Stratix II GX SSN Design Guidelines

August 2007, ver. 1.0 Application Note 472

Introduction

Today's FPGA devices exhibit sub-nanosecond edge rates to meet the critical timing requirements needed to run high-speed memory interfaces, and communicate over high-speed serial links. The downside of fast edge rates coupled with large parallel I/O buses (DDR) can cause a variety of signal integrity problems like crosstalk, including simultaneous switching noise (SSN), which can result in degradation of system performance if it is not accounted for during the initial design phase.

This application note provides information on the major mechanisms of simultaneous switching noise, along with various techniques to mitigate the observed noise in the Stratix® II GX device family. In addition, this application note explains how Altera® defines signal margin, covers the impact of multiple I/Os toggling on transceiver performance, and concludes by providing a high-level overview of best practices for designing a board.

It is critical to follow the best practices described in this document to ensure the best performance from your system.

The following key terms are used in this document:

Victim: Any I/O that has noise induced on it can be considered a victim. An example of the victim signal can be CLOCK, RESET, and STATUS signals switching asynchronously with respect to data signals in a DDR interface.

Aggressor: Any I/O that induces noise can be considered an aggressor.

Quiet Low (QL): The quiet low voltage level on a victim net is the observed voltage level at the far end of the victim net when driven to logic zero (low) in the absence of any aggressor I/Os toggling simultaneously.

Quiet Low Noise (QLN): The quiet low noise is the noise that is coupled onto the victim driven static low sharing the same VCCIO/GND return path with multiple aggressors switching simultaneously.

Quiet High (QH): The quiet high voltage level on a victim net is the observed voltage level at the far end of the victim net when driven to logic one (high) in the absence of any aggressor I/Os toggling simultaneously.

Quiet High Noise (QHN): The quiet high noise is the noise that is coupled onto the victim driven static high sharing the same VCCIO / GND return path with multiple aggressors switching simultaneously.

VIL_{Max}/VIH_{Min} (DC): For a given receiver, the VIL / VIH (DC) is defined such that the final logic state is unambiguously defined, that is once the receiver input has crossed this value, the receiver changes to and maintains the new logic state.

Figure 1 gives a schematic representation of the definitions listed above.

Figure 1. QHN and QLN Definitions

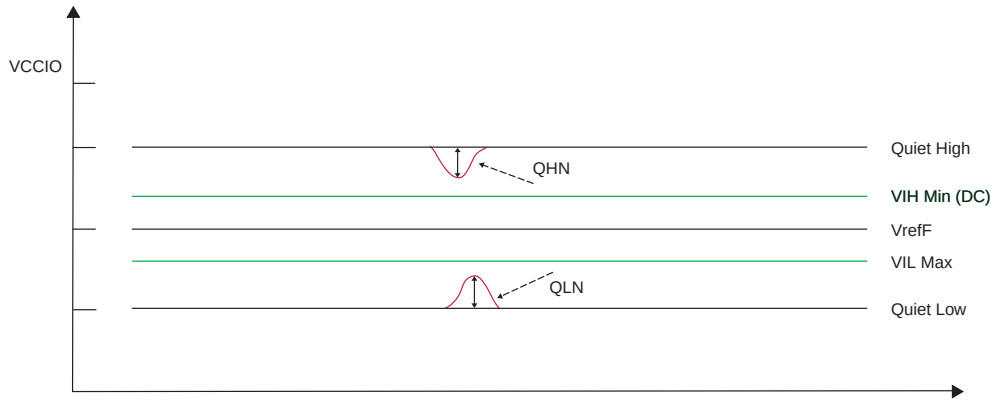

The quiet high and quiet low voltage levels observed at the far end of the transmission line are dependent on the output buffer characteristics along with the termination scheme used for the interface.

Simultaneous Switching I/O Noise

Introduction

Simultaneous switching noise is the inductive noise caused by multiple outputs switching at the same time. A victim net switching by itself may exhibit good signal quality. However, when all the signals in the bus are switching simultaneously, noise generated from the other signals can corrupt the signal quality of the victim net. This can be observed by probing the quiet I / O pin driven high or low.

Simultaneous switching noise is broken down into three distinct mechanisms:

- Inductive crosstalk

- Delta-I
- Power supply compression

This noise can result in the degradation of circuit performance at the near or far ends of the bus. Figure 2 captures the important circuit components for discussion of the mechanisms and the resulting simultaneous switching noise.

Figure 2. Schematic Diagram of FPGA and PCB Topology

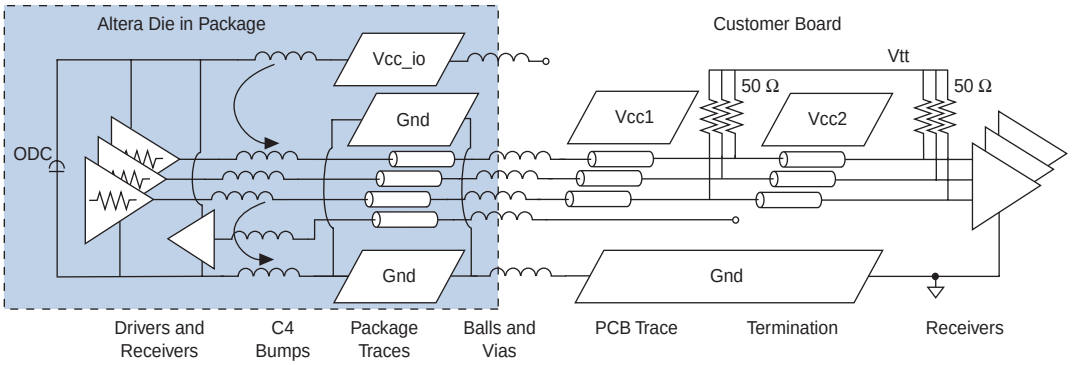

The blue box represents the Stratix II GX silicon and the surrounding electronic package. There may be up to 85 drivers in an I/O bank within a die that can switch simultaneously, thus giving rise to the SSN. The die is connected to the package through die bumps (C4), which are represented by an inductance. Signals traverse from the die to the ball locations through transmission lines. Package power planes bring power into the chip, and ground planes provide a return path for both the power and signals. Finally, all signals, power and ground of the FPGA package, are connected to the printed circuit board (PCB) with solder balls. The remainder of the system (PCB) is outside the blue box. HSTL class II is the signaling technology represented in Figure 2.

