Cyclone III Simultaneous Switching Noise (SSN) Design Guidelines

December 2007, ver. 1.0 Application Note 508

Introduction

Low-cost FPGAs designed on 90-nm and 65-nm process technologies are made to support high performance applications with system clocks and interfaces such as DDR2 pushing 200 MHz and LVDS applications up to 840 Mbps. The device is designed with fast edge rates which can cause signal integrity problems such as simultaneous switching noise (SSN) which may limit system performance and affect circuit operation. These problems can exist even at low data rates as the signal edge rate remains the same regardless of the data rate. For high performance applications such as DDR2 interfaces, Altera® has largely done the work by pre-assigning DQ and DQS signals and accounting for SSN when setting performance limits. For other applications, it is up to the designer to determine the pin placement and I/O settings to optimize performance. By following some simple guidelines and best practices for device settings, pin-out selection, and PCB design, you can avoid many of the signal integrity problems for your designs.

This application note provides a framework to describe SSN and understand the sources of SSN, discusses ways to mitigate SSN for Cyclone® III FPGAs by using I/O settings and selecting proper I/O standards, and provides guidelines on PCB design that are good practice for general high speed digital designs. To demonstrate the effects of the various recommendations, characterization data measured on Cyclone III devices is shown throughout.

SSN Terminology

Definitions of Key Terms used in this document:

Victim: The victim pin is the pin that is affected by SSN noise, where the SSN measurements are taken.

Aggressors: The aggressors are the switching I/Os that cause SSN noise on the victim pin.

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 driven to logic zero (low), in the absence of any aggressor I/Os toggling.

Quiet Low Noise (QLN): The Quiet Low Noise is the noise coupled onto the victim driven quiet low, which is sharing the VCCIO/GND return path with aggressor(s) switching.

Quiet High (QH): The voltage level on a victim net is the observed voltage level at the far end of the victim net 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 coupled onto the victim driven quiet high, which is sharing the VCCIO/GND return path with aggressor(s) switching simultaneously.

VIL_{Max}/VIH_{Min} (DC): For a given receiver, the VIL/VIH (DC) are the values that determine the final logic state unambiguously. When the receiver input crosses these values, the receiver logic state will change and maintain the new value.

di/dt: di/dt is the rate of current change over time, also referred to as the current slew rate.

Figure 1 describes the definitions for QH/QL, QHN/QLN, and VIL/VIH.

Figure 1. QH / QL / QHN / QLN Definitions

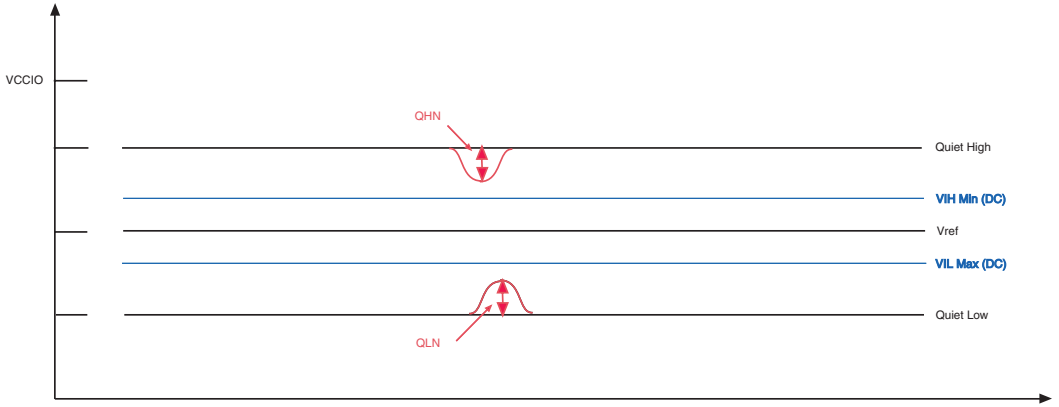

SSN Mechanisms and Metrics

A high level of signal and power integrity performance is necessary for FPGAs required to run DDR memory interfaces and communicate over high-speed serial links. Performance goals for simultaneous switch noise (SSN) are the primary metrics. There are three major mechanisms that contribute to SSN: inductive crosstalk, delta-I noise and power supply compression[1].

For more information about signal and power integrity refer to the document *FPGA Design for Signal and Power Integrity*.

SSN noise is generated when all drivers switch concurrently. Mutual coupling from aggressor signals to victim and delta-I noise associated with the inductance of power and ground paths are the primary mechanisms that cause noise during the rise/fall time of the aggressors. The di/dt of the aggressors is responsible for this noise. Both horizontal structures (transmission lines and planes) and vertical structures (wires, balls, and vias) contribute to SSN crosstalk. Figure 2 shows an example of the important circuit components for discussion of the mechanisms and the resulting SSN.

