DesignCon 2010

Simulation Techniques for 6+ Gbps Serial Links

Donald Telian, Siguys
telian@siguys.com

Sergio Camerlo, Ericsson
sergio.camerlo@ericsson.com

Brian Kirk, Amphenol TCS
brian.kirk@amphenol-tcs.com

Abstract

This paper details new simulation techniques for serial link Signal Integrity analysis at data rates above 6 Gbps. Proper modeling techniques for both the active (SerDes, AMI, transistor-level, etc) and passive (3D extraction, S-parameters, parasitics, etc) devices are described, along with solutions for making progress when desired models are not available. Sensitivity analysis reveals variables that limit system performance and guides design choices to optimize the link – adding as much as 50% more system margin. With variables and tolerances understood, exhaustive worst-case analysis is performed to quantify anticipated design margin.

Authors Biographies

Donald Telian is an independent Signal Integrity Consultant. Building on over 25 years of SI experience at Intel, Cadence, HP, and others, his recent focus has been on helping customers correctly implement today's Multi-Gigabit serial links. He has published numerous works on this and other topics that are available at his website siguys.com. Donald is widely known as the SI designer of the PCI bus and the originator of IBIS modeling and has taught SI techniques to thousands of engineers in more than 15 countries.

Sergio Camerlo is an Engineering Director with Ericsson Silicon Valley, which he joined through the Redback Networks acquisition. His responsibilities include the New Platforms Channel Design and Backplane Technology, including Signal and Power Integrity. He has also served on the company Patent Committee. In his previous assignment, Sergio was VP, Systems Engineering at MetaRAM, a local startup, where he dealt with die stacking and 3D integration of memory structures. Before, Sergio spent close to a decade at Cisco Systems, where he served in different management capacities. Sergio holds thirteen US Patents on signal and power distribution, interconnects and packaging.

Brian Kirk is a signal integrity engineer at Amphenol TCS. He has previously worked for Digital Equipment Corporation, Compaq, Hewlett Packard and Teradyne. His previous experiences include a variety of signal integrity tasks, module designs and FPGA designs for servers and routers. He is currently involved in connector development, simulation and correlation for high-speed interconnects. He received his PHD in Electrical Engineering from the University of New Hampshire.

1. Introduction

Though proprietary solutions have been available longer, open-market impulse response based channel simulators have now been available for over half a decade [1,2]. During that time, industry standards have begun to require them [3,4] and repeatable processes have emerged for their use [4,5].

This paper describes simulation techniques that deploy these maturing technologies in new ways to solve problems related to serial links operating at data rates over 6 Gbps. Though these techniques have been used for lower-speed Signal Integrity (SI) analysis for many years [6,7], practices such as pre-route sensitivity analysis and exhaustive worst-case analysis are now becoming feasible with channel simulators. Considering that both operating frequencies and the number of bits being simulated have increased three to four orders of magnitude, this has been a considerable challenge to overcome. In addition, simulation can quantify the types and amounts of equalization required to achieve robust performance across a variety of channel lengths and manufacturing tolerances.

Analysis is only meaningful when it is based on accurate models. For passive elements, S-parameter models [8] capture the behavior of 3D structures such as vias, connectors, and their associated grounds while 2D trace models must factor in how their Printed Circuit Board (PCB) materials change with frequency [9]. Since serial links are terminated on both ends, simulations provide meaningful results as long as accurate models of the losses and discontinuities in the passive interconnect are in place – even when using more approximate models of the active devices, or SerDes (Serializer-Deserializer). This is a relevant and important observation, since industry standard formats for SerDes models that connect to channel simulators are just now coming into place [10,4]. Indeed the analog behaviors of such models are quite simple, particularly when compared to the equalization schemes required for data rates over 6 Gbps [11,12]. Simplistic SerDes models function well provided there is a mathematical way to overlay anticipated equalization, as will be demonstrated.

This paper will follow a typical design process from modeling to optimizing to verifying a 6+ Gbps channel. Once appropriate models are in place, the simulation effort can be divided into the three analysis phases shown in Figure 1: Feasibility, Sensitivity, and Exhaustive. Each phase builds on the preceding work, and serves to further understand, optimize, and quantify the performance of the channel.

The Figure illustrates how the process begins by applying approximate models to perform Feasibility Analysis to comprehend the boundaries of the Design Space. The items in the top portion of the design space are challenges imposed by the system architecture, while the items in the lower portion are variables resolved by the design team. Over time, the design space narrows as Sensitivity Analysis guides design choices until a final “channel design” is defined. Once defined the design space widens slightly again; this time due to fully comprehending manufacturing tolerances during Exhaustive Analysis. Engineering choices made during this phase guide the fine tuning of equalization settings and finalize the simulated pre-hardware link performance. Ideally, this process is completed prior to the fabrication of hardware.

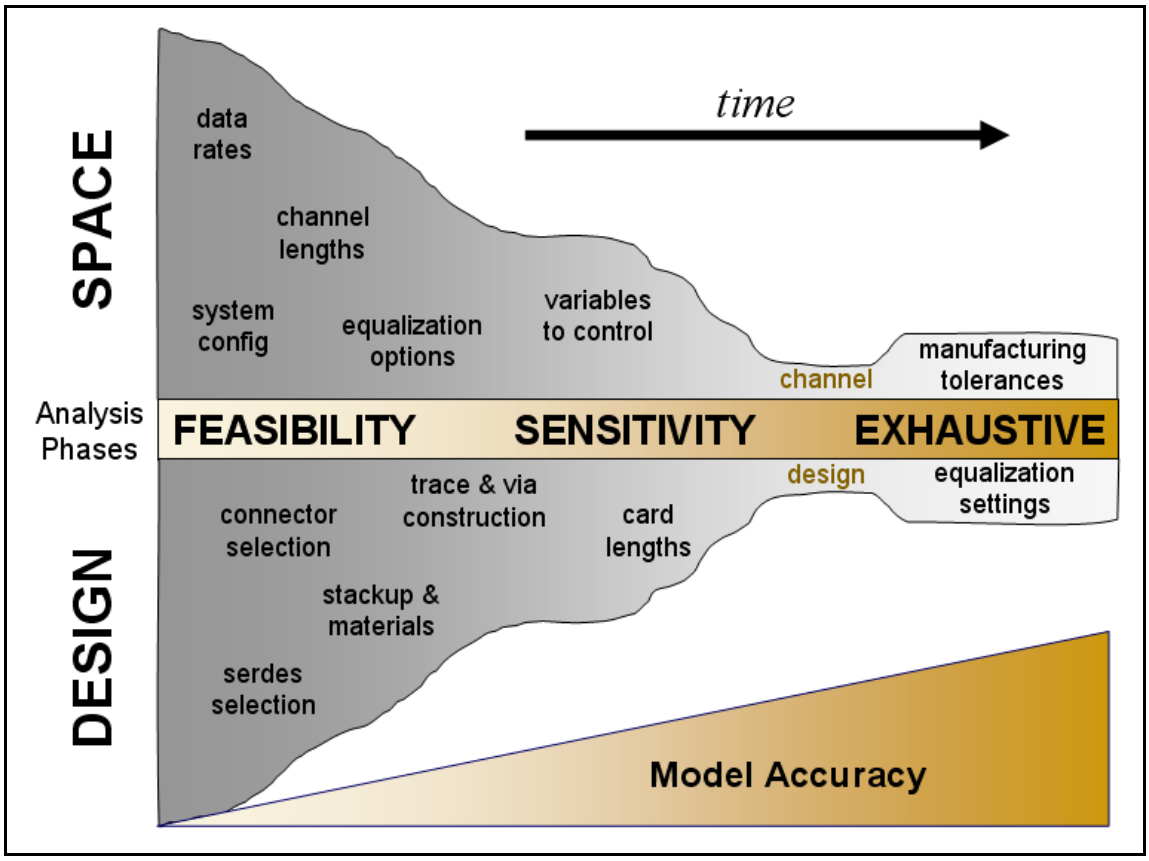

Figure 1: Channel Simulation Phases vs. Design Space

Figure 1 also illustrates the simulation engineer's responsibility to continually refine and improve the quality of the link models throughout the process. It's important to proactively work on this task with component and tool vendors, as well as your own design team, as further detailed in the next section.

2. Modeling a Channel

Though measured post-hardware channels might typically be represented by a single S-Parameter file, during system design it is necessary to represent a channel using cascaded elements. In this way, contributors to system loss and discontinuities can be understood, quantified, and tuned. Figure 2 shows the elements typically found in a 6+ Gbps channel, along with the likely source and type of each model. This is a single (non-crosstalk) channel Transmitter (Tx) to Receiver (Rx) connection, and will serve as our reference channel throughout this paper.

From left to right, the example channel consists of three PCBs: the Tx card, the backplane, and the Rx card. On the Tx card is found an active SerDes Tx, a model of the SerDes package, a via down to a stripline layer, the stripline route, and a via up to the backplane connector. The backplane consists of the connector model, a via down to the routing layer, a stripline route, a via up to the connector, and another connector. The Rx card has a connector via down to a stripline route, a via up to an AC capacitor, a microstrip route to the Rx which includes models of both the IC package and the Rx SerDes.

