

Advanced FPGAs present tremendous time-to-market advantages for end users by incorporating the latest silicon fabrication technologies. However, the very nature of these devices presents challenges to the FPGA vendor in presenting accurate timing models for end users to design-in new devices into systems. This paper describes the timing models available for end-user FPGA design, in advance of production FPGAs.

Introduction

Users of leading-edge FPGAs need detailed timing characteristics well before production-ready silicon is available. However, FPGAs are inherently difficult to characterize due to their use of advanced process nodes and extreme configurability. This leads to a longer silicon characterization cycle than is typical in non-programmable devices.

To meet the needs of FPGA users well before the availability of volume-produced silicon, timing models are made available in several stages. Each stage is increasingly accurate as the FPGA device development cycle proceeds, while sample silicon is characterized and correlated with timing models. The availability of timing models allows you to develop designs concurrently with new FPGA device development.

This paper explains the development cycle of FPGAs, along with the timing models available at each stage. This paper also describes the best practices for use of each timing model stage.

FPGAs Tackle Unique Timing Model Challenges

Advanced process nodes present many challenges as process nodes push the physical limits of what can be done in silicon. Because of the enormous complexity, cost, and risks associated with working at the most advanced process geometries, few IC designers work in this rarefied environment; the return on investment does not justify the large outlay of engineering time. FPGA vendors are among the few who see the economies of scale that justify the expense of finding solutions to the device modeling problems. Because of this, FPGA vendors work at the leading edge of silicon design and manufacture.

The challenges stem from both basic device physics at the nanometer scale and from the unique configurability of FPGAs to meet RTL designer needs.



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

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or 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.



Timing Model Challenges at Advanced Silicon Process Geometries

In advanced microelectronics fabrication, circuit behavior is difficult to predict and model. In past generations, the behavior of silicon at previous process nodes could be used as the basis of predicting the behavior of the next process node through extrapolation. Some of these effects could be ignored at 90 nm and even 65 nm. But with each new process node, physical effects are increasingly infeasible to extrapolate—new simulation techniques need to be devised, and measurements must be made on actual silicon samples and calibrated to simulation results. The Altera® white paper, *Guaranteeing Silicon Performance with FPGA Timing Models*, discusses some of the challenges inherent in today's leading-edge silicon. These challenges include:

- On-die process variation
- Asymmetric N- and P- channel transistor speed
- Clock uncertainty and jitter
- Metal variation
- End-of-life degradation effects
- Global variation
- On-chip local variation
- Signal crosstalk
- Distributed parasitic RC wire effects
- Power distribution effects

Characterization and Correlation of Silicon Measurements vs. Circuit Simulations

When silicon devices are fabricated and early silicon is available for measurement, the sampled product is measured according to a carefully planned, thorough, and time-consuming process called characterization. These tests determine the actual speed of the fabricated devices, and the results are correlated with the original behavior from circuit simulations. Differences between simulation and actual measurement result in timing model calibration so end-user timing models are effectively tuned to actual silicon measurements.

The process of characterization is a complex task in itself. A modern chip contains billions of transistors that cannot be exhaustively tested. Characterization tests are carefully designed to determine the actual operating characteristics of key elements on the device, over process, voltage, and temperature (PVT) variation extremes.

In FPGAs, the already complex characterization task is exacerbated by the very feature that makes FPGAs attractive to users in the first place—their almost infinite programmability.

An FPGA comprises many types of elements with differing timing, loading, and power needs:

- Internal memory
- Transceivers

- Standard-cell blocks
- System-on-chip features (for example, embedded processors)
- General logic elements (LEs)
- Local and global routing elements

These can all be included into a user design in a vast configuration space, interacting in ways that cannot be exhaustively enumerated and tested. Because of this, a characterization plan must be carefully considered and executed to account for all possible ways a device may be configured in the field.

Fabricated device behavior is measured over the full range of process, operating voltage, and operating temperature variations to guarantee device timing is met for the binned speed grade.

Characterization of a device continues until all known differences are accounted for and fixed in the timing model, which is eventually encoded into the timing engine used by the TimeQuest Timing Analyzer and in the Quartus® II software fitter place-and-route tool.

Timing Analysis at All Operating Conditions

An FPGA must operate over a range of temperatures, supply voltages, and manufacturing process variations. There will even be variations within a particular speed grade from device to device. All of these variables affect device timing in different ways. At advanced process geometries of 28 nm and beyond, timing delays may not vary in direct relationship with temperature or voltage, due to nanometer-scale effects like temperature inversion. To fully account for worst- and best-case timing situations, timing must be analyzed at multiple operating conditions.