Figure 2. Diagram of Circuit Components Including FPGA and PCB Load Topology

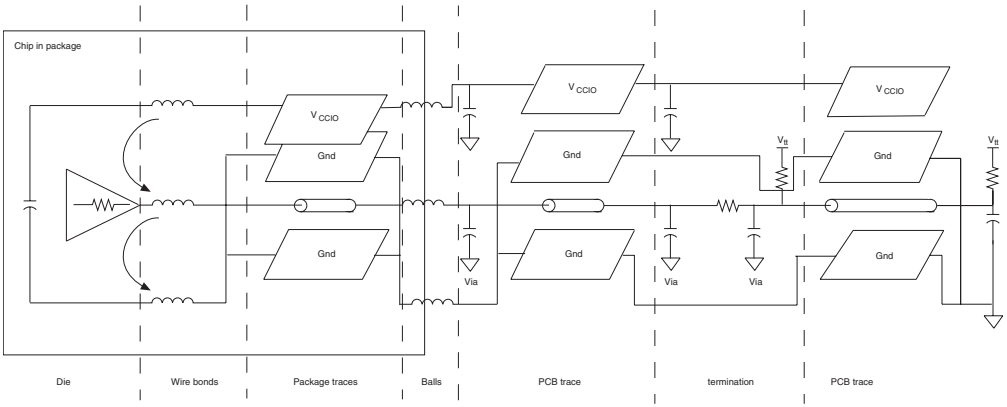

Traditionally, SSN is quantified by absolute voltage of noise. It works for any individual I/O standard. However, FPGAs are programmable. It is difficult to predetermine how the circuits will be used. I/O drivers may be utilized in different I/O standards, such as SSTL, HSTL, LVTTL, CMOS, and LVDS. Therefore, an alternative metric is necessary to specify performance goals for any bus interface that the FPGA device might drive. In this document, QHN and QLN are quantified by the percentage of the signal margin, as well as the traditional absolute voltage of noise when necessary.

Inductive Crosstalk

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

$$V = \frac{m \times di}{dt}$$

Where *m* is the mutual inductance between the aggressors and the victim. Inductive coupling does not happen when the driver current is constant. It only happens when 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. The magnitude of inductive coupling is proportional to the parallel length of the aggressor and the victim signals.

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 conductors. When many drivers switch from high to low, signal current enters the signal pins and must exit through the ground pins. Similarly, when these drivers switch from low to high, the current must come in through the power pins. The governing equation is:

$$V = \frac{L \times di}{dt}$$

Where L is inductance associated with the current path, usually dominated by the self inductance of the power or ground path. 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 when the driver current is constant because there is no di/dt to generate the noise.

Power Supply Compression

Power supply compression noise is the variation of the voltage between local power and ground within the device. This voltage difference between power and ground enables the circuits to perform their tasks and the variation can affect the silicon performance.

When an output buffer changes its state, the output structure will momentarily construct a low impedance path from power supply rail to ground. The output transition causes the output to charge or discharge, requiring the current be immediately available on the output load to reach the required voltage level. Local capacitance provides the stored energy required for this current transient. This lowers the local voltage and brings in current from outside. Current flowing through the capacitance generates the voltage variation. The governing equation is:

$$V = \frac{1}{C} \times \int Idt$$

Power supply compression does not happen instantaneously but occurs long after the driver has made the transition, as long as current flows through the capacitance.

Descriptive SSN Metrics

The traditional way to quantify SSN is "mV of noise". It works for one signaling technology. For different I/O standards, an absolute voltage cannot indicate how much the noise impacts the signal. For example, a 2.5V LVTTL bus can tolerate more mV of noise than a 1.8V SSTL bus.

FPGAs are programmable. The I/O interface that a driver can be used for is unpredictable before the device is programmed. Therefore, an FPGA needs a set of SSN metrics that can apply to different bus technologies, power supply voltages and drive strengths. One possibility is to quantify SSN as a percentage of the signal margin.

The signal margin is the difference between Quiet High and VIH min (for QHN) or the difference between VIL max and Quiet Low (for QLN). Incoming signals to the inputs of receivers are guaranteed to be correctly interpreted as a 0 or 1 when the voltage is within the signal margin. Figure 3 shows a signal eye diagram with QLN/QHN quantified in percentage of the signal margin for a typical SSTL class II signal. Systems should be designed such that a substantial amount of signal margin is left for other requirements after accounting for SSN.

Figure 3. Signal Eye Diagram with QLN/QHN Quantified in Percentage of Signal Margin