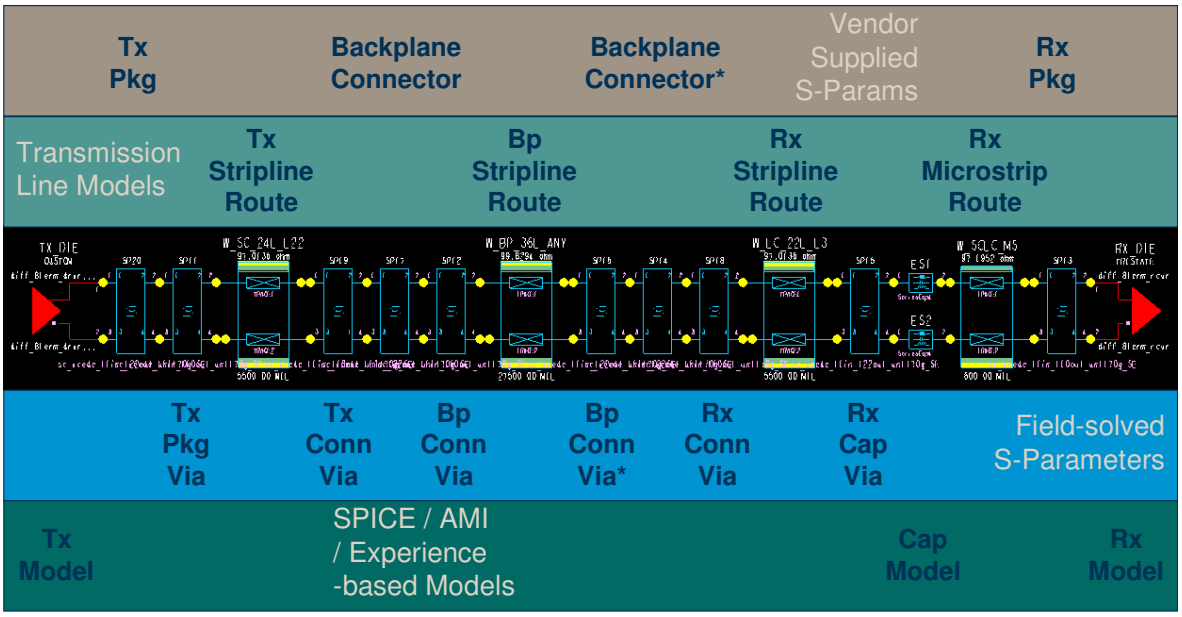

Figure 2: Cascading Models to Form a Typical Serial Link Simulation

Figure 2 also shows the typical source for each model in four horizontal rows. The top two rows contain models that will be supplied by either component vendors (top row) or trace models derived by the simulation tool of choice (second row). Since the design process has little impact on these items, other than component and tool selection, they will be dealt with minimally here. The lower two rows are more interesting, and will be covered in more detail. Aside from connectors, vias represent the primary discontinuities in a 6+ Gbps channel. As such, field solvers are deployed to develop accurate S-parameter models of these 3D structures while providing the ability to understand and optimize their performance.

The SerDes Tx and Rx are found in the bottom row of Figure 2. While these models might be supplied by a component vendor, the majority of models in existence today do not simulate fast enough for the types of analysis described herein – particularly if they are implemented at the transistor-level. As such, the model must typically be abstracted into its relevant behaviors capturing datasheet items such as edge rate, impedance, and return loss. If the vendor does not provide a transportable description of the device's equalization capabilities, generic Algorithmic Modeling Interface (AMI) models [10, 13] can be deployed to approximate device behavior. Until the day when AMI models are as readily available as IBIS models, experience working with numerous SerDes devices is required to implement the simulation environment correctly. The important point to remember is that *approximate SerDes models provide the ability to understand, adapt, and tune a passive channel.* Indeed, even a generic simulated TDR can provide insight into a passive channel. However, an approximate SerDes model additionally provides the ability to quantify eye performance and derive Bit Error Rate (BER) estimates.

2.1 Via Models

It seems odd that something as simple as a via would get so much attention during channel design and simulation. Indeed, not too long ago they were simply represented by a 0.5pF capacitor – if they were even acknowledged as "in the circuit" at all. However, as previously

stated, in the design of 6+ Gbps channels the lowly via has become one of the primary discontinuities. Make a mistake here, and 90% of your signal disappears. As will be shown, subtle changes can improve a via 50-100% while not-so-subtle changes can yield a 7x improvement in performance. In general, the biggest problem with vias is stubs [14] (a problem found not only in vias, but also in other structures such as testpoints), while the subtler problem involves balancing the metal cross-sections and hence impedances in the signal path.

So how different is the behavior of an actual via when compared with a 0.5pF capacitor? The answer is: it depends. Figure 3 is a comparison of the time domain (differential TDR, at left) and frequency domain (insertion loss, at right) response of a complex 3D differential via model (in red) versus two 0.5 pF capacitors (in blue). For this particular via, the response of the simple capacitor model is surprisingly accurate in both domains – particularly when you consider the complexity and the cost of derivation of the 3D model. Albeit the time domain response rolls off a bit differently and the frequency domain response is less accurate as you move away from 6 Gbps (3 GHz). Not too surprisingly, the via looks more like a transmission line and can be approximated even closer with a 40pS 70 Ohm differential transmission line shown in dashed gray at left. So if vias can be so well represented with such simple models, why all the fuss about 3D field solutions?

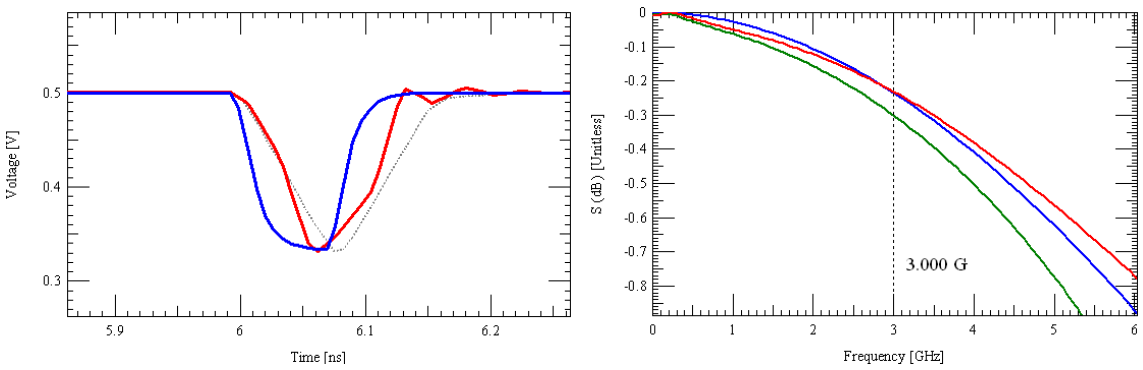

Figure 3: Time & Frequency Domain Responses of 0.5pF and 3D Via Models

The answer to that question is two-fold: (1) 3D field solutions give you the ability to connect design decisions with via performance, and (2) there is a significant range in the performance of vias. The example given above was contrived because we were able to search through a library and find a differential via with characteristics similar to 0.5 pF capacitors. Even so the single-ended insertion loss (green curve in Figure 3) of the vias is 30% greater than both the capacitor model and the differential insertion loss. This variation is typically not true of trace models, and hence ability to see and model this difference is another reason for using 3D solvers. And it goes without saying that the proper modeling of a 40pS discontinuity is quite relevant now that it is ¼ of the cycle time and on par with the edge rates.

The actual impedance of vias can vary significantly depending on their construction. Figure 4 shows the simulated TDR response of various 3D via models from a recent channel project superimposed on the same blue and red responses shown in Figure 3. Note that not only do actual differential impedances range from 60 to 110 Ohms, but the time delays can vary significantly as well; variations that require the use of a 3D solver to quantify and understand.

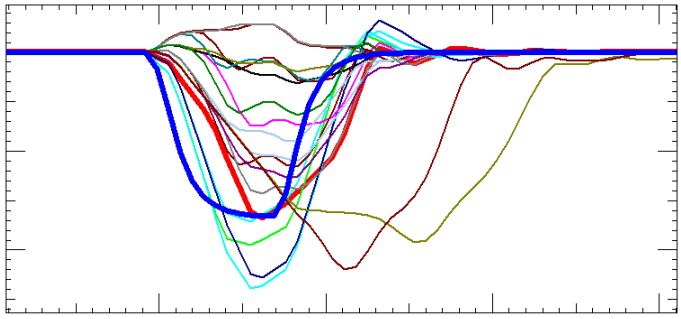

Figure 4: Range of via time domain characteristics

While differential via characteristics vary significantly by construction, an even wider variation might be seen if the 3D field solution is not performed correctly. Figure 5 demonstrates potential impedance and loss variations for 10 mil (drill size) differential vias in a 120 mil thick PCB. The variations shown are caused primarily by including (or not including) associated signal and ground vias in the solution. The correct solution includes both signal vias and their associated stitching grounds for a total of four vias, and is shown in red in both plots. The other vias were solved by using subsets of the four vias, such as a single via with no ground, one via with one ground, and both signal vias with no grounds. At left the differential impedance is shown to vary by up to 50%, erring on the side of higher impedance primarily due to the missing grounds. At right, the loss seen by a differential signal is shown to vary by more than 1 dB for the various solutions. As the total loss in our reference channel is around 12 dB, incorrect via field solutions on the six vias in the channel can cause a 50% error in total loss (and hence simulated signal amplitude).