Inductive Coupling

Inductive coupling is often the dominant mechanism for SSN, and occurs when the current from one conductor generates a magnetic field that is coupled to another conductor and generates a voltage across it.

The inductive coupling mechanism can be represented by the following equation:

$$V = \sum_{k=1}^{N} M_k \times \frac{di_k}{dt}$$

Where:

V = Total voltage induced onto the victim

k = Number of aggressors

M_k = Mutual inductance between the victim and the k^{th} aggressor I/O.

di_k / dt - Rate of change of current over time for the k^{th} aggressor

Inductive coupling only occurs when the current changes as a function of time. Conductors associated with aggressor drivers generate a change in magnetic field patterns during the rise and fall time of the signal waveform and couple noise voltage to victim conductors.

Most inductive crosstalk occurs in vertical structures rather than in horizontal transmission line structures. Examples of vertical coupling structures include the die bumps (C4), package vias, solder balls (package pins), PCB vias, and pins in a connector. The magnitude of inductive coupling is proportional to the parallel length of the aggressor and victim signals. All vertical structures contribute some amount of inductive coupling. However, most of the coupling occurs at the interface between the FPGA package and the PCB in the package balls, and the PCB vias, where the parallel path is the longest between aggressors and victims. Noise is inductively coupled from aggressor to victim conductors during the aggressor rise and fall time, and is not coupled at any other time.

Delta-I Noise

Delta-I noise occurs when a high amount of current tries to enter or exit the package through a small number of power or ground pins. When many drivers switch from high to low, the signal current enters the signal pins and must exit through the ground pins. When these drivers switch from low to high, the current must come in through the power pins.

The delta-I noise mechanism can be represented by the following equation:

$$V = \sum_{k=1}^{N} L_k \times \frac{di_k}{dt}$$

Where:

V = Ground bounce or power sag

k = Number of aggressor I/Os

L_k = Loop inductance associated with the current path (signal to ground loop, signal to power loop) usually dominated by the self-inductance of the power or ground path.

di_k / dt - Rate of change of current over time for the k^{th} aggressor.

Delta-I noise results in the movement of the on-die supply rails (V_{CC} and ground) with respect to PCB supply rails, which is the basis for industry-standard terms, such as “ground bounce" and "V_{CC} sag". Like inductive coupling, delta-I noise only occurs during the signal transition, as this is the only time where the current changes as a function of time. Delta-I noise does not occur in time frames where the driver current is constant because there is no di/dt to generate the noise.

Power Supply Compression

Power supply compression noise is the difference between power and ground voltages. It is the voltage that appears across the power supply terminals in silicon circuits on the die. This is the voltage that enables the circuits to perform their intended tasks.

For the low to high transition case, power supply compression becomes an important factor, as drivers continue to source steady state current after the SSN event has taken place. This current initially comes from the ODC (on die capacitance), which lowers the on-die voltage, and in turn brings in current from the outside through the package inductances. The local on-die power supply voltage may take several nanoseconds after the SSN event to reach its minimum or maximum value. This occurs long after the drivers have made their simultaneous transitions, and is more dependent upon delta-I noise than it is upon di/dt.

Quiet low noise and its quiet high noise equivalent are voltage fluctuations on a victim pin driven low or high resulting from inductive crosstalk, delta-I, or power supply compression. The effect of quiet low noise can be seen on a victim driven low and vice versa. Enough quiet low noise on a quiet low signal level can cause its value to be mistaken for logic high when clocked during the sampling window. The same can be said about quiet high noise on a victim pin to be mistaken for logic low, depending on the magnitude of the noise when clocked during the sampling window.

Signal Margin Terminology

This section introduces the concept of signal margin to represent the observed QLN / QHN on a victim net caused by multiple I / Os toggling simultaneously. This section also provides an easy way to breakup the total noise budget between the FPGA and other components that make up the link, thereby helping you to make good design tradeoffs to meet the total noise budget.

Scenario I - Victim Net Driven Low:

$$SML_{(K)} = \frac{QLN_{(K)} - QL}{VIL_{Max(DC)} - QL}$$

Where:

K = 0, 1, 2, 3 .... (Maximum number of I / Os in an I / O Bank – 1)

$SML_{(K)}$ = Signal Margin Low when K aggressors are switching simultaneously

$QLN_{(K)}$ = quiet low noise when K aggressors are switching simultaneously

QL = Quiet Low Voltage (No aggressors are switching)

$VIL_{Max\,(DC)}$ = Receiver Maximum DC Input Low Voltage

Scenario II - Victim Net Driven High:

$$SMH_{(K)} = \frac{QH - QHN_{(K)}}{QH - VIH_{Min(DC)}}$$

Where:

K = 0, 1, 2, 3 .... (Maximum number of I/Os in an I/O Bank – 1)

$SMH_{(K)}$ = Signal Margin High when K aggressors are switching simultaneously

$QHN_{(K)}$ = quiet high noise when K aggressors are switching simultaneously

QH = Quiet High Voltage (No aggressors are switching)

$VIH_{Min\ (DC)}$ = Receiver Minimum DC Input High Voltage

A schematic view of the SML/SMH definitions is shown in Figure 3.

Figure 3. SML and SMH Definitions

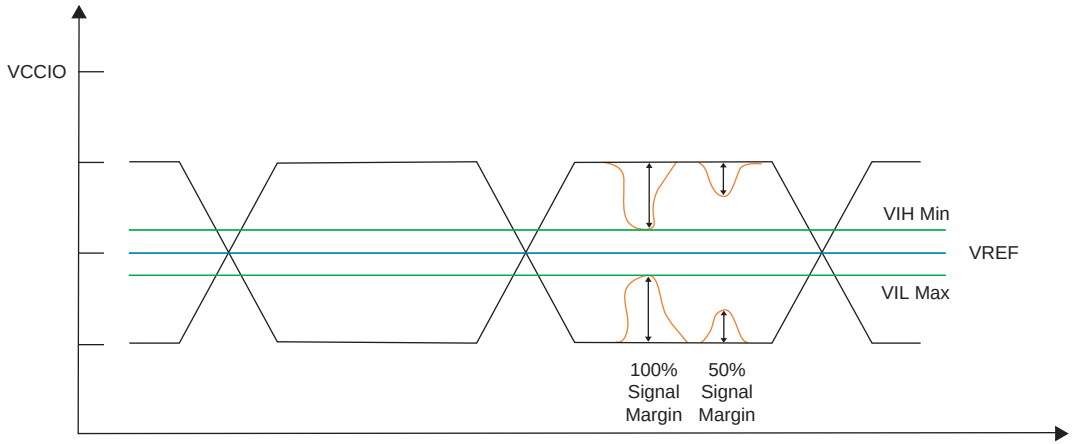

For example, based on Altera's SSN characterization data listed in Table 1 for Stratix II GX, SSTL2 Class II 25-Ω OCT, the observed signal margin for a quiet low and quiet high scenario is as follows:

$VIL_{Max\ (DC)}$ for SSTL2 = Vref – 180 mV = 1070 mV

$VIH_{Min\ (DC)}$ for SSTL2 = Vref + 180 mV = 1430 mV

Where:

Vref = VCCIO/2; VCCIO = 2.5 V

From Table 1, QL = 565 mV; QH = 1890 mV

$QLN_{(K = 4)}$ = 636 mV; $QLN_{(K = 87)}$ = 813 mV

$QHN_{(K = 4)}$ = 1800 mV; $QHN_{(K = 87)}$ = 1619 mV

$$SML_{(K = 4)} = \frac{(636 - 565)}{(1070 - 565)} = 14\%;\ SML_{(K = 87)} = \frac{(813 - 565)}{(1070 - 565)} = 49\%$$

$$SMH_{(K = 4)} = \frac{(1890 - 1800)}{(1890 - 1430)} = 20\%;$$

$$SMH_{(K = 87)} = \frac{(1890 - 1619)}{(1890 - 1430)} = 59\%$$

Table 1. Characterization Data for QLN/QHN for a SSTL2 CII 25 Ω OCT Stratix II GX Device

K (No. of Aggressors)	QLN (mV)	QHN (mV)
0	565 (QL)	1890 (QH)
4	636	1800
87	813	1619

Based on the data from Table 1, for the SSTL2 Class II 25-ΩOCT interface, you have a 41% (100% - 59%, from worst case scenario) margin left to account for other components that make up the memory interface.

Methods to Mitigate the Simultaneous Switching Noise

You can employ various methods to mitigate the observed SSN noise due to fast edge rates. The amount of noise observed on the victim is significantly lower when the aggressors are using a terminated standard like SSTL or HSTL, rather than unterminated standards like LVTTL or LVCMOS. Most of the techniques discussed in the following sections are device dependent, with the "High-Speed Board Design Guidelines" section going over the best practices for board design. The amount of observed noise on a victim is highly dependent on the choices made during the board layout from a termination scheme, signal breakout, and power delivery network design.

Use the following methods to mitigate observed simultaneous switching noise on the device:

- Lower drive strength

- Using I/O standards with termination
- Using lower voltage I/O standards
- Adding return paths based on pin placement

Lower Drive Strength

A simple and very effective method of reducing SSN is by directly reducing the amount of current utilized by the aggressor I/Os by reducing the I/O drive strength settings.

Figure 4 shows the dependence of the quiet low noise for the victim pin held quiet low for the 2.5 V LVTTL interface as a function of the number of aggressor I/Os and the current strength.

Figure 4. Impact of Drive Strength on Signal Margin

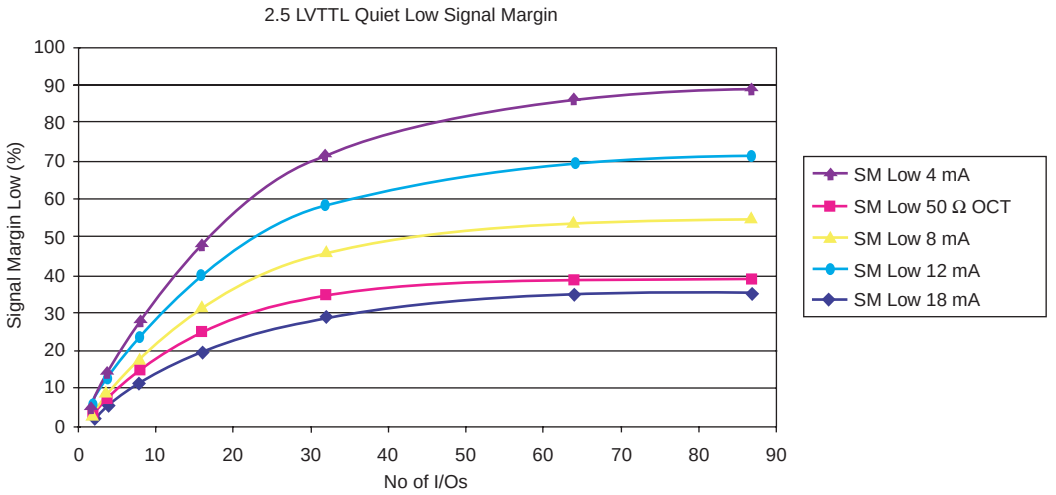

Modifying the drive strength for all SSO pins produces significantly different results in the observed quiet low noise. Quiet high noise is similarly sensitive to the drive strength.

Reducing the drive strength from 16 mA to 4 mA results in a ~54% improvement in signal margin for the quiet low noise and a similar improvement in the signal margin for quiet high noise when 87 I/Os are toggling simultaneously. This approach has the greatest benefit for interfaces that make use of unterminated standards like LVTTL or LVCMOS, compared to HSTL and SSTL.

Take care to ensure that the SSN gains from reducing the drive strength are not completely offset by the reduced signal swing and reduced quiet margin.

In practical designs, system speed requirements constrain how low a current drive strength can be utilized. Timing margins on the fast-switching data interfaces may prevent the use of low current drive strengths.

Whenever possible, do not use more I/O drive strength than is required. Perform signal integrity and timing analysis to determine the optimum drive strength setting of the I/O.

I/O Standards with Termination

Terminated I/O standards help in optimizing the signal transmission on a high-speed trace by reducing reflections. In addition, series termination resistors also act as current limiters by reducing the output voltage swing at the far end of the line.