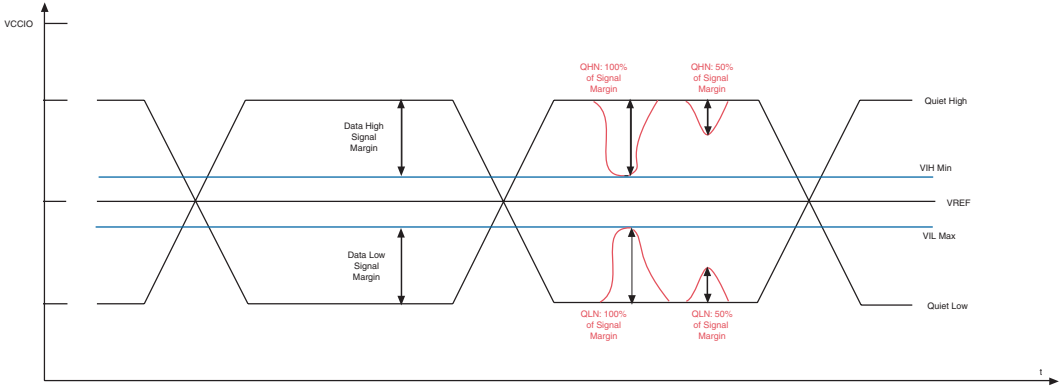

Synchronous and Asynchronous Signals and SSN Glitches

For synchronous signals, there is a period of time called the sampling window, usually a setup and hold time around the sampling clock, where incoming signals are expected to be stable within the signal margin zone. During this sampling window, a functional error may occur if SSN causes the incoming signal to exceed the signal margin. SSN glitches out of the sampling window do not cause functional failures because the receiver is not sensitive to noise glitches during this time period.

For asynchronous signals, there is no sampling window. Any SSN glitches that causes the signal to break the signal margin, will affect the circuitry controlled by this asynchronous signal. To guarantee the circuitry performance controlled by the asynchronous signals, it is critical to shield the asynchronous control signals from SSN glitches.

Figure 4 shows the effect of SSN glitches on synchronous signals and asynchronous signals.

Figure 4. Synchronous and Asynchronous Signals and SSN Glitches

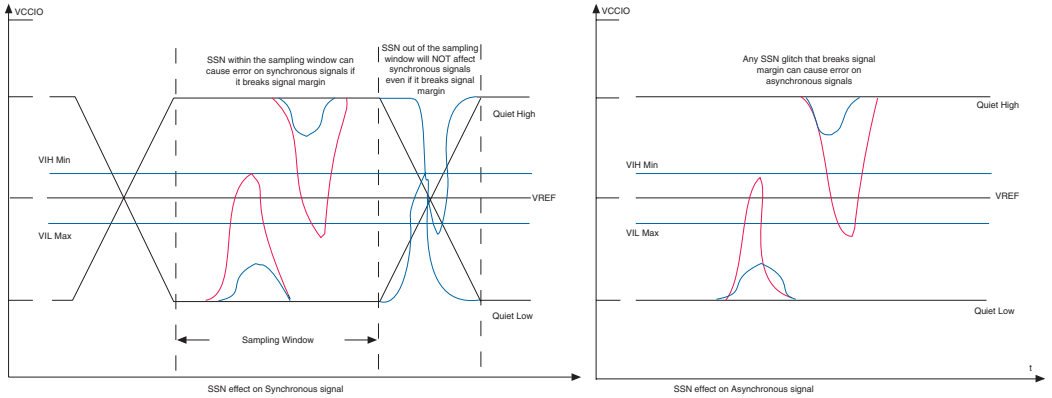

Forward and Reverse Spiral Pattern

When measuring SSN vs. the number of aggressors, we can describe the noise in a forward or reverse spiral pattern. The shape of the forward and reverse spiral aids in identifying the dominant mechanism of SSN, and in determining the sensitivity of victim location with respect to the switching aggressors. The patterns add or remove switching I/Os around the victim in a specific order. Figure 5 shows an example of the on-die locations of the spiral binary distribution pattern of unused I/Os and switching I/Os.

Figure 5. Spiral Binary Switching Pattern of the Switching I/Os and Unused I/Os (On Die)

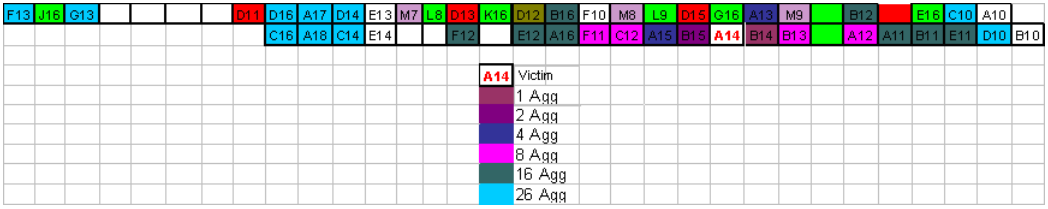

Figure 6 shows the on-package locations of the same spiral binary distribution pattern.

Figure 6. Spiral Binary Switching Pattern of the Switching I/Os and Unused I/Os (On Package)

	10	11	12	13	14	15	16	17	18		A14	victim
A		A11	A12	A13	A14	A15	A16	A17	A18			Aggressors = 1
B		B11	B12	B13	B14	B15	B16					Aggressors = 2
C	C10		C12		C14		C16					Aggressors = 4
D	D10				D14		D16					Aggressors = 8
E		E11	E12									Aggressors = 16
F		F11	F12	F13								Aggressors = 26
G				G13								
H												

The forward spiral pattern starts with no aggressor and gradually increases to the maximum number of aggressors, starting from the pins closest to the victim pin. The reverse spiral pattern starts with the maximum number of aggressors and gradually decreases to no aggressor, turning off the pins closest to victim pin first.

