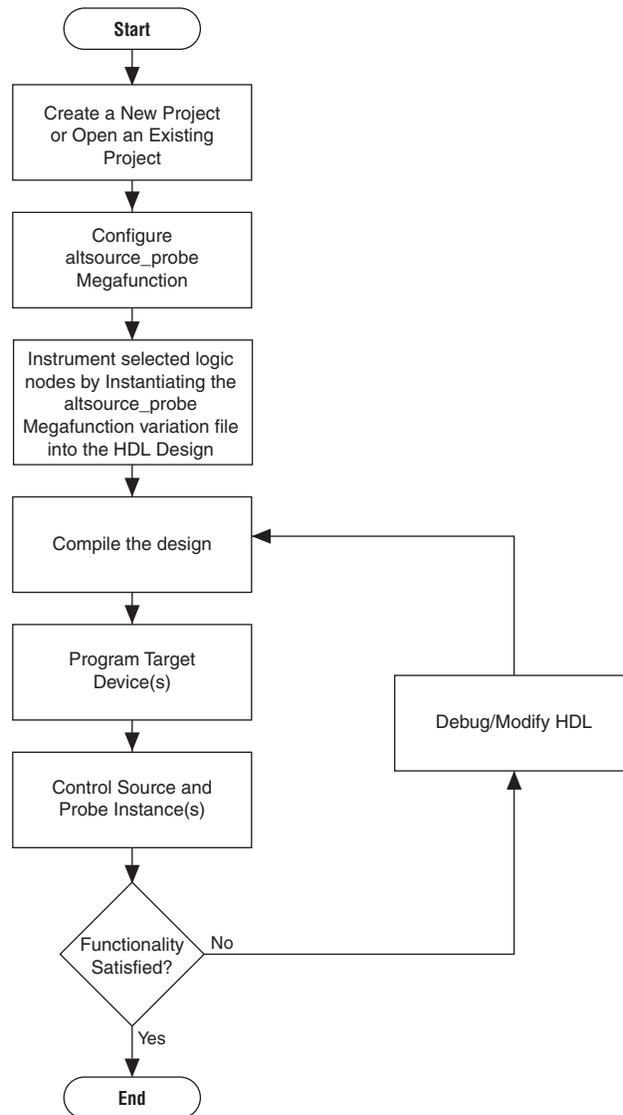
The In-System Sources and Probes Editor supports the following device families:

- Arria® GX
- Stratix® series
- HardCopy® II
- Cyclone® series
- MAX® II

Design Flow Using the In-System Sources and Probes Editor

The In-System Sources and Probes Editor supports an RTL flow. Signals that you want to view in the In-System Sources and Probes editor are connected to an instance of the ALTSOURCE_PROBE megafunction. After you compile the design, you can control each ALTSOURCE_PROBE instance via the **In-System Sources and Probes Editor** pane or via a Tcl interface. The complete design flow is shown in [Figure 18-2](#).

Figure 18-2. FPGA Design Flow Using the In-System Sources and Probes Editor



Configuring the ALTSOURCE_PROBE Megafunction

To use the In-System Sources and Probes Editor in your design, you must first instantiate the ALTSOURCE_PROBE megafunction variation file. You can configure the ALTSOURCE_PROBE megafunction with the MegaWizard™ Plug-In Manager. Each source or probe port can be up to 256 bits. You can have up to 128 instances of the ALTSOURCE_PROBE megafunction in your design.

To configure the ALTSOURCE_PROBE megafunction, performing the following steps:

1. On the Tools menu, click **MegaWizard Plug-In Manager**.
 2. Select **Create a new custom megafunction variation**.
 3. Click **Next**.
 4. On page 2a of the MegaWizard Plug-In Manager, make the following selections:
 - a. In the **Installed Plug-Ins** list, expand the **JTAG-accessible Extensions** folder and select **In-System Sources and Probes**.
-  Verify that the currently selected device family matches the device you are targeting.
- b. Select an output file type and enter the name of the ALTSOURCE_PROBE megafunction. You can choose AHDL (.tdf), VHDL (.vhd), or Verilog HDL (.v) as the output file type.
 5. Click **Next**.
 6. On page 3 of the MegaWizard Plug-In Manager, make the following selections:
 - a. Under **Do you want to specify an Instance Index?**, turn on **Yes**.
 - b. Specify the '**Instance ID**' of this instance.
 - c. Specify the width of the probe port. The width can be from 0 bit to 256 bits.
 - d. Specify the width of the source port. The width can be from 0 bit to 256 bits.
 7. On page 3 of the MegaWizard Plug-In Manager, you can click **Advanced Options** and specify other options, including the following:
 - **What is the initial value of the source port, in hexadecimal?**—Allows you to specify the initial value driven on the source port at run time.
 - **Write data to the source port synchronously to the source clock**—Allows you to synchronize your source port write transactions with the clock domain of your choice.
 - **Create an enable signal for the registered source port**—When turned on, creates a clock enable input for the synchronization registers. You can turn on this option only when the **Write data to the source port synchronously to the source clock** option is turned on.

 The In-System Sources and Probes Editor does not support simulation. You must remove the ALTSOURCE_PROBE megafunction instantiation before you create a simulation netlist.

Instantiating the ALTSOURCE_PROBE Megafunction

The MegaWizard Plug-In Manager produces the necessary variation file and the instantiation template based on your inputs to the MegaWizard. Use the template to instantiate the ALTSOURCE_PROBE megafunction variation file in your design. The port information is shown in [Table 18-1](#).

Table 18-1. ALTSOURCE_PROBE Megafunction Port Information

Port Name	Required?	Direction	Comments
probe[]	No	Input	The outputs from your design.
source_clk	No	Input	Source Data is written synchronously to this clock. This input is required if you turn on Source Clock in the Advanced Options box in the MegaWizard Plug-In Manager.
source_ena	No	Input	Clock enable signal for source_clk. This input is required if specified in the Advanced Options box in the MegaWizard Plug-In Manager.
source[]	No	Output	Used to drive inputs to user design.

You can include up to 128 instances of the ALTSOURCE_PROBE megafunction in your design, if your device has available resources. Each instance of the ALTSOURCE_PROBE megafunction uses a pair of registers per signal for the width of the widest port in the megafunction. Additionally, there is some fixed overhead logic to accommodate communication between the ALTSOURCE_PROBE instances and the JTAG controller. You can also specify an additional pair of registers per source port for synchronization.

Compiling the Design

When you compile your design with the In-System Sources and Probes megafunction instantiated, an instance of the ALTSOURCE_PROBE and SLD_HUB instances are added to your compilation hierarchy automatically. These instances provide communication between the JTAG controller and your instrumented logic.

You can modify the number of connections to your design by editing the ALTSOURCE_PROBE megafunction. To open the design instance you want to modify in the MegaWizard Plug-In Manager, double-click the instance in the Project Navigator. You can then modify the connections in the HDL source file. You must recompile your design after you make changes.

You can use the Quartus II incremental compilation feature to reduce compilation time. Incremental compilation allows you to organize your design into logical partitions. During recompilation of a design, incremental compilation preserves the compilation results and performance of unchanged partitions and reduces design iteration time by compiling only modified design partitions.



For more information about the Quartus II incremental compilation feature, refer to the *Quartus II Incremental Compilation for Hierarchical and Team-Based Design* chapter in volume 1 of the *Quartus II Handbook*.

Running the In-System Sources and Probes Editor

The In-System Sources and Probes Editor gives you control over all ALTSOURCE_PROBE megafunction instances within your design. The editor allows you to view all available run time controllable instances of the ALTSOURCE_PROBE megafunction in your design, provides a push-button interface to drive all your source nodes, and provides a logging feature to store your probe and source data.

To run the In-System Sources and Probes Editor, on the **Tools** menu, click **In-System Sources and Probes Editor**.

The In-System Sources and Probes Editor contains three panes:

- **JTAG Chain Configuration**—Allows you to specify programming hardware, device, and file settings that the In-System Sources and Probes Editor uses to program and acquire data from a device.
- **Instance Manager**—Displays information about the instances generated when you compile a design, and allows you to control data that the In-System Sources and Probes Editor acquires.
- **In-System Sources and Probes Editor**—Logs all data read from the selected instance and allows you to modify source data that is written to your device.

When you use the In-System Sources and Probes Editor, you do not need to open a Quartus II software project. The In-System Sources and Probes Editor retrieves all instances of the ALTSOURCE_PROBE megafunction by scanning the JTAG chain and sending a query to the device selected in the **JTAG Chain Configuration** pane. You can also use a previously saved configuration to run the In-System Sources and Probes Editor.

Each **In-System Sources and Probes Editor** pane can access the ALTSOURCE_PROBE megafunction instances in a single device. If you have more than one device containing megafunction instances in a JTAG chain, you can launch multiple **In-System Sources and Probes Editor** panes to access the megafunction instances in each device.

Programming Your Device With JTAG Chain Configuration

After you compile your project, you must configure your FPGA before you use the In-System Sources and Probes Editor. To configure a device to use with the In-System Sources and Probes Editor, perform the following steps:

1. Open the In-System Sources and Probes Editor.
2. In the **JTAG Chain Configuration** pane, point to **Hardware**, and then select the hardware communications device. You may be prompted to configure your hardware; in this case, click **Setup**.
3. From the **Device** list, select the FPGA device to which you want to download the design (the device may be automatically detected). You may need to click **Scan Chain** to detect your target device.
4. In the **JTAG Chain Configuration** pane, click to browse for the SRAM Object File (.sof) that includes the In-System Sources and Probes instance or instances. (The .sof may be automatically detected).
5. Click **Program Device** to program the target device.

Instance Manager

The **Instance Manager** pane provides a list of all ALTSOURCE_PROBE instances in the design and allows you to configure how data is acquired from or written to those instances.

The following buttons and sub-panes are provided in the **Instance Manager** pane:

- **Read Probe Data**—Samples the probe data in the selected instance and displays the probe data in the **In-System Sources and Probes Editor** pane.
- **Continuously Read Probe Data**—Continuously samples the probe data of the selected instance and displays the probe data in the **In-System Sources and Probes Editor** pane; you can modify the sample rate via the **Probe read interval** setting.
- **Stop Continuously Reading Probe Data**—Cancels continuous sampling of the probe of the selected instance.
- **Write Source Data**—Writes data to all source nodes of the selected instance.
- **Probe Read Interval**—Displays the sample interval of all the In-System Sources and Probe instances in your design; you can modify the sample interval by clicking **Manual**.
- **Event Log**—Controls the event log in the **In-System Sources and Probes Editor** pane.
- **Write Source Data**—Allows you to manually or continuously write data to the system.

The status of each instance is also displayed beside each entry in the **Instance Manager** pane. The status indicates if the instance is **Not running Offloading data**, **Updating data**, or if an **Unexpected JTAG communication error** occurs. This status indicator provides information about the sources and probes instances in your design.

In-System Sources and Probes Editor Pane

The **In-System Sources and Probes Editor** pane allows you to view data from all sources and probes in your design. The data is organized according to the index number of the instance. The editor provides an easy way to manage your signals, and allows you to rename signals or group them into buses. All data collected from in-system source and probe nodes is recorded in the event log and you can view the data as a timing diagram.

Reading Probe Data

You can read data by selecting the ALTSOURCE_PROBE instance in the **Instance Manager** pane and clicking **Read Probe Data**. This action produces a single sample of the probe data and updates the data column of the selected index in the **In-System Sources and Probes Editor** pane. You can save the data to an event log by turning on the **Save data to event log** option in the **Instance Manager** pane.

If you want to sample data from your probe instance continuously, in the **Instance Manager** pane, click the instance you want to read, and then click **Continuously read probe data**. While reading, the status of the active instance shows **Unloading**. You can read continuously from multiple instances.

You can access read data with the shortcut menus in the **Instance Manager** pane.

To adjust the probe read interval, in the **Instance Manager** pane, turn on the **Manual** option in the **Probe read interval** sub-pane, and specify the sample rate in the text field next to the **Manual** option. The maximum sample rate depends on your computer setup. The actual sample rate is shown in the **Current interval** box. You can adjust the event log window buffer size in the **Maximum Size** box.

Writing Data

To modify the source data you want to write into the ALTSOURCE_PROBE instance, click the name field of the signal you want to change. For buses of signals, you can double-click the data field and type the value you want to drive out to the ALTSOURCE_PROBE instance. The In-System Sources and Probes Editor stores the modified source data values in a temporary buffer. Modified values that are not written out to the ALTSOURCE_PROBE instances appear in red. To update the ALTSOURCE_PROBE instance, highlight the instance in the **Instance Manager** pane and click **Write source data**. The **Write source data** function is also available via the shortcut menus in the **Instance Manager** pane.

The In-System Sources and Probes Editor provides the option to continuously update each ALTSOURCE_PROBE instance. Continuous updating allows any modifications you make to the source data buffer to also write immediately to the ALTSOURCE_PROBE instances. To continuously update the ALTSOURCE_PROBE instances, change the **Write source data** field from **Manually** to **Continuously**.

