DesignCon 2007

FPGA Design for Signal and Power Integrity

Larry Smith, Altera Corporation

Hong Shi, Altera Corporation

CP-01023-1.0
January 2007

Abstract

FPGAs have traditionally been optimized for low-cost environments where signal and power integrity are minor considerations. With today's requirements for high-speed memory and serial interfaces, FPGA silicon and packages must be designed to provide good signal and power integrity while still maintaining cost objectives. Performance goals for near- and far-end simultaneous switch noise (SSN) noise as well as power supply quality are the primary metrics.

Both near- and far-end noise is generated when all drivers switch concurrently creating SSN. Mutual coupling from aggressor signals to victims and delta-I noise associated with the inductance of power and ground paths are the primary mechanisms that cause noise during the rise 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 (bumps, balls, and vias) contribute to SSN crosstalk. The vertical structures are responsible for most of the noise. Neither near- nor far-end noise should exceed 50 percent of the signal noise margin. This goal is achieved by controlling the signal I/O to power and ground ratios.

The quality of the power supply, seen by the circuits on the die, is important for proper circuit performance and the ability to meet timing and jitter specifications. The power distribution network (PDN) for the package die combination is an important consideration in determining power supply quality. Through the use of on-package decoupling (OPD) capacitance and on-die capacitance (ODC), the power supply voltage tolerance is held to +/-10 percent of nominal throughout the difficult die/package resonance frequency band.

This paper describes the design trade-offs made for the Altera® Stratix® III family of FPGAs to achieve performance while maintaining costs. Measured data on design prototypes as well as simulation predictions are presented.

Author Biographies

Larry D. Smith is a principal engineer at Altera Corporation concentrating on SSN noise, power and signal integrity. Prior to joining Altera in 2005, he worked at Sun Microsystems from 1996 to 2005 where he did development work in the field of signal and power integrity. Before this, he worked at IBM in the areas of reliability, characterization, failure analysis, power supply and analog circuit design, packaging and signal integrity. Mr. Smith received a BSEE degree from the Rose-Hulman Institute of Technology and a MS degree in material science from the University of Vermont. He has 13 patents and has authored numerous journal and conference papers.

Dr. Hong Shi is technical lead for the Packaging Technology Department at Altera Corporation. As manager of the Packaging group's electrical design team, Dr. Shi's responsibilities include developing the strategy for high-density and high-performance FPGA packaging, simulating system level electrical performance and establishing chip-package-board interconnect co-design capability. Before joining Altera, Dr. Shi worked

at HP and Agilent Technologies where he was project leader for Agilent's first 40-Gbps digital communications analyzer module. He has published over 30 technical papers in areas of optoelectronics, microwave circuits and digital circuit packages. Dr. Shi obtained his BSEE from Xi'an Jiaotong University, MS Physics from DePaul University, and PhD in electrical engineering from CREOL College of Optics at University of Central Florida.

Introduction

In the last decade, tremendous advances have been made in signal and power integrity, primarily when considering microprocessors, ASICs, and memory. A high level of signal and power integrity performance is now being demanded from FPGAs as they are required to run DDR memory interfaces and communicate over high-speed serial links. The FPGA industry has traditionally been very cost sensitive, but FPGA-based products are now being upgraded to deliver improved signal and power integrity performance while continuing to closely monitor costs.

This paper describes the transition from a Stratix II FPGA to the Stratix III FPGA family. Careful measurements were made on all aspects of signal and power integrity to understand how the previous product performed and the most effective ways to improve upon that performance for the next generation. Software simulations were correlated to hardware measurements and formed the basis for predicting future performance.

Emphasis is placed on understanding signal and power integrity mechanisms, which are the underlying reasons for performance. Three major mechanisms contribute to SSN, including inductive crosstalk, delta-I and power supply compression.

By their very nature, FPGAs are programmable, so it is difficult to predetermine how the circuits will be used. I/O driver circuits may be required to service many different bus applications, including SSTL, HSTL, LVTTL, CMOS, and LVDS. For this reason, a new method for specifying performance goals was developed that applies to any bus interface that the FPGA might drive. SSN is quantified by the percentage of the signal margin as well as the traditional millivolt (mV) of noise.

After identification of underlying mechanisms and careful measurement and simulation, a design point was established for the new generation of Stratix III FPGAs. It involves an increase in the number of power and ground return paths for I/O, and improvement in power quality by adding capacitance in critical positions of the circuit.

Simultaneous Switching Noise Mechanisms

SSN is broken down into three distinct mechanisms: inductive crosstalk, delta-I, and power supply compression. This noise can result in degradation of circuit performance at the near or far ends of the bus, and can create power supply quality problems that impact all circuits on that supply. Figure 1 captures the important circuit components for discussion of the mechanisms and the resulting SSN.

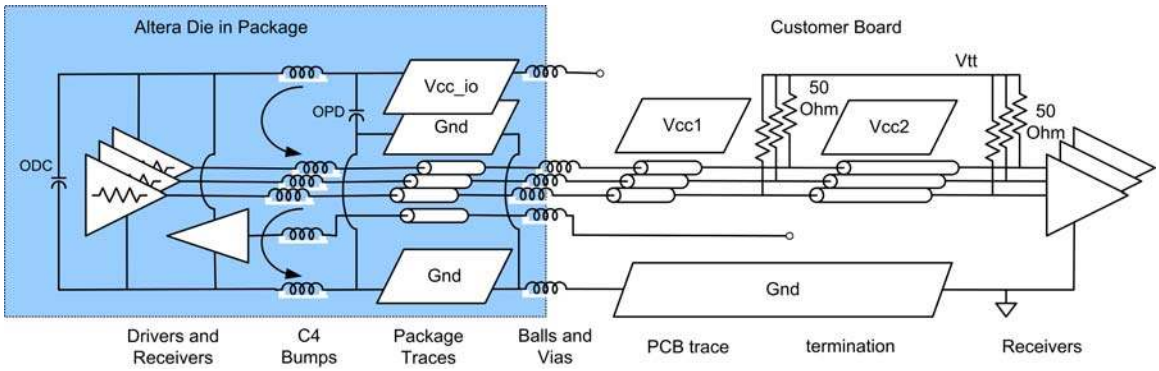

Figure 1. Schematic Diagram of System including FPGA and PCB Load Topology

The blue box represents the FPGA silicon and surrounding electronic package. There may be up to 100 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 flip chip C4 bumps, 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 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 is outside the blue box. Beginning at the PCB pads for surface mount, each has an associated via. PCB vias may connect with signal transmission lines, power, or ground planes within the PCB stack. Typically, transmission lines may be at any level within the PCB. The PDN that supplies power to the I/O circuits on die may also be at any level of the PCB. SSTL class II is the signaling technology represented in Figure 1. This technology places a 50Ω termination resistor near the FPGA driver as well as near the receiver. PCB transmission lines may be referenced to power or ground planes within the stackup.

Inductive Coupling

Inductive coupling is often the dominant mechanism for SSN. This 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 $V = m\, di / dt$. Inductive coupling 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.

Most of the inductive crosstalk occurs in vertical structures rather than horizontal transmission line structures. Horizontal transmission line impedance is controlled by associated reference planes. The distance that significant inductive coupling can travel is limited to approximately three times the height of the trace above the plane, which usually limits the effect of inductive coupling only to the first adjacent signal. However, in the vertical dimension, there are often multiple signal conductors for each return path that provide opportunities for aggressor signals to strongly couple to victims. Examples of vertical coupling structures include the C4 solder bumps, 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 conductors. When many drivers switch from high to low, 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 governing equation is $V = L\, di/dt$ 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 in time frames where the driver current is constant because there is no *di/dt* to generate the noise.