Figure 7 shows a measurement sample of forward and reverse spiral SSN in percentage of signal margin. The blue curve is the SSN measurement for the forward spiral pattern. The blue arrow shows the direction in which the number of aggressors increases from 1 to 26. The pink curve is the SSN measurement for the reverse spiral pattern. The pink arrow shows the direction in which the number of aggressors decreases from 26 to 1. The figure shows that the noise measured on the victim pin is very sensitive to the location or proximity of the switching aggressors. Four switching I/Os in the forward spiral pattern consume as much signal margin as 22 switching I/Os in the reverse spiral pattern. In this example, by placing a victim pin 4 pads away from the switching aggressors, you are able to switch 18 more I/Os while maintaining the same signal margin at the victim pin.

Figure 7. Forward and Reverse Spiral SSN in Percentage of Signal Margin

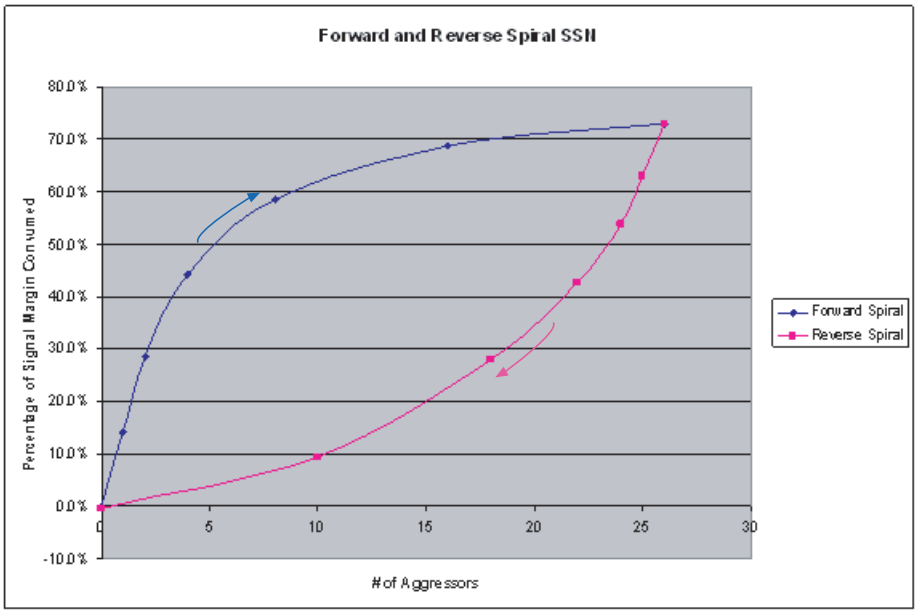

p

Package effect on SSN

The Cyclone III device family uses a wire-bond package, a very common low-cost chip-to-substrate interconnection technology. The wires connecting on-die pad and substrate pad are long while the space between them is very small. These result in large self inductance along the wires and mutual inductance between the wires, which are the major contributors to SSN. Therefore, cross-talk and delta-I noise are the dominant SSN factors and power supply compression is the secondary SSN factor. Figure 8 shows a typical cross-section and side view diagrams of wire bond assembly.

Figure 8. Wire Bond Assembly Diagram

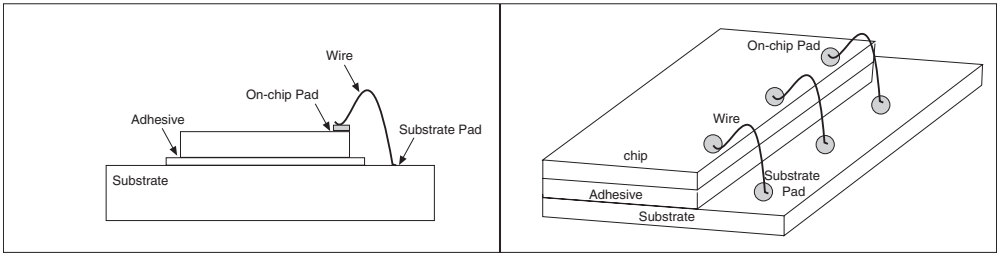

SSN Mitigation with Cyclone III

While SSN is an ever-increasing design challenge, there are several options designers can use to reduce SSN. Due to the package effect, inductive coupling and delta I noise dominates the SSN for Cyclone III devices. This section provides several methods to mitigate SSN sourced from these two mechanisms.

Inductive coupling and delta-I noise occur during the signal transition, where di/dt is the cause of SSN. The amount of SSN is determined by effective inductance, including self inductance and mutual inductance. Consider the following equation:

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

Where:

- V is the SSN voltage.
- N is the number of effective aggressors.
- L_k is the effective inductance from the kth aggressor. For delta-I noise, it is the self inductance. For cross-talk, it is the mutual inductance.
- di_k/dt is the slew rate of the kth aggressor.

Therefore, di/dt, effective inductance, and N are the main concerns for SSN caused by these two mechanisms. The following sections discuss methods that work on each of them to minimize SSN.

Reduce di/dt

di/dt is the change in current over time, which represents how fast the current in the I/O driver switches. To reduce *di/dt*, the designer can either reduce *di*, the amount of current switching, or increase *dt*, the transition time, or affect both *di* and *dt*.