Organizing Data

The **In-System Sources and Probes Editor** pane allows you to group signals into buses, and also allows you to modify the display options of the data buffer.

To create a group of signals, select the node names you want to group, right-click and select **Group**. You can modify the display format in the Bus Display Format and the Bus Bit order shortcut menus.

The **In-System Sources and Probes Editor** pane allows you to rename any signal. To rename a signal, double-click the name of the signal and type the new name.

The event log contains a record of the most recent samples. The buffer size is adjustable up to 128k samples. The time stamp for each sample is logged and is displayed above the event log of the active instance as you move your pointer over the data samples.

You can save the changes that you make and the recorded data to a Sources and Probes File (**.spf**). To save changes, on the File menu, click **Save**. The file contains all the modifications you made to the signal groups, as well as the current data event log.

Tcl interface for the In-System Sources and Probes Editor

To support automation, the In-System Sources and Probes Editor supports the procedures described in this chapter in the form of Tcl commands. The Tcl package for the In-System Sources and Probes Editor is included by default when you run `quartus_stp`.

The Tcl interface for the In-System Sources and Probes Editor provides a powerful platform to help you debug your design. The Tcl interface is especially helpful for debugging designs that require toggling multiple sets of control inputs. You can combine multiple commands with a Tcl script to define a custom command set.

 For more information about Tcl scripting, refer to the *Tcl Scripting* chapter in volume 2 of the *Quartus II Handbook*. For more information about settings and constraints in the Quartus II software, refer to the *Quartus II Settings File Manual*. For more information about command-line scripting, refer to the *Command-Line Scripting* chapter in volume 2 of the *Quartus II Handbook*.

Table 18-2 shows the Tcl commands you can use instead of the In-System Sources and Probes Editor.

Table 18-2. In-System Sources and Probes Tcl Commands

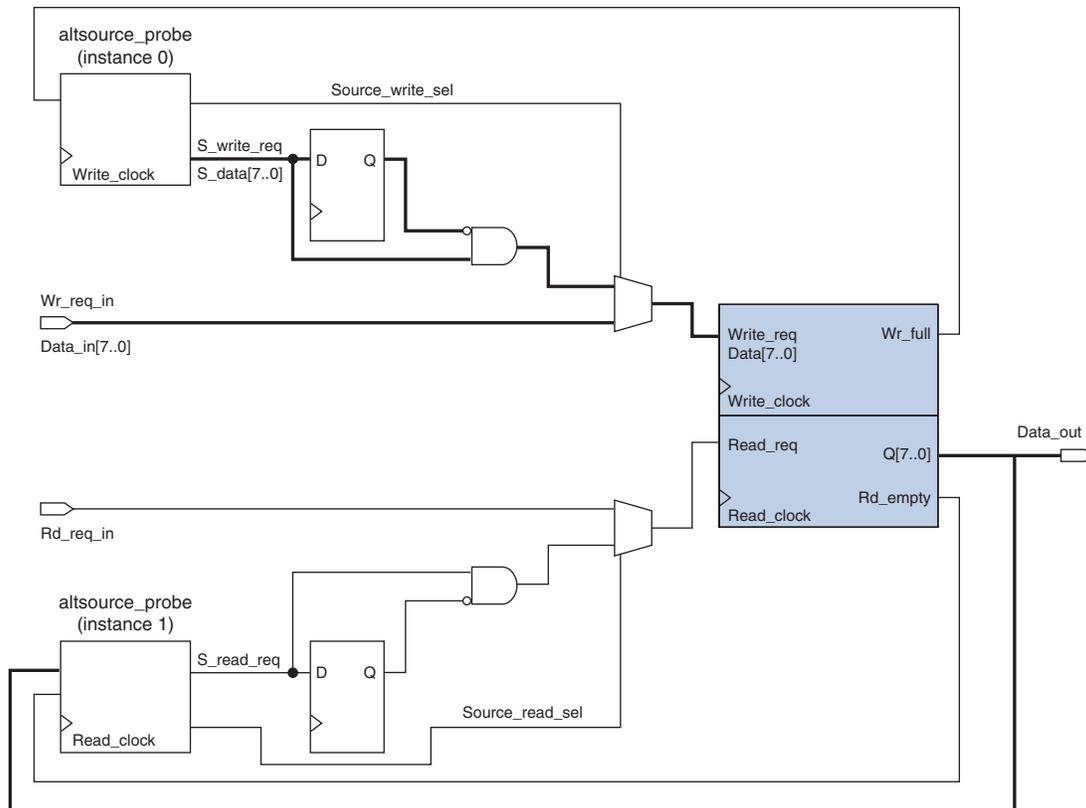
Command	Argument	Description
start_insystem_source_probe	-device_name <device name> -hardware_name <hardware name>	Opens a handle to a device with the specified hardware. Call this command before starting any transactions.
get_insystem_source_probe_instance_info	-device_name <device name> -hardware_name <hardware name>	Returns a list of all ALTSOURCE_PROBE instances in your design. Each record returned is in the following format: {<instance index>, <source width>, <probe width>, <instance name>}
read_probe_data	-instance_index <instance_index> -value_in_hex (optional)	Retrieves the current value of the probe. A string is returned that specifies the status of each probe, with the MSB as the left-most bit.
read_source_data	-instance_index <instance_index> -value_in_hex (optional)	Retrieves the current value of the sources. A string is returned that specifies the status of each source, with the MSB as the left-most bit.
write_source_data	-instance_index <instance_index> -value <value> -value_in_hex (optional)	Sets the value of the sources. A binary string is sent to the source ports, with the MSB as the left-most bit.
end_interactive_probe	None	Releases the JTAG chain. Issue this command when all transactions are finished.

Example 18-1 shows an excerpt from a Tcl script with procedures that control the ALTSOURCE_PROBE instances of the design as shown in Figure 18-3. The example design contains a DCFIFO with ALTSOURCE_PROBE instances to read from and write to the DCFIFO. A set of control muxes are added to the design to control the flow of data to the DCFIFO between the input pins and the ALTSOURCE_PROBE

instances. A pulse generator is added to the read request and write request control lines to guarantee a single sample read or write. The ALTSOURCE_PROBE instances, when used with the script in [Example 18-1](#), provide visibility into the contents of the FIFO by performing single sample write and read operations and reporting the state of the full and empty status flags.

Use the Tcl script in debugging situations to either empty or preload the FIFO in your design. For example, you can use this feature to preload the FIFO to match a trigger condition you have set up within the SignalTap II Logic Analyzer.

Figure 18-3. A DCFIFO Example Design Controlled by the Tcl Script in [Example 18-1](#)



Example 18-1. Tcl Script Procedures for Reading and Writing to the DCFIFO in Figure 18-3 (Part 1 of 2)

```
## Setup USB hardware - assumes only USB Blaster is installed and
## an FPGA is the only device in the JTAG chain

set usb [lindex [get_hardware_names] 0]
set device_name [lindex [get_device_names -hardware_name $usb] 0]
## write procedure : argument value is integer

proc write {value} {

    global device_name usb
    variable full

    start_insystem_source_probe -device_name $device_name -hardware_name $usb

    #read full flag
    set full [read_probe_data -instance_index 0]

    if {$full == 1} {end_insystem_source_probe
return "Write Buffer Full"
}
}
```

Example 18-1. Tcl Script Procedures for Reading and Writing to the DCFIFO in Figure 18-3 (Part 2 of 2)

```

##toggle select line, drive value onto port, toggle enable
##bits 7:0 of instance 0 is S_data[7:0]; bit 8 = S_write_req;
##bit 9 = Source_write_sel

##int2bits is custom procedure that returns a bitstring from an integer
## argument

write_source_data -instance_index 0 -value /[int2bits [expr 0x200 | $value]]
write_source_data -instance_index 0 -value [int2bits [expr 0x300 | $value]]

##clear transaction

write_source_data -instance_index 0 -value 0

end_insystem_source_probe
}

proc read {} {

    global device_name usb
    variable empty
    start_insystem_source_probe -device_name $device_name -hardware_name $usb

    ##read empty flag : probe port[7:0] reads FIFO output; bit 8 reads empty_flag
    set empty [read_probe_data -instance_index 1]

    if {[regexp {1.....} $empty]} { end_insystem_source_probe
    return "FIFO empty" }

    ## toggle select line for read transaction
    ## Source_read_sel = bit 0; s_read_reg = bit 1

    ## pulse read enable on DC FIFO
    write_source_data -instance_index 1 -value 0x1 -value_in_hex
    write_source_data -instance_index 1 -value 0x3 -value_in_hex

    set x [read_probe_data -instance_index 1 ]

    end_insystem_source_probe

    return $x
}

```

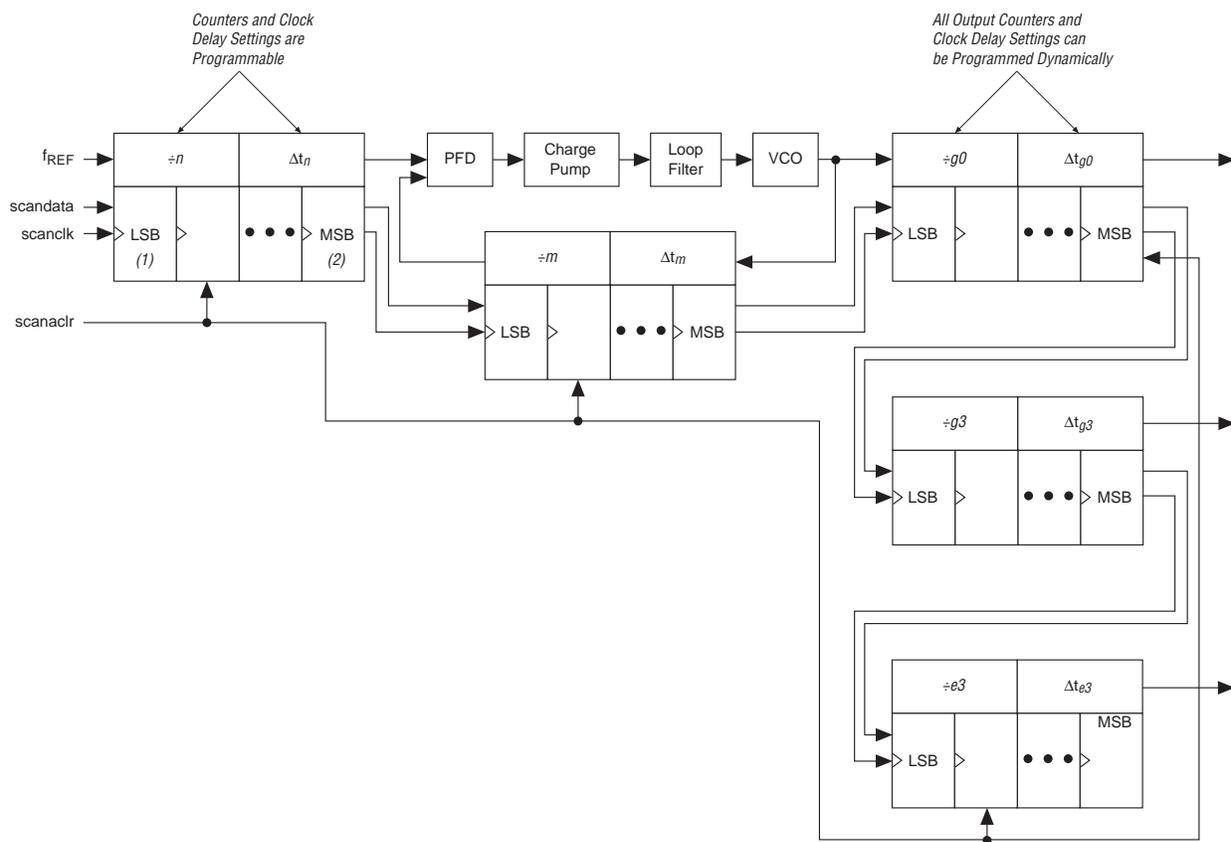
Design Example: Dynamic PLL Reconfiguration

The In-System Sources and Probes Editor can help you create a virtual front panel during the prototyping phase of your design. You can create relatively simple, high functioning designs of in a short amount of time. The following PLL reconfiguration example demonstrates how to use the In-System Sources and Probes Editor to provide a GUI to dynamically reconfigure a Stratix PLL.

Stratix PLLs allow you to dynamically update PLL coefficients during run time. Each enhanced PLL within the Stratix device contains a register chain that allows you to modify the pre-scale counters (m and n values), output divide counters, and delay counters. In addition, the ALTPLL_RECONFIG megafunction provides an easy interface to access the register chain counters. The ALTPLL_RECONFIG megafunction provides a cache that contains all modifiable PLL parameters. After you update all the PLL parameters in the cache, the ALTPLL_RECONFIG megafunction drives the PLL register chain to update the PLL with the updated parameters. Figure 18-4 shows a Stratix-enhanced PLL with reconfigurable coefficients.