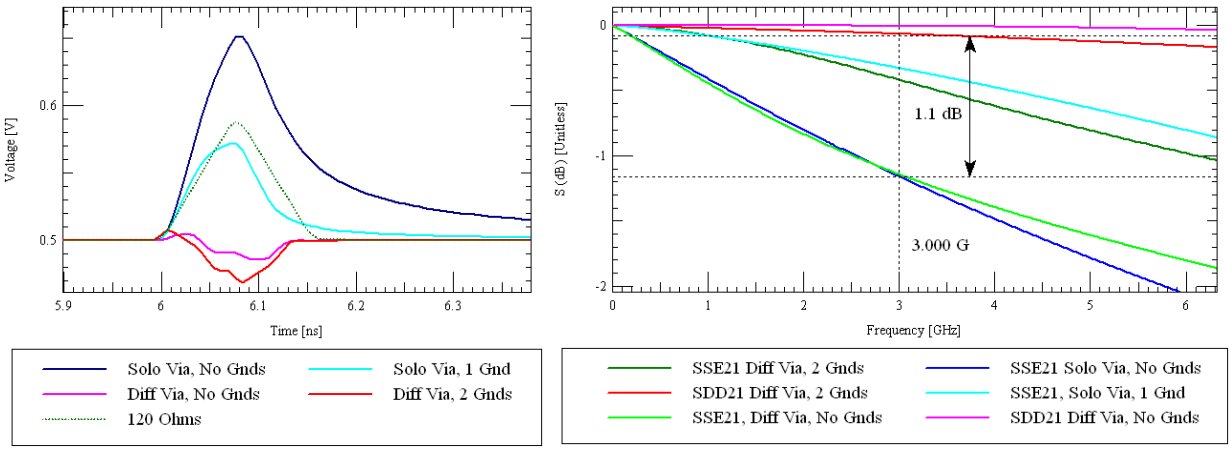

Figure 5: Impedance and Loss Variations in the Same Via Solved Different Ways

Correct setup of the 3D geometry and solution space is also critical for correct characterization of via structures. Special attention needs to be paid to the boundary conditions and the excitation ports. Since all channels need to be divided into smaller subsections that can be efficiently modeled in 3D field solvers, the excitation port represents the point of concatenation where the models are joined. Incorrect excitation and port geometries can lead to overly pessimistic or optimistic results. The excitation of the via structure needs to closely resemble the geometry of the interfaces that define the model. In Figure 6 at right, the excitation on the top of the PCB

model contains a portion of the connector pin. This is required to correctly model the ground path transitions from the PCB model to connector model. It also correctly models the impedance around the surface pads of the PCB.

Each sub-model extracted from the interconnect path contains artificial boundaries used to control the numerical solution space. The boundaries provide a way to limit the calculation and contain the 3D geometry, yet have no true physical structure. It is critical to understand how different boundary conditions affect the overall 3D solution. For the model below, the boundary condition on the perimeter of the model was varied in distance with respect to signal pins. As the volume of the model is reduced by 1mm, the model shows up to 20% more insertion loss in the graph at left. This is a result of the radiation boundary condition used in the model being in close proximity to the field strength of the excited signal path. This energy is absorbed by the boundary condition in the model as radiation. Radiation boundaries are commonly used for via models to effectively make the planes infinite in size. Perfect conducting walls can produce cavity resonance within the 3D structure, since they can change the structure into a sealed box. Since 3D electromagnetic modeling is a computationally intensive task, the model is generally reduced to minimal size for increased simulation throughput. This reduction is size can lead to erroneous results when the boundary conditions are not defined correctly. Correct field solution of differential vias with their associated grounds is essential when performing channel simulation.

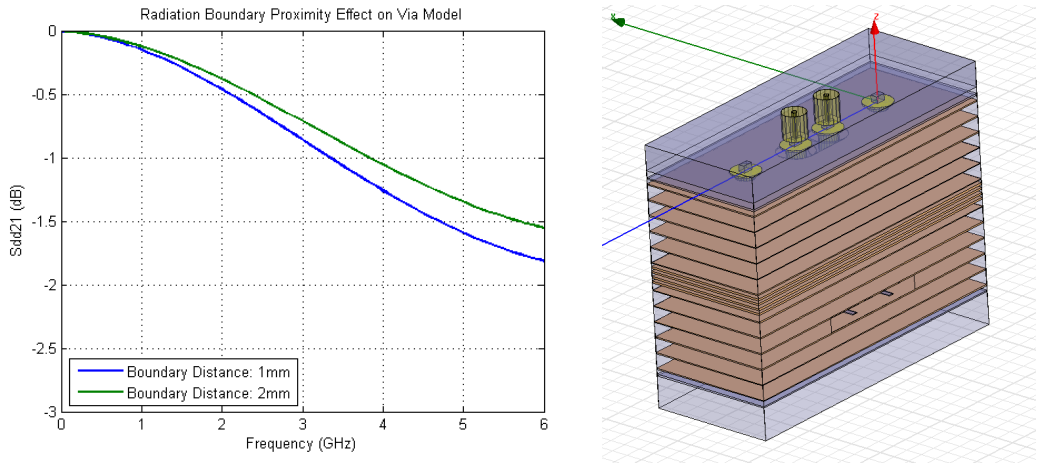

Figure 6: Boundary Condition Effects on Via Models

Via stubs cause a rapid increase in loss due to reflected energy that cancels out the signal at a frequency related to the length of the stub. Bogatin derives a rule of thumb [14] that that shows this frequency to be about equal to 1.5/stub (stub in inches, frequency in GHz). This typically causes vias with shorter layer transitions (yet longer stubs) to have more loss than longer vias, which might be considered not intuitive. Figure 7 shows the differential insertion loss in another set of vias from our 120 mil PCB. They represent the same via construction connecting the top layer to the lowest routing layer (in red), the highest routing layer (in blue), and various layers in between these (greens). This causes various lengths of stubs, the longest of which causing a null at ~12.5 GHz (=1.5/0.12, using Bogatin's rule of thumb above) as seen at left. At right we see the differential loss can vary by over 2x at 6 Gbps, and even more at higher data rates.

Consequently, though they have more total length, vias to lower layers are typically a better choice than those to upper layers.

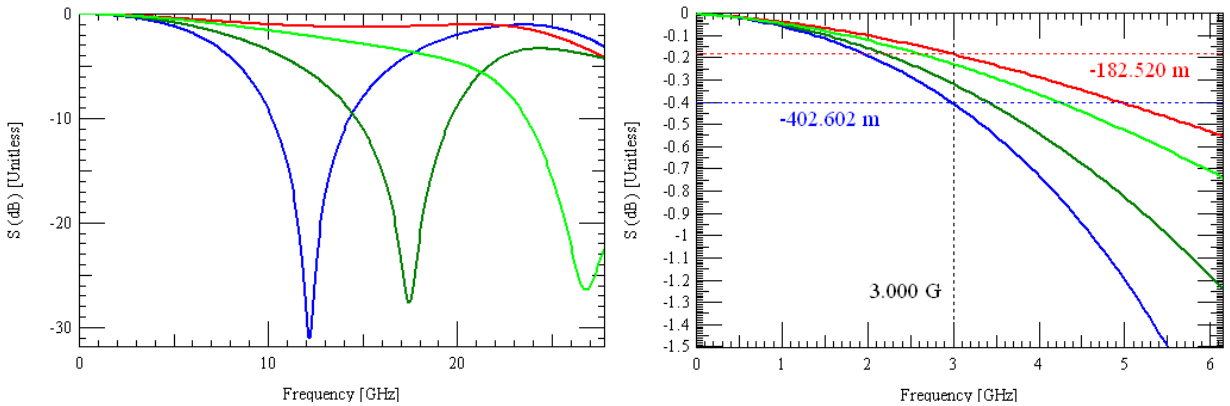

Figure 7: Differential Loss Variation in Vias Due to Stub Length

In thicker backplanes the loss created by stubs is typically problematic, so back-drilling is used to remove them. To demonstrate the performance gained by back-drilling, Figure 8 shows an upper layer via in a ~1/4" backplane with and without back-drilling. At left, the loss is shown to improve by ~7x (single-ended and differential insertion losses shown, lighter shades are back-drilled). At right the impedance improves by more than 60% from 45 to 73 Ohms; significant improvements.

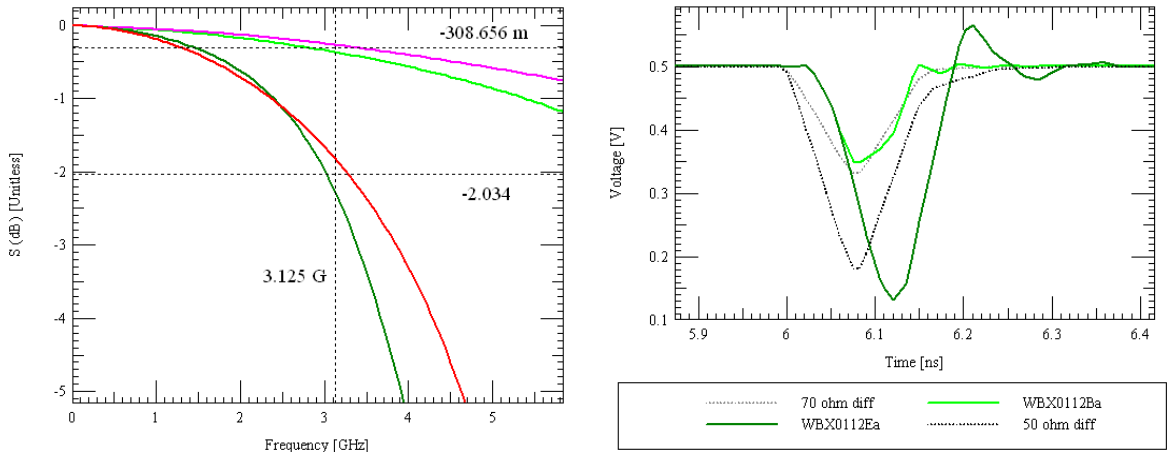

Figure 8: Via Loss and Impedance Improvements, with and without Back-drilling

Figure 9 shows how the performance of a via can be optimized through process and design improvements. In red - a decently accurate field solution with two vias and their two grounds - shows the starting point. Removing metal (and hence capacitance) from the system by elongating antipads, voiding unused power planes, and de-embedding trace length gained ~5 Ohms as shown in green. After further refining material frequency dependence, connection mechanism, and error tolerances an additional 5 Ohm improvement is shown in blue. Note that the plot includes 70 and 80 Ohm reference traces, which illustrate a ~10 Ohm improvement overall along with a general shortening of the time delay of the discontinuity.