Figure 2 compares LVTTL and SSTL class II bus topologies. For parallel-terminated bus topologies (SSTL), DC current is still flowing long after the SSN event. The SSN event involves two processes: turning off the current in the old (pull-down) field effect transistor (FET) and turning on the current in the new (pull-up) FET. The precondition current causes the voltage on the driver output to switch, simply by turning off the old FET. The new FET does not have to source or sink a significant amount of current until two transmission line time-of-flight delays from the driver to the parallel termination resistors have occurred. Therefore there is a significant difference between the delta-I noise for the series and parallel-terminated busses with LVTTL producing much more delta-I noise than SSTL class II type busses.

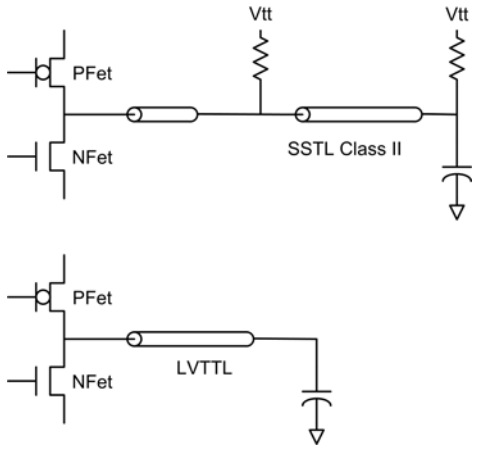

Figure 2. SSTL Class II and LVTTL Circuit Topology

Delta-I noise results in movement of the on-die supply rails (V_{CC} and ground) with respect to PCB supply rails. This is the basis for the "ground bounce" and "V_{CC} sag" terminology that is often used in conjunction with the measurement of quiet-high and quiet-low signals. This language has roots back in the wire bond and lead frame package days when inductance in the ground or power lead was the primary culprit in SSN noise. The terminology still applies to LVTTL bus topology because current is directed into the ground lead (or drawn from the power lead) immediately when the driver switches, but it does not accurately describe the inductive coupling noise found in both SSTL and LVTTL topologies. "Quiet low" and "quiet high" are the preferred terms because they accurately describe the measured noise. Also, ground bounce is a dangerous concept because it cannot be measured and is not unique. In his book [1], Brian Young discusses the non-uniqueness of voltages on opposite sides of inductive loops as well as the non-uniqueness of inductance matrices. Rather than refer to ground bounce and V_{CC} sag anywhere in the system, it is better to discuss quiet-low and quiet-high noise on signals with respect to local ground.

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 die. This is the voltage that enables the circuits to perform their intended tasks. Some amount of on-die decoupling capacitance, either intentionally designed or naturally occurring, happens within the circuits. This capacitance prevents instantaneous voltage changes across the power terminals of the circuits. The governing equations are $I = C\,dv/dt$ and $V = \frac{1}{C}\int I\,dt$. Current must flow through the capacitance and be integrated to produce power supply compression. It does not happen instantaneously and only has a minor effect during the rise and fall times of the drivers.

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, which lowers the on-die voltage, and in turn brings in current from the outside through the package inductances. Current is integrated in the ODC resulting in power supply compression. 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 *dI* than it is upon *di/dt*.

Note that the power supply compression may be negative. When drivers are initially in the high state, they must drive steady state current out to the network for parallel-terminated bus technologies (i.e., SSTL class II). This steady state current enters the package though the inductive power path. When the drivers turn off (ceases to draw current from the power supply), package inductance continues to drive current into the die where it is integrated in the ODC. This causes a voltage peak (negative compression) on the local on-die power supply.

The three major SSN mechanisms—inductive crosstalk, delta-I, and power supply compression—are present to some degree in most applications. The FPGA must be designed to reasonably address each of these mechanisms and support very diverse design possibilities without prior information on how the FPGA will be used.

Performance Consequences of SSN Mechanisms

The SSN mechanisms described above may result in degraded system performance in three different areas: near-end noise, far-end noise [2], and power-supply quality. Far-end noise is well recognized in the industry. This is the noise observed and easily measured at the far end of the transmission line where the receiver resides. Near-end noise affects receivers on the same die where the noise is generated. All three SSN mechanisms contribute to near-end noise but inductive coupling is the most important. Mutual inductance from aggressor signals couples noise to near-end victims in the same manner as it does to far-end victims. A noise voltage appears across the victim conductor (usually a ball or via), and is conducted to the input of a receiver.

SSN Figure of Merit

The traditional figure of merit used to evaluate the extent of SSN is "mV of noise" and may be published for the near- and far-end. This makes sense in the context of one signaling technology, but may not be the best figure of merit for the many bus technologies. For example, a 2.5V LVTTL bus can tolerate more mV of noise than a 1.8V SSTL bus. Also, more mV of signal margin is gained with greater drive strength. It makes sense to generate a set of SSN metrics for FPGAs that is a good figure of merit for many different bus technologies, power supply voltages, drive strengths, etc. One possibility is expressing SSN as a percentage of the signal margin.

In all signal technologies the signal voltage reaches steady state when the driver does not switch for a long period of time. There is usually a specification for V_{il} and V_{ih}, the highest voltage that will be interpreted as a "0" state and the lowest voltage that will be interpreted as a "1", respectively. The signal margin is the difference between quiet high and V_{ih} or the difference between V_{il} and quiet low. Incoming signals applied 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 the signal margin as part of the eye diagram for a typical SSTL class II signal.

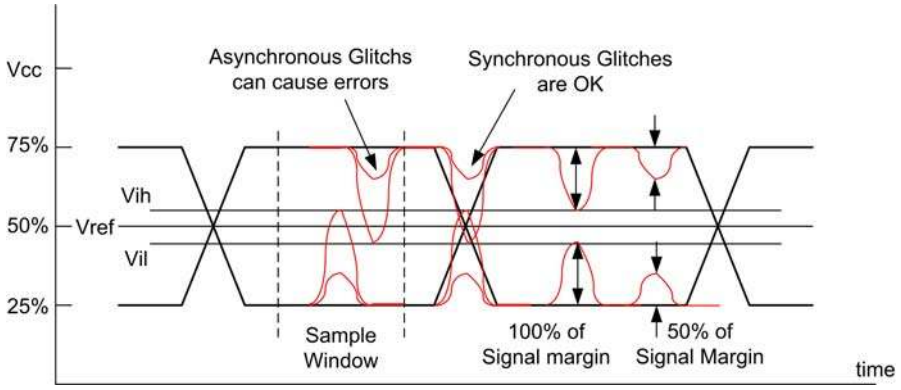

Figure 3. Signal eye diagram. SSN noise glitches can cause functional failures if they occur within the sample window or timing issues outside the sample window. 100% of signal margin is the difference between quiet high and V_{ih} or V_{il} and quiet low. A reasonable goal for SSN noise glitches is 50% of signal margin.

In the eye diagram, the vertical dimension is in units of volts and the horizontal dimension is in units of time. A key concept is the sampling window. There is a period of time or sampling window, often a setup and hold time around a clock, where incoming signals are expected to be stable-high or stable-low and remain 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 and penetrate into the V_{ih} to V_{il} switching zone.

For synchronous signals, meaning all signals synchronized to the same clock, SSN glitches outside of the sampling window will not cause functional failures because the receiver is not sensitive to noise glitches during this time period. This is true for both near- and far-end synchronous signals. However, SSN glitches outside of the sampling window may lead to timing pull-in and push-out.

The design goal for the Stratix III FPGA is to keep the SSN glitch height to less than 50 percent of the signal margin for both the high and low states. The 50 percent goal allows half of the signal margin to be taken up by the driving device, leaving the other half of the signal margin for system level noise that may occur on the PCB, connectors, receiver package, power supply, and V_{REF} tolerance, etc. This represents a substantial improvement in signal quality over competitive FPGA offerings. Changes in the die and package are required to meet this goal including additional power and ground return paths and additional capacitance. These changes are costly but must occur in order to meet the high performance demands of many of the FPGA customers.