 Stratix II and Stratix III devices also allow you to dynamically reconfigure PLL parameters. For more information about these families, refer to the appropriate data sheet. For more information about dynamic PLL reconfiguration, refer to [AN 282: Implementing PLL Reconfiguration in Stratix & Stratix GX Devices](#) or [AN 367: Implementing PLL Reconfiguration in Stratix II Devices](#).

Figure 18-4. Stratix-Enhanced PLL with Reconfigurable Coefficients



The following design example uses an ALTSOURCE_PROBE instance to update the PLL parameters in the ALTPLL_RECONFIG megafunction cache. The ALTPLL_RECONFIG megafunction connects to an enhanced PLL in a Stratix FPGA to drive the register chain containing the PLL reconfigurable coefficients. This design example uses a Tcl/Tk script to generate a GUI where you can enter in new m and n values for the enhanced PLL. The Tcl script extracts the m and n values from the GUI, shifts the values out to the ALTSOURCE_PROBE instances to update the values in the

ALTPLL_RECONFIG megafunction cache, and asserts the reconfiguration signal on the ALTPLL_RECONFIG megafunction. The reconfiguration signal on the ALTPLL_RECONFIG megafunction starts the register chain transaction to update all PLL reconfigurable coefficients. A block diagram of a design example is shown in Figure 18-5. The Tk GUI is shown in Figure 18-6.

Figure 18-5. Block Diagram of Dynamic PLL Reconfiguration Design Example

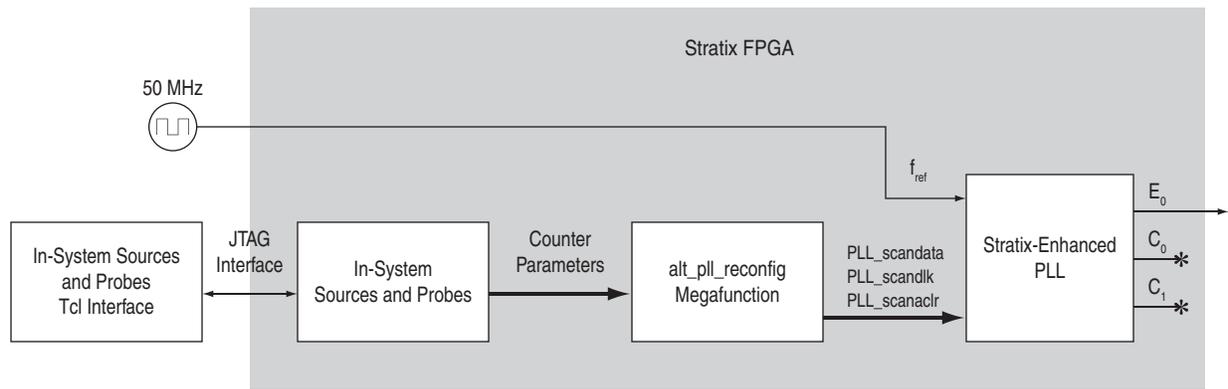
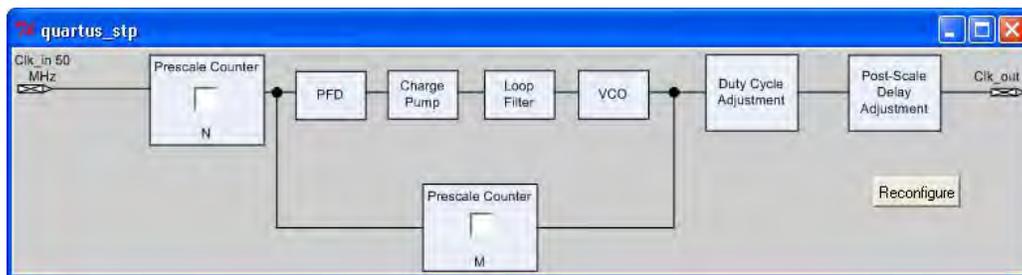


Figure 18-6. Interactive PLL Reconfiguration GUI Created with Tk and In-System Sources and Probes Tcl Package



This design example was created using a Nios® II Development Kit, Stratix Edition. The file **sourceprobe_DE_dynamic_pll.zip** contains all the necessary files for running this design example, including the following:

- **Readme.txt**—A text file that describes the files contained in the design example and provides instructions about running the Tk GUI shown in Figure 18-6.
- **Interactive_Reconfig.qar**—The archived Quartus II project for this design example.

Download the **sourceprobe_DE_dynamic_pll.zip** file from the [Literature: Quartus II Handbook](#) page of the Altera website.

Conclusion

The In-System Sources and Probes Editor provides stimuli and receives responses from the target design during run time. With the simple and intuitive interface, you can add virtual inputs to your design during run time without using external equipment. When used in conjunction with the SignalTap II Logic Analyzer, you can use the In-System Sources and Probes Editor to obtain greater control of the signals in your design, and thus help shorten the verification cycle.

Document Revision History

Table 18-3 shows the revision history for this chapter.

Table 18-3. Document Revision History

Date	Version	Changes
December 2010	10.1.0	Minor corrections. Changed to new document template.
July 2010	10.0.0	Minor corrections.
November 2009	9.1.0	<ul style="list-style-type: none"> ■ Removed references to obsolete devices. ■ Style changes.
March 2009	9.0.0	No change to content.
November 2008	8.1.0	Changed to 8-1/2 x 11 page size. No change to content.
May 2008	8.0.0	<ul style="list-style-type: none"> ■ Documented that this feature does not support simulation on page 17-5 ■ Updated Figure 17-8 for Interactive PLL reconfiguration manager ■ Added hyperlinks to referenced documents throughout the chapter ■ Minor editorial updates

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

The Quartus® II software easily interfaces with EDA formal design verification tools such as the Cadence Encounter Conformal and Synopsys Synplify software. In addition, the Quartus II software has built-in support for verifying the logical equivalence between the synthesized netlist from Synopsys Synplify and the post-fit Verilog Quartus Mapped (**.vqm**) files using Cadence Encounter Conformal software.

This section discusses formal verification, how to set-up the Quartus II software to generate the **.vqm** file and Cadence Encounter Conformal script, and how to compare designs using Cadence Encounter Conformal software.

This section includes the following chapter:

- [Chapter 19, Cadence Encounter Conformal Support](#)

The Quartus® II software provides formal verification support for Altera® designs through interfaces with a formal verification EDA tool, the Cadence Encounter Conformal Logic Equivalence Check (LEC) software.

The two types of formal verification are equivalence checking and model checking. This chapter discusses equivalence checking with the Conformal LEC software.

Use the Conformal LEC software to verify the functional equivalence of a post-synthesis Verilog Quartus Mapping (.vqm) netlist file from Synopsys Synplify Pro software, a post-fit Verilog Output File (.vo) from the Quartus II software, or both. You can also use the Conformal LEC software to verify the functional equivalence of the register transfer level (RTL) source code and post-fit .vo file with the Quartus II software when using Quartus II integrated synthesis. These formal verification flows support designs for the Arria® GX, Cyclone®, Cyclone II, Cyclone III, HardCopy® II, Stratix®, Stratix II, Stratix II GX, Stratix III, and Stratix IV device families.

This chapter contains the following sections:

- “Formal Verification Design Flow” on page 19–2
- “RTL Coding Guidelines for Quartus II Integrated Synthesis” on page 19–4
- “Black Boxes in the Conformal LEC Flow” on page 19–8
- “Generating the Post-Fit Netlist Output File and the Conformal LEC Setup Files” on page 19–10
- “Understanding the Formal Verification Scripts for Conformal LEC” on page 19–12
- “Comparing Designs Using Conformal LEC” on page 19–15
- “Known Issues and Limitations” on page 19–16
- “Black Box Models” on page 19–18
- “Conformal Dofile/Script Example” on page 19–19

Equivalence checking uses mathematical techniques to compare the logical equivalence of two versions of the same design rather than using test vectors to perform simulation. The two compared versions can be post-map design and post-fit design, or RTL design and post-fit design. Equivalence checking greatly shortens the verification cycle of the design.

Formal Verification Versus Simulation

Formal verification cannot be considered as a replacement for vector-based simulation. Formal verification only complements the existing vector-based simulation techniques to speed up the verification cycle. Vector-based simulation techniques of gate-level designs can take a considerable amount of time.

You can use Vector-based simulation techniques to perform the following functions:

- Verify design functionality
- Verify timing specifications
- Debug designs

Formal Verification: What You Need to Know

If you use formal verification techniques to verify logic equivalence of your design, you can save time by foregoing a comprehensive vector-based simulation of the gate-level design. However, there might be an impact on area and performance during recompilation of your design with the Quartus II software if you choose to use formal verification flow for Conformal LEC software. The area and performance of your design might be affected by the following factors:

- Hierarchy preservation
- ROM implementation by logic elements (LEs)
- Disabled retiming is disabled

Refer to [“Known Issues and Limitations” on page 19-16](#) before you consider using the formal verification flow in your design methodology.

Formal Verification Design Flow

Altera supports formal verification using the Conformal LEC software for the following two synthesis tools:

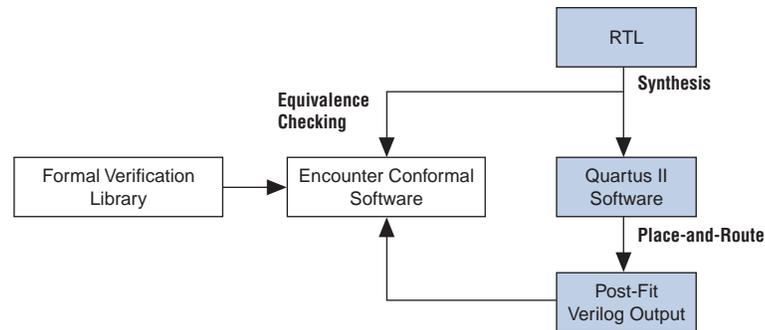
- [“Quartus II Integrated Synthesis” on page 19-3](#)
- [“Synplify Pro” on page 19-3](#)

The following sections describe the supported design flows for these synthesis tools.

Quartus II Integrated Synthesis

The design flow for formal verification using the Quartus II integrated synthesis is shown in Figure 19-1. This flow performs equivalency checking for the RTL source code and the post-fit netlist generated by the Quartus II software. The RTL source code can be in Verilog HDL or VHDL format. The post-fit netlist generated by the Quartus II software is always in Verilog HDL format.

Figure 19-1. Formal Verification Using Quartus II Integrated Synthesis and the Conformal LEC Software



EDA Tool Support for Quartus II Integrated Synthesis

The formal verification flow using the Quartus II software and Conformal LEC software supports the following software versions and operating systems:

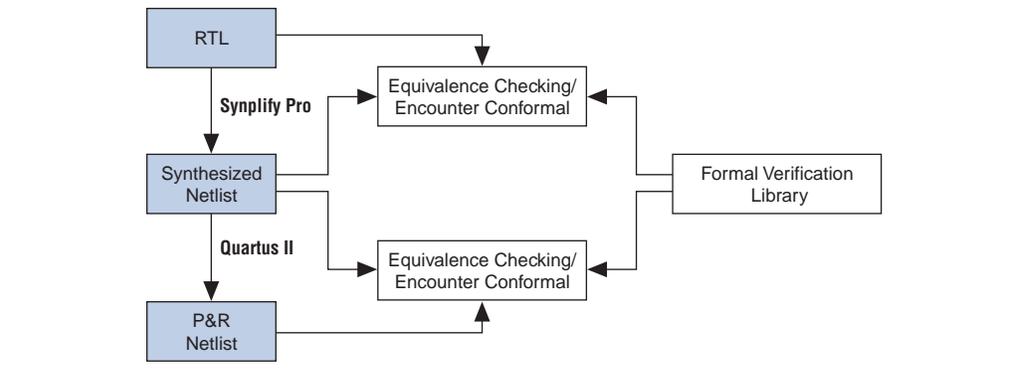
- Quartus II software beginning with version 4.2
- Cadence Conformal LEC software beginning with version 4.3.5A
- Solaris and Linux operating systems

Synplify Pro

The design flow for formal verification using Synplify Pro Synthesis performs equivalency checking for the post-synthesis netlist from Synplify Pro and the post-fit netlist generated by Quartus II software, as shown in Figure 19-2 on page 19-4.

- For additional information about performing equivalency checking between RTL and post-synthesis netlists generated from Synplify Pro software, refer to the Synplify Pro documentation.

Figure 19-2. Formal Verification Flow Using Synplify Pro and the Conformal LEC Software



RTL Coding Guidelines for Quartus II Integrated Synthesis

The Conformal LEC software can compare the RTL code against the post-fit netlist generated by the Quartus II software. The Conformal LEC software and the Quartus II integrated synthesis parse and compile the RTL description in slightly different ways. The Quartus II software supports some RTL features that the Conformal LEC software does not support and vice versa. The style of the RTL code is of particular concern because neither tool supports some constructs, leading to potential formal verification mismatches; for example, state machine extraction, wherein different encoding mechanisms can result in different structures. Therefore, for successful verification, both tools must interpret the RTL code in the same manner.