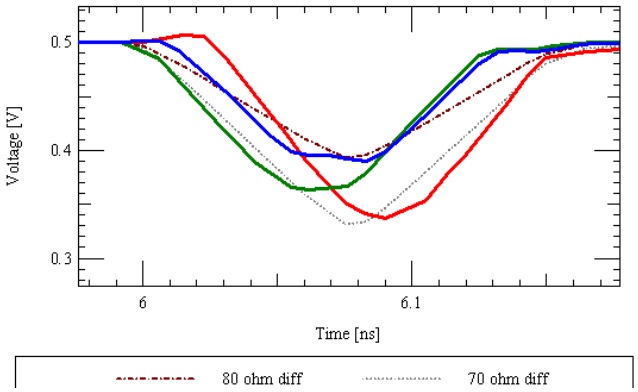

Figure 9: Better Via Impedance Through Design and Process Improvements

This section has detailed the proper modeling and use of vias which must be achieved for accurate 6+ Gbps channel simulation.

2.2 SerDes Modeling

This section provides more detail on the construction of approximate SerDes models for rapid simulation in a channel simulator. Note that for the types of analysis described in this paper, it is important to have a SerDes model that simulates rapidly for both impulse response characterization and subsequent channel simulation. A typical SerDes behavioral model has two sections: analog and equalization. The analog portion is typically relevant for impulse response characterization, while the equalization portion is deployed during channel simulation. Both portions of the model are detailed separately below.

SerDes deployed in 6+ Gbps links typically implement Feed-Forward Equalization (FFE) in the Tx and Decision-Feedback Equalization (DFE) in the Rx. This paper assumes that either generic or vendor-provided AMI models are available to overlay these types of equalization or signal processing. If Rx pin-level specs are available and a signal can be extracted at that location, the Rx DFE model may not be required. An example AMI call in a Tx model is shown next. Six taps are configured: 2 Pre-cursors, 1 Main cursor, and 3 Post-cursors. Offset sets two pre-cursors, and the values configured in this example are: Pre1=-2%, Pre2=-10%, Main=100%, Post1=-50%, Post2=-25%, Post3=-10%. They can be set to any values allowed by the actual device and the voltage swing can be adjusted to amplitude settings supported by the device by scaling the Main cursor.

```
(ami_ffe <path to AMI DLL file>
 (fwd 6 ) (offset 2 )
 (fwdtaps "160e-12 6 -0.02 -0.10 1.0 -0.50 -0.25 -0.10" ) )
```

The analog portion of the SerDes model implements the various datasheet specs shown Table 1, and can be implemented with only a few elements. As some SerDes specify “return loss” (a measure of how much energy reflects off the component) instead of die capacitance, it may be necessary to combine the analog model with the vendor’s package model and adjust model parameters until the return loss is within limits. Both higher die capacitance and larger termination values cause greater return loss. Note also that rise/fall times are typically specified at the pin, and also need to be adjusted with the package model in place.

Spec	Fast	Typ	Slow	Unit
Vswing	500	500	500	mV
Tr, Tf	30	50	70	pS
p/n skew	0	2	5	pS
Z_term	80	100	120	Ohms
C_die	0.5	0.75	1	pF

Table 1: Typical 6+ Gbps SerDes Behaviors

The Rx model would implement only Z_term and C_die. While Vswing can be varied in the analog model, it is held constant here anticipating that it will be adjusted using the AMI model instead for greater flexibility. Other SerDes specs such as Tx jitter (Rj, Dj, DCD) are typically implemented in other places in the channel simulator. Figure 10 shows waveforms that might be obtained by connecting the analog portions of a 6+ Gbps Serdes Tx model to an Rx across the various device corners such as fast (red), typical (green), and slow (blue).

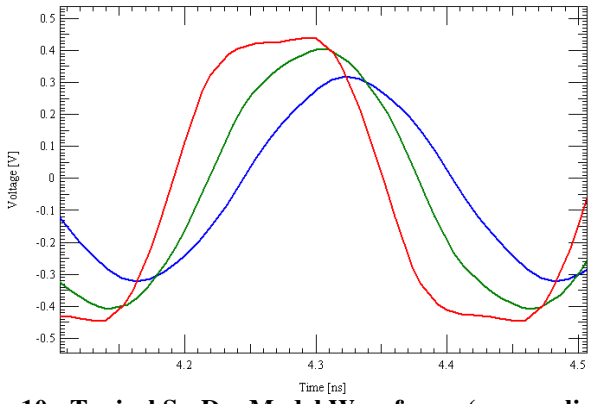

Figure 10: Typical SerDes Model Waveforms (no equalization)

With the relevant parameters implemented in a behavioral SerDes model, it is now possible to quantify the channel's anticipated eye and BER performance. Note that if a more complex SerDes model is provided by the vendor, it can also be used in the channel simulator to verify and correlate important conclusions formed by using the behavioral SerDes model. Either way, it is important to understand that valuable channel simulation can be performed even though there is a present lack of efficient and transportable vendor-supplied SerDes models.

2.3 Other Items in the Channel Model

Traces are another important element of the channel. Their geometry affects the interconnect electrical characteristics, like insertion loss and crosstalk. When choosing the trace width, trace to trace and pair to pair spacing, dielectric material type, thickness and weave style complex trade offs must be made to optimize the stackup for both cost and performance. Parameters like routing density, board thickness, loss and crosstalk figures, to name just a few, need to be quantified. This aspect of system design, indeed very important, is not the focus of this paper as it has been addressed from a variety of perspectives elsewhere [17–21].

In addition, the modeling, impedance control, and tuning of the AC coupling capacitors and their associated break out pattern can be easily overlooked. As with all via and pad structures, this

region needs to be tuned for optimal performance. The location of the AC coupling capacitors in relation to the SerDes package makes this of critical importance. Reflections between the Serdes receiver, package, BGA footprint, and AC coupling capacitor have very little dissipation between the elements due to the tight spacing, and can become problematic. As will be shown later, the physical structures close to receiver have the highest level of sensitivity in relation to to link performance.

3. Loss Analysis

With all the passive components of the link modeled and assembled, it's important to quantify the end-to-end loss. This can be done both by summing the individual elements, or extracting the end-to-end S-parameters of the channel.

Figure 11 shows the cascaded elements in the reference channel, followed by a differential loss summation of the individual elements to 12.11 dB. Note that the Tx and Rx packages are not included since they are not considered part of the channel loss under control of the system designer.

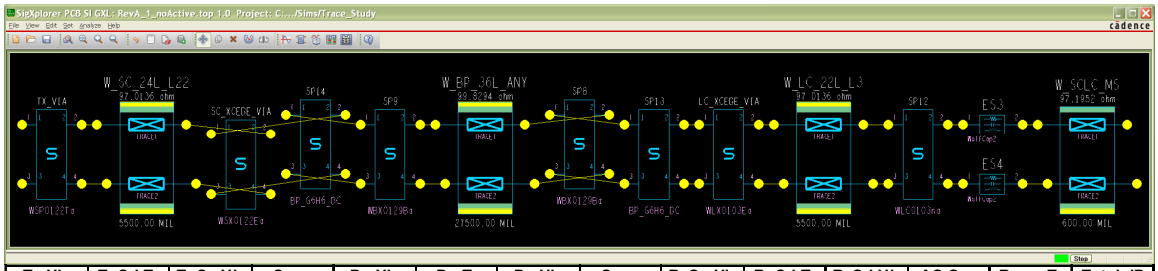

Tx Via	TxCd Tr	TxCn Via	Conn	Bp Via	Bp Tr	Bp Via	Conn	RxCn Via	RxCd Tr	RxCd Via	AC Cap	Rx us Tr	Total dB
0.12	1.97	0.21	0.51	0.40	5.38	0.40	0.51	0.31	1.97	0.12	0.01	0.20	12.11

Figure 11: Differential Loss Summation of Models in a Reference Channel

The summation above agrees with an insertion loss plot of the same circuit in Figure 12. Red is SDD21, green is SSE21 (SE=Single-Ended), blue is SCC21. This also shows 12.1 dB for our reference channel. This raises confidence that the channel model is functioning correctly, and also suggests that we should only see 25% of the Tx signal at the Rx. Note that although most trace models propagate differential and common signals equivalently, other components in the channel in our example such as vias and connectors present less loss to the differential signal than the common signal. As such a variation is seen in the red and blue plots.