The 50 percent of signal margin goal represents a new metric for discussing SSN, and applies to all drive strengths and bus topologies. The traditional figure of merit worked nicely in specific cases, but the mV of SSN measured at either the near or far end is highly dependent upon the power supply voltage, drive strength, bus termination type, temperature and process condition (fast, typical, or slow). The new metric works for all applications where an FPGA may be used. The measured mV of SSN varies greatly according to the application, but the SSN expressed as a percentage of signal margin is relatively constant for a given die and package combination for all applications.

An FPGA product that meets the 50 percent of signal margin goal will have a very low probability of producing functional SSN failures and a minimal impact on timing for synchronous busses.

Measured Results

To help project the SSN performance of Stratix III FPGAs, Stratix II GX test chips were measured. The most straightforward SSN measurements are made at the far end of a transmission line, approximately in the receiver position of the circuit of Figure 1. A victim driver was chosen in the load bank and set to be either quiet high or quiet low. Aggressor drivers were then switched in the vicinity of the victim, and noise on the victim was measured with a high bandwidth oscilloscope. Figure 4 shows an overview of far-end noise waveforms for 25Ω SSTL class II drivers. The blue trace represents the aggressors, the green trace is quiet high and the yellow trace is quiet low. All traces are on the same voltage scale.

Figure 4. Typical SSN waveforms with 86 drivers switching. The aggressor waveform is in blue, quiet high in green and quiet low in yellow. Inductive coupling and delta I noise cause quiet waveforms to spike up with the falling edge of aggressors and dip with rising edge of aggressors. Power supply compression causes ringing at the PDN resonant frequency, particularly on the aggressor falling edge.

Saturation Curves

A substantial amount of information can be gained by analyzing the forward and reverse saturation curves [3]. SSN measurements usually saturate at some maximum value as demonstrated in Figures 6, 7, and 8. The "forward" saturation curve gives the amplitude of SSN when the aggressor drivers associated with package pins nearest the victim are switched first, followed by the aggressor pins farthest from the victim in an outward spiral pattern. The "reverse" saturation curve gives the amplitude of SSN when drivers with pin positions farthest away from the victim are switched first, with drivers closer to the victim then progressively switched in an inward spiral pattern.

The measurements begin with zero drivers switching. All drivers are in high-impedance, tri-state mode except for the victim, which is programmed to be in either the high or low state. This establishes the quiet voltage levels with no SSN present. A binary spiral pattern of 1, 2, 4, 8, 16, 32, 64, and 87 aggressors are then activated and noise measurements are made for each pattern. Then the victim is changed to the opposite state and the pattern is repeated. Figure 5 shows a typical pin (solder ball) position pattern on the bottom of the FPGA (in this case, the victim is pin AL13). The aggressors switched in the binary pattern are also identified in Figure 5.

U1	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
AA	GND	GND						VCC	GND	VCC	GND	VCC	GND	VCC	GND	VCC
AB	GND	GND		GND	GND	GND	GND	GND	GND	GND	AB16			GND		
AC	GND	GND	GND		GND		GND	AC13	AC14	AC15	AC16	AC17				
AD	GND	GND	GND	GND	GND	GND	GND	AD13	AD14	GND	VCCPD	VCCPD	GND	VCCPD	VCCPD	GND
AE	GND					AE11	VCCN	AE13	AE14	VCCN	VCC		VCC	VCCN	VCCN	VCCN
AF	GND				AF10	AF11	VCCN	AF13	AF14	VCCN	VCC	GND	GND	GND	VCC	VCCN
AG	GND	GND	AG8	GND	AG10	AG11	GND	AG13	AG14	GND	VCCN	GND	GND	VCCN		GND
AH	GND	AH7	AH8	AH9	AH10	AH11	AH12	AH13	AH14	AH15	AH16	AH17				
AJ	GND	AJ7	AJ8	AJ9	AJ10	AJ11	AJ12	AJ13	AJ14	AJ15	AJ16	AJ17				
AK	GND	GND	AK8	GND	VREF	AK11	GND	VREF	AK14	GND	AK16	VREF	GND	AK19		GND
AL	GND	AL7	AL8	AL9	AL10	AL11	AL12	AL13 -	AL14 -	AL15	AL16	AL17	AL18	AL19		
AM	GND	AM7	AM8	AM9	AM10	AM11	AM12	AM13 -	AM14 -	AM15	AM16	AM17	AM18			
AN	GND	GND	AN8	GND	AN10	AN11	GND	AN13	AN14	GND	AN16	AN17	GND			GND
AP	GND	AP7	AP8	AP9	AP10	AP11	AP12	AP13	AP14	AP15	AP16	AP17	AP18	AP19		

victim
∑ Aggressors = 1
∑ Aggressors = 2
∑ Aggressors = 4
∑ Aggressors = 8
∑ Aggressors = 16
∑ Aggressors = 32
∑ Aggressors = 48
∑ Aggressors = 64
∑ Aggressors = 87 (All)

Figure 5. Pin position for binary aggressor patterns. The AL13 victim pin is red. Forward and reverse binary patterns of aggressors are switched with 1, 2, 4, 8, 16, 32, 64 and 86 drivers being switched from the inside out and the outside in respectively.

Figure 6 shows the far-end SSN voltage measured for 25Ω drivers with an SSTL class II load topology with 1.8V power supplies. The curves show the SSN glitch height measured for quiet-high and quiet-low victims for the forward and reverse binary spiral patterns. The SSN glitch on quiet low spikes up on an aggressor falling edge and the SSN glitch on quiet high spikes down on an aggressor rising edge, as seen in Figure 4. The magnitude of the glitch height with respect to far-end local ground is measured by a high-bandwidth probe and high-bandwidth oscilloscope, and graphed in Figure 6. Also shown are the V_{il} and V_{ih} levels. The forward and reverse saturation curves reveal much about the underlying SSN mechanisms.

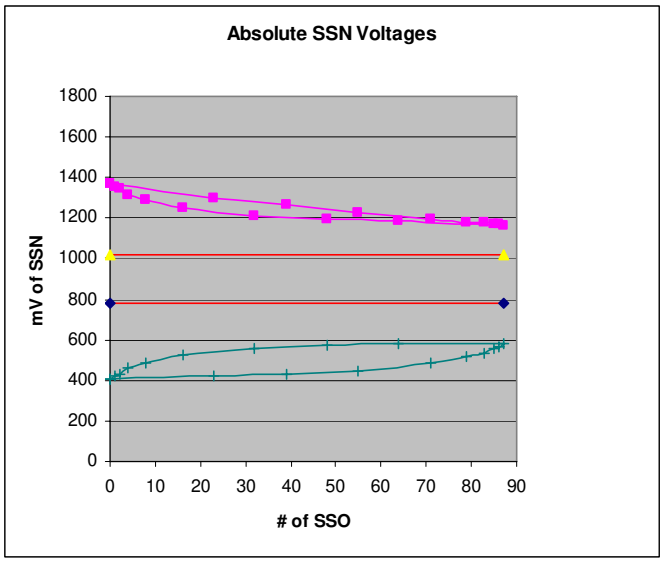

Figure 6. Absolute measured SSN voltages for forward and reverse spiral patterns with various number of 25 Ohm SSTL class II drivers switching. Power supply was 1.8V.

Figure 7 shows the subtracted SSN. The initial voltage on the victim, with all drivers in tri-state mode, is subtracted from the highest peak measured on quiet signals for each number of drivers defined by the binary pattern. Both forward and reverse spiral patterns are shown, as is the subtracted data for quiet-high and quiet-low victims. Figure 8 shows the subtracted data presented as a percentage of signal margin. As discussed above, the signal margin is the difference between the quiet-high or quiet-low voltage, with no drivers switching and V_{ih} or V_{il} respectively.