The following section provides information about recognizing and preventing problems that can arise in the formal verification flow.

- For more details about RTL coding styles for Quartus II integrated synthesis, refer to the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook*.

Some of the coding guidelines apply to both the Quartus II integrated synthesis and Synplify Pro flow, as indicated in each of the guidelines in the following sections.

Synthesis Directives and Attributes

Synthesis directives, also known as pragmas, play an important role in successful verification of RTL against the post-fit `.vo` netlist file from the Quartus II software.

Pragmas and trigger keywords that are supported in Quartus II integrated synthesis and the Conformal LEC software are also supported in the formal verification flow. The Quartus II integrated synthesis and Conformal LEC both support the trigger keywords “`synthesis`” and “`synopsys`.” When the Quartus II software does not recognize a keyword (such as “`verplex`”), the keyword is disabled in the formal verification scripts produced for use with the Conformal LEC software. Therefore, it is important to use caution with unsupported pragmas because they can lead to verification mismatches.

For example, you can use Quartus II integrated synthesis to synthesize RTL code with the synthesis directive `read_comments_as_HDL` (Example 19-1 and Example 19-2).

Example 19-1. Verilog HDL Example of Read Comments as HDL

```
// synthesis read_comments_as_HDL on
// my_rom lpm_rom (.address (address),
// .data (data));
// synthesis read_comments_as_HDL off
```

Example 19-2. VHDL Example of Read Comments as HDL

```
-- synthesis read_comments_as_HDL on
-- my_rom : entity lpm_rom
-- port map (
-- address => address,
-- data => data, );
-- synthesis read_comments_as_HDL off
```

 The Conformal LEC software does not support the synthesis directive `read_comments_as_HDL`, and the directive has no affect on the Conformal LEC software.

Table 19-1 lists supported pragmas and trigger keywords for formal verification.

Table 19-1. Supported Pragmas and Trigger Keywords for Formal Verification

Pragmas (1)	Trigger Keywords
full_case parallel_case pragma synthesis_off synthesis_on translate_off translate_on	synthesis synopsys

Note to Table 19-1:

(1) Do not use Verilog 2001-style pragma declarations. The Quartus II software and the Conformal LEC software support this style of pragma in different manners.

Stuck-at Registers

Quartus II integrated synthesis eliminates registers that have their output stuck at a constant value. Quartus II integrated synthesis gives a warning message and adds an entry to the corresponding report panel in the formal verification folder of the **Analysis & Synthesis** section of the Compilation Report. If the Conformal LEC software does not find the same optimizations, it can lead to unmapped points in the golden netlist. Example 19-3 on page 19-6 illustrates the issue.

Example 19-3. Verilog HDL Example Showing Stuck at Registers

```

module stuck_at_example {clk, a,b,c,d,out};
input a,b,c,d,clk;
output out;
reg e,f,g;
  always @(posedge clk) begin
    e <= a and g;// e is stuck at 0
    g <= c and e;// g is stuck at 0
    f <= e | b;
  end
assign out = f and d;
endmodule

```

In this module description, registers `e` and `g` are tied to logic 0. In this example, the Quartus II software generates the following warning message:

```

Warning: Reduced register "g" with stuck data_in port to stuck value GND
Warning: Reduced register "e" with stuck data_in port to stuck value GND

```

Quartus II integrated synthesis then adds a command to the formal verification scripts telling Conformal LEC that a register is stuck at a constant value, as shown in [Example 19-4](#).

Example 19-4. Conformal LEC Script Showing Commands for Instance Equivalence

```

// report floating signals
// Instance-constraints commands for constant-value registers removed
// during compilation
// add instance constraints 0 e -golden
// add instance constraints 0 g -golden

```

The command is commented in the formal verification script, forcing the Conformal LEC software to treat the register as stuck at a constant value and potentially hiding a compilation error. You must verify that input to the `e` and `g` registers is constant in the design and uncomment the command to obtain accurate results.



Altera recommends recoding your design to eliminate “stuck-at” registers.

The stuck-at register information in this section also applies to the Synplify Pro flow.

ROM, LPM_DIVIDE, and Shift Register Inference

For the purpose of formal verification, the Quartus II integrated synthesis implements both ROM and shift registers in the form of LEs instead of using dedicated on-chip memory resources. Using LEs can be less area-efficient than inferring a megafunction that can be implemented in a RAM block. However, the Quartus II software generates a warning message indicating that the megafunction was not inferred. Quartus II integrated synthesis also reports a suggested ROM or shift register instantiation that enables you to either use the MegaWizard™ Plug-In Manager to create the appropriate megafunction explicitly, or to isolate the corresponding logic in a separate entity that you can set as a black box. By setting black box properties on a particular module or

entity, you are telling the formal verification tool not to look inside the module or entity for formal verification. If the black box properties are set on the corresponding megafunction before synthesis, you can verify the megafunction with the Conformal LEC software. For details about setting black box properties on a particular module, refer to [Table 19-2 on page 19-9](#).

If the design contains division functionality, the Quartus II software infers an LPM_DIVIDE megafunction, which is treated as a black box for the purpose of formal verification.

RAM Inference

When the Quartus II software infers the ALTSYNCRAM megafunction from the RTL code, the Quartus II software generates the following warning message:

```
Created node "<mem_block_name>" as a RAM by generating altsyncram megafunction to implement register logic with M512 or M4K memory block or M-RAM. Expect to get an error or a mismatch for this block in the formal verification tool.
```

This warning is generated because the memory block (altsyncram) is a new instance in the post-fit netlist that is handled as a black box by the formal verification tool. However, no such instance exists in the original RTL design, resulting in mismatch or error reporting in the formal verification tool.

Latch Inference

A latch is implemented in the Quartus II integrated synthesis using a combinational feedback loop. The Conformal LEC software infers a latch primitive in the Conformal LEC library (DLAT) to implement a latch. This results in having a DLAT on the golden side and a combinational loop with a cut point on the revised side, leading to verification mismatches. The Quartus II software issues a warning message whenever a latch is inferred, and the Quartus II software adds an entry to the report panel in the Formal Verification folder of the Analysis & Synthesis report. Altera recommends that you avoid latches in your design; however, if latches are necessary, Altera recommends using the corresponding LPM_LATCH megafunction.

 For more information about the problems related to latches, refer to the [Recommended HDL Coding Styles](#) chapter in volume 1 of the *Quartus II Handbook*.

Combinational Loops

If the design consists of an intended combinational loop, you must define an appropriate cut point for both the RTL and the post-fit .vo netlist file. A warning that a combinational loop exists in the design is found in the **Formal Verification** subfolder of the Quartus II software Analysis & Synthesis report.

For more information about issues with combinational loops, refer to [“Known Issues and Limitations” on page 19-16](#).

Finite State Machine Coding Styles

When a state machine is inferred by the Conformal LEC software, it uses sequential encoding as the default encoding when no user encoding is present. The Quartus II software selects the encoding most suited for the inferred state machine if the **State Machine Processing** setting is set to the default value **Auto**. To do this, perform the following steps:

1. On the Assignments menu, click **Settings**. The **Settings** dialog box appears.
2. In the **Category** list, select **Analysis & Synthesis Settings**. The **Analysis & Synthesis Settings** page appears.
3. Click **More Settings**. The **More Analysis & Synthesis Settings** dialog box appears.
4. Under **Option**, in the **Name** list, select **State Machine Processing**. In the **Setting** list, select **Auto**.
5. Click **OK**.
6. Click **OK**.

Use the coding style described in the *Recommended HDL Coding Styles* chapter in volume 1 of the *Quartus II Handbook* when writing finite state machines (FSMs). This allows the Quartus II integrated synthesis and the Conformal LEC software to infer a similar state machine for the same RTL code.

Black Boxes in the Conformal LEC Flow

The Quartus II software usually generates a flattened netlist; however, you must treat some modules in the design as black boxes. The following is a list of some of these modules:

- LPMs and megafunctions without formal verification models
- Encrypted IP functions
- Entities not implemented in Verilog HDL or VHDL

To perform equivalence checking of a design between its version, which consists of the modules listed above and its implemented version, the modules must be treated as black boxes by the Conformal LEC software. To facilitate the formal verification flow, the Quartus II software reconstructs the hierarchy on the black boxes with a port interface that is identical to the module on the golden side of the design.

If your golden netlist (**.vqm** netlist file from Synplify Pro or RTL) includes any design entity not having a corresponding formal verification model, that entity is handled as a black box with its boundary interface preserved. There are three types of black boxes with their required user actions described in [Table 19-2 on page 19-9](#).

Verilog Output Files (**.vo**) written by the Quartus II software contain the black box hierarchy when you make an EDA Formal Verification Hierarchy assignment with the value **BLACKBOX**.

If this assignment is not made for a module, the Quartus II software implements that module with logic cells. When this happens, the .vo netlist file no longer contains the black box hierarchy and does not preserve the port interface, resulting in a mismatch within the Conformal LEC software.

Table 19-2. Black Boxes and Required User Action

Type of Black Box	Required User Action
Altera library of parameterized modules (LPMs) and megafunctions.	No action required. The Quartus II software automatically creates a black box list of components and preserves the hierarchy.
Any parameterized entity other than those listed in the <i>Guidelines for Creating a Design for Use with the Encounter Conformal and Quartus II Software</i> topic in Quartus II Help.	User must designate the wrapper that instantiates the parameterized entity as a black box.
Non-parameterized entities that the user wants to designate as a black box.	User can designate the entity itself as a black box.

You can also use Tcl commands or the Quartus II GUI to set the black box property on the entities, which the formal verification tool does not compare.

Tcl Command

Use the Tcl command shown in [Example 19-5](#) to preserve the boundary interface of a black box entity: dram.

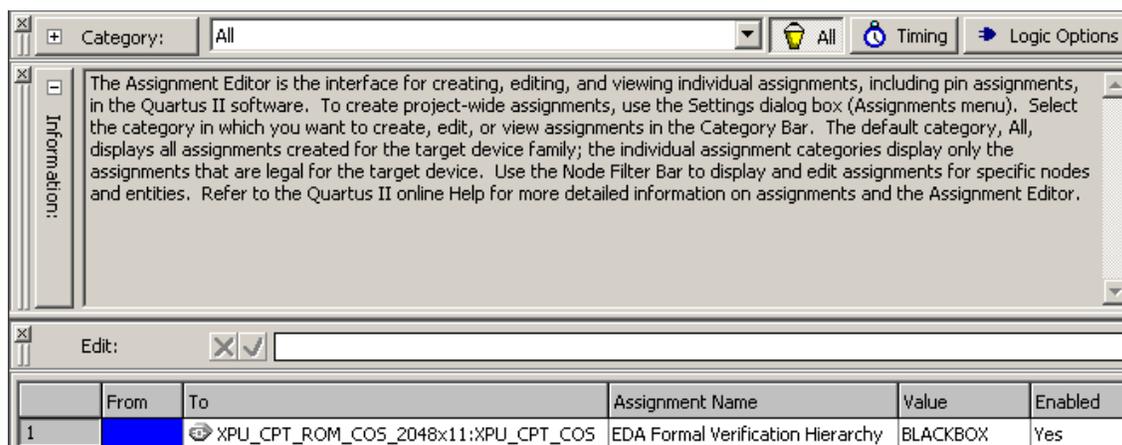
Example 19-5. Tcl Command to Create a Black Box

```
set_instance_assignment -name EDA_FV_HIERARCHY BLACKBOX -to | -entity dram
```

GUI

To preserve the boundary interface of an entity using the GUI, make an EDA Formal Verification Hierarchy assignment to the entity with the value BLACKBOX as shown in [Figure 19-3](#).

Figure 19-3. Setting the Black-Box Property on a Module



Generating the Post-Fit Netlist Output File and the Conformal LEC Setup Files

The following steps describe how to set up the Quartus II software environment to generate the post-fit .vo netlist file and the Conformal LEC script for use in formal verification. With the exception of step 2, the steps are identical for both of the Synthesis tools:

To create a new Quartus II project or open an existing project, perform the following steps:

1. On the Assignments menu, click **Settings**. The **Settings** dialog box appears.
2. In the **Category** list, click **EDA Tool Settings**.

If you are using Quartus II integrated synthesis, perform the following steps:

- a. In the **Category** list, under **EDA Tool Settings**, select **Design Entry/Synthesis**. Select **<None>** from the **Tool name** list.
- b. In the **Category** list, under **EDA Tool Settings**, select **Formal Verification**. Select **Conformal LEC** from the **Tool name** list.