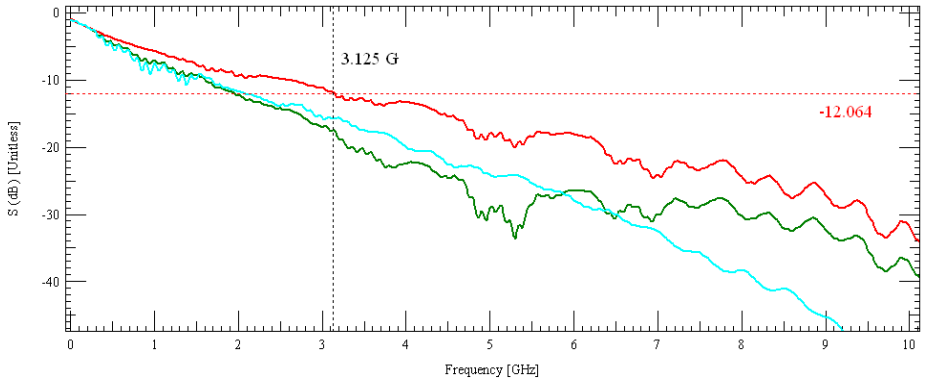

Figure 12: Loss in Reference Channel

It is critical to analyze the differential SDD21 response of channel as opposed to the SSE21 response. An ideal interconnect could be described with highly linear low loss SDD21 and strongly dissipative SCC21; such a system would be immune to common mode influences and skew along the path. The response of SSE21 alone doesn't accurately describe these elements and sometimes is mistakenly the focus of system designers. While stripline and other trace geometries show common performance between SSE21 and SDD21 responses, connectors and vias typically have significant difference in the responses of these different modes. More complex 3D structures, such as connector and vias, have strong coupling between the halves of the differential pair, while strip-line geometries strongly to couple to reference planes.

Depending on the targeted link type, it may also be relevant to compare the channel's loss against various pre-defined masks and standards. Although it is not a direct measure of performance, it is becoming typical to compare channel loss against masks found in Annex 69B of the 802-3ap standard [15]. Examples of these comparisons for the reference channel are shown in Figure 13 for the channel's Max Insertion Loss (left), Insertion Loss Deviation (center) and Return Loss (right). Defined mathematical post-processing of the loss data is required to generate these plots.

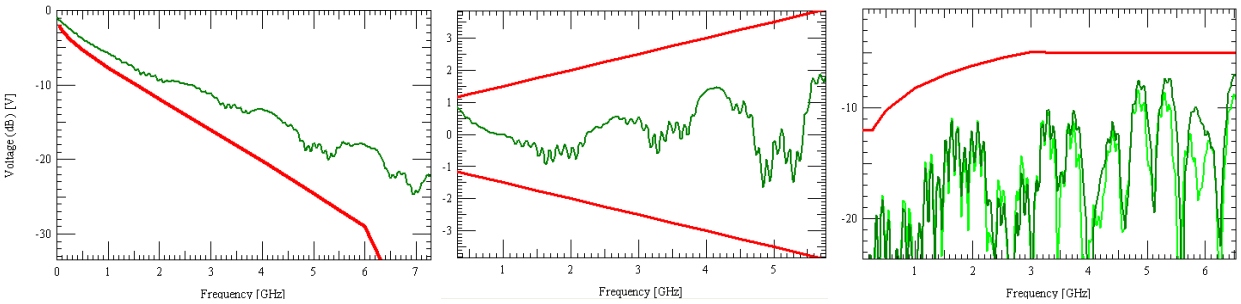

Figure 13: Reference Channel Loss Compared to Defined Masks

4. Channel Simulation

With an accurate channel model in place, and loss cross-correlated, eye and BER performance can be measured and compared against various permutations of the channel. In this way, design choices, materials, and component trade-offs can be made to influence and minimize loss and discontinuities.

Standard metrics for quantifying channel performance are signal eye height and width at the Rx, as shown in Figure 14. Eye height is a bit less than the 250 mV that might be expected in a channel with 12 dB of loss due to various reflections from channel discontinuities (a source of ISI, or Inter-Symbol Interference) as well as extra loss in the SerDes package. Furthermore, the hour-glass shape of the outer eye suggests that the signal arrives at the Rx somewhat over-equalized. Eye width is often assessed with statistical jitter sources applied (as shown at right), from which the width is extracted at different quantities of bits that imply a certain BER. In the Figure the width is determined at a typical value of 1e12 bits, but other values could be used as well. The values shown here apply simple FFE equalization, and will provide a baseline measure of performance against which system variables can be tested.

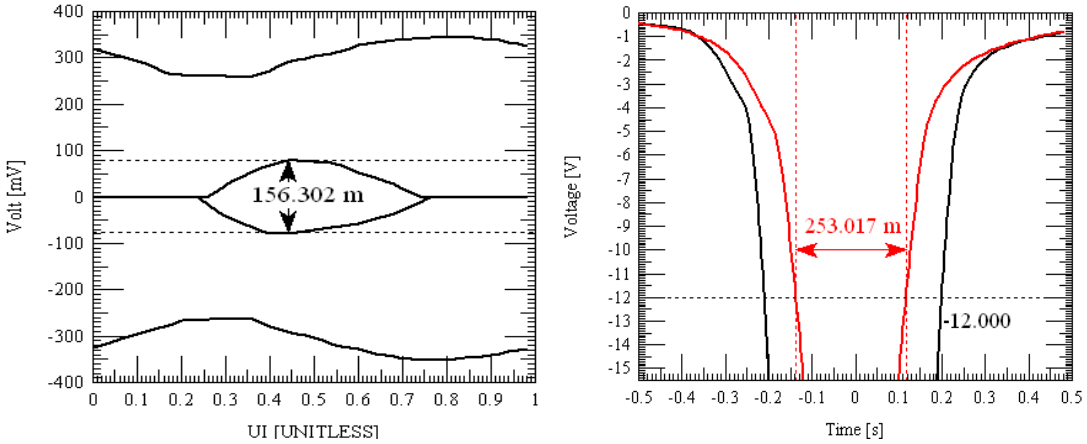

Figure 14: Eye Height and Width are Standard Measures of Channel Performance

Before detailing the various types of analysis, it should be noted that – while we typically use millions of bits for analysis – variation in the types of data patterns used also has a significant effect on eye performance. For a given link, the types of data encoding and/or the patterns against which performance should be measured is typically well-defined. Figure 15 shows that while eye height may vary up to 25% with change in common data patterns, no substantial change in eye width was noted for our reference channel. In the Figure, eye height is normalized to the compliance jitter pattern (CJTPAT) and declines with larger PRBS values, as might be expected. The performance of 64b/66b encoding is shown to lie somewhere between PRBS11 and PRBS23.

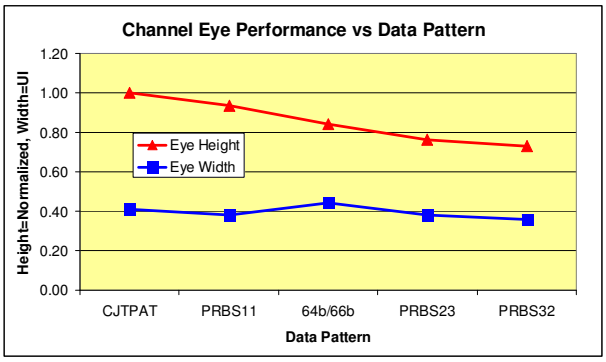

Figure 15: Eye Performance vs. Common Data Patterns

4.1 Sensitivity Analysis

There are many variables in a manufactured channel, all of which are subject to certain tolerances. To gain an understanding of which variables might be more relevant than others, "Sensitivity Analysis" identifies the primary variables in a channel and then varies them one at a time to better understand the system's sensitivity to each.

Table 2 lists the primary variables associated with each PCB in the reference channel. For example, the routes on the Tx Card have an impedance tolerance, route length range, and might be on upper or lower layers (this primarily affects the via stub length). Furthermore, the Tx SerDes might be operating at various corners (fast/typ/slow). The Rx Card and the Backplane have similar variables, while connector pin variation is associated with the Backplane.

Tx Card				Backplane					Rx Card			
Tx	Card	Length	Layer	TxConn	Bp	Layer	Length	RxConn	Card	Length	Layer	Rx
Typ	NomZ	Max	Lower	Conn1	NomZ	Lower	Max	Conn1	NomZ	Max	Upper	Typ
Fast	HighZ	Min	Upper	Conn2	HighZ	Upper	Min	Conn2	HighZ	Min	Lower	Fast
Slow	LowZ			Cpnn3	LowZ		Typ	Conn3	LowZ			Slow

Table 2: Primary Variables in Reference Channel, and Baseline Configuration

The row of variables highlighted in yellow establishes a "baseline" configuration against which all variable corners will be compared. As such, there are six variations associated with both the Tx and Rx Cards, and seven for the backplane since the Conn values track each other.

The eye performance (height and width) for each simulation are graphed in Figure 16. The points furthest to the left marked as "Baseline" provide the reference simulation to which all other variations are compared. The rows in the graph are configured so that the middle band is centered on the Baseline values for both height and width. This allows us to quickly observe which variations stay within a 10% variation (the center band) and which ones cause data in the outer bands. The outer bands help determine which variables cause the most improvement (upper bands) or degradation (lower bands) to the eye.