Select Lower Drive Strength

The drive strength setting determines the amount of current the driver will draw. Higher drive strength setting means faster current change when I/O switches. Therefore, limiting the drive strength helps minimize the I/O switching current change and reduce SSN.

Figure 9 shows the QLN and QHN measurements taken on LVTTL 3.0V at different drive strength: 4mA and 16mA. From 16mA to 4mA, the QLN in percentage of signal margin gains up to 60% and the QHN in percentage of signal margin gains up to 37%.

Figure 9. QLN and QHN in Percentage of Signal Margin on 3.0V LVTTL at Different Drive Strength

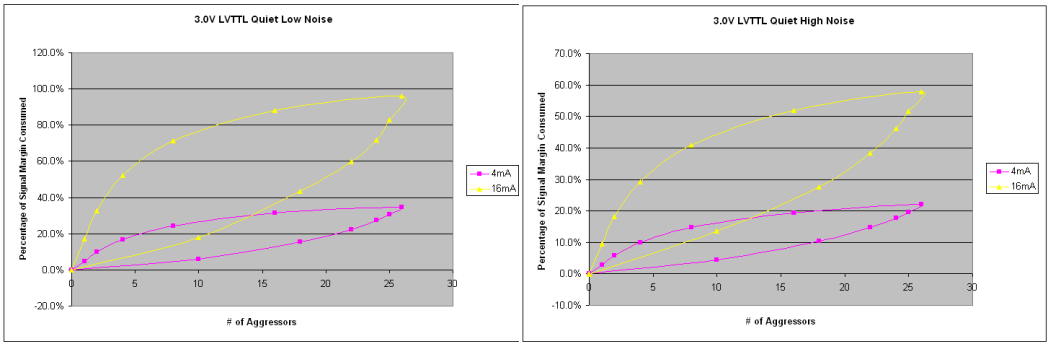

Figure 9 shows how sensitive SSN is to the drive strength. Reducing the drive strength can significantly reduce the noise and improve the signal margin.

In practical designs, the system speed and timing margin usually determine the minimum drive strength. To minimize the current change at I/O switching, select the lowest drive strength that meets the performance requirements.

Utilize Programmable Slew Rate Control Settings

Slew rate is the rate of change at output, or the amount of time it takes to switch from 10% to 90% of its final value. Figure 10 shows the definition of slew rate.

Figure 10. Definition of Slew Rate

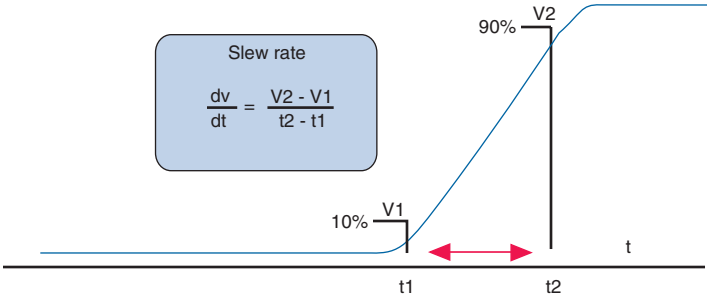

The Cyclone III device family supports programmable slew rate control. The output buffer for each I/O pin provides three output slew-rate settings for each supported I/O standard: slow, medium, or fast. A faster slew rate provides high-speed transitions for high-performance systems, but may introduce noise into the system. A slower slew rate reduces system noise, but adds a nominal delay to rising and falling edges.

Figure 11 shows the SSTL-18 CI 12mA QLN and QHN at different slew rates. By switching slew rate from fast to slow, the QLN in percentage of signal margin gains up to 20% and the QHN in percentage of signal margin gains up to 40%.

Figure 11. SSTL-18 CI 12mA QLN and QHN at Different Slew Rates: Fast, Medium, and Slow

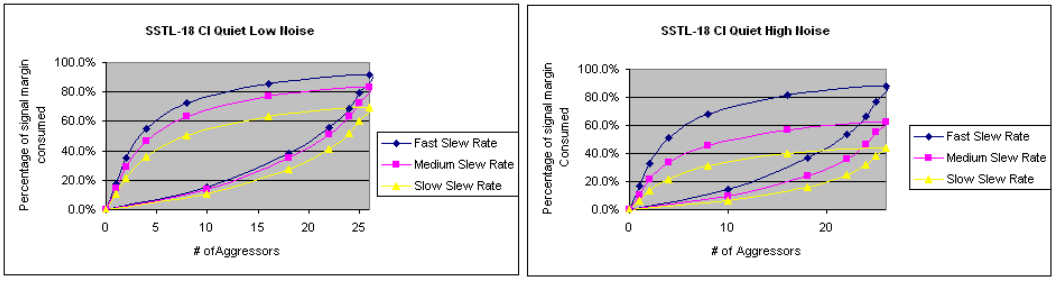

Since each I/O pin has an individual slew-rate control, designers can specify the slew rate on a pin-by-pin basis. In practical design, select the slowest slew rate that meets the system speed requirement.