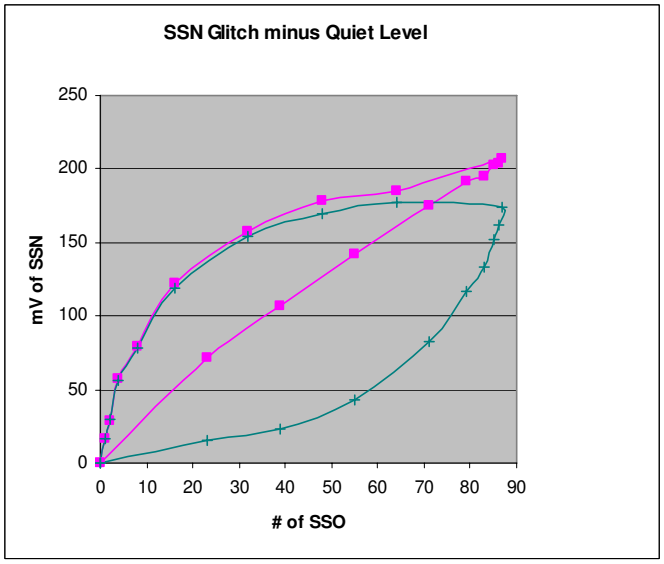

Figure 7. Subtracted SSN voltages. The quiet low voltage with zero switching aggressors is subtracted from the quiet low noise spikes and the quiet high noise spikes are subtracted from the quiet high voltage with no aggressors switching.Reverse saturation curve indicates PDN noise.

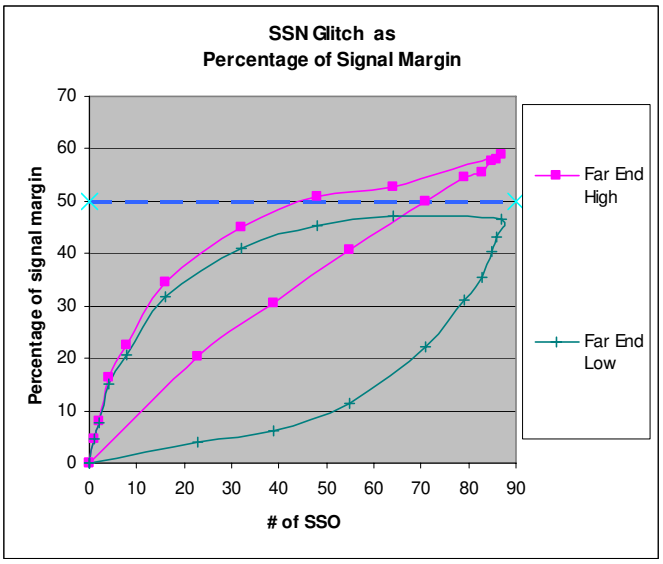

Figure 8. SSN expressed as a percentage of signal margin. Signal margin is determined by Vih – quiet low and quiet high – Vih when no drivers are switching.

Waveform Analysis

Figures 9, 10, and 11 show the far-end SSN waveforms on an expanded scale. Figure 9 shows the SSN glitch on quiet low as 87 aggressors make a high to low transition. The sharp spike at the beginning of the waveform is due to inductive coupling and delta-I noise ($m\,di/dt$ and $L\,di/dt$) followed by reflection activity from the transmission lines. The initial spike is the biggest because this is where the drivers make their transition. The source function for inductive coupling and delta-I noise is the *di/dt* of the aggressors, as this type of noise is only generated when the drivers are in transition. When the drivers finish the transition, the forcing function is finished. The remainder of the noise waveform is generated from transmission line reflections from the initial spike.

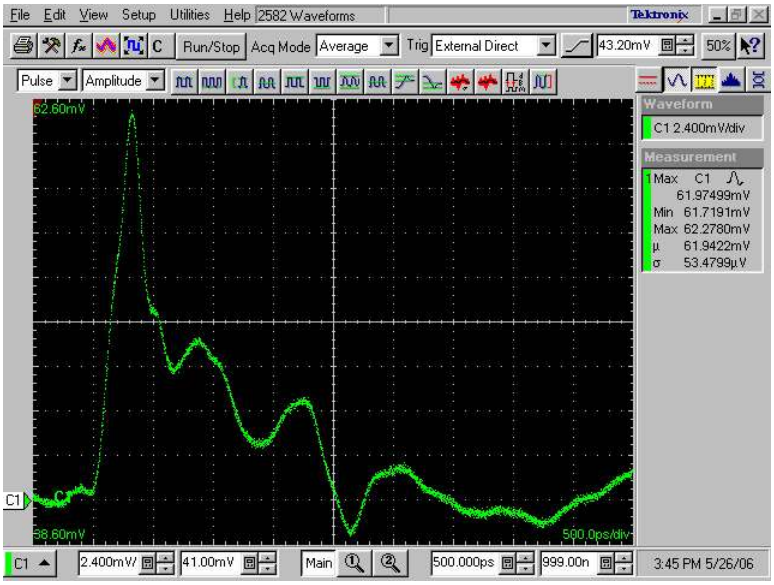

Figure 9. Quiet Low SSN Waveform with 87 Drivers Switching

Figure 10 shows the SSN glitch on quiet high as 87 aggressors make the transition from low to high. The initial spike in the waveform is due to inductive coupling and delta-I noise similar to that described for the quiet-low line. But the noise lasts several nanoseconds longer compared to quiet-low noise. The extended noise is due to power supply compression, and is related to the ODC and the PDN inductance as described below. The 87 aggressors create *di/dt* noise during their transitions, but must continue driving current out into the transmission lines for SSTL class II loads. This current depletes the charge in the on-die decoupling capacitance for a period of time until the current is able to rush in through the PDN. The power supply voltage sags momentarily, then recovers. For 87 aggressors, the magnitude of the power supply compression noise is nearly the same as the *di/dt* noise.

Figure 10. Quiet High SSN Waveform with 87 Drivers Switching

Figure 11 shows the quiet-high waveform for 55 aggressors switching in the reverse spiral pattern. Aggressors far away from the victim are active, but those nearer are not. Inductive coupling is greatly diminished because the 31 nearest aggressors, where inductive coupling is the strongest, are not switched. The spike observed at the beginning of the waveform is small compared to that of Figure 10, and is due to delta-I noise plus the small amount of inductive coupling that can traverse the distance from the aggressor pins to the victim pin. The large dip that occurs several nanoseconds after the initial spike is due to power supply compression. When the on-die power supply droops, all drivers in the I/O bank are affected, including the quiet-high driver being measured. Position in the I/O bank is not important for power supply compression noise.

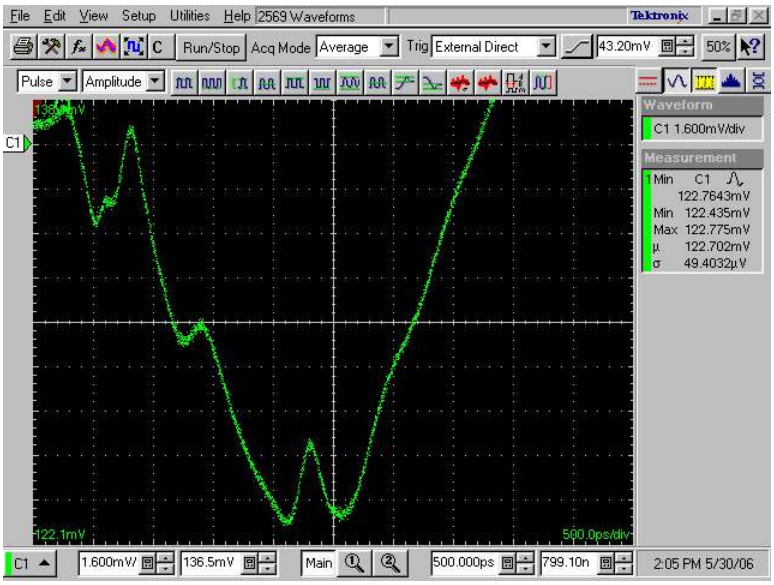

Figure 11. Quiet low SSN with 55 aggressors switching in the reverse saturation binary pattern. The initial dip due to inductive coupling has almost vanished because aggressors near the victim are not being switched. SSN noise in this case is dominated by the power supply compression mechanism.