Far-end parallel terminations limit the amount of noise observed at the far end of the line since without them the noise launched onto the transmission line would be reflected at the far end due to the impedance mismatch. The amplitude of noise measured at the far end is higher due to multiple reflections in the absence of terminations. Figure 5 shows the difference in the observed quiet low noise and quiet high noise for the victim pin held low or high for 2.5 V LVTTL 16 mA and SSTL2 CII 16 mA current drive strength, respectively. The observed signal margin for quiet low noise for SSTL2 CII 16 mA is greater by ~54% compared to that of 2.5 V LVTTL 16 mA. There is an improvement of ~47% in the observed signal margin for quiet high noise.

Figure 5. Signal Margin Comparison Between 2.5 V LVTTL 16 mA and SSTL2 CII 16 mA

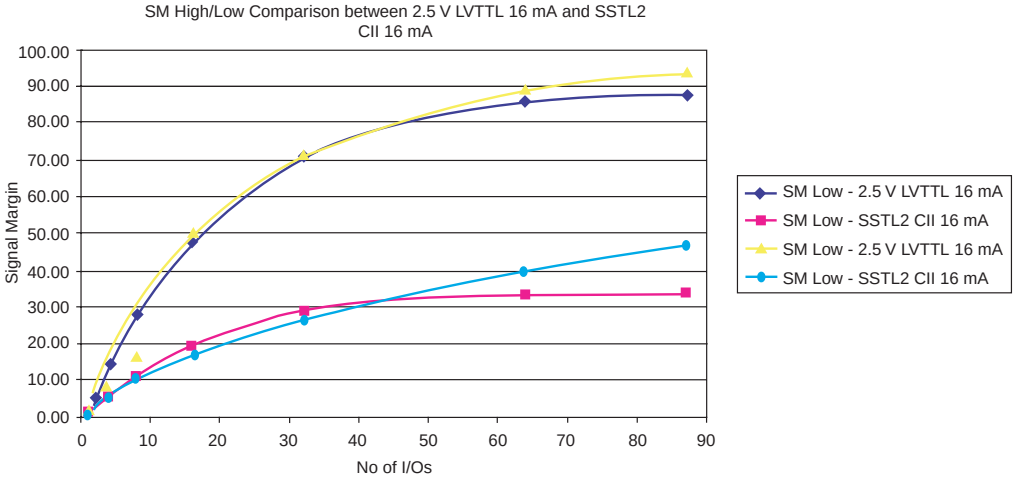

Lower Voltage I/O Standards

Using low-voltage signaling standards helps reduce noise on the quiet victim pin located near numerous aggressor pins.

Based on the following equation:

$I = V / R$

The transmission line impedance, together with the output drive strength and the power supply voltage, determines the current through the driver, and therefore the di/dt for the aggressor I/Os.

The caveat here is that with lower voltage swings come lower voltage noise margins. Figure 6 and Figure 7 show the signal margin difference between HSTL15 CII 16 mA and HSTL18 CII 16 mA for the quiet low or high scenario. The signal margin low with 87 I/Os toggling for the quiet low noise is similar for both HSTL15 CII 16 mA and HSTL18 CII 16 mA interfaces. The signal margin high for quiet high noise is greater for HSTL15 (52.7%) compared to the HSTL18 (45.2%) interface, though the noise numbers in absolute millivolts is higher for HSTL18 (190 mV) interface compared to HSTL15 (173 mV) interface. You need to take into account additional design criteria, for example, meeting EMI/EMC specifications, before deciding to go with either a low noise/low noise margin (low voltage signaling) compared to a high noise/high noise margin (higher voltage signaling) interface.

Figure 6. Signal Margin Low Comparison Between HSTL18CII 16 mA and HSTL15 CII 16 mA

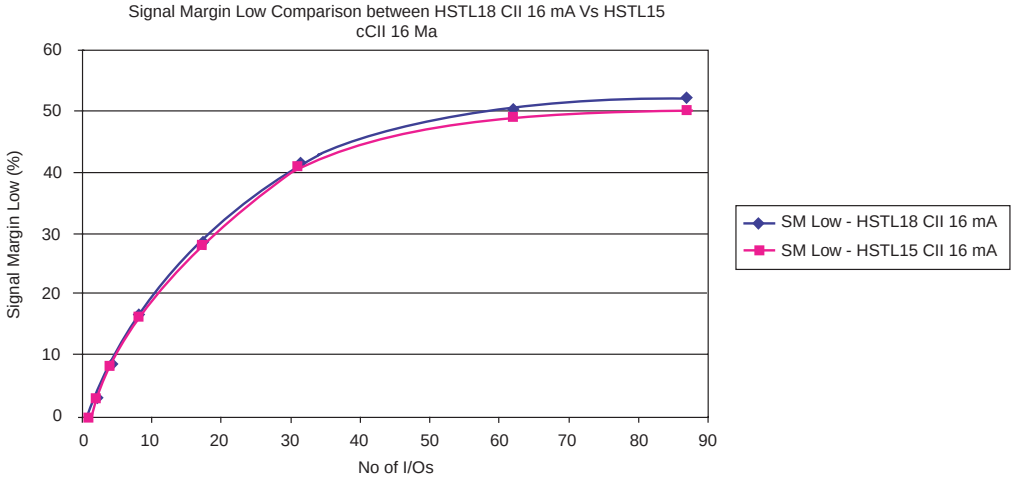

Figure 7. Signal Margin High comparison Between HSTL18 CII 16 mA and HSTL15 CII 16 mA

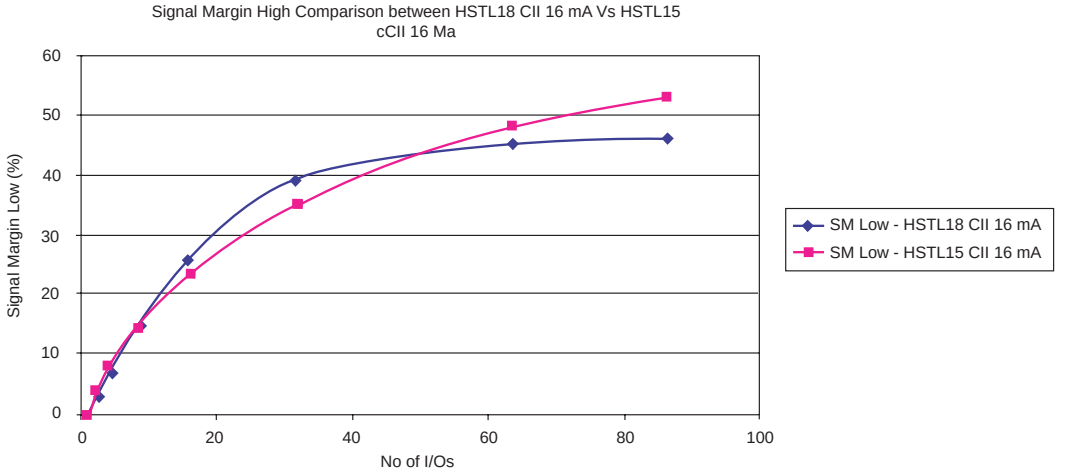

Adding Return Paths

The purpose of providing high speed return paths is to reduce the signal / ground, signal / power, and power / ground loop inductance, since any signal launched from an output buffer has to return to the source to

complete the loop. At high frequencies, the signal will do so through the path of least impedance which in majority of cases is the path of least inductance. This can be achieved as follows:

- Add return paths with unused I/O pins set as programmable power and ground pins
- Utilize separate return paths by locating offender pins away from victim pins
- Add return paths for victims and offenders by locating them closer to power and ground pins

Add Return Paths with Unused I/O Pins Set as Programmable Power and Ground Pins

If possible, program all unused I/O pins to drive either low or high, and connect them to board ground and power accordingly. Always use as many programmable grounds and power as possible. Ideally, set half of all unused outputs to ground, and the other half to VCCIO. Then spread out the pairs of programmable VCCIO and programmable GND pins evenly throughout the bank, as shown in Figure 8. The yellow-colored I/O cells in the configuration below are all unused I/Os distributed across the combined banks 4 and 9 set to programmable power and ground in the EP2SGX90EF1152 device. Doing so reduces the signal-to-return loop inductance along with the power-to-ground loop inductance.

Figure 8. Distributed Unused I/O Configuration Across the Bank

	7	8	9	10	11	12	13	14	15	16	17	18	19
A	A7	A8	A9	A10	A11	A12	A13	A14	A15	A16	A17	PVSS	A19
B	VSS	PVCCIO	VSS	PVCCIO	B11	VSS	PVCCIO	B14	VSS	PVCCIO	B17		
C	C7	C8	C9	C10	C11	PVSS	C13	PVSS	C15	C16	C17	PVSS	
D	D7	D8	PVSS	D10	D11	D12	D13	D14	D15	D16	PVSS	D18	
E	VSS	PVCCIO	VSS	VREF	E11	VSS	VREF	E14	VSS	VREF	E17		
F	F7	F8	F9	PVCCIO	F11	F12	PVCCIO	F14	F15	PVCCIO	F17	F18	F19
G	G7	G8	PVSS	G10	G11	G12	G13	PVSS	G15	G16		PVSS	G19
H	VSS	PVCCIO	VSS			VSS	H13	H14	VSS				
J						VCCIO	J13		VCCIO				
K						VCCIO	K13		VCCIO				
L							L13	L14					
M							M13	PVSS	M15	M16	PVCCIO	PVSS	M19
N											N17	N18	N19

VSS	Dedicated VSS Pin
VCCIO	Dedicated VCCIO Pin
PVCCIO	Programmable VCCIO Pin
PVSS	Programmable VSS Pin
C13	Victim Pin
	Aggressor I/Os
VREF	Dedicated VREF Pins

Since the vast amount of the quiet low and quiet high noise is produced by simultaneous switching outputs, Altera recommends giving preference to the outputs and bidirectional I/O pins as opposed to the inputs, as far as the programmable VCCIO and GND distribution in the bank is concerned. In memory interfaces, such as DDR or QDR, distributing programmable references among the data pins (writes) is more important than among the address and control pins, since the latter does not switch as frequently.

Spreading programmable VCCIO and GND in pairs throughout the bank is more effective than programming all unused pins to just one logic state (high or low) and/or clustering them in the bank. This approach helps to minimize the observed quiet low and quiet high noise on any victim across the bank, since any aggressor pin can become a victim pin during different clock cycles, depending on the data pattern. The aggressor I/Os for the configuration above are programmed to 1.8 V LVTTL standard, 12 mA drive strength.

Table 2 below gives the observed quite low noise and quite high noise when the unused I/Os for the configuration above are set to VCCIO/VSS, respectively.

Table 2. QLN/QHN Results for Distributed I/O Configuration

Scenario – I	Quiet Low Noise (mV)	Quiet High Noise (mV)
Unused I/Os distributed across the Bank (Figure 8)	286	405

Utilize Separate Return Paths by Locating Aggressor Pins Away From Victim Pins

When the victim I/O (for example, clock, reset, or control signals) is known early in the design cycle, you can achieve further reduction in the SSN by moving aggressor pins away from victim pins whenever possible, in addition to providing programmable power and grounds. Consider Figure 9 with aggressor I/Os (programmed to 1.8 V LVTTL standard, 12-mA drive strength) clustered around the quiet victim pin `C13` in the combined banks 4 and 9 in one case, and aggressor I/Os located as far away as possible from `C13` in the other case (Figure 10).

