



High-Speed PCB Design: Issues, Tools, and Methodologies

In this series on signal integrity, *Xcell* explores tools and methods you can use to combat signal and power integrity distortions throughout product development.

Table of Contents:

Ten Reasons Why Performing SI Simulations is a Good Idea	11
Interfacing SMA Connectors to Virtex-II Pro MGTs	12
For Synchronous Signals, Timing Is Everything	16
Designing High-Speed Interconnects for High-Bandwidth FPGAs	20
Accurate Multi-Gigabit Link Simulation with HSPICE	24
Eyes Wide Open	28
Backplane Characterization Techniques	31
A Low-Cost Solution for Debugging MGT Designs	36
Tolerance Calculations in Power Distribution Networks	40
High-Speed PCB Design Resources	44

by Suresh Sivasubramaniam
Senior Design Engineer
Xilinx, Inc.
suresh.subramaniam@xilinx.com

Philippe Garrault
Technical Marketing Engineer
Xilinx, Inc.
philippe.garrault@xilinx.com

RocketIO™ transceivers are a standard feature on the most advanced Xilinx FPGAs. Our first-generation transceivers operate in the 1-3.125 Gbps bandwidth, while the latest generation transceivers have an operating bandwidth of up to 10 Gbps.

These data rates mean that bit period and signal rise and fall times are extremely small. Designing physical links/channels on a PCB for these high-speed devices with small amplitude and timing budgets necessitates a careful analysis of the possible signal integrity (SI) and power integrity (PI) distortions.

SI/PI issues and reduced amplitude and timing budgets are not limited to just high-speed serial links. In recent years, the amount of logic cells inside Xilinx devices has grown tremendously. Additionally, the pin count on device packages has gone from a few to more than a thousand. Increased I/O performance, in conjunction with the large number of I/Os available, means that for each of your new designs, a lot more transistors are switching more often.

The common denominator in these problems is poor management of the three “bad boys” of electric circuits: resistance, inductance, and capacitance.

Additional requirements for the efficient design of high-speed buses may be dictated by the needs of a specific application. For example, a 266 MHz, 64-bit DDR RAM interface will be sensitive to skew between the different byte lanes. Large parallel buses also have the potential to generate simultaneous switching output (SSO) noise and voltage droop. All of these factors translate into the need to manage the transient current demands of a particular application through proper design of the power distribution system (PDS).

Resistance, Inductance, and Capacitance Pull the Strings

In general, SI and PI issues arise when designers pay inadequate attention to these broad categories:

- Termination schemes
- Skin effect (frequency-dependent attenuation)
- Dielectric losses
- Impedance discontinuities/reflections
- Data coding (DC balanced codes, run length, channel memory)
- Equalization/pre-emphasis
- Inter-symbol interference
- Crosstalk
- Decoupling/bypassing in power distribution
- Board stack-up
- Signal edge rates.

The common denominator in these problems is poor management of the three “bad boys” of electric circuits: resistance, inductance, and capacitance (Figure 1). In addition, you must understand and employ the right measurement techniques in the lab to accurately measure and validate designs against simulations or design specifications.

The objective is to build systems right the first time.

Minimize SI/PI Effects

In this special series on signal integrity, we have assembled articles that will provide you with practical and technical resources towards achieving that goal. From characterization and model extraction techniques in the lab to methods for simulating signal degradations of synchronous parallel/asynchronous serial systems to case studies, this series covers many aspects of SI.

In a sidebar to this article, Xilinx Principal Engineer Austin Lesea lists “Ten Reasons Why Performing SI Simulations is a Good Idea.” Although this may sound very familiar to some of you, understanding the benefits of performing SI analysis throughout the design cycle can help you achieve your performance, reliability, and time-to-market goals.

“Interfacing SMA Connectors to Virtex-II Pro MGTs” details Warren Miller and Vince Gavagan’s experience designing the interface between Virtex-II Pro™ multi-gigabit transceivers (MGTs) and Sub Miniature version A (SMA) connectors for the Virtex-II Pro Aurora Design Kit. Through prototyping and time domain reflectometry (TDR) measurements, they illustrate how SMA connector choice influences signal quality.




Figure 1 – The three bad boys of electric circuits (Courtesy of Educator’s Corner/Agilent Technologies)

Bill Hargin believes that “For Synchronous Signals, Timing Is Everything.” His article outlines a method for extracting correction values that can be applied to the clock-to-out and flight time numbers. The resulting timing values in the datasheet are representative of the actual load and topology of your design. This technique specifically applies to source-synchronous links.

Predicting the interconnect performance of high-speed links made of complex via, connector, and trace structures is no easy task. However, as Ansoft’s Lawrence Williams explains in “Designing High-Speed Interconnects for High-Bandwidth FPGAs,” combining electromagnetic, circuit, and system simulations greatly helps in the design of reliable and fast data transmission channels.

When designing multi-gigabit asynchronous channels, you must carefully analyze the link’s physical and electrical properties. In his article, “Accurate Multi-Gigabit Link Simulation with HSPICE,” Dr. Scott Wedge explains how the combination of an EM solver, coupled transmission lines, S-parameter support, and SPICE and IBIS modeling to the HSPICE® circuit simulator helps accurately account for high-speed signal distortions.