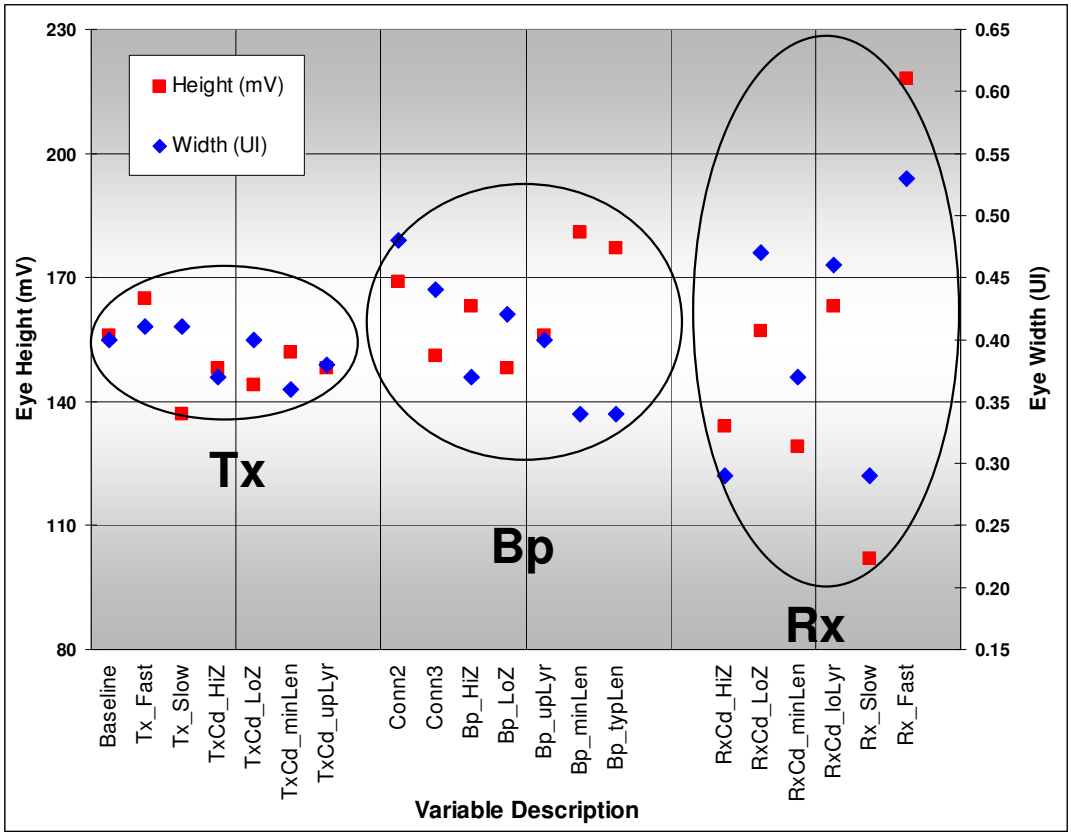

Figure 16: Eye Performance Results from Sensitivity Analysis

The ovals in the Figure help quickly identify that the widest variations are caused by variations and discontinuities on the Rx card. As such, optimizing and controlling the Rx design has the greatest potential for contribution to the overall system design. This will be discussed in more detail later.

Before discussing the variables that cause the widest variation, it's instructive to examine which variables remain within the center band and hence do NOT cause much variation. These are not all intuitive. Most notably (from left to right):

a. **Tx Discontinuities.** All discontinuities on the Tx card (higher/lower impedance traces, less desirable upper-layer vias) have a slight negative effect – even the minimum length trace to the connector. As such, there is likely no preferred direction to intentionally skew manufacturing variables to improve performance. Though not a major contributor to system-level performance, it's worth our effort to reduce Tx card discontinuities whenever possible.
b. **Backplane Impedance.** Backplane trace impedance variations show a measurable yet not too significant change from the baseline, conversely affecting height and width (ie, height gains are at the expense of more jitter while height loss improves jitter). These can likely be tolerated without paying extra for greater impedance control.
c. **Backplane Layers.** Changing backplane routing layers had no measureable effect on our reference ~1/4" thick backplane. Due to a consistent stackup and via back-drilling only the via length changes, and that change (very slight additional delay and loss) is not perceptible in the Rx_input eye shape. The behavior we observe in this example simplifies the backplane routing task since any layer can be selected.

All other variations cause data excursions beyond the center band, and hence need to be understood and managed more carefully. Observations from these variables include:

1. **Tx Swing.** Tx corners (impedance, edge rate, etc.) affect the eye as expected (voltage scales across channel) yet have little effect on jitter. Figure 17 shows that scaling the Tx was found to linearly improve eye height at the Rx without introducing jitter. This raises confidence that the Tx amplitude can be scaled in the SerDes without introducing adverse effects - such as non-linear eye shape or excess jitter - in order to gain more Rx eye opening if necessary. The down-side to this scaling is primarily increased power.

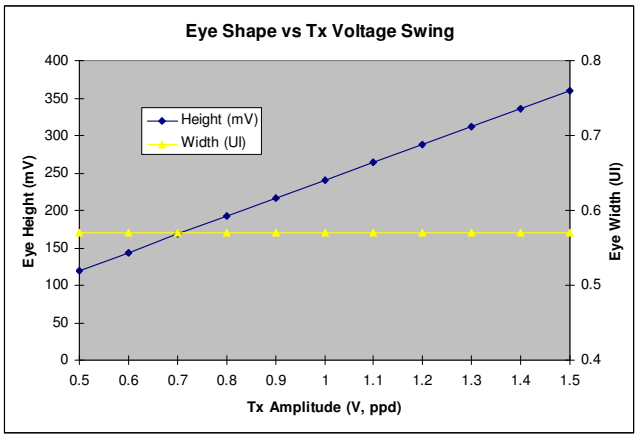

Figure 17: Eye Performance vs Tx Amplitude

2. **Connector Impact.** Changing connector pins has more effect on performance than expected, and is worth examining in more detail. Figure 18 shows little difference

between the SDD21 insertion loss of two different length pins for connector only simulations, especially around 3GHz. However, when the vias are added, the insertion loss of the two pins can vary by up to 1dB. The connector is attached to the PCB through via structures which are typically not an ideal impedance match. The ability to tune the impedance of these vias is often constrained by manufacturing and cost constraints. Reflections from these impedance discontinuities can create a cavity resonance or standing wave. The frequency at which this resonance occurs is dependent upon the delay of the connector pins. Longer pins have a lower frequency resonance. It is important to sweep the different length pins in the overall channel analysis to understand the interaction of the connector system and its associated vias.

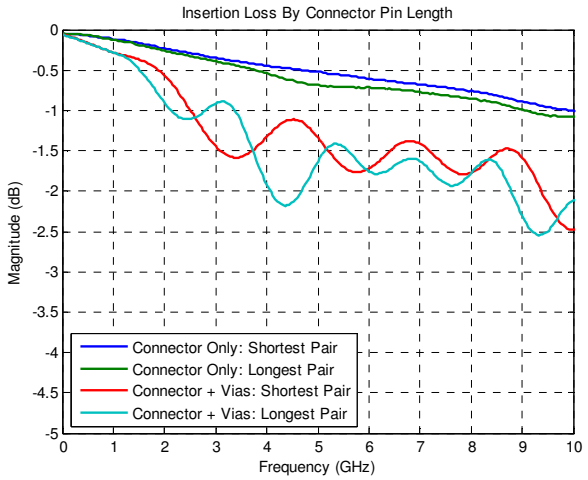

Figure 18: Insertion Loss Effect by Pin Length, with and without Vias

3. **Backplane Lengths.** Shortening backplane lengths causes an anticipated increase in eye height (less loss) at the expense of increased jitter (over-equalized signal, hence greater ringing). There is benefit from defining route_length / Tx_tap_setting signal bundles, which will be discussed more detail later.
4. **Rx Discontinuities.** Discontinuities on the Rx card are more significant than those on the Tx. Routing on the lower layers (less via discontinuity) improves performance as shown. Short lengths on the Rx card adversely affects performance, and controlling the minimum length will be shown in the next section to add over 40% more margin to the system. This is not intuitive, since we might expect that minimizing lengths would reduce loss and provide for a better eye. However, this is not true on the Rx card where the relationship between eye performance and route length is quite complex. To illustrate this, consider the impulse responses of two channels in Figure 19. The total Tx and Rx card lengths remains constant so the overall lengths of the channels are the same, however eye performance in the channel with the shorter Rx card length (in red) is 20% worse. The impulse responses reinforce that the variation is due to Rx discontinuities since the two plots are identical until the roundtrip time to the Rx connector (~1.5nS later) where the primary Rx discontinuities are found. Since the channels have the same overall length, the noise after a complete roundtrip to the Tx is similar and in phase (but note that the noise from the Tx connector interface is out of phase due to the length variations).

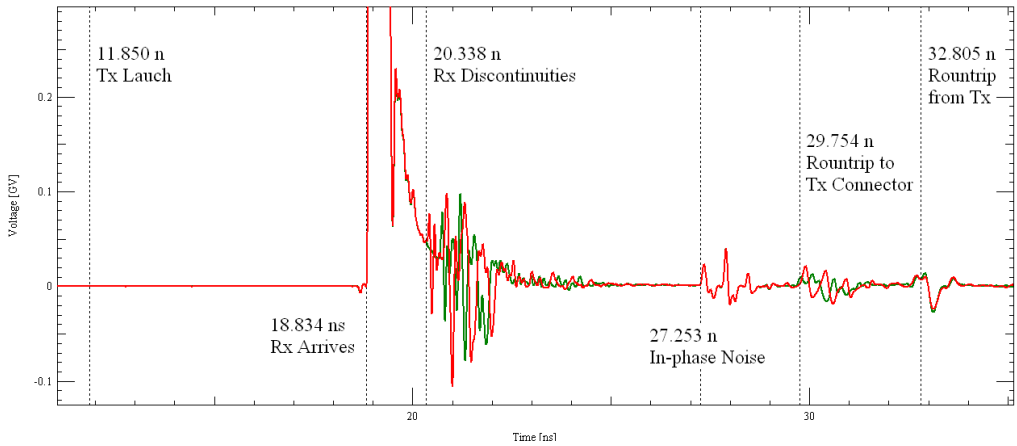

Figure 19: Impulse Responses for Channels with Same Length but Different Performance

5. **Rx Analog Parameters.** Rx die analog corners (impedance and die capacitance) have the most significant effect on performance when compared to all the other variables - ~35% increase/decrease from baseline for both height and width. Further analysis reveals that the variation is linear and hence somewhat predictable, with the impedance change having the greater influence (~3x in height yet similar in width) as shown in Table 4.