Mechanism Identification

The waveforms at the far end of the transmission line have been interpreted to identify the relative importance of the underlying SSN mechanisms. The saturation curves shown in Figures 7 and 8 reveal similar information. The quiet-low curve builds up SSN quickly in the forward spiral, but does not show much SSN until the final drivers are switched for the reverse spiral. This is typical of inductive coupling where most of the SSN comes from the nearest aggressor pins. The forward spiral for quiet high shows a similar pattern, but the reverse spiral does not. The amount of SSN is almost linearly proportional to the number of drivers switched, indicating that the pin positions of victim and aggressors are unimportant. This is an indication of either delta-I noise or power supply compression. From the far-end voltage waveform, it is determined that power supply compression dominates the SSN for SSTL class II drivers rather than delta-I noise. Mechanism identification is important in determining the most effective ways to reduce SSN.

Near-End Noise

Far-end noise voltages are easily observed with an oscilloscope, but near-end noise measurements are more challenging as it is usually impossible to place oscilloscope probes near a victim receiver on an aggressing die. Near-end SSN can be measured by applying a DC voltage to the input of the victim receiver and noting voltages where the receiver is stable and where it is not. Figure 12 shows a schematic diagram for near-end SSN measurement where the same forward and reverse spiral patterns are used for the aggressors as were used for far-end SSN measurements, but the victim pin is changed from a driver output to a receiver input. The receiver compares the voltage on the input to V_{REF}. V_{REF} is typically stable on die because the amount of capacitance on this node prevents it from responding to SSN. While there is only a small amount of capacitance on the receiver input, it is vulnerable to noise coupled to the input pin.

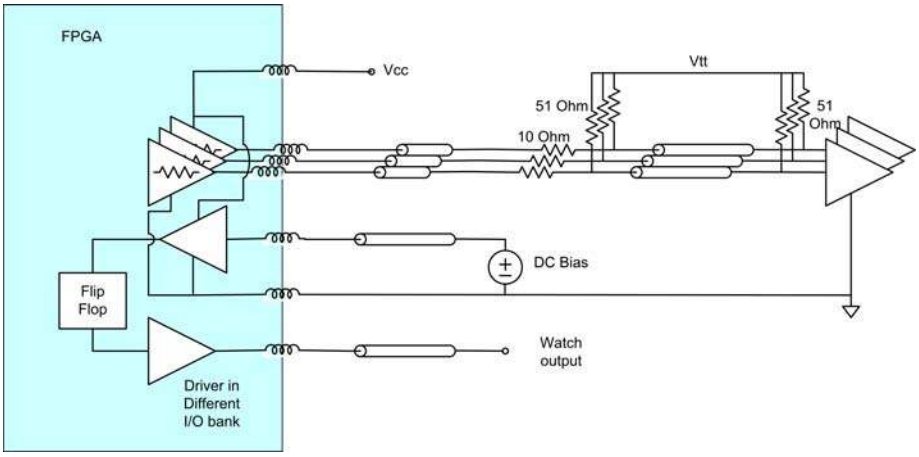

Figure 12. Near end SSN noise measurement. DC bias is adjusted to cause the input receiver to toggle or not toggle in the presence of aggressor SSN.

The DC voltage source is initially set to a low voltage, allowing the receiver to detect a "0" with no aggressors switching. The voltage is gradually turned up until the receiver toggles, generally when the input voltage is between V_{il} and V_{REF}. There is an input voltage range near V_{REF}, where the victim receiver is unstable and may toggle up and down due to random noise. Eventually, as the input voltage approaches V_{ih}, the receiver will consistently interpret the input as a high condition. It is important that the flipflop is edge sensitive as opposed to level sensitive, as this allows a fast glitch from the victim receiver to trip and toggle. The output of the flip flop leads to a driver in a different I/O bank and out to the PCB for observation with an oscilloscope. This process determines the initial receiver sensitivity without SSN.

The aggressors are then activated in the binary pattern. When one aggressor is switched next to the victim pin, the DC voltage ranges that are interpreted as either high or low will be slightly diminished. As more aggressors switch, the stable voltage ranges are further diminished. This technique is used to obtain the curves shown in Figure 13, as are the far-end SSN measurements, making for easy comparison. Figure 13a shows the near-end voltages measured at the edge of the toggle zone. Figure 13b shows the difference between the toggle voltage and the initial non-toggle voltage when zero aggressors are switching. This curve can be interpreted as the amount of SSN detected by the victim receiver as a function of the number of aggressors switching. Figure 13c shows the near-end noise as a percentage of signal margin.

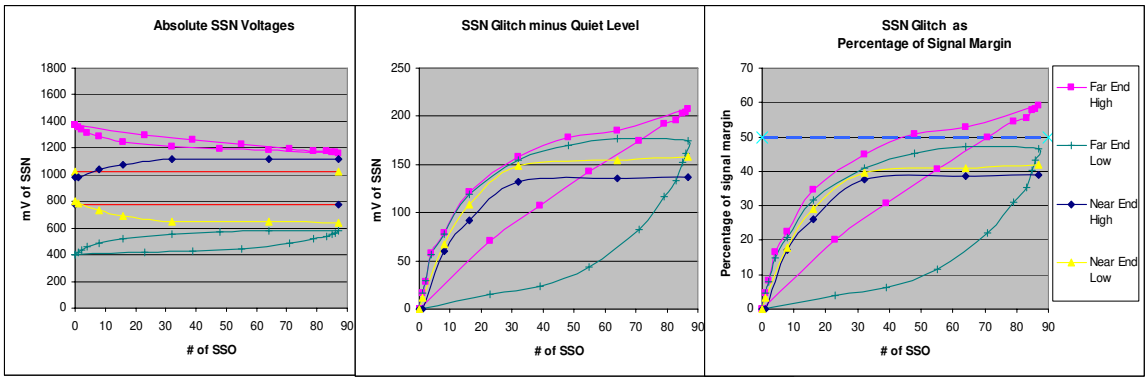

Figures 13a, 13b, 13c. Near end SSN measurements compared to far end SSN measurements. (a) measured voltages. (b) subtracted voltages showing the difference between zero aggressors and the binary pattern of aggressors switching. (c) SSN expressed as a percentage of signal margin.

Power Supply Compression

Power supply compression is described above as the difference between local power and ground voltages on-die. It is impossible to get probes into the die to measure this voltage, but it is possible to observe the quiet-high and quiet-low waveforms at the far end. These voltage waveforms can be subtracted on an oscilloscope to give a reasonably accurate representation of the on-die power supply compression. Figure 14 shows the quiet-high and quiet-low signal waveforms at the far end with the subtracted result. Figure 14a shows the timing relationship between an aggressor (yellow), quiet high (blue), quiet low (purple), power supply compression (orange) subtracted at the far end, and finally the V_{CCIO} supply on the PCB near the FPGA (green). Figure 14b shows the same waveforms with similar scale factors so the relative magnitudes are apparent.

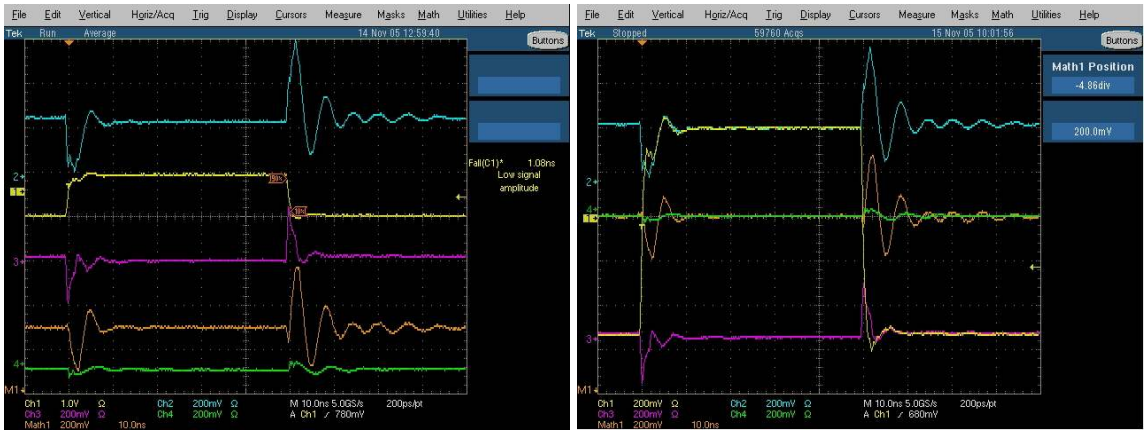

Figure 14. Power supply compression measured at far end together with aggressor, quiet high and quiet low waveforms.