Select Low Voltage I/O Standards

Using low-voltage signaling standards can help reduce noise on the quiet victim pin located near numerous aggressor pins. Consider the relationship between the current and output voltage:

Resistive load:

$$I = \frac{V_{OUT}}{R}$$

Capacitive load:

$$I = \frac{C_L \times dV_{OUT}}{d_t}$$

Reducing the output voltage swing at the aggressors can help reduce the current required in the output transistor, thus reduce di/dt at aggressors and eventually SSN voltage at the victim pin. The caveat is that the voltage noise margin will also decrease as signal level decreases.

Figure 12 shows the QLN comparison between SSTL-18 CII 16mA and SSTL-2 CII 16mA. By moving from SSTL-2 CII to SSTL-18 CII, the QLN in percentage of Signal Margin gains no more than 5%, even though the absolute QLN voltage can be different by more than 50mv.

Figure 12. QLN Difference Between SSTL-2 CII 16mA and SSTL-18 CII 16mA

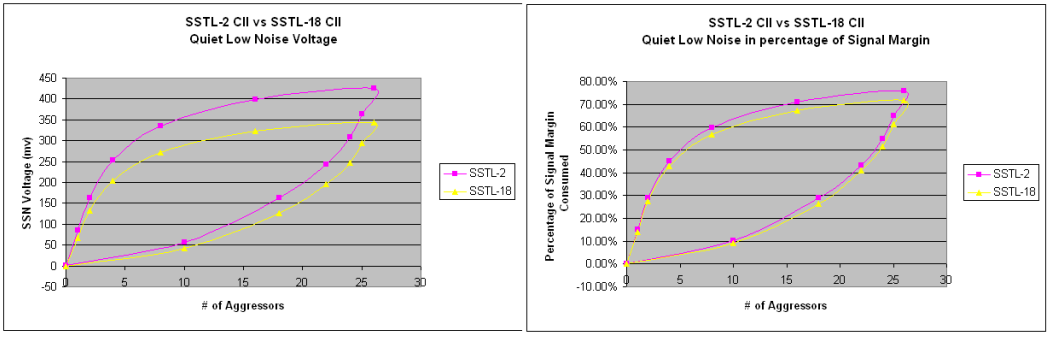

Figure 13 shows the QHN comparison between SSTL-18 CII 16mA and SSTL-2 CII 16mA. By moving from SSTL-2 CII to SSTL-18 CII, the QHN in percentage of signal margin gains no more than 9%, even though the absolute QHN voltage can be different by ~100mv.

Figure 13. QHN Difference Between SSTL-2 CII 16mA and SSTL-18 CII 16mA

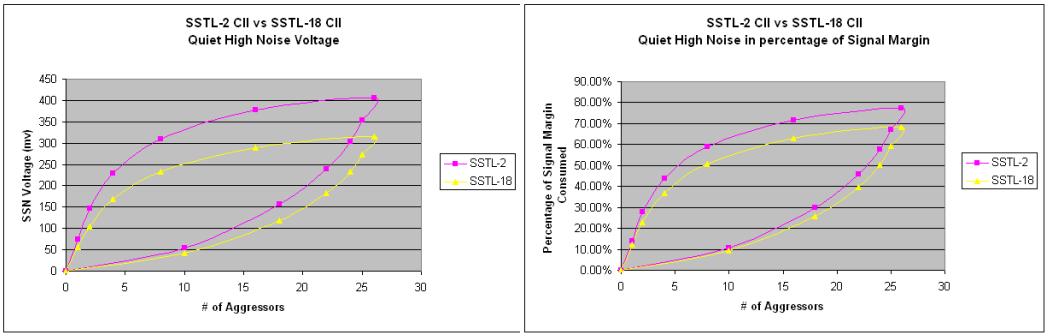

In practical designs, the system designer should also consider additional design criteria such as meeting EMI/EMC specifications before selecting the voltage for certain I/O standards.

Select 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. Without these terminations, the noise launched onto the transmission line will be reflected at the far end due to the impedance mismatch. The amplitude of noise measured at the far end will be higher due to multiple reflections in the absence of terminations. Figure 14 shows the QLN and QHN difference between 1.8V LVTTL 16mA and SSTL-18 CII 16mA. The observed QLN and QHN in percentage of signal margin for SSTL-18 CII are better by up to 60% than those of 1.8V LVTTL