Simulation	Height (mV)	% Change	Width (UI)	% Change
SlowCorner	102	-35%	0.3	-25%
SlowImpOnly	114	-27%	0.34	-15%
SlowCapOnly	142	-10%	0.36	-10%
Baseline	157	0%	0.4	0%
FastCapOnly	173	10%	0.46	15%
FastImpOnly	207	32%	0.45	13%
FastCorner	221	41%	0.52	30%

Table 3: Eye Performance Variation vs Rx Analog Parameters

4.2 Net/Equalization Bundles

The previous section determined the system's sensitivity to assorted system variables, and found that performance improvements can be achieved by defining net_length/equalization_setting signal bundles while controlling lengths on the Rx card. It is not surprising that different channel lengths have different optimal equalization settings, since more length brings more loss – and loss is the primary phenomenon equalization is meant to overcome. This section illustrates a process for bundling varying channel lengths with equalization settings.

It's not uncommon for a system to have a variety of backplane trace lengths ranging from the minimal slot to next-slot lengths to much longer lengths in the order of 20" or even 30". Even in a backplane using good PCB materials with trace loss at ~0.2 dB/inch such length difference may easily result in a variation of 5 dB, or up to 40% of our system loss. In addition, the length/loss variations on the cards of ~1.5 dB means the SerDeses may encounter a 70% change in loss throughout all the nets in the fabric. With a variation this wide it is difficult to pick a single tap setting that will perform well for all nets. As such, it's likely we can improve design margin by defining "net bundles" that have different tap settings in the Tx.

For the discussion that follows we will limit ourselves to a single post tap, or the ability to "de-emphasize" a non-transition bit by a certain percentage of the transition bit. To understand the relationship between this post tap setting and backplane length (card lengths held constant), consider Figure 20. The ideal setting for the 1st post tap as a percentage of the main cursor (most simulators include a tap optimizer that calculates this value) is plotted on the y axis. From the plot, we see the increase in the tap's percentage is fairly linear with increasing backplane length.

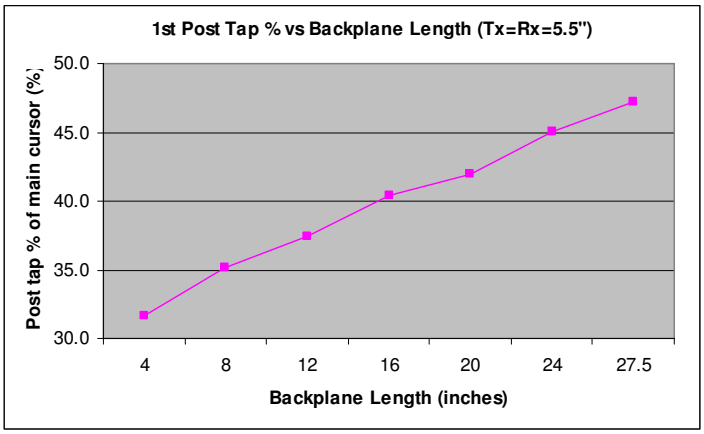

Figure 20: Optimal Post Tap Values vs Backplane Length

While the plot above assists with a conceptual understanding of the relationship of tap weight versus trace length, it is overly simplistic since there are many other variables in the system. Figure 21 plots 22 channel simulations of the reference channel using optimal post-tap settings with variations in the primary lengths. All three PCB lengths are varying, and the pink line "Tot_Len" is a summation of BP_len (purple), Tx_len (dark green), and Rx_Len (light green). The plot at left is sorted by total channel length increase, from 8 to 39 inches. The plot reveals that, among other things, for this channel optimal post taps (dark blue) might range from 22% to 47% - more than a 2x increase due to variation in length.

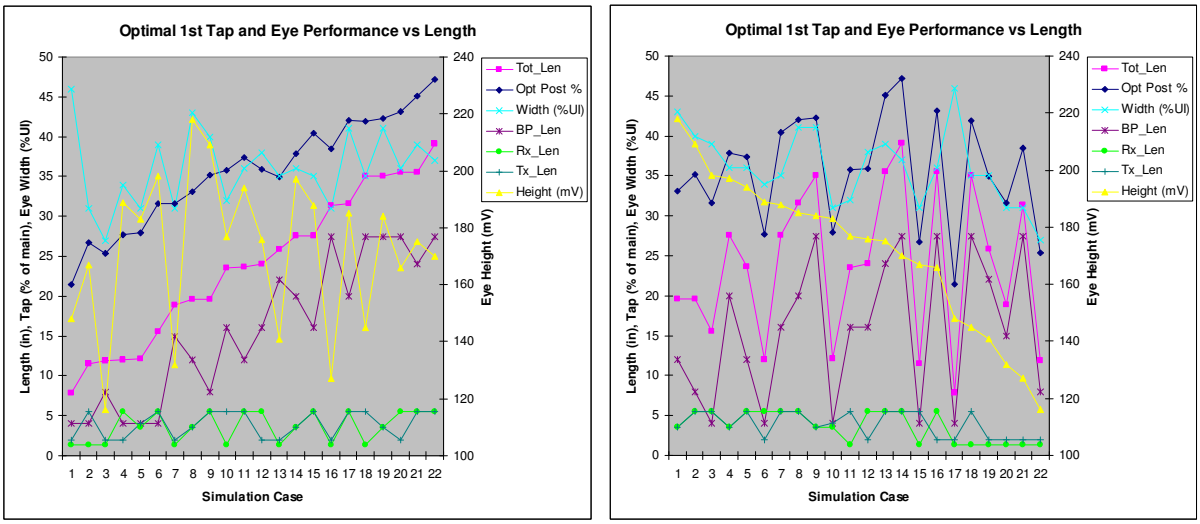

Figure 21: Eye Performance vs Varying Channel Lengths and Post Tap Values

Aside from getting a glimpse into the range of optimal post tap settings, the plot at left also reveals that eye height is swinging inconsistently +/- 50mV around 170 mV – and the variation is not related to total length or optimal tap values. Interestingly, all the minimum eye height values are related to the minimum lengths (primarily on the Rx card). This can be seen in the plot at right which sorts the simulation data by eye height. This again underlines the relevance of Rx discontinuities on eye performance.

Knowing the range of optimal tap settings (22 to 47%), the first step at deriving net bundles would be to test two tap ranges: 22-34, and 35-47. If two bundles are sufficient, we have greatly simplified system configuration. Since knowing the total length of a channel is difficult in practice, configuration is further simplified if those ranges can be based solely on the backplane length. Averaging the optimal settings in these ranges suggests we should test the shorter bundle (backplane lengths up to 12") at 28% and the longer bundle (lengths over 12") at 40%.

Simulating the channels with these settings yields the results shown in Figure 22, sorted by eye height margin. This figure illustrates that – even with non-optimal taps – the lowest margins still are associated with minimum lengths on the Rx card. From this we can hypothesize that if the Rx minimum length is raised to 3.5" the low margins in cases 15 to 22 would be removed and we'd gain at least 50 mV of margin. If more margin is desired, observe that – with cases 15 to 22 removed – the lowest margins now correspond to the longest channels. This is a "desirable" problem that can be fixed by either raising the Tx amplitude or creating a third net bundle with a larger post tap (~47%) for the longest channels.

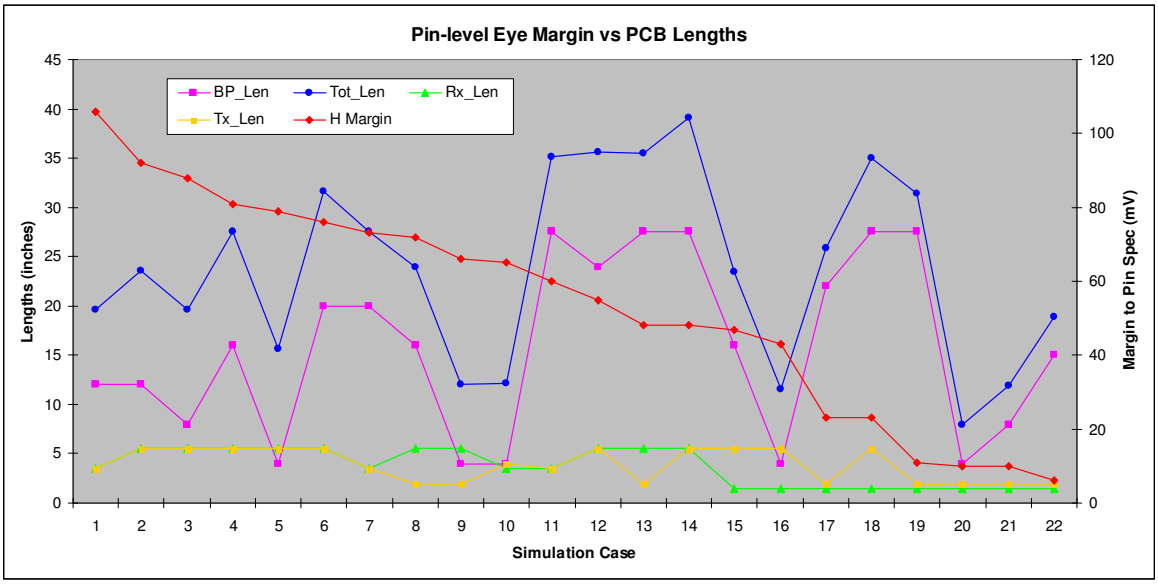

Figure 22: Eye Height Margin Using Two Net/Tap Bundles

This section has conjectured that good margin can be found by using backplane route lengths to define two tap weight net bundles, assuming the minimum length on the Rx card can be better controlled. These are likely good assumptions, yet they need to be further verified by exhaustive analysis that stresses all system variables as described next.