If you are using Synplify Pro, perform the following steps:

- a. In the **Category** list, under **EDA Tool Settings**, select **Design Entry/Synthesis**. Select **Synplify Pro** from the **Tool name** list.
 - b. In the **Category** list, under **EDA Tool Settings**, select **Formal Verification**. Select **Conformal LEC** from the **Tool name** list.
3. In the **Category** list, click the “+” icon to expand **Compilation Process Settings**, and select **Incremental Compilation**. The **Incremental Compilation** page appears.
 4. Select **Full Incremental Compilation** to turn on Incremental Compilation.

or

Turn on Incremental Compilation by typing the following Tcl command in the Quartus II software Tcl console:

```
set_global_assignment -name INCREMENTAL_COMPILATION FULL_INCREMENTAL_COMPILATION
```



Altera requires that Incremental Compilation be turned on for Formal Verification, and that your design does not contain any user-created partitions. The incremental compilation feature is on by default.

5. In the **Category** list, click the “+” icon to expand **Compilation Process Settings** and click **Physical Synthesis Optimizations**. The **Physical Synthesis Optimizations** page appears.
6. Turn off **Perform register retiming**.



If **Perform register retiming** is not turned off, an error occurs during compilation: “Physical Netlist Optimization Register retiming is not supported by Formal Verification tool Conformal LEC”.

- Under **Optimize for fitting (physical synthesis for density)**, turn off both **Perform physical synthesis for combinational logic** and **Perform logic to memory mapping** to prevent logic from being mapped to RAMs.

Retiming a design, either during the synthesis step or during the fitting step, usually results in moving and merging registers along the critical path and is not well-supported by the equivalence checking tools. Because equivalence checkers compare the cone of logic terminating at registers, do not use retiming to move the registers during optimization in the Quartus II software.

 For more information about physical synthesis, refer to the *Netlist Optimizations and Physical Synthesis* chapter in volume 2 of the *Quartus II Handbook*.

- Perform a full compilation of the design. On the Processing menu, click **Start Compilation**, or click the **Start Compilation** icon on the Toolbar.

Quartus II Software Generated Files, Formal Verification Scripts, and Directories

After successful compilation, the Quartus II software generates a list of files, formal verification scripts, and directories in the `<project_directory>/fv/conformal/` directory (Table 19-3).

Table 19-3. Quartus II Software Compiler-Generated Files and Directories (Part 1 of 2)

File or Directory	Name	Details
.vo file	<code><proj rev>.vo</code>	The Quartus II software-generated netlist for formal verification.
Script file	<code><proj rev>.ctc</code>	The <code><proj rev>.ctc</code> file references <code><proj rev>.clg</code> and <code><proj rev>.clr</code> that read the library files and black box descriptions. The <code><proj rev>.ctc</code> file also references the <code><proj rev>.cmc</code> file containing information about the mapped points. (1)
	<code><proj rev>.cec</code>	The <code><proj rev>.cec</code> file contains information for instance equivalences.
	<code><proj rev>.cep</code>	The <code><proj rev>.cep</code> file contains information for black box pin equivalences in the design.
	<code><proj rev>.cmp</code>	The <code><proj rev>.cmp</code> file contains information for the black box pin mapping between the golden and revised sides. (2)
	<code><proj rev>.cmc</code>	The <code><proj rev>.cmc</code> file contains information about the additional points to be mapped in addition to the points selected by the tool.
	<code><proj rev>_trivial.cmc</code>	This <code><proj rev>_trivial.cmc</code> file contains mapping information for all the key points in the design. (3)
	<code><proj rev>.clr</code>	The <code><proj rev>.clr</code> file contains information about the macros and libraries for the revised design.
	<code><proj rev>.clg</code>	The <code><proj rev>.clg</code> file contains information about the macros and libraries for the golden design.

Table 19-3. Quartus II Software Compiler-Generated Files and Directories (Part 2 of 2)

File or Directory	Name	Details
blackboxes directory	<code><project directory>/fv/conformal/<project rev>_blackboxes</code>	This directory contains top-level module descriptions for all the user-defined black box entities and contains modules with definitions other than Verilog HDL or VHDL, for example, Block Design File (.bdf) in the design directory <code><project directory>/fv/conformal/<project rev>_blackboxes</code>

Notes to Table 19-3:

- (1) This file is used with the Conformal LEC software.
- (2) This file is called from the `<proj rev>.ctc` script file. By default, the line where this file is called is commented out. These files are only useful for HardCopy II device families.
- (3) In some cases, Conformal LEC software performs incorrect key point mapping, resulting in formal verification mismatches. To overcome the verification mismatches, the Quartus II software writes out the `<proj rev>_trivial.cmc` file that contains mapping information for all the key points in the design. Reading this file during the formal verification setup can result in increased run time. Therefore, the Quartus II software writes out the top-level script file `<proj rev>.ctc` with the command to read the `<proj rev>_trivial.cmc` file commented out. If the formal verification results are not acceptable, the user can uncomment the command and read the `<proj rev>_trivial.cmc` file. The command in the `<proj rev>.ctc` file is:


```
//Trivial mappings with same name registers
//read mapped points $PROJECT/fv/conformal/<proj rev>_trivial.cmc
```

The script file contains the setup and constraints information to use with the formal verification tool. The file `<entity>.v` in the **blackboxes** directory contains the module description of entities that are not defined in the formal verification library. The file also contains entities that you specify as black boxes. For example, if there is a reference to a black box for an instance of the ALTDPRAM megafunction in the design, the **blackboxes** directory does not contain a module description for the ALTDPRAM megafunction because it is defined in the `altdpram.v` file of the formal verification library. When a module does not have an RTL description, or the description exists only in the formal verification library and you do not want to compare the module using formal verification, a file containing only the top-level module description with port declaration is written out to the **blackboxes** directory and read into the Conformal LEC software. To learn more about black boxes, refer to [“Black Boxes in the Conformal LEC Flow” on page 19-8](#).

Understanding the Formal Verification Scripts for Conformal LEC

The Quartus II software generates scripts to use with the Conformal LEC software. This section elaborates on the details of the Conformal LEC commands used within the scripts to help you compare the revised netlist with the golden netlist. In most cases, you do not have to add any more Conformal LEC constraints to verify your netlists.

A sample script generated by the Quartus II software is provided in [“Conformal Dofile/Script Example” on page 19-19](#).

Conformal LEC Commands within the Quartus II Software-Generated Scripts

The value for the variable `QUARTUS` is the path to the Quartus II software installation directory:

```
setenv QUARTUS <Quartus Installation Directory>
```

The Quartus II software assigns the current working directory of the project to the `PROJECT` variable. Use this variable to change the project directory to the directory where the design files are installed when moving from a UNIX to a Windows environment, or vice versa:

```
setenv PROJECT <Quartus Project Directory>
```

The following command reads both the golden and the revised netlists, along with the appropriate library models:

```
read design <design files>
```

 You must update the project location when the files are moved from the Windows environment to the UNIX environment.

The post place-and-route netlist from the Quartus II software might contain net and instance names that are slightly different from those of the golden netlist. By using the following command, the Quartus II software defines temporary substitute string patterns enabling the Conformal LEC software to automatically map key points when the names are not the same:

```
add renaming rule <rule>
```

The Conformal LEC software employs three name-based methods to map key points to compare the revised netlist with the golden netlist. Scripts set the correct method to get the best results.

```
set mapping method <mapping_rule>
```

The Quartus II software performs several optimizations, including optimizing the registers whose input is driven by a constant. Under these circumstances, for the formal verification software to compare the netlists properly, the command `set flatten model` is used with the option `seq_constant`.

```
set flatten model <flattening_rule>
```

When you use the command `report black box`, verify that the following modules are listed as black boxes, along with any of the modules black boxed by the user, in both the golden and revised netlists:

- LPMs and megafunctions without the formal verification models
- Encrypted IP functions
- Entities not implemented in Verilog HDL or VHDL

Use the following command to set the same implementation on multipliers for both the golden and revised netlists:

```
set multiplier implementation <implementation_name>
```

If there are any combinational loops or instances of `LPM_LATCH`, the Quartus II software cuts the loop at the same point using the following command on both the golden and revised netlists:

```
add cut point
```

The Conformal LEC software does not always automatically map all of the keypoints, or can incorrectly map some keypoints. To help the Conformal LEC software successfully complete the mapping process, the Quartus II software records optimizations performed on the netlist as a series of add mapped points in the Conformal LEC `<file_name>.cmc` script.

```
add mapped points <key_points>
```

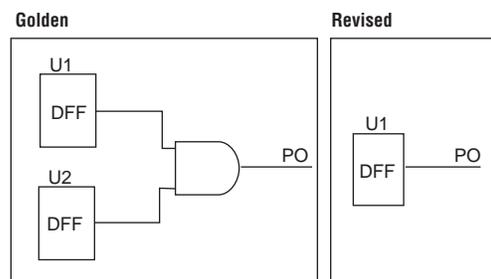
There are situations in which the inverter in front of the register is moved after the register. In this situation, the following command is used:

```
add mapped points <key_points> -invert
```

The following command reads in the mapped point information from the specified file:

```
read mapped points <file_name>.cmc
```

Figure 19-4. Instance Equivalence



During the process of optimization, the Quartus II software might merge two registers into one (Figure 19-4). The Quartus II software informs the formal verification tool that the U1 and U2 registers are equivalent to each other using the following command:

```
add instance equivalence <instance_pathname ..> [-Golden]
```

If the register duplication takes place, the following command is used:

```
add instance equivalence <instance_pathname ..> [-revised]
```

The following command is used when the inverter is moved beyond the register along with either register duplication or merging:

```
add instance equivalences <instance_pathname>
[-invert <instance_pathname>]
```

At times, the register output is driven to a constant, either logic 0 or logic 1. The Quartus II software sets the value of the register to a constraint using the add instance constraint command. For more information about this command, refer to “Stuck-at Registers” on page 19-5.

```
add instance constraint <constraint_value>
```

Comparing Designs Using Conformal LEC

This section addresses using the Conformal LEC software to compare designs, and to prove logical equivalence between two versions of the design.

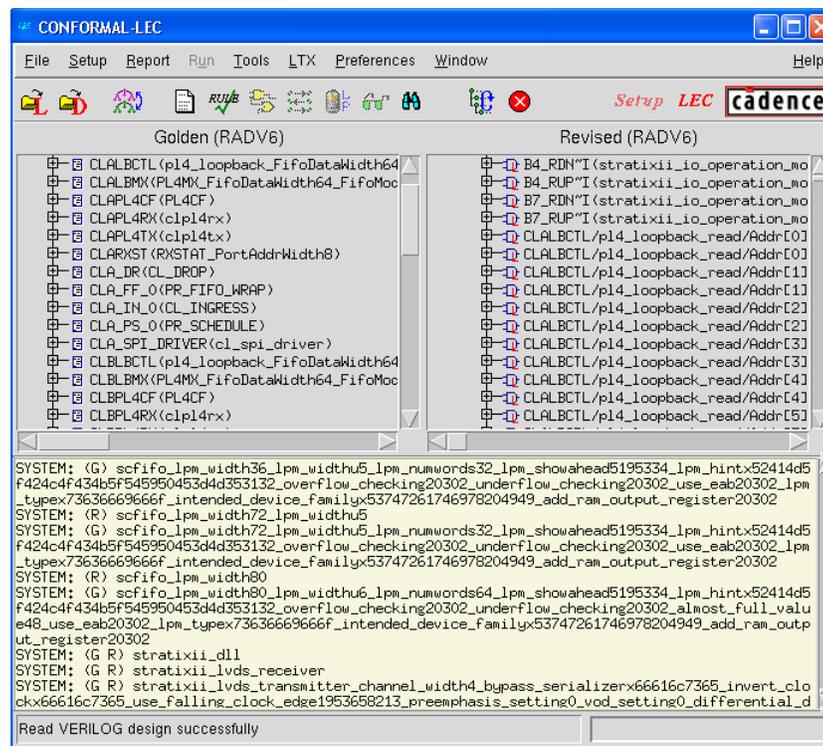
Running the Conformal LEC Software from the GUI

To run the Conformal LEC software from the GUI, follow these steps:

1. Open the Conformal LEC software.
2. On the File menu, click **Do Dofile**.
3. Select the file `<path to project directory>/fv/conformal/<proj rev>.ctc`.

The Conformal LEC software GUI displays the comparison results (Figure 19-5). The Golden window displays the original RTL description or the post synthesis .vqm netlist file from Synplify Pro, and the Revised window displays the information from the post-fit netlist generated by the Quartus II software. The message section at the bottom of the window reports the verification results and the number of unmapped and non-equivalent points found in the design.

Figure 19-5. Conformal LEC Software GUI Display of Functional Comparisons



To investigate the verification results, click the **Mapping Manager** icon in the toolbar, or on the Tools menu, click **Mapping Manager**. The Conformal LEC software reports the mapped, unmapped, and compared points in the **Mapped Points**, **Unmapped Points**, and **Compared Points** windows, respectively.

 For more information about how to diagnose non-equivalent points, refer to the Conformal LEC software user documentation.