Figure 15 shows the measured power supply compression for a Stratix II GX test chip (a) without and (b) with OPD capacitance. OPD greatly reduces power supply compression and changes the resonant frequency as discussed below.

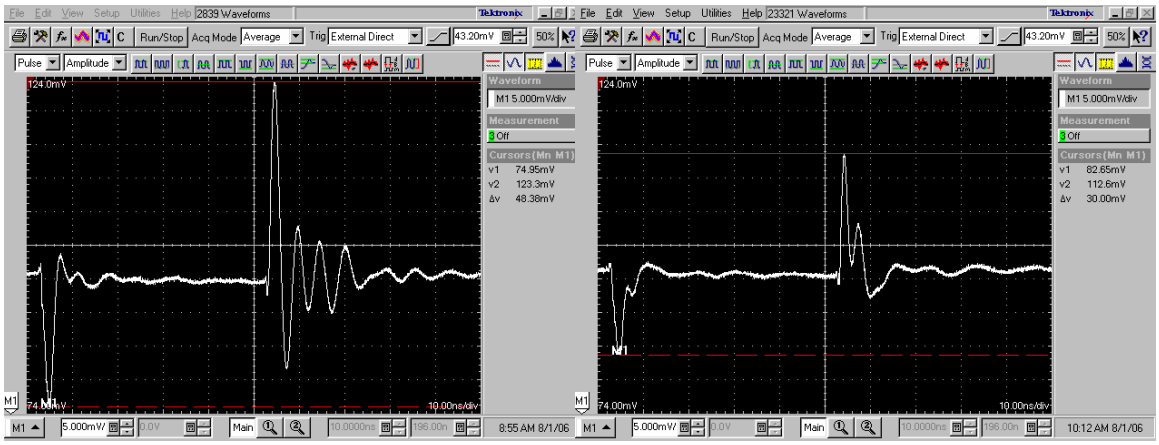

Figure 15. Power supply compression at far end for S2GX test chip (a) without and (b) with On-Package Decoupling (OPD) capacitance. A 10:1 attenuation probe was used for these measurements.

Power Supply Quality

The third metric for SSN is power supply quality, expressed in terms of percentage of the power supply voltage. The power supply quality metric is a distinct contributor to SSN as it is related to SSN, but is not unique for single-ended I/O driver circuits. All die-level power supplies have a resonant frequency associated with the die capacitance and the mounted package inductance [4]. The inductance of the package resonates with the ODC at the characteristic frequency $1/(2\pi\sqrt{LC})$. This resonance is stimulated anytime a fast transient current is consumed by circuits on die and drawn from the on-die decoupling capacitor, thus causing power supply compression. This is the case for LVTTL circuits that drive current into the load long enough for two transmission line time-of-flights to transpire and for SSTL class II circuits that deliver current continuously into the parallel terminations.

In Figures 14 and 15, the initial power supply compression dip damps out quickly, but the peak is under-damped. This is the ringing out of the chip/package resonance. When drivers switch from low to high, resonance is damped out by the parallel resistance of signal drivers and transmission lines. However, when drivers switch from high to low, the transmission lines are not in parallel with the resonance and the circuit has a higher Q, thus the power supply compression noise is often much more pronounced as drivers switch from high to low.

Figure 16 shows an equivalent circuit schematic for SSTL class II drivers, ODC, package and mount inductance and the PCB PDN. The VNA measurement techniques of [4] are used to determine some of the component values and others are determined by observation of the time domain resonant frequency. The circuit is simulated in SPICE-generating voltage waveforms at the far end and the V_{CCIO} voltage at the near end.

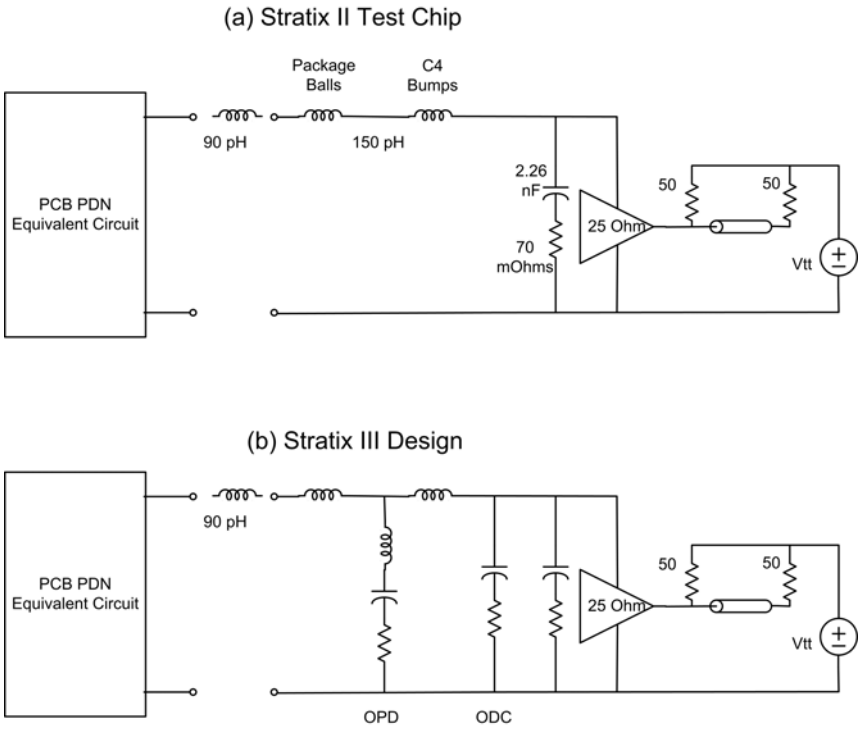

Figure 16. Schematic diagrams of PDN environment. (a) Stratix IIGX test chip and (b) Stratix III with on-package decoupling (OPD) and on-die capacitance (ODC).

Simulation results are displayed in Figure 17. Good matching is found for both the amplitude and resonant frequencies. For SSTL class II topology with 25Ω drivers, the power supply compression observed at the far end is about half of the V_{CCIO} noise voltage on die, consistent with voltage division. The simulated far-end power supply compression is 426 mV p-p and the on-die V_{CCIO} voltage is 814 mV p-p resulting in a power supply quality of 1.8V +/-23 percent.