Altera provides timing models for four extreme operating conditions. Timing analysis is done at all operating conditions to guarantee correct operation over temperature, voltage, and manufacturing process variations.

FPGA Timing Analysis Accuracy Requires Circuit Simulation

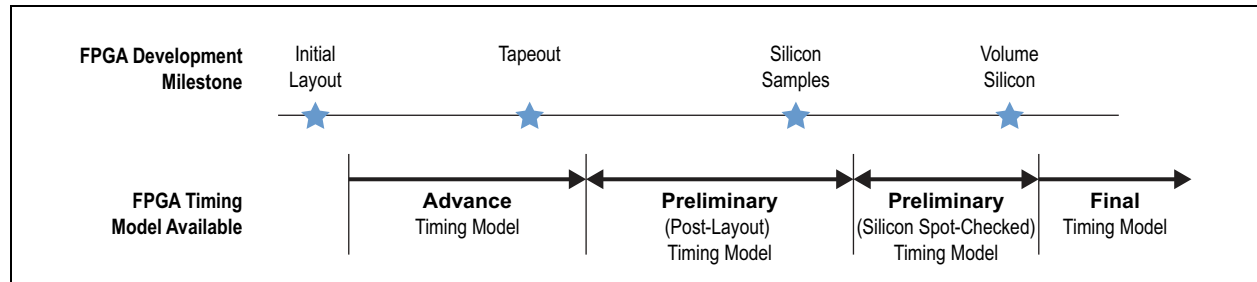
If a device's logic topology is not known until it is programmed by the RTL designer, how can we deliver accurate timing information to a device?

Altera's TimeQuest Timing Analyzer performs its analysis on a user's design with two distinct approaches—a static timing-delay database and a custom circuit simulator. The delay database is used for elements with limited and well-defined configurability such as logic and hard intellectual property (IP) blocks. The circuit simulator does detailed simulation, taking into account the particular configuration of the user design. The circuit simulation handles routing interconnect delays that cannot be accurately modeled with static delay numbers due to the highly configurable routing fabric. Circuit simulation accounts for the effects of capacitive loading, listening position on the routing element, delays for the selected metal routing track, the input waveform applied to the circuit, and other effects. A detailed SPICE circuit simulation is impractical for an entire end-user design. However, the Quartus II software circuit simulator is optimized for the FPGA architecture so it runs several orders of magnitude faster than a general purpose SPICE simulation.

The FPGA Development Process and Timing Model Availability

To understand the availability of user timing models, it is useful to relate the available models to the FPGA development process. See [Figure 1](#).

Figure 1. FPGA Development Models



Initial Layout

At this stage of design, the architecture of a new device is finalized. The device's mixture of internal components is settled upon (a percentage of die devoted to routing, general logic, memory, transceivers, and so on) and a floorplan is created. At this point, detailed layout has not been done. However, a good estimate of block-to-block delay is available and extracted from the EDA tools. Detailed timing simulation has not been done at this stage. The design is ready for detailed layout and simulation at completion of this stage. The advance timing model is available upon completion of this stage.

Tapeout

Detailed layout is done, and circuit simulations are completed. Circuit delays are derived from post-layout simulations which include the latest information on physical device effects in partnership with the foundry. Once this design phase is completed, final parameter extraction is done from the EDA tools and the layout and simulations are at a level of detail ready for tapeout to the foundry. The preliminary post-layout timing model is available upon completion of this stage.

Spot-Check Silicon

Once the foundry has fabricated sample devices, these are spot-checked according to correlation plans, and the results are correlated with the original simulation results. Any differences are traced to its source. During this stage, timing models are updated as silicon continues to be spot-checked over time. The preliminary silicon spot-checked timing model is available upon completion of this stage.