Figure 9. Simultaneous Switching Aggressors Next to Victim Pin

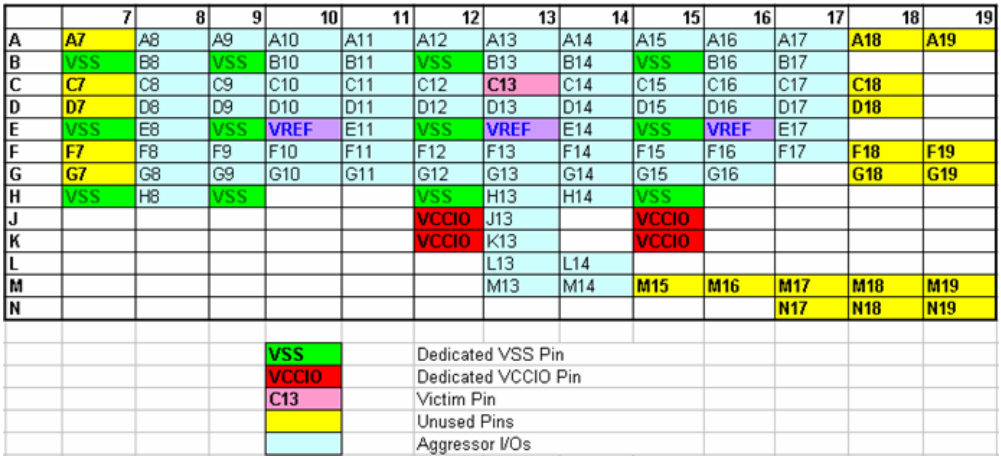

	7	8	9	10	11	12	13	14	15	16	17	18	19
A	A7	A8	A9	A10	A11	A12	A13	A14	A15	A16	A17	A18	A19
B	VSS	B8	VSS	B10	B11	VSS	B13	B14	VSS	B16	B17		
C	C7	C8	C9	C10	C11	C12	C13	C14	C15	C16	C17	C18	
D	D7	D8	D9	D10	D11	D12	D13	D14	D15	D16	D17	D18	
E	VSS	E8	VSS	VREF	E11	VSS	VREF	E14	VSS	VREF	E17		
F	F7	F8	F9	F10	F11	F12	F13	F14	F15	F16	F17	F18	F19
G	G7	G8	G9	G10	G11	G12	G13	G14	G15	G16		G18	G19
H	VSS	H8	VSS			VSS	H13	H14	VSS				
J						VCCIO	J13		VCCIO				
K						VCCIO	K13		VCCIO				
L							L13	L14					
M							M13	M14	M15	M16	M17	M18	M19
N											N17	N18	N19

VSS	Dedicated VSS Pin
VCCIO	Dedicated VCCIO Pin
C13	Victim Pin
	Unused Pins
	Aggressor I/Os

Figure 10. Simultaneous Switching Aggressors Away from Victim Pin

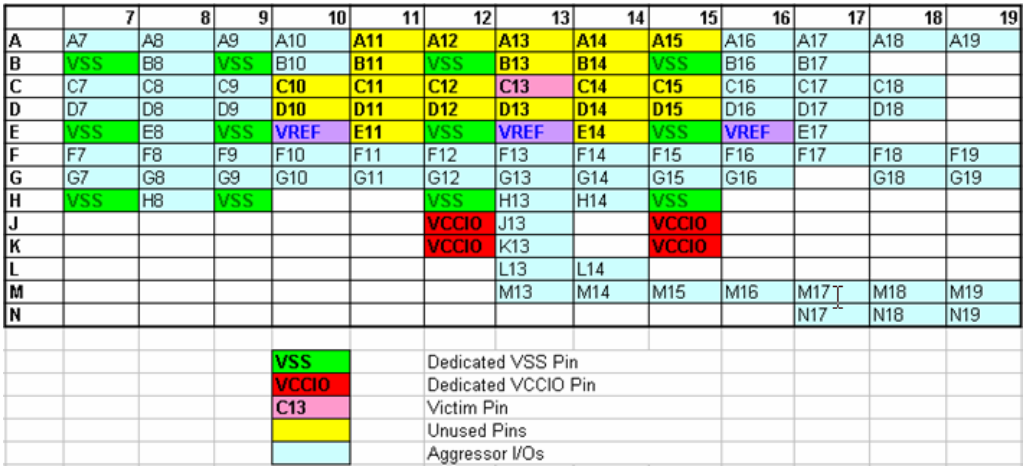

	7	8	9	10	11	12	13	14	15	16	17	18	19
A	A7	A8	A9	A10	A11	A12	A13	A14	A15	A16	A17	A18	A19
B	VSS	B8	VSS	B10	B11	VSS	B13	B14	VSS	B16	B17		
C	C7	C8	C9	C10	C11	C12	C13	C14	C15	C16	C17	C18	
D	D7	D8	D9	D10	D11	D12	D13	D14	D15	D16	D17	D18	
E	VSS	E8	VSS	VREF	E11	VSS	VREF	E14	VSS	VREF	E17		
F	F7	F8	F9	F10	F11	F12	F13	F14	F15	F16	F17	F18	F19
G	G7	G8	G9	G10	G11	G12	G13	G14	G15	G16		G18	G19
H	VSS	H8	VSS			VSS	H13	H14	VSS				
J						VCCIO	J13		VCCIO				
K						VCCIO	K13		VCCIO				
L							L13	L14					
M							M13	M14	M15	M16	M17	M18	M19
N											N17	N18	N19

VSS	Dedicated VSS Pin
VCCIO	Dedicated VCCIO Pin
C13	Victim Pin
	Unused Pins
	Aggressor I/Os

As a result of locating aggressor pins away from the victim pin C13, different common return paths are utilized, lowering the level of quiet low noise and quiet high noise coupled onto C13. Table 3 gives the observed quiet low noise and quiet high noise when the unused I/Os for the configurations above were set to VSS, respectively.

Table 3. QLN/QHN for SSO Aggressors Near To and Away From Victim Pin

Scenarios II and III	Quiet Low Noise (mV)	Quiet High Noise (mV)
Aggressors next to the victim (Figure 9)	386	488
Aggressors far away from the victim (Figure 10)	96	393

Locate Victim and Aggressor Pins Closer to Power and Ground Pins

Locating the aggressor I/Os near the power and ground pins adds more localized return paths, thereby reducing the signal-to-return inductance loop and the overall effective loop inductance. The aggressor pin locations are critical for the current return path; the closer to power and ground pins, the better.

Impact of Multiple I/Os Switching on Jitter

Under specific conditions, there is a negligible impact on TX output jitter due to the aggressor I/O toggling. The difference in the output jitter on TX0, TX10 is negligible when measured under the two following scenarios.

Scenario I

In Scenario I, the output jitter on TX0 and TX10 was measured when the nearest neighboring single ended I/O banks (Bank4, Bank 7) are quiet.

Scenario II

In Scenario II, the output jitter on TX0 and TX10 was measured when the nearest neighboring single-ended I/O banks (Bank4, Bank 7) are toggling with 2.5 V LVTTL 16 mA drive strength.

Figure 11 and Figure 12 show the victim and aggressor I/O locations in the EP2SGX90EF1152 device.