With “Eyes Wide Open,” Steve Baker shows you how to use the RocketIO Design Kit for ICX™ to evaluate pre-layout options (such as placement, connectors, or stackup) as well as post-layout options (such as detailed routing structures) to achieve high-speed serial link performance.

As much as Xilinx recommends SI simulation and analysis before manufacturing a PCB, there are two very valuable lab instruments that you can use on prototype/exploration boards. With these instruments, you can characterize interconnect properties and high-speed signal behavior, explore different topology performances, or extract simulation models. In his article, “Backplane



Characterization Techniques,” Eric Bogatin explains the need for making measurements, illustrating the concept of measurement and model bandwidth. He also discusses SMA launches, information contained in TDR traces, and differential S-parameters.

In “A Low-Cost Solution for Debugging MGT Designs,” Joel Tan presents a solution comprising a bit-error rate testing module connected to a flexible on-chip logic analyzer core, both implemented in FPGA fabric. Together with the ChipScope™ Pro software suite, these two

components allow you to perform diagnostic testing, debugging, and development of an MGT system without the use of expensive lab equipment such as logic analyzers and BERT testers.

And in “Tolerance Calculations in Power Distribution Networks,” Sun Microsystems’ Istvan Novak walks you through different scenarios of bypass capacitor configurations to demonstrate the importance and influence of the capacitors’ technology, value, and number in designing a decoupling/power distribution network.

Conclusion

We hope you will find in this series instructive material on the sources of SI/PI effects, along with practical information about the resources and tools available to you. Our experience tells us that careful simulations, analysis, and measurements of PI and SI effects early in the design process guarantees first-time success more often than not.

If you’d like to send us feedback about the topics discussed, please e-mail us at si_xcell@xilinx.com ✉

Ten Reasons Why Performing SI Simulations is a Good Idea

by Austin Lesea

Principal Engineer, Advanced Products Group
Xilinx, Inc.

Not so long ago, the rise and fall times of signals, the coupling from one trace to another, and the de-coupling of power distribution on a PCB were tasks that were routinely handled by a few simple rules. Occasionally, you might use the back of an envelope, scribbling down a few equations to make sure that the design would work.

Those days are gone forever. Sub-nanosecond, single-ended I/O rise and fall times, 3 to 10 Gb transceivers, and tens of ampere power needs at around 1V have all led to increased engineering requirements.

Your choice is simple: simulate now and have a working result on the first PCB, or simulate later after a series of failed boards. The cost of signal integrity tools more than outweighs the cost of making the board over and over with successive failures.

In keeping with the theme of this special issue, here are my 10 best reasons why signal integrity engineering is a good idea:

1. You’re tired of making PCBs over and over and still not having them work.

Seriously, without simulating all signals, as well as power and ground, you risk making a PCB that will just not work. IR (voltage) drop, inadequate bypassing or de-coupling, crosstalk, and ground bounce are just a few of the possible problems.

2. You’re tired of being late to market and watching your competition succeed.

Every time you have to fix a problem with a PCB, it necessitates a new or changed layout, a new fabrication, and another assembly cycle. It also requires the re-verification of all parameters. Taking the time to do these things right has both monetary and competitive advantages.

3. You’re tired of spending all this money, only to scrap the first three versions of PCBs and all of the components that went with them.

See reason number two.

4. Your eye pattern is winking at you.

If the eye pattern of a high-speed serial link is closing, or closed, it’s likely that the link has a serious problem and will have dribbling errors – or worse, will be unable to synchronize at all. You must simulate every element of the design to assure an error-free channel.

5. All 1s or all 0s suddenly breaks the system.

Unfortunately, many systems do not have a choice of what data may be processed. Often the data pattern will create conditions that, if not simulated a priori, will cause errors in the system.

6. Hot and cold, fast and slow, and high and low voltages cause failures.

Without simulating the “corners” of the silicon used as well as the environmental factors, you’re playing Russian Roulette with five of the six chambers loaded.

7. You cannot meet timing, and you are unable to find out why.

Poor signal integrity is the primary cause of adding jitter to all signals in a design. Ground bounce, crosstalk, and reflections all conspire to add jitter. And once added, jitter is virtually impossible to remove.

8. The FCC Part 15 or VDE EMI/RFI test fails every time you test a board.

Radiated and conducted radio frequency emissions, as well as susceptibility to radio frequency sources, is a sign of poor SI practices. Fixing the problem by shielding increases the system cost substantially, and may not even be possible in extreme cases.

9. Customers complain, but when you get the boards back, you don’t find any problems.

One of the biggest problems with SI is that the errors and failures observed are difficult to correlate and sometimes impossible to find. Was it a problem with voltage, temperature, or with the data pattern itself? It might have been someone turning lights on and off (ground disturbance). Don’t risk a return that cannot be fixed.

And last, but certainly not the least:

10. Your manager has suggested that you look for other employment.

Do not let this happen to you. Stay current, educated, and productive. Get the right tools to do the job. Realize that signal integrity engineering is a valuable and irreplaceable skill in great demand in today’s design environments.