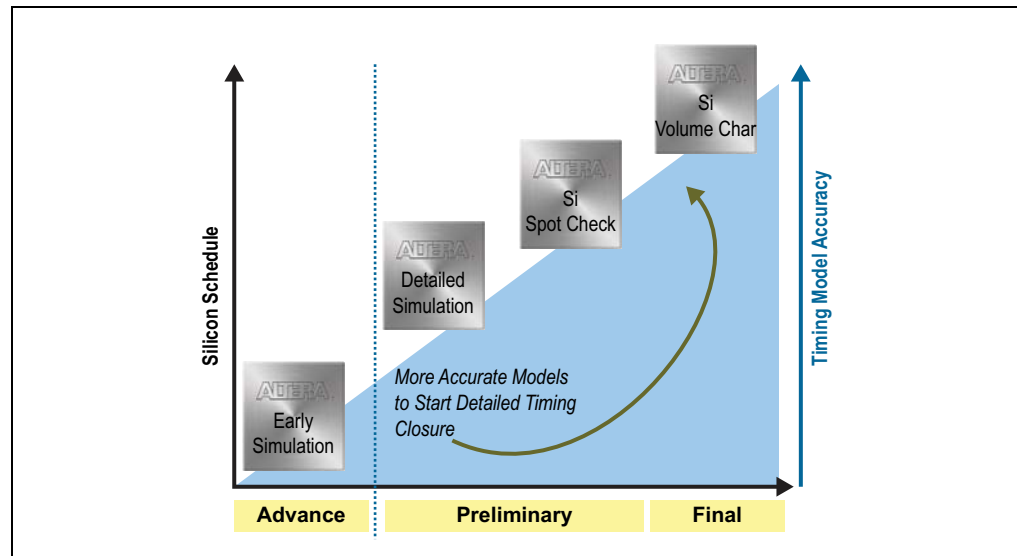
Volume Silicon

Eventually, volume production of devices is attained by the foundry, with large samples of silicon available to correlate to timing models. Many characterization goals depend on a statistically significant sample size to characterize behavior over process variation corners, so volume measurements are an important milestone. At this stage, high confidence in characterization measurements is reached, and correlation work is completed. The final timing model is available upon completion of this stage.

Best Practices for Using Various Timing Models

To allow end users to design systems around FPGAs that are not yet available, Altera supplies increasingly accurate timing models at each stage of the FPGA development. See [Figure 2](#).

Figure 2. Timing Model Types



Advance Timing Models

Timing models at this stage include initial estimate of delays based on pre-layout information and engineering estimates, which are subject to change after layout and simulation is complete. Because of this, the absolute delay numbers are subject to change so the Advance model is to be used for system architecture studies, and not to be used for detailed timing closure work. Advance timing models should be used with the following expectations:

- System latency assessments
- Basic timing assessments (pipeline budgeting)
- I/O transfer strategy (datapath width, burst depth, I/O standards tradeoffs)



Advance timing models should not be used for timing closure work. Because of this, FPGA programming file (.pof) output is not supported for this stage of timing model development.

Preliminary Models

Preliminary models incorporate timing data that is subject to change as characterization and correlation work is yet to be done. These models are more accurate than advance models because they are based on detailed simulation results, and calibrated with silicon measurements as characterization work continues from early, spot-check production sampling through high-volume production sampling.

Preliminary Post-Layout Timing Models

First few models designated as *Preliminary* are those that incorporate the delays extracted from circuit simulation of the final device layout. These delays are extracted from the tapeout-quality representation of the implemented FPGA. However, the initial preliminary models do not include measured results from physical silicon. For this reason, these models are not to be used for system design signoff, although initial timing closure efforts can start.

Preliminary Silicon Spot-Checked Timing Models

Later preliminary models contain timing information correlated to actual early sample silicon measurements. At this stage, high-volume characterization work has not yet been completed--these preliminary models are updated as characterization work proceeds at Altera, and silicon results are correlated with timing models and transferred to Quartus II software tools. Preliminary timing models should be used with the following expectations:

- Timing closure efforts can begin and continue with more accurate results as improved preliminary models are available.
- Programming file (.pof) support for early access customer, for system bring-up in lab



Designers should expect to use a newer build of software with final models before releasing their product.

Final Models

At this stage, silicon has been characterized at sufficient volume to fully characterize timing behavior. These models are approved for use in production systems.

- Final timing closure
- Suitable for customer production systems

Conclusion

System designers working with FPGAs face competitive pressures in their markets to release systems as soon as component devices are physically available. Design must start well in advance of the availability of physical FPGAs. Rigorous system design demands the use of accurate timing models for FPGAs.

However, leading-edge FPGAs present difficult physical modeling challenges, which delay the availability of final timing models until volume silicon characterization has been completed. Compounding the characterization process is the enormous configurability of current high-end FPGAs.

Altera enables users to design and integrate FPGAs into their systems well in advance of available silicon. This is done by providing several tiers of timing models during the FPGA development process, each encoding the most accurate timing information from the earliest stages of development, through to the final, fully correlated timing models.

Further Information

- White paper: *Guaranteeing Silicon Performance with FPGA Timing Models*
www.altera.com/literature/wp/wp-01139-timing-model.pdf

Acknowledgements

- Stanley Ng, Applications Engineer, Technical Services, Altera Corporation
- Adrian Drury, Applications Engineer, Technical Services, Altera Corporation

Document Revision History

Table 1 shows the revision history for this document.

Table 1. Document Revision History

Date	Version	Changes
February 2014	1.0	Initial release.