Figure 11. Bank 4 Configuration for Jitter Measurement

	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
A	TX1P	TX1N	VSS	A7	A8	A9	A10	A11	A12	A13	A14	A15	A16	A17	A18	A19
B			VSS	VSS	B8	VSS	B10	B11	VSS	B13	B14	VSS	B16	B17		
C	TX0P	TX0N	VSS	C7	C8	C9	C10	C11	C12	C13	C14	C15	C16	C17	C18	
D			VSS	D7	D8	D9	D10	D11	D12	D13	D14	D15	D16	D17	D18	
E			VSS	VSS	E8	VSS	VREF	E11	VSS	VREF	E14	VSS	VREF	E17		
F			VSS	F7	F8	F9	F10	F11	F12	F13	F14	F15	F16	F17	F18	F19
G			VSS	G7	G8	G9	G10	G11	G12	G13	G14	G15	G16		G18	G19
H			VSS	VSS	H8	VSS			VSS	H13	H14	VSS				
J	TX2P	TX2N	VSS						VCCIO	J13		VCCIO				
K			VSS						VCCIO	K13		VCCIO				
L	TX3P	TX3N	VSS							L13	L14					
M			VSS							M13	M14	M15	M16	M17	M18	M19
N			VSS											N17	N18	N19

Victim
Aggressor I/Os

Figure 12. Bank 7 Configuration for Jitter Measurement

AA	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
AB			VSS													
AC			VSS										AB16			
AD	TX9P	TX9N	VSS							AC13	AC14	AC15	AC16	AC17		
AE			VSS							AD13	AD14	VSS				
AF	TX8P	TX8N	VSS						VCCIO	AE13		VCCIO				
AG			VSS	VSS			AF10		VCCIO	AF13		VCCIO				
AH			VSS	VSS	AG8	VSS	AG10	AG11	VSS	AG13	AG14	VSS				
AJ			VSS	AH7	AH8	AH9	AH10	AH11	AH12	AH13	AH14	AH15	AH16	AH17		
AK			VSS	AJ7	AJ8	AJ9	AJ10	AJ11	AJ12	AJ13	AJ14	AJ15	AJ16	AJ17		
AL			VSS	VSS	AK8	VSS	VREF	AK11	VSS	VREF	AK14	VSS	AK16	VREF		AK19
AM	TX10P	TX10N	VSS	AL7	AL8	AL9	AL10	AL11	AL12	AL13	AL14	AL15	AL16	AL17	AL18	AL19
AN			VSS	AM7	AM8	AM9	AM10	AM11	AM12	AM13	AM14	AM15	AM16	AM17	AM18	
AP	TX11P	TX11N	VSS	VSS	AN8	VSS	AN10	AN11	VSS	AN13	AN14	VSS	AN16	AN17		
			VSS	AP7	AP8	AP9	AP10	AP11	AP12	AP13	AP14	AP15	AP16	AP17	AP18	AP19

Victim
Aggressor I/Os

Guidelines for I/O Placement When Using Transceivers

A simple way to minimize the impact of aggressor I/O toggling on TX output jitter is to adhere to the following two rules:

- Do not use the closest I/Os next to the TX pins for signals toggling at high frequencies. These I/Os can be used for static signals.
- Follow the board design practices listed in "High-Speed Board Design Guidelines".

Guidelines to Minimize Impact of I/O Toggling on PLL Jitter

The PLL power pin VCCA_PLL can be affected by the surrounding I/Os when they switch simultaneously. The noise coupling between VCCA_PLL and the adjacent I/Os is through mutual inductance. Each VCCA_PLL pin has eight neighboring pins, as shown in Figure 13. You can assign a maximum of four neighboring I/Os while the remaining pins are dedicated GND or PWR pins. Some VCCA_PLL pins have fewer than four I/Os around them.

Figure 13. Typical VCCA_PLL Pin Placement

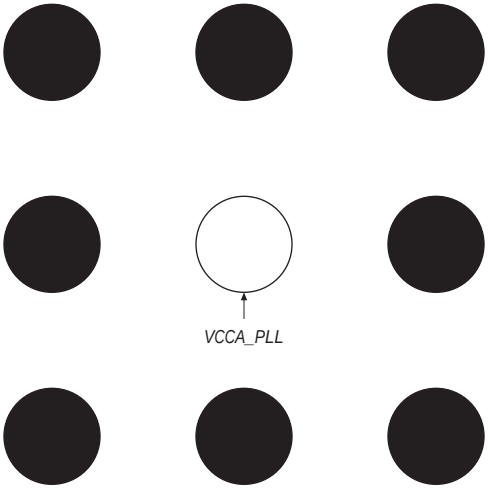

Table 4 shows the list of user I/Os around the analog PLL power pin in an EP2SGX90EF1152 device.

Table 4. EP2SGX90EF1152 VCCA_PLL Pin Neighboring I/O Pin Map (Part 1 of 2)

PLL Supply		Neighboring User I/Os			
Pin Name	Pin Number				
VCCA_PLL1	T25	U24	-	-	-
VCCA_PLL2	V25	U24	-	-	-
VCCA_PLL5	K16	-	-	-	-
VCCA_PLL6	AE16	-	-	-	-
VCCA_PLL7	J26	H26	H25	J27	K27
VCCA_PLL8	AF26	AF27	AE27	AG26	AG25

Table 4. EP2SGX90EF1152 VCCA_PLL Pin Neighboring I/O Pin Map (Part 2 of 2)					
PLL Supply		Neighboring User I/Os			
Pin Name	Pin Number				
VCCA_PLL12	K17	-	-	-	-
VCCA_PLL12	AE18	-	-	-	-

PLL I/O Placement Recommendation

A set of three simple placement rules can help minimize the noise coupling onto the VCCA_PLL pins. These rules are:

- Do not use any neighboring user I/Os for switching signals except used as LVDS I/Os or inputs.
- Always program unused I/Os to GND and connect them to board GND.
- Use lowest current strength for aggressor I/O pins as appropriate.

High-Speed Board Design Guidelines

The following sections provide a brief overview of some common high-speed board design practices. PCB design plays a key role in the amount of noise that is observed on the victim pin. A good PCB stackup that provides controlled impedance traces along with high-speed return paths, coupled with a robust power delivery solution, helps in minimizing the observed noise on the victim pin. Refer to "References" for more in-depth coverage on the topics covered below.

Stackup Design

Stackup design refers to the layer arrangement in a printed circuit board. The placement of power and ground planes in the PCB stackup (determined by layer order) has a significant impact on the loop inductances and interplaner capacitance between the power and ground planes. PCB designers need to consider the stackup design in the early stages of the design cycle, placing difficult to bypass power rails on the top half of the board and low-transient current rails in the middle to lower part of the board with the assumption that the FPGA device is on the top layer.