Figure 14. QLN and QHN Difference Between SSTL-18 CII 16mA and 1.8V LVTTL 16mA

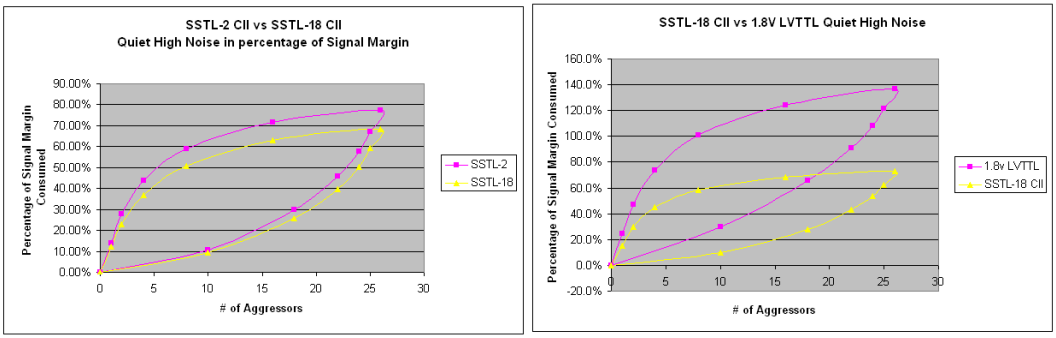

It should be noted that the data shown in Figure 14 was measured on a setup with no far end load which will result in higher SSN compared to real designs. In real systems the far-end load or receiver loads will filter the noise seen at the victim pin, reducing the SSN.

Reduce SSN by Limiting the Number of SSO pins

Another method to reduce SSN is to limit *N*, the number of effective aggressors switching simultaneously in a given bank. Additional aggressors bring additional SSN contributions. Thus limiting *N* reduces overall effective inductance between all aggressors and victim. It also reduces the total amount of current in the inductance loop, which results in reducing overall di/dt. Figure 15 shows the relationship between QLN and QHN and the number of aggressor I/Os, measured on CIII for SSTL-18 CII. Both QLN and QHN increase as the number of SSO pins increase. The saturation in the curve shows additional SSO pins will have negligible effect on the noise.

In practical designs, avoid assigning the highly switching I/Os in the same bank. Spread them into different banks.

Figure 15. SSTL-18 CII QLN and QHN vs. the Number of SSOs

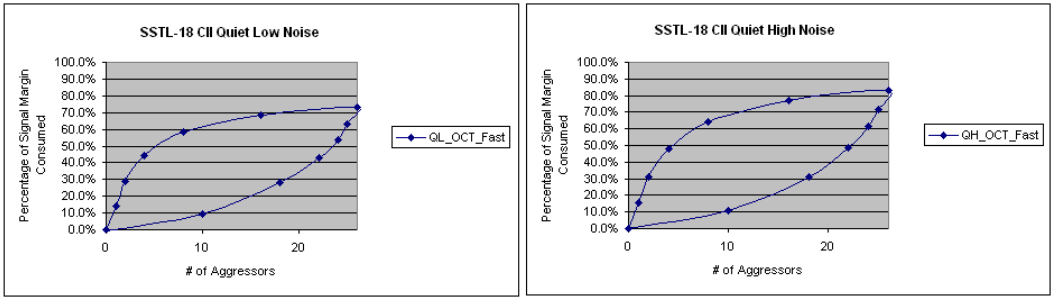

Reduce Effective Inductance

When implementing high speed memory interfaces like DDR/DDR2, the wide memory bus and tight timing budgets will limit the flexibility to reduce the number of SSOs. Another option is to spread out the SSO pins across the bank if possible, when I/O usage in the bank is substantially below 100%. By spreading out the switching I/Os, there are less switching aggressor I/Os in the vicinity of a potential victim I/O. The effective inductance between the aggressor I/Os and the victim I/O decreases and results in reduced SSN.

Note that in Cyclone III devices, when implementing DDR/DDR2, DQ and DQS pins are pre-assigned. Thus spreading I/Os might not be an option for DQ/DQS but only for command or control signals.

Once the unused I/Os are interspersed among the switching I/Os, the status of the unused I/Os will also affect SSN. Hardwiring these unused I/Os to the ground or power plane helps reduce SSN.

Critical asynchronous control signals, e.g., `clock/reset/enable` signals, are vulnerable to SSN as they usually have direct impact on the system performance. To guarantee the system operation is not interfered by SSN, shielding the critical signals away from aggressors is highly recommended. The further the aggressors are from the victim pin, the less the effective inductance is between the aggressor I/Os and the victim I/O. Figure 16 shows the SSN measurements of forward-spiral and reverse-spiral patterns on SSTL-2 CII (16mA with OCT on). In Figure 16, with the same number of SSO pins, there is a significant margin of noise from forward-spiral pattern curve than from reverse-spiral pattern curve. When 10 aggressors are switching, QLN from the forward-spiral pattern

in percentage of signal margin is about 50% more than the reverse-spiral pattern and QHN from the forward spiral pattern in percentage of signal margin is about 65% more.

Note that when the aggressor number is the same, aggressors in the forward-spiral pattern are closer to the victim than those in reverse-spiral pattern. In conclusion, the closer the aggressors are to the victim I/O, the worse the SSN is. Therefore, whenever possible, assign critical and asynchronous control signals away from switching I/Os.