Running the Conformal LEC Software From a System Command Prompt

To run the Conformal LEC software without using the GUI, type the command shown in [Example 19-6](#) at a system command prompt.

Example 19-6. Conformal LEC Command to Run Formal Verification

```
lec -dofile /<path to project directory>/fv/conformal/<proj rev>.ctc -nogui
```

To get a downloadable design example showing the formal verification flow with Quartus II software, refer to the [Formal Verification Design Example](#) page of the Altera website.



For more information about the latest debugging tips and solutions for formal verification flow between the Conformal LEC software and the Quartus II software, go to www.altera.com and perform an advanced search with keywords “formal verification”.

Known Issues and Limitations

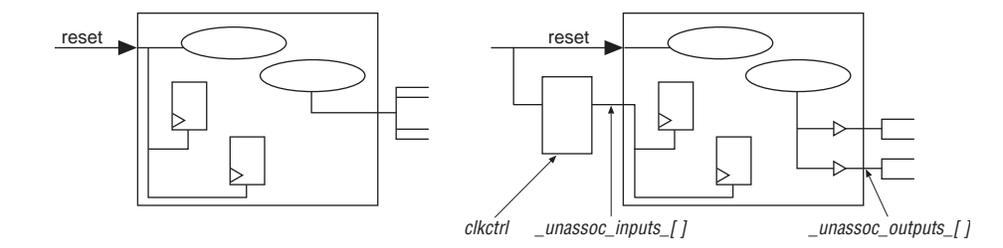
The following known issues and limitations can occur when using the formal verification flow described in this chapter:

- When a port on a black box entity drives two or more signals within the black box, the Quartus II software pushes the connections outside of the black box, and creates that many ports on the black box. This problem is only associated with Stratix II and HardCopy II designs.

The additional ports on the black box are named `_unassoc_inputs_[]` and `_unassoc_outputs_[]` ([Figure 19-6](#)). This issue is generally associated with reset and enable signals. [Figure 19-6](#) shows an example in which the reset pin is split into two ports outside of the black box and the `_unassoc_inputs_[]` port is driven by the `clkctrl` block. In such situations, the `.vo` netlist file generated by the Quartus II software has signals driving these black box ports, but golden RTL does not contain any signals to drive the `_unassoc_inputs_[]` port, resulting in a formal verification mismatch of the black box. The black box module definition generated by the Quartus II software in the directory `<Quartus_project>\fv\conformal*_blackboxes` contains these additional `_unassoc_inputs_[]` and `_unassoc_outputs_[]` ports. This black box module is read on both the golden and revised sides of the design, which results in unconnected ports on the golden side and formal verification mismatches.

[Figure 19-6](#) shows the creation of the `_unassoc_inputs_[]` and `_unassoc_outputs_[]` ports for the reset signal.

Figure 19-6. Creation of `_unassoc_inputs_[]` and `_unassoc_outputs_[]`

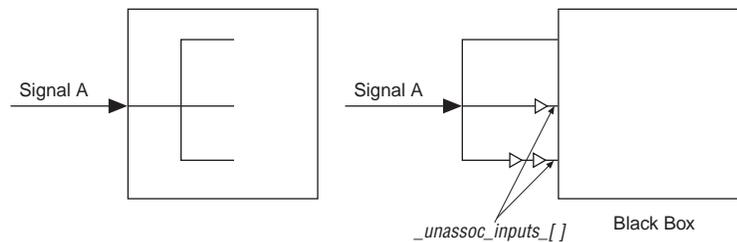


Another common occurrence of this issue is in HardCopy II designs. Whenever a port drives large fan-out within the black box, the Quartus II software inserts a buffer on the net and moves the logic outside of the black box (Figure 19-7).

To fix the problem of `_unassoc_input_[]` ports causing black box mismatches, use Conformal LEC commands to change the type of the black box `_unassoc_input_[]` keypoint to a primary output keypoint, and then mark the appropriate pin equivalences. Similarly, to fix the problem of register mismatches due to `_unassoc_output_[]` pins from black boxes, use Conformal LEC commands to change the type of the blackbox `_unassoc_output_[]` keypoint to a primary input, and then mark the equivalent pins as such. The commands to perform these actions are written in the `<proj rev>.cep` file.

Figure 19-7 shows the creation of `_unassoc_inputs_[]` for a signal with large fan-out.

Figure 19-7. Creation of `_unassoc_inputs_[]` for a Signal with Large Fan-out



- In designs with combinational feedback loops, the Conformal LEC software can insert extra cut points in the revised netlist, causing unmapped points and ultimately verification mismatches.
- For Cyclone II designs, Conformal LEC might report non-equivalent flipflops and extra cut points for the revised (post-fit) design under the following conditions:
 - when your HDL source code instantiates the `lpm_ff` primitive with an asynchronous load signal `aload` (with or without any other asynchronous control signals) and,
 - when the asynchronous clear signal `aclr` and asynchronous set signal `aset` are used together.

To avoid this problem, ensure that there is a wrapper module or entity around the `lpm_ff` instantiation, and black box the module or entity that instantiates the `lpm_ff` primitive.

- For Stratix III designs, the Conformal LEC software creates cut points for the combinational loops on the golden side and might fail equivalence checking due to improper mapping. The combinational loops are due to logic around the registers emulating multiple sets, resets, or both. These cut points are also reported during the mapping step in Quartus II software with warning messages. You can add Conformal LEC commands manually to add cut points, which can result in proper mapping and formal verification.

- To perform formal verification, certain synthesis optimization options (such as register retiming, optimization through black box hierarchy boundaries, and disabling the ROM and shift register inference) are turned off, which can have an impact on the area resource and performance.
- RAM and ROM instantiations, inferences, or both are not verified using formal verification.
- Incremental compilation for the purpose of formal verification does not support user-created design partitions.
- Formal verification does not support clear box netlists due to unconnected ports on its WYSIWYG instances.
- Formal verification does not support VHDL megafunction variations due to undriven ports on the megafunctions.
- When a black box contains bidirectional ports, the Quartus II software fails to reconstruct the hierarchy. For this reason, the black box is represented by a flat netlist, resulting in formal verification mismatches.
- ROMs in the design must be black boxed before compilation using Quartus II integrated synthesis, because the Quartus II software might perform some optimizations on the ROM, resulting in formal verification mismatches.
- The Conformal LEC software might report mismatches or abort comparisons of some key points when a DSP megafunction is implemented in LEs by the Quartus II software, due to implicit optimizations within the DSP and the complexity of the multiplier logic in terms of LEs.
- Unused logic optimized within and around a black box by the Quartus II software can result in a black-box interface different from the interface in the synthesized .vqm netlist file.

Black Box Models

The black box models are interface definitions of entities, such as primitives, atoms, LPMs, and megafunctions. These models have a parameterized interface, and do not contain any definition of behavior. They are designed and tested specifically to work with the Conformal LEC software, which uses these models along with your design to generate black boxes for instances of the entity with varying sets of parameters in the design.

- ④ For a complete list of supported black box models, refer to *Guidelines for Creating a Design for Use with the Encounter Conformal and Quartus II Software* in Quartus II Help.

Conformal Dofile/Script Example

The following example script (Example 19-7), generated by the Quartus II software, lists some of the setup commands used in Conformal LEC software.

Example 19-7. Conformal LEC Script (Part 1 of 2)

```
// Copyright (C) 1991-2008 Altera Corporation
// Your use of Altera Corporation's design tools, logic functions
// and other software and tools, and its AMPP partner logi
// functions, and any output files from any of the foregoing
// (including device programming or simulation files), and any
// associated documentation or information are expressly subject
// to the terms and conditions of the Altera Program License
// Subscription Agreement, Altera MegaCore Function License
// Agreement, or other applicable license agreement, including,
// without limitation, that your use is for the sole purpose of
// programming logic devices manufactured by Altera and sold by
// Altera or its authorized distributors. Please refer to the
// applicable agreement for further details.

// Script generated by the Quartus II software

reset
set system mode setup
set log file mfs_3prm_1a.fv.log -replace
set naming rule "%s" -register -golden
set naming rule "%s" -register -revised
// Naming rules for Verilog
set naming rule "%L.%s" "%L[%d].%s" "%s" -instance
set naming rule "%L.%s" "%L[%d].%s" "%s" -variable
// Naming rules for VHDL
// set naming rule "%L:%s" "%L:%d:%s" "%s" -instance
// set naming rule "%L:%s" "%L:%d:%s" "%s" -variable
// set undefined cell black_box -both
// These are the directives that are not supported by the QIS RTL to gates FV flow
set directive off verplex ambit
set directive off assertion_library black_box clock_hold compile_off compile_on
set directive off dc_script_begin dc_script_end divider enum infer_latch
set directive off mem_rowselect multi_port multiplier operand state_vector template
add notranslate module alt3pram -golden
add notranslate module alt3pram -revised
setenv QUARTUS /data/quark/build/ajaishan/quartus
setenv PROJECT /net/quark/build/ajaishan/quartus_regtest/eda/fv/conformal/synplify/
stratix/mfs_3prm_1a_v1_/mfs_3prm_1a/qu_allopt
```

Example 19-7. Conformal LEC Script (Part 2 of 2)

```

read design \
  $QUARTUS/eda/fv_lib/vhdl/dummy.vhd \
  -map lpm $QUARTUS/eda/fv_lib/vhdl/lpms \
  -map altera_mf $QUARTUS/eda/fv_lib/vhdl/mfs \
  -map stratix $QUARTUS/eda/fv_lib/vhdl/stratix \
  -vhdl -noelaborate -golden
read design \
  -file $PROJECT/fv/conformal/mfs_3prm_1a.clg \
  $PROJECT/p3rm_block.v \
  $PROJECT/mfs_3prm_1a.v \
  -verilog2k -merge none -golden
read design \
  $QUARTUS/eda/fv_lib/vhdl/dummy.vhd \
  -map lpm $QUARTUS/eda/fv_lib/vhdl/lpms \
  -map altera_mf $QUARTUS/eda/fv_lib/vhdl/mfs \
  -map stratix $QUARTUS/eda/fv_lib/vhdl/stratix \
  -vhdl -noelaborate -revised
read design \
  -file $PROJECT/fv/conformal/mfs_3prm_1a.clr \
  $PROJECT/fv/conformal/mfs_3prm_1a.vo \
  -verilog2k -merge none -revised
// add ignored inputs _unassoc_inputs_* -all -revised
add renaming rule r1 "~I\" "/" -revised
add renaming rule r2 "_I\" "/" -revised
set multiplier implementation rca -golden
set multiplier implementation rca -revised
set mapping method -name first
set mapping method -nounreach
set mapping method -noreport_unreach
set mapping method -nobbox_name_match
set flatten model -seq_constant
set flatten model -nodff_to_dlat_zero
set flatten model -nodff_to_dlat_feedback
set flatten model -nooutput_z
set root module mfs_3prm_1a -golden
set root module mfs_3prm_1a -revised
report messages
report black box
report design data
// report floating signals
dofile $PROJECT/fv/conformal/mfs_3prm_1a.cec
// dofile $PROJECT/fv/conformal/mfs_3prm_1a.cep
// Instance-constraints commands for constant-value registers removed
// during compilation
set system mode lec -nomap
read mapped points $PROJECT/fv/conformal/mfs_3prm_1a.cmc

// Trivial mappings with same name registers
// read mapped points $PROJECT/fv/conformal/mfs_3prm_1a_trivial.cmc
// dofile $PROJECT/fv/conformal/mfs_3prm_1a.cmp
map key points
remodel -seq_constant -repeat
add compare points -all
compare
usage
// exit -f

```

Conclusion

Formal verification software enables verification of the design during all stages from RTL to placement and routing. Verifying designs takes more time as designs increase in size. Formal verification is a technique that helps reduce the time needed for your design verification cycle.

Document Revision History

Table 19-4 shows the revision history for this chapter.

Table 19-4. Document Revision History

Date	Version	Changes
December 2010	10.1.0	Changed to new document template. Removed Table 21-1.
July 2010	10.0.0	Updates for new GUI changes, and added link to Help.
November 2009	9.1.0	Updated “Black Boxes in the Encounter Conformal Flow” section.
March 2009	9.0.0	Updated Table 21-1.
November 2008	8.1.0	<ul style="list-style-type: none">■ Changed to 8-1/2 x 11 page size.■ Added support for Stratix IV devices.■ Added support for Cadence Conformal LEC version 7.2 and Synplify Pro version 9.6.2.
May 2008	8.0.0	<ul style="list-style-type: none">■ Added support for Cyclone III devices.■ Updated “Black Boxes in the Encounter Conformal Flow” section.■ Updated Table 18-1 and Table 18-5.

 For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this chapter.

The Quartus® II software offers a complete software solution for system designers who design with Altera® FPGA and CPLD devices, including device programming. The Quartus II Programmer is part of the Quartus II software package that allows you to program Altera CPLD and configuration devices, and configure Altera FPGA devices. This section describes how you can use the Quartus II Programmer to program or configure your device after you successfully compile your design.