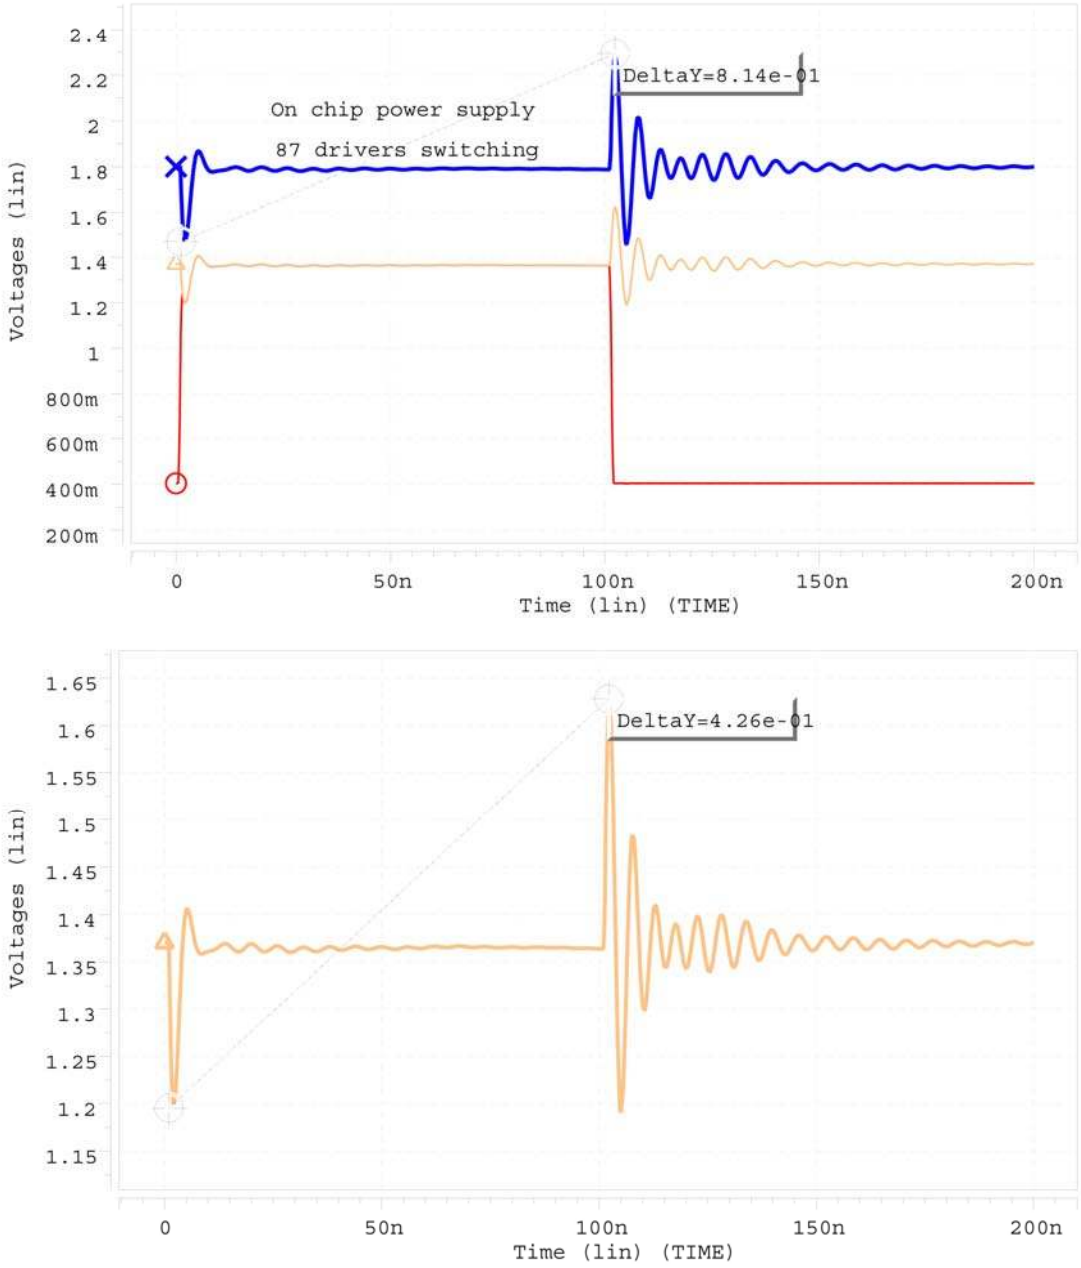

Figure 17. Simulation of the circuit of Figure 16 to show the on-die power supply compression compared to the far end measurement.

Figure 18 shows frequency and time domain simulation results with various amounts of capacitance in the PDN, similar to the simulations performed in [5]. The initial frequency domain simulation shows an impedance peak at 230 MHz. Subsequent simulations show the effect of ODC and OPD. The impedance peak is reduced with additional capacitance. OPD is effective at lower frequency and ODC at higher frequency. Time domain simulations show the improvement of on-die power supply quality with the additional capacitance, which closely match the measured results of Figure 15. The projected power supply compression for Stratix III FPGAs is 328 mV p-p on die, resulting in 1.8V +/-9 percent for the power supply quality.

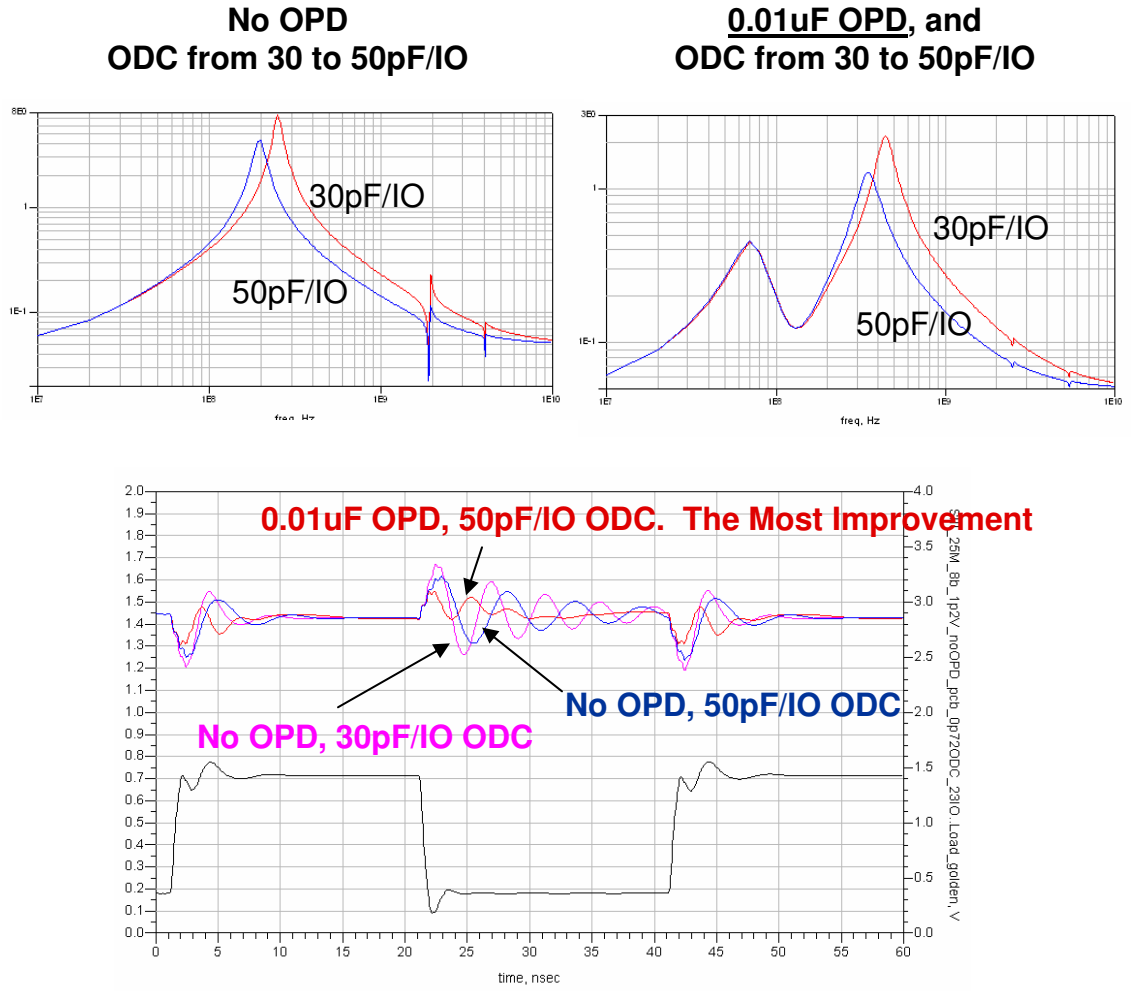

Figure 18. Simulation of Stratix III PDN including OPD and ODC. Power Supply Quality is greatly improved.