Figure 16. SSTL-2 CII QLN and QHN with Forward / Reverse Binary Switching Patterns

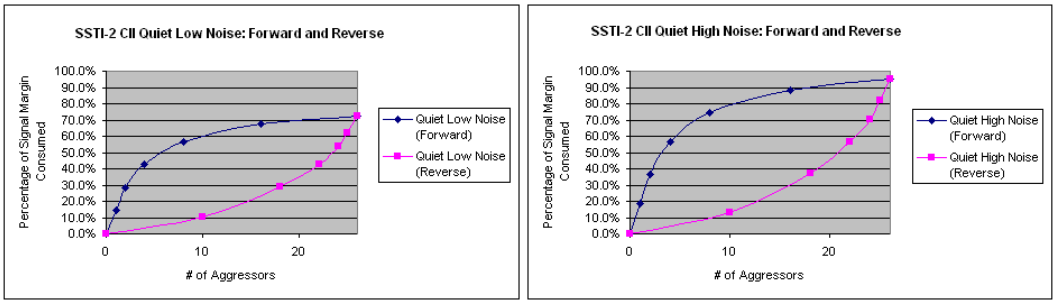

High-Speed Board Design Guidelines

The following sections provide a brief overview of several common high speed board design practices. PCB design plays a key role in the amount of noise 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 minimize noise on the victim pin.

Refer to "Documents Referenced" for more in-depth coverage on stack up design, signal breakout, high-speed return paths, and power delivery network design.

Stackup Design and Signal Breakout

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 inter-plane capacitance between power and ground planes. PCB designers need to consider the layer order in the early stages of the design cycle. Place the difficult-to-bypass power rails on the top half of the stackup and low-transient current rails in the middle to lower part of the board, assuming the FPGA device is on the top layer.

The power planes for power supplies with high transient currents, such as SSTL or HSTL, should be placed close to the device to reduce the Via inductance connecting to the plane. If possible, each power plane should have a ground plane placed adjacent to minimize the loop inductance. The power and ground plane pairs that carry high frequency current transients have higher priority over the ones that carry low frequency current transients. The PCB designer needs to identify the priority of each pair when designing the stackup.

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 is to have a ground plane above with a VCCIO (of the signal) plane below or vice versa. Plane cutouts should be done in a manner that provides good return path for high-speed traces.

High Speed Return Paths

The physical characteristics of the current return path are as important as the signal trace. A very common mistake in PCB design is to put extra emphasis in designing a controlled impedance signal trace without considering the current return path. At high frequencies, any signal launched onto a trace from an output buffer will return 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 increase in series inductance. The extra inductance can filter out high frequency components, degrade the edge rate and round the corners. Another effect is a very high coupling coefficient between the traces traversing the same gap.

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 is inevitable, place decoupling capacitors on both sides of the signal line to provide an ac short across the gap to minimize the effect.

Changing reference planes should always be avoided for high speed signals. In the event that it happens, sufficient decoupling should be provided near the layer change to minimize the impedance of 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 the devices on the PCB. It is important to have a stable supply voltage, since any fluctuations in the reference and/or supply voltages will significantly affect the timing and signal integrity of the individual components[2].

There are essentially two levels of power delivery that must be considered 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 close to the device. The bandwidth of the second tier capacitors does 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[1]. When designing a PDN, remember to minimize the inductive path between the decoupling capacitors and the device (load).

The reason to have several tiers of decoupling capacitors is to meet the target impedance across a wide range of frequency. Once the maximum transient current[3] and maximum percent voltage ripple for the power supply are available, the target impedance can be calculated using the following equation:

$$Z_{TARGET} = \left[\frac{VoltageRail \bullet PercentofRipple}{MaxTransientCurrent}\right]$$

The target impedance for any given power rail is calculated to maintain power integrity throughout the entire frequency range, as the voltage regulators are only effective at the lower frequencies. The PDN relies on the high/mid/low frequency decoupling of the power rail using decoupling capacitors and inter-plane capacitances, i.e. capacitance from the power-ground sandwich in the board stackup.

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 with layout techniques such as via diameter selection, via location, power plane distance and via-to-pad distance.

Apply this technique to any size capacitor regardless of its physical dimensions or electrical characteristics. In practice, small capacitors (0402, 0603, 0805 and 1206) 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.

Conclusion

This application note described 3 SSN mechanisms: inductive coupling, delta-I noise and power supply compression, and presented the metrics to quantify SSN with apercentage of signal margin consumed. Based on the dominant SSN mechanisms, it provides some recommendations for designers to mitigate SSN on Cyclone III designs. Following is a summary of these methods:

- Reduce *di/dt*
- Select lowest drive strength that meets the system speed and timing margin
- Select slowest slew rate that meets the system speed requirement
- Select low voltage I/O standards that meets the design criteria
- Select I/O Standards with termination
- Limit the number of SSO pins
- Reduce effective inductance by spreading SSOs

It also provides several high speed board design guidelines for stackup design, high speed return paths and PDN designs.

By applying the SSN mitigation methods and following the high speed board design guidelines, you can significantlyreduce the risk of SSN issues at the early stage of the design cycle.

Documents Referenced

This application note references the following documents:

- *FPGA Design for Signal and Power Integrity*
- *High-Speed Digital System Design*
- *Power Play Early Power Estimator User Guides*
- *RIGHT THE FIRST TIME, A Practical Handbook on High Speed PCB and System Design*

Revision History

Table 1 shows the revision history for this document.

Table 1. Document Revision History

Date and Document Version	Changes Made	Summary of Changes
December 2007 v1.0	Initial Release	—

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