This section includes the following chapter:

- [Chapter 20, Quartus II Programmer](#)

This chapter describes how to program and configure Altera® CPLD, FPGA, and configuration devices with the Quartus® II Programmer.

The Quartus II software offers a complete software solution for system designers who design with Altera FPGA and CPLD devices. After you compile your design, you can use the Quartus II Programmer to program or configure your device.

This chapter contains the following sections:

- “Programming Flow”
- “Quartus II Programmer GUI” on page 20–3
- “Programming and Configuration Modes” on page 20–5
- “Scripting Support” on page 20–10

- ❓ For more information about how to use the Quartus II Programmer GUI to program and configure your device, refer to *Programming Devices* in Quartus II Help.

Programming Flow

The programming flow is as follows:

1. Compile your design, such that the Quartus II Assembler generates the programming or configuration file.
2. Perform programming or configuration file conversion for different configuration devices, or optional programming and configuration file creation.
3. Configure or program the FPGA, CPLD, or configuration devices using the programming or configuration file with the Quartus II Programmer.

Figure 20-1 shows the programming file generation flow.

Figure 20-1. Programming File Generation Flow

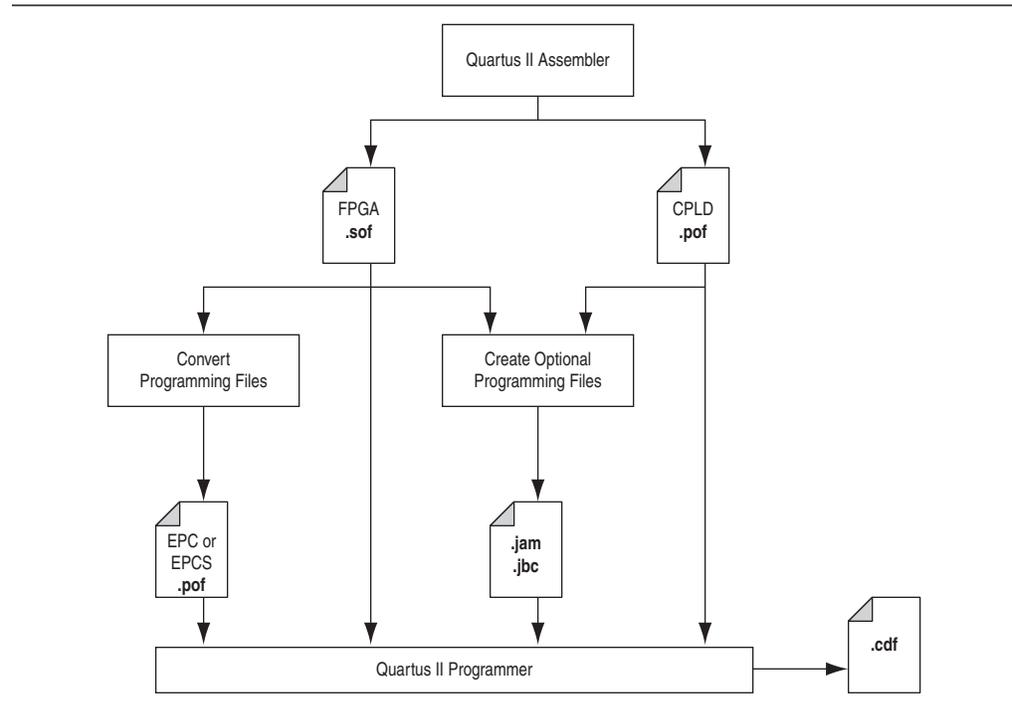


Table 20-1 shows the programming and configuration file formats supported by Altera FPGAs, CPLDs, and configuration devices.

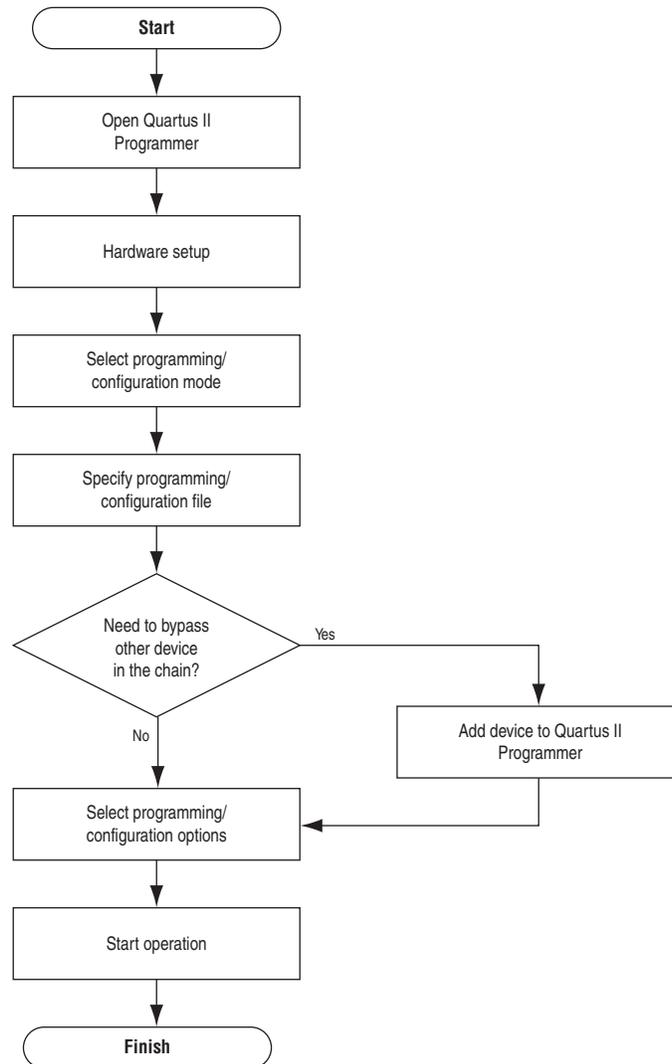
Table 20-1. Programming and Configuration File Format

File Format	FPGA	CPLD	Configuration Device	Serial Configuration Device
SRAM Object File (.sof)	✓	—	—	—
Programmer Object File (.pof)	—	✓	✓	✓
JEDEC JESD71 STAPL Format File (.jam)	✓	✓	✓	—
Jam Byte Code File (.jbc)	✓	✓	✓	—

❓ For more information about Chain Description Files (.cdf), refer to *About Programming* in Quartus II Help.

Figure 20-2 shows the programming flow using the Quartus II Programmer.

Figure 20-2. Programming Flow



Quartus II Programmer GUI

The Quartus II Programmer is a window in which you can add your programming and configuration files, specify the programming options and hardware, and then proceed with the programming or configuration of the device.

To open the Programmer window, on the Tools menu, click **Programmer**. The Quartus II message window reports the status of each operation, whether programming is successful or not.

- ❓ For a description of the Programmer window, refer to *Programmer Window* in Quartus II Help. For a description of options in the Tools menu, refer to *Programmer Page (Options Dialog Box)* in Quartus II Help.

Hardware Setup

The Quartus II Programmer provides the flexibility to choose a download cable or the programming hardware. Before you can program or configure your device, you must have the correct hardware setup.

 For hardware settings, refer to *Setting Up Programming Hardware* in Quartus II Help.

 For more information about programming hardware driver installation, refer to the *Design Software Support* page on the Altera website.

JTAG Settings

The JTAG server allows programs such as the Quartus II Programmer to access the JTAG hardware. You can also access the JTAG download cable or programming hardware connected to a remote computer through the JTAG server of that computer. With the JTAG server, you can control the programming or configuration of devices from a single computer through other computers at remote locations. The JTAG server uses the TCP/IP communications protocol.

 For more information about JTAG settings, refer to *Using the JTAG Server* in Quartus II Help.

JTAG Chain Debugger Tool

The JTAG Chain Debugger tool is a Quartus II Programmer feature that allows you to test the JTAG chain integrity and detect intermittent failures of the JTAG chain. In addition, the tool allows you to shift in JTAG instructions and data through the JTAG interface and step through the test access port (TAP) controller state machine for debugging purposes.

JTAG Chain Debugger Example

To set and send the USERCODE JTAG instruction in your device with the JTAG Chain Debugger, follow these steps:

1. On Assignment menu, click **Device**.
2. Click **Device and Pin Option**.
3. In the General list, type abcdefcb in the **JTAG user code (32-bit hexadecimal)** field.
4. Click **OK**.
5. Click **OK**.
6. On the Tools menu, click **JTAG Chain Debugger**.
7. In the **JTAG Chain Debugger** dialog box, click the **JTAG Chain Debugging** tab.
8. In the **Command** pulldown list, select **Scan Instruction Register**.
9. In the **TAP State** pulldown list, select **End scan in RTI state**.
10. In the **Clocks** field, type 10.

11. In the **TDI** field, type 0x07, and then click **Run**. The **TDO** field displays 0x155.

To read the **USERCODE** JTAG instruction with the JTAG Chain Debugger, follow these steps:

1. In the **Command** pulldown list, select **Scan Data Register**.
2. In the **TAP State** pulldown list, select **End scan in RTI state**.
3. In the **Clocks** field, type 40.
4. In the **TDI** field, type 0000 0000 or any value, and then click **Run**. The **TDO** field displays 0x00ABCDEFAB.

 For more information, refer to *Using the JTAG Chain Debugger* in Quartus II Help.

Other Programming Tools

The following section describes other programming tools in more detail.

Stand-Alone Quartus II Programmer

If you do not have the full version of the Quartus II software, Altera offers the free Quartus II Stand-Alone Programmer. The stand-alone programmer has the same full functionality as the Quartus II Programmer in the Quartus II software. You can download the Quartus II Stand-Alone Programmer from the [Download Center](#) on the Altera website.

Programming and Configuration Modes

The following section describes the Quartus II Programmer and the Programmer configuration modes in more detail.

Configuration Modes

The Quartus II Programmer supports four configuration modes, including JTAG, passive serial (PS), active serial (AS), and in-socket modes (ISM).

[Table 20–2](#) shows the programming and configuration modes supported by Altera devices.

Table 20–2. Programming and Configuration Modes

Mode	FPGA	CPLD	Configuration Device	Serial Configuration Device
JTAG	✓	✓	✓	—
PS	✓	—	—	—
AS	—	—	—	✓
In-Socket Programming	—	✓ (1)	✓	✓

Notes to Table 20–2:

- (1) MAX II CPLDs do not support in-socket programming mode.

With the Quartus II Programmer, you can program and configure one or more devices. This section provides references to Quartus II Help procedures that describe how to program or configure Altera devices, and how to bypass Altera and non-Altera devices in a JTAG chain.

- ① For more information about programming or configuring a single device in JTAG and AS programming modes, refer to *Programming Devices* in Quartus II Help. For more information about how to use the different configuration modes, refer to *About Programming* in Quartus II Help.
- For more information about JTAG configuration or programming mode and JTAG pin connections, refer to the *Configuration Handbook*, or the device handbook or data sheet for the respective FPGA, CPLD, or configuration device.
- For more information about PS configuration mode and PS pin connection, refer to the *Configuration Handbook* or the chapter about configuration in the appropriate FPGA device handbook.
- For more information about programming the serial configuration device, configuring the FPGA with the serial configuration device through AS mode, or the AS pin connections, refer to the *Serial Configuration Data Sheet* in the *Configuration Handbook* or the chapter about configuration in the appropriate FPGA device handbook.
- For a list of programming adapters available for Altera devices, refer to www.altera.com.

Design Security Key Programming

The Quartus II Programmer supports the generation of encryption key programming files and encrypted configuration files for Altera FPGAs that support the design security feature. You can also use the Quartus II Programmer to program the encryption key into the FPGA.

- For more information about using the design security feature with the Quartus II software, refer to *AN 341: Using the Design Security Feature in Stratix II and Stratix II GX Devices* and *AN 512: Using the Design Security Feature in Stratix III Devices*.

The Quartus II software can generate optional programming or configuration files in various formats that you can use with programming tools other than the Quartus II Programmer. When you compile a design in the Quartus II software, the Assembler automatically generates either a **.sof** or **.pof**. The Assembler also allows you to convert FPGA configuration files to programming files for configuration devices.

- ① For more information, refer to *About Optional Programming Files* in Quartus II Help.

- For more information about the programming and configuration file formats, refer to file format topics in the Quartus II Help or the *Configuration File Formats* chapter of the *Configuration Handbook*. For more information about using the **.jam** and **.jbc** programming files with the Jam STAPL Player, Jam STAPL Byte-Code Player, and the `quartus_jli` command-line executable, refer to *AN 425: Using Command-Line Jam STAPL Solution for Device Programming*.

Generating Secondary Programming Files

The Quartus II software generates programming files of various formats for use with different programming tools.

Table 20-3 lists the file types generated by the Quartus II software and supported by the Quartus II Programmer.