Where possible, VCCIO should have a ground plane adjacent to it in the stackup to minimize spreading (loop) inductance. The PCB designer must determine which VCCIO and ground plane pairs have high or low priority based on the transient current requirements.

Ideally, the board stackup should allow the PCB designer to route high-speed traces referenced to AC ground on each side, above and below. A more realistic configuration would be to have a ground plane above with a VCCIO (of the signal) plane below or vice versa. Plane cutouts should be done in such a manner that provides good return path for high-speed traces.

High-Speed Return Paths

When designing a PCB, it is important to keep in mind that the physical characteristics of the current return path are just as important as that of the signal trace. A very common mistake in PCB design is to put extra emphasis in designing a controlled impedance signal trace with no thought whatsoever on the current return path. At high frequencies, any signal launched onto a trace from an output buffer returns to the source through the path of least impedance, which in most cases is the path of least inductance. The most fundamental effect of a discontinuity is an effective increase in series inductance. The extra inductance filters out some of the high frequency components of the signal and degrades the edge rate and round the corners. Another effect of this specific non-ideal return path is a very high coupling coefficient between the traces traversing the same gap.

When designing a PCB, it is always a good design practice to provide a continuous reference plane for all the high-speed traces. If traversing a gap in the reference plane becomes unavoidable, you can minimize the effect by placing decoupling capacitors on both sides of the signal line to provide an AC short across the gap. If the high-speed trace is routed on multiple layers, and is referencing the same reference plane (either power or ground plane) on multiple layers, take care to provide enough stitching vias between the multiple layers of the reference plane, where the high-speed trace changes layers to provide a reliable return path.

Note that changing the reference planes, the high speed trace referencing power plane on one layer, and the ground plane on an other layer, should always be avoided. If it is absolutely necessary, provide sufficient decoupling near the layer change to minimize the return current path.

Brief Overview of Power Delivery Network Design

The goal of the PDN (Power Delivery Network) design is to provide stable and uniform voltages for all devices on the PCB. It is important to have a stable supply voltage since any fluctuations in the reference and or supply voltages significantly affects the timing and signal integrity of the individual components.

There are essentially two levels of power delivery that must be accounted for when designing a PDN. The high frequency component of the power delivery system must supply the instantaneous current demanded by the device. This can be achieved by providing maximum on-chip capacitance along with placing capacitors on package or close to the device on the PCB. The second tier of decoupling is required to replenish the charge of the high frequency caps that are placed closed to the device. The bandwidth requirements of the second tier caps do not need to support the full di / dt requirements of the device; they only have to support a bandwidth high enough to recharge the high frequency capacitors that are near the device before they are required to supply current to the device.

When designing the PDN, remember to minimize the inductive path between the decoupling capacitors and the device (load). When designing a PDN, it may be required to have several tiers of decoupling capacitors to hit the target impedance across a wide range of frequency. This involves first determining the maximum transient current and maximum percent voltage ripple for the power supply. Since the voltage level on that power supply is already known, the target impedance can be calculated using the equation below:

$$Z_{(TARGET)} = \frac{VoltageRail \times \%Ripple}{MaxTransientCurrent}$$

The target impedance for any given power rail is calculated because the voltage regulators are only effective at the lower frequencies. To maintain power integrity throughout the entire frequency range of operation, the PDN relies on the high / mid / low frequency decoupling of the power rail using decoupling capacitors and inter-plane capacitances (capacitance from the power-ground sandwich in the board stackup). By selecting the correct components, you can achieve the impedance profile that meets the target impedance over the desired frequency range of operation.

The effectiveness of any decoupling capacitor is limited by the ESL (Equivalent Series Inductance) of the cap, along with the mounting (PCB via) inductance and the lateral inductance, also known as “spreading” or “routing” inductance (distance of the cap with respect to the device). Minimize the mounting inductance by using layout techniques which includes via diameter selection, via location, power plane distance, and via-to-pad distance.

Apply these techniques to any size capacitor regardless of its physical dimensions and electrical characteristics. In practice, small capacitors (0402, 0603, 0805, and 1206) which are located close to the point of load do not require more than one via per pad. Adding more capacitors reduces the number of routing channels out of the device. The larger decoupling capacitors are located further from the load and can have multiple vias per pad in order to minimize the mounting inductance.

Summary

This application note outlines some of the strategies you can follow to maximize signal margin when multiple I/Os are toggling simultaneously. To summarize:

- Select the minimum drive strength that meets the signal quality and timing requirements. Signal integrity and timing analysis are necessary since the output edge degradation of drivers with reduced current strength reduces the timing margins.
- Provide programmable power and ground wherever possible since they provide additional return paths.
- Provide terminations as dictated by memory interface design guidelines or I/O standard termination requirements.
- Use low voltage signaling standards since they reduce the current di/dt requirement in the output buffer, thereby reducing the quiet low noise and quiet high noise.
- Minimize the impact on transceiver output jitter by designing the pinout in neighboring banks in such a way so as to minimize having the aggressors next to the transceiver channel.
- Designing a PCB with a continuous return path for every high-speed signal.
- Avoid changing routing layers when routing high-speed signals. If the layer change must occur, provide the ground stitching vias in close proximity to the signal via so as to minimize the signal-to-return loop inductance.
- Minimize the mounting inductance of the decoupling caps during PDN design for effective decoupling.

References

1. S.H. Hall, G.W. Hall, J.A. McCall, *High-Speed Digital System Design*.

2. “FPGA Design for Signal and Power Integrity” DesignCon 2007 Paper by Larry Smith and Hong Shi.

Document Revision History

Table 5 shows the revision history for this application note.

Table 5. Document Revision History

Date and Document Version	Changes Made	Summary of Changes
August 2007 v1.0	Initial Release	N/A

101 Innovation Drive
San Jose, CA 95134
www.altera.com
Technical Support:
www.altera.com/support
Literature Services:
www.altera.com/literature

Copyright © 2007 Altera Corporation. All rights reserved. Altera, The Programmable Solutions Company, the stylized Altera logo, specific device designations, and all other words and logos that are identified as trademarks and/or service marks are, unless noted otherwise, the trademarks and service marks of Altera Corporation in the U.S. and other countries. All other product or service names are the property of their respective holders. Altera products are protected under numerous U.S. and foreign patents and pending applications, maskwork rights, and copyrights. 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 Corporation. 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.

I.S. EN ISO 9001