Stratix III FPGA Signal and Power Integrity

Stratix III FPGAs build upon the successes of Stratix II devices for signal and power integrity. Stratix II FPGAs have approximately 20 signals for each ground and power pin, resulting in SSN that exceeds 50 percent of the signal margin. Stratix II FPGAs were optimized for I/O density. Figure 19 shows the relative far-end noise for several ratios of signal to ground. Stratix II GX devices use an 8:1 ratio because it is an optimal balance

between inductive coupling performance and I/O density at the knee of the curve, and substantially improves the SSN performance in measured results as shown above. Stratix III FPGAs have power pins located among their I/O pins, which give an 8:1:1 ratio for signal : ground : power, further improving SSN performance. The signal to return-current pin ratio is 4:1.

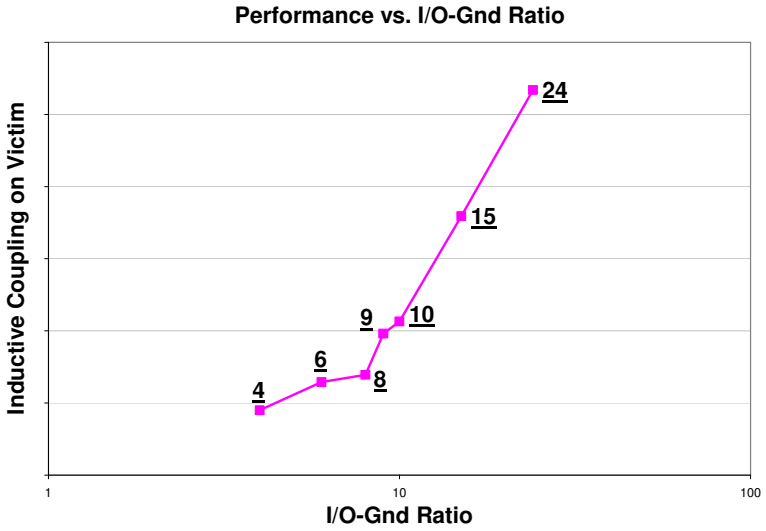

Figure 19. Inductive coupling is improved by improving the ratio of signals to return path. Only marginal improvement is obtained by making the ratio less than 8:1.

Figure 20 shows the typical relationship between signal, power, and ground balls for Stratix III FPGAs. This pattern will surpass the effectiveness of the Stratix II GX pattern by reducing inductive coupling and delta-I mechanisms that generate SSN. The pattern enables a Stratix III device to constrain SSN within 50 percent of the signal margin for most drive strengths and termination topologies over all process, voltage, and temperature (PVT) conditions.

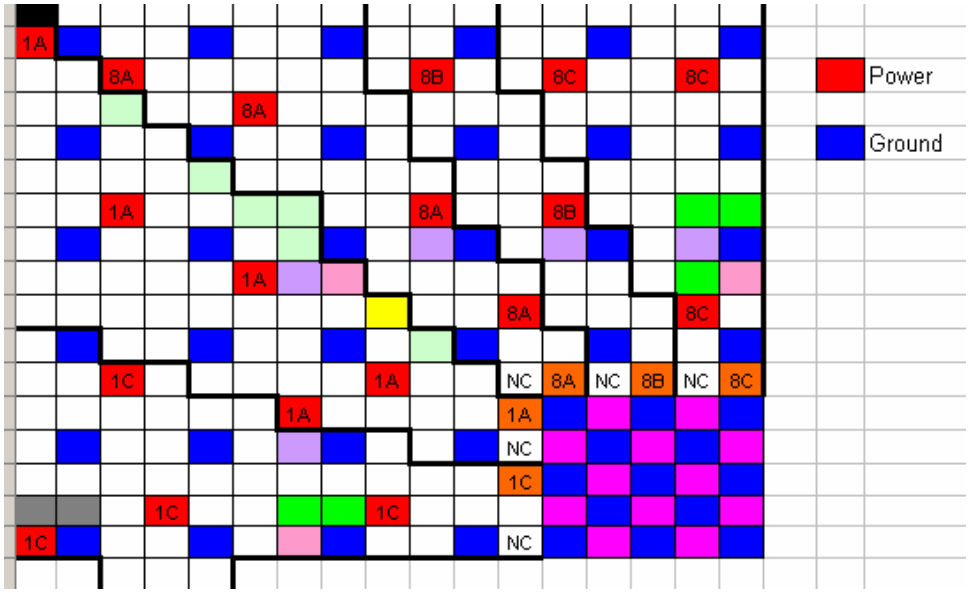

Figure 20. Pin pattern for typical Stratix III I/O bank including signals, power and ground.

OPD and ODC capacitance are incorporated into Stratix III FPGAs to improve the power supply quality, with on-die supplies remaining within 10 percent of the nominal voltage. The Stratix III FPGA family has been carefully architected to maximize SSN performance and give the best signal and power integrity results possible without significantly increasing the cost of the device.

Conclusions

The fundamental signal and power integrity mechanisms that contribute to SSN noise have been identified in this paper. Inductive coupling, which occurs mostly in vertical structures, including the package balls and PCB vias, enables aggressor signals to affect nearby quiet victims through mutual inductance ($m\,di/dt$). Delta-I noise occurs when the signal current loop is completed and voltage builds up in either the power or ground path ($L\,di/dt$). Power supply compression occurs when transient current either is drawn from or forced into the OPD capacitance and causes fluctuations in the power supply voltage. These mechanisms create SSN at both the near and far ends (driver and receiver die) and may cause signal-integrity errors. The mechanisms are identified by observing the far end voltage wave forms and by the shape of the forward and reverse saturation curves with aggressor drivers switching in a binary spiral pattern.

The traditional way of describing SSN is in mV of noise. This metric is not descriptive enough for FPGA devices that typically run a variety of signal technologies at various voltage levels. A more descriptive metric would express SSN in terms of signal margin. Signal margin in mV is calculated from the quiet-high voltage minus V_{ih} or V_{il} minus the quiet-low voltage. The height of the SSN glitch is then expressed as a percentage of signal margin.

The Stratix III FPGA family has been architected to carefully manage the signal and power integrity mechanisms. A low ratio of signal-to-return path pins is most effective in reducing inductive coupling and delta-I noise, with die and package capacitance added to improve power supply quality, thereby greatly improving SSN without adding excessive cost to the product line.

Acknowledgements

The authors would like to thank Kundan Chand, Mark Flanigan, Zhe Li and Raymond Wibowo for their many hours of careful measurements and the ICD, PE, Applications and Packaging departments for supporting this work.

References

[1] B. Young, *"Digital Signal Integrity",* Prentice Hall PTR, Upper Saddle River, NJ, 2001, p138, pp. 208-218.

[2] T. Budell, J. Audet, D. Kent, J. Libous, D. O'Connor, S. Rosser, and E. Tremble, *"Comparison of Multilayer Organic and Ceramic Package Simultaneous Switching Noise Measurements using a 0.16 um CMOS Test Chip,"* Electronic Components and Technology Conference, May 2001, pp. 1087-1094.

[3] J.P. Libous, D.P. O'Connor*, "Measurement, Modeling, and Simulation of Flip-Chip CMOS ASIC Simultaneous Switching Noise on a Multilayer Ceramic BGA,"* IEEE Transactions on Components, Packaging, and Manufacturing Technology – Part B, Vol. 20, No. 3, Aug 1997, pp. 266-271.

[4] L.D. Smith, R.E. Anderson, T. Roy*, "Chip-Package Resonance in Core Power Supply Structures for a High Power Microprocessor",* Proc. Interpack'01 Conference, July 2001.

[5] Hong Shi, Geping Liu, Alan Liu, *"Analysis of FPGA Simultaneous Switching Noise in Three Domains: Time, Frequency, and Spectrum",* DesignCon 2005.

101 Innovation Drive
San Jose, CA 95134
(408) 544-7000
http://www.altera.com