Table 20-3. File Types Generated by the Quartus II Software and Supported by the Quartus II Programmer

File Type	Generated by the Quartus II Software	Supported by the Quartus II Programmer
.sof	✓	✓
.pof	✓	✓
.jam	✓	✓
.jbc	✓	✓
JTAG Indirect Configuration File (.jic)	✓	✓
.svf	✓	—
In System Configuration File (.isc)	✓	—
Hexadecimal (Intel-Format) Output File (.hexout)	✓	—
Raw Binary File (.rbf)	✓	—
Tabular Text File (.ttf)	✓	—
Raw Programming Data File (.rpd)	✓	—

- For more information, refer to *Generating Secondary Programming Files* in Quartus II Help.

Convert Programming Files Dialog Box

The **Convert Programming Files** dialog box contains the following sections:

- **Conversion setup file**—Allows you to import conversion setup information from a file and save the conversion setup information for future use.
- **Output programming file**—Allows you to specify programming file types, configuration device types and configuration modes, file names, and Remote/Local update difference files. You can also add new files.
- **Input programming file**—Allows you to add file for the conversion.

- **Advanced Options**—Allows you to debug configuration issues.



Use the **Advanced** option only when you encounter an issue with device configuration.

The **Convert Programming Files** dialog box in the Programmer allows you to convert programming files from one file format to another. For example, to store the FPGA data in configuration devices, you can convert the **.sof** data to another format and then program the configuration device. You can also convert FPGA data into other file formats, including **.pof**, **.hexout**, **.rbf**, **.rpd**, or **.jic**.

You must have the Programmer open to access the **Convert Programming Files** dialog box. On the Quartus II main menu, click **File**, and then click **Convert Programming Files**. You can then perform the following tasks:

- Configure multiple devices, such as combining multiple **.sof** files into one **.pof**.
- Configure multiple devices with an external host, such as a microprocessor or CPLD. For example, you can combine multiple **.sof** files into one configuration file.



For more information about converting programming files with the Quartus II software, refer to the *Configuration File Formats* chapter of the *Configuration Handbook*.

You can use the **Advanced** option in the **Convert Programming Files** dialog box to debug your configuration.



When you change settings in the **Advanced** option, the change affects **.pof**, **.jic**, **.rpd**, and **.rbf** files.

You must choose the setting that applies to your Altera device. You can force the Quartus II software to enable the option by turning the option on or off in the **Advanced Options** dialog box.

Table 20-4 describes the **Advanced Options** settings in more detail.

Table 20-4. Advanced Options Settings

Option Setting	Description
Disable EPCS ID check	<ul style="list-style-type: none"> ■ FPGA skips the EPCS silicon ID verification. ■ Default setting check box is grayed out (EPCS ID check is enabled). ■ Applies to the single- and multi-device AS configuration modes on all FPGA devices.
Disable AS mode CONF_DONE error check	<ul style="list-style-type: none"> ■ FPGA skips the CONF_DONE error check. ■ Default setting check box is grayed out (EPCS ID check is enabled). ■ Applies to single- and multi-device (AS) configuration modes on all FPGA devices.
Program Length Count (PLC) settings	<ul style="list-style-type: none"> ■ Specifies the offset you can apply to the computed PLC of the entire bitstream. ■ Default setting is 0. The value should be an integer. ■ Applies to single- and multi-device (AS) configuration modes on all FPGA devices.
Post-chain bitstream pad bytes	<ul style="list-style-type: none"> ■ Specifies the number of pad bytes appended to the end of an entire bitstream. ■ Default value is set to 0 if the bitstream of the last device is uncompressed. Set to 2 if the bitstream of the last device is compressed.

Table 20-4. Advanced Options Settings

Option Setting	Description
Post-device bitstream pad bytes	<ul style="list-style-type: none"> ■ Specifies the number of pad bytes appended to the end of the bitstream of a device. ■ Default value is 0. No negative integer. ■ Applies to all single-device configuration modes on all FPGA devices.
Bitslice padding value	<ul style="list-style-type: none"> ■ Specifies the padding value used to prepare bitslice configuration bitstreams, such that all bitslice configuration chains simultaneously receive their final configuration data bit. ■ Default value is 1. Valid setting is 0 or 1. ■ Use only in 2, 4, and 8-bit PS configuration mode, when you use an EPC device with the decompression feature enabled. ■ Applies to all FPGA devices that support enhanced configuration devices.

Table 20-5 lists symptoms you may encounter if a configuration fails, and describes the **Advanced Options** you must use to debug your configuration.

Table 20-5. Failure Symptoms and Options Settings

Failure Symptoms	Disable EPCS ID Check	Disable AS Mode CONF_DONE Error Check	PLC Settings	Post-Chain Bitstream Pad Bytes	Post-Device Bitstream Pad Bytes	Bitslice Padding Value
Configuration failure occurs at the end of a configuration cycle. Decompression feature enabled. Encryption feature enabled.	—	—	—	Use only for multi-device chain.	Use only for single-device chain.	—
CONF_DONE stays Low at the end of a configuration cycle.	—	—	Start with positive offset to the PLC settings.	Use only for multi-device chain.	Use only for single-device chain.	—
CONF_DONE goes High momentarily at the end of a configuration cycle.	—	—	Start with negative offset to the PLC settings.	—	—	—
FPGA does not enter user mode even though CONF_DONE goes High.	—	—	—	Use only for multi-device chain.	Use only for single-device chain.	—
Configuration failure occurs at the beginning of a configuration cycle.	—	—	—	—	—	—
Newly introduced EPCS, such as EPCS128.	—	—	—	—	—	—
Failure in .pof generation for EPC device using Quartus II Convert Programming File Utility when the decompression feature is enabled.	—	—	—	—	—	—

The **Input files to convert** option in the **Convert Programming Files** dialog box allows you to add one or more files for the conversion process after specifying the the output file types in the **Output Programming Files** option.

-  For more information about converting programming files with the Quartus II software, refer to the *Configuration File Formats* chapter of the *Configuration Handbook*.
-  For more information about the **Convert Programming Files** dialog box, refer to *Convert Programming Files Dialog Box* in Quartus II Help.

Flash Loaders

Parallel and serial configuration devices do not support the JTAG interface. You are not able to program parallel and serial configuration devices directly through normal JTAG programming. You can use a flash loader to program serial configuration devices in-system via the JTAG interface. To do so, use an FPGA as a bridge between the JTAG interface and the serial configuration device. Altera supports parallel and serial flash loaders.

-  For more information, refer to *About Flash Loaders* in Quartus II Help.

Scripting Support

In addition to the Quartus II Programmer GUI, you can use the Quartus II command-line programmer (`quartus_pgm`) to enter commands. The programmer accepts the **.pof**, **.sof**, and **.jic** programming or configuration files and Chain Description Files (**.cdf**).

Example 20-1 shows a command that programs a device:

Example 20-1. Programming a Device

```
quartus_pgm -c byteblasterII -m jtag -o bpv:design.pof ↵
```

Where:

- `-c byteblasterII` specifies the ByteBlaster™ II download cable
- `-m jtag` specifies the JTAG programming mode
- `-o bpv` represents the blank-check, program, and verify operations
- `design.pof` represents the **.pof** used for the programming

The Programmer automatically executes the erase operation before programming the device.

-  You can also run the procedures for the programmer in a Tcl script. For more information about scripting command options, refer to *About Quartus II Scripting* in Quartus II Help.

The jtagconfig Debugging Tool

You can use the `jtagconfig` command-line utility (which is similar to the auto detect operation in the Quartus II Programmer) to check the devices in a JTAG chain and the user-defined devices.

For more information about the `jtagconfig` utility, type one of the following commands at the command prompt:

Example 20-2.

```
jtagconfig -h ←
jtagconfig --help ←
```



The help switch does not reference the `-n` switch, which you may find to be the most useful switch. The `jtagconfig -n` command shows each node for each jtag device.



For more information about Tcl scripting, refer to the *Tcl Scripting* chapter in volume 2 of the *Quartus II Handbook*. For information about all settings and constraints in the Quartus II software, refer to the *Quartus II Settings File Manual*. For more information about command-line scripting, refer to the *Command-Line Scripting* chapter in volume 2 of the *Quartus II Handbook*.

Conclusion

The Quartus II Programmer offers you a wide variety of options to program and configure your Altera devices. With the Quartus II Programmer, the Quartus II software provides you with a complete solution for your FPGA or CPLD design prototyping, which can also be performed in the production environment.

Document Revision History

Table 20-6 shows the revision history for this chapter.

Table 20-6. Document Revision History

Date	Version	Changes
December 2010	10.1.0	<ul style="list-style-type: none"> ■ Changed to new document template. ■ Updated “JTAG Chain Debugger Example” on page 20-4. ■ Added links to Quartus II Help. ■ Reorganized chapter.
July 2010	10.0.0	<ul style="list-style-type: none"> ■ Added links to Quartus II Help. ■ Deleted screen shots.
November 2009	9.1.0	No change to content.
March 2009	9.0.0	<ul style="list-style-type: none"> ■ Added a row to Table 21-4. ■ Changed references from “JTAG Chain Debug” to “JTAG Chain Debugger”. ■ Updated figures.



For previous versions of the *Quartus II Handbook*, refer to the [Quartus II Handbook Archive](#).

 Take an [online survey](#) to provide feedback about this handbook chapter.

This chapter provides additional information about the document and Altera.

About this Handbook

This handbook provides comprehensive information about the Altera® Quartus® II design software, version 10.1.

How to Contact Altera

To locate the most up-to-date information about Altera products, refer to the following table.

Contact (1)	Contact Method	Address
Technical support	Website	www.altera.com/support
Technical training	Website	www.altera.com/training
	Email	custrain@altera.com
Product literature	Website	www.altera.com/literature
Non-technical support (General) (Software Licensing)	Email	nacomp@altera.com
	Email	authorization@altera.com

Note to Table:

(1) You can also contact your local Altera sales office or sales representative.

Third-Party Software Product Information

Third-party software products described in this handbook are not Altera products, are licensed by Altera from third parties, and are subject to change without notice. Updates to these third-party software products may not be concurrent with Quartus II software releases. Altera has assumed responsibility for the selection of such third-party software products and its use in the Quartus II 10.1 software release. To the extent that the software products described in this handbook are derived from third-party software, no third party warrants the software, assumes any liability regarding use of the software, or undertakes to furnish you any support or information relating to the software. EXCEPT AS EXPRESSLY SET FORTH IN THE APPLICABLE ALTERA PROGRAM LICENSE SUBSCRIPTION AGREEMENT UNDER WHICH THIS SOFTWARE WAS PROVIDED TO YOU, ALTERA AND THIRD-PARTY LICENSORS DISCLAIM ALL WARRANTIES WITH RESPECT TO THE USE OF SUCH THIRD-PARTY SOFTWARE CODE OR DOCUMENTATION IN THE SOFTWARE, INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, AND NONINFRINGEMENT. For more information, including the latest available version of specific third-party software products, refer to the documentation for the software in question.

Typographic Conventions

The following table shows the typographic conventions this document uses.

Visual Cue	Meaning
Bold Type with Initial Capital Letters	Indicate command names, dialog box titles, dialog box options, and other GUI labels. For example, Save As dialog box. For GUI elements, capitalization matches the GUI.
bold type	Indicates directory names, project names, disk drive names, file names, file name extensions, software utility names, and GUI labels. For example, \qdesigns directory, D: drive, and chiptrip.gdf file.
<i>Italic Type with Initial Capital Letters</i>	Indicate document titles. For example, <i>Stratix IV Design Guidelines</i> .
<i>italic type</i>	Indicates variables. For example, $n + 1$. Variable names are enclosed in angle brackets (< >). For example, <file name> and <project name>.pdf file.
Initial Capital Letters	Indicate keyboard keys and menu names. For example, the Delete key and the Options menu.
“Subheading Title”	Quotation marks indicate references to sections within a document and titles of Quartus II Help topics. For example, “Typographic Conventions.”
Courier type	Indicates signal, port, register, bit, block, and primitive names. For example, data1, tdi, and input. The suffix n denotes an active-low signal. For example, resetn. Indicates command line commands and anything that must be typed exactly as it appears. For example, c:\qdesigns\tutorial\chiptrip.gdf. Also indicates sections of an actual file, such as a Report File, references to parts of files (for example, the AHDL keyword SUBDESIGN), and logic function names (for example, TRI).
	An angled arrow instructs you to press the Enter key.
1., 2., 3., and a., b., c., and so on	Numbered steps indicate a list of items when the sequence of the items is important, such as the steps listed in a procedure.
	Bullets indicate a list of items when the sequence of the items is not important.
	The hand points to information that requires special attention.
	A question mark directs you to a software help system with related information.
	The feet direct you to another document or website with related information.
	A caution calls attention to a condition or possible situation that can damage or destroy the product or your work.
	A warning calls attention to a condition or possible situation that can cause you injury.
	The envelope links to the Email Subscription Management Center page of the Altera website, where you can sign up to receive update notifications for Altera documents.