4.3 Exhaustive Analysis

It has recently become possible to perform exhaustive analysis of 6+ Gbps channels that comprehend system-wide channel variations across manufacturing tolerances and process corners. This has been achieved by reducing the time to simulate 100k bits in a channel permutation to ~2 minutes (including impulse response characterization) and automating the data extraction from batch simulations. More details on how to achieve this with various tools is explained elsewhere [16].

While the analysis in the previous section (Figure 22) guided us to design decisions that create margin against length variation using typical impedances and process corners, it's necessary to test the robustness of our design against the range of manufacturing tolerances anticipated in the final product design. Figure 23 plots the eye height margin against a random collection of manufacturing variations for the derived design at five process corners: FastTx/FastRx (FF), FastTx/SlowRx (FS), SlowTx/FastRx (SF), SlowTx/SlowRx (SS), and TypTx/TypRx (TT). This particular scenario utilized 27 channel characterizations with 5 jitter/data_pattern scenarios each, for a total of 135 simulations per process corner. While the Figure shows that the majority of the typical (TT) simulations do indeed have 50mV or more of eye height margin, there are a few permutations that do not. Furthermore there are numerous permutations at the FS and SS corners with even worse margins, as might be expected.

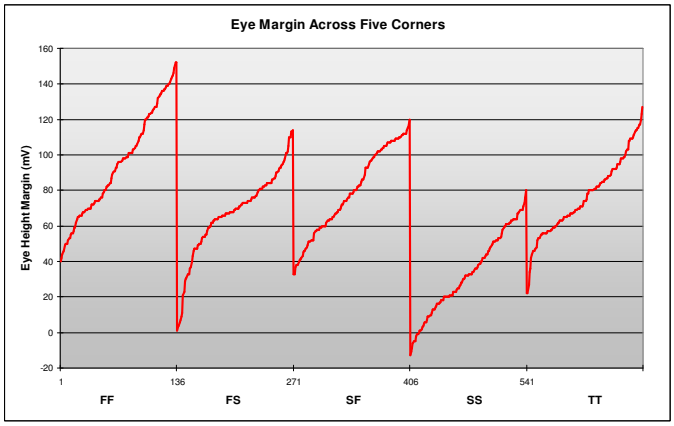

Figure 23: Eye Margin vs ~700 Simulations at Five Process Corners

This data guides us to focus on the FS and SS corners to find ways to create margin across an even more robust set of channels. As shown previously various solutions can be utilized such as:

- increase Tx amplitude to create height margin without increasing jitter
- adjust equalization to compensate for additional loss due to manufacturing tolerances
- move shorter channels into the shorter net bundles to avoid over-equalization.

After implementing various changes and testing them against a larger set of channels, the revised eye performance results are shown in Figure 24. The new set of channels systematically captures five impedance corners against five card length scenarios, superimposed upon twelve backplane lengths for a total of 300 (=5x5x12) channels processed at each corner with worst-case jitter. Measuring the results in two ways across three process corners results in 1,800 channel simulations requiring ~2.5 days to complete. The results show the desired improvement in height margin and the width margin is found to be fairly stable across process corners.

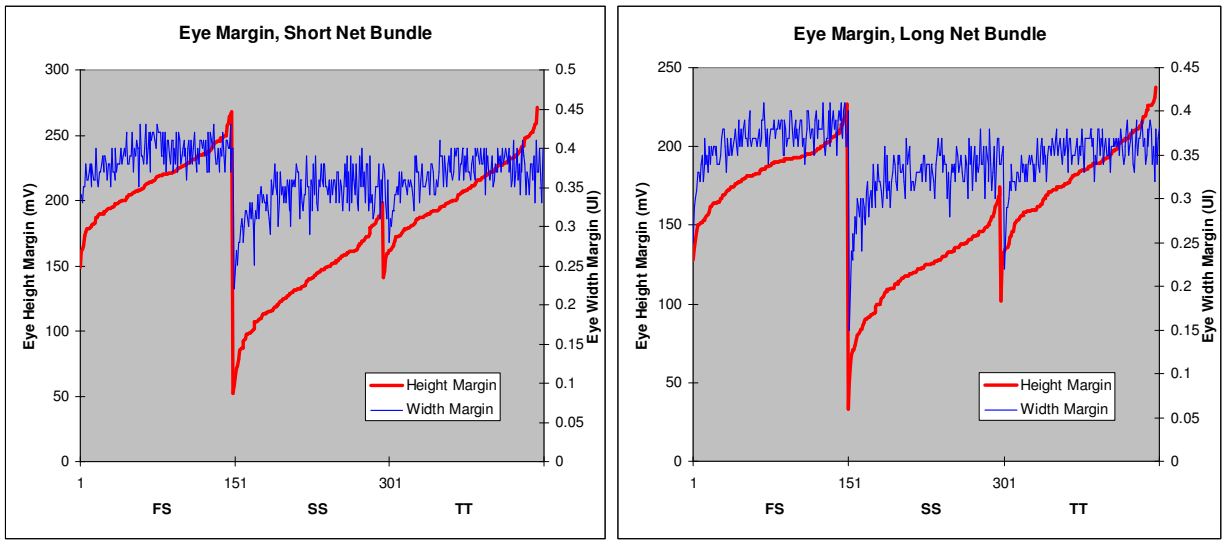

Figure 24: Exhaustive Analysis of Eye Margin vs 1,800 Simulations

The tails/range on the eye heights and the width variation can be narrowed further as desired by creating more than two net bundles, at the cost of increasing system configuration complexity. In some cases, it may be necessary to additionally bound the maximum eye height, which – though not addressed here – could be pursued in a similar fashion.

When balancing these trade-offs against increasing system complexity and power, the best answer becomes a matter of sound engineering judgment and experience. Thankfully, with a properly implemented simulation environment, the tools are in place to quantify performance and add the necessary data to the engineering process.

5. Summary

This paper has detailed simulation and modeling techniques relevant for 6+ Gbps channel design. While accurate trace modeling is taken as a given, the relevance of careful via modeling is discussed in detail. Nimble and efficient SerDes models are chosen in order better understand and constrain variables under the control of the system designer, trading model complexity for design insight. The design of a reference channel is used to illustrate ways to use sensitivity analysis to isolate important variables and improve design margins by 10 to 50%. With an accurate and well-tuned simulation environment in place, exhaustive analysis provides the ability to balance equalization options and design margin against design and manufacturing variations.

Acknowledgements

The authors wish to thank Steve Tam and Kusumakumari Matta at Ericsson for their support in pioneering new simulation methodologies as described here. Also to Chris Heard and Jose Paniagua at Amphenol-TCS for continued help in advancing 6+ Gbps signaling. Additional thanks to Brad Griffin and Aaron Tang of Cadence for their continued assistance with Allegro PCB SI GXL, and to Deana Woodrow and Oscar Wang of Ansoft-Ansys. Without the efforts of these and others this work would not have been possible.

References

[1] "Introducing Channel Analysis for PCB Systems", 2004 Cadence webinar, http://www.siguys.com/resources/2004_Webinar_Introducing_Channel_Analysis.pdf

[2] "New Techniques for Designing and Analyzing Multi-GigaHertz Serial Links", Telian, Wang, Maramis, Chung – DesignCon 2005, http://www.siguys.com/resources/2005_DesignCon_New_MGH_Techniques_ISP_CA_PCIe_SATA.pdf

[3] SAS Specification, SAS-2, Project T10/1760-D, Revision 16, 18 April 2009, http://www.t10.org/cgi-bin/ac.pl?t=f&f=sas2r16.pdf

[4] "New Serial Link Simulation Process, 6 Gbps SAS Case Study", Telian, Larson, Ajmani, Dramstad, Hawes, DesignCon 2009 Paper Award, http://www.siguys.com/resources/2009_DesignCon_6Gbps_Simulation_Paper.pdf

[5] "A Process for Serial Link Signal Integrity Analysis", Telian, Xrosstalk Magazine Jan09, http://www.siguys.com/resources/2009_XrosstalkMagazine_SerialLinkSI_Process.pdf

[6] "Signal Integrity Engineering in High-Speed Digital Systems", Telian, DCon 97 Best Paper, http://www.siguys.com/resources/1997_DesignCon_HighSpeedDesign.pdf

[7] "An Optimized Methodology for High-Speed Design" Telian, Perger, DesignCon '98 http://www.siguys.com/resources/1998_DesignCon_Optimized_SI_Methodology.pdf

[8] "Understanding and Using S-Parameters for PCB Signal Integrity", Cadence webinar, http://www.siguys.com/resources/2004_Webinar_Understanding&Using_S-parameters_for_PCB_SI.pdf

[9] "Frequency Dependent Material Properties, So What?", Bogatin, DeGroot, Warwick, DesignCon 2010 paper 7-WA3, http://www.designcon.com/2010/attendees/7_wa3/index.asp

[10] IBIS ATM Working Group, and the IBIS-AMI Standard, http://www.eda.org/pub/ibis/macromodel_wip/

[11] "Adapting SI Tools and Techniques for 6 Gbps and Beyond", Telian, CDNLive! 2007 http://www.siguys.com/resources/2007_CDNLive_Adapting_SI_Tools_for_6Gbps+.pdf

[12] "Signals on Serial Links: Now you see 'em, now you don't." DesignCon07 Article, http://www.siguys.com/resources/2007_Article_SignalsOnSerialLinks.pdf

[13] Much of the work described here (as in [4]) was performed using Cadence Allegro PCB SI GXL version 16.2. http://www.cadence.com/products/pcb/pcb_si/pages/default.aspx Consult your EDA and component vendors regarding their support of AMI models.

[14] "Signal and Power Integrity Simplified" by Eric Bogatin. ISBN 978-0132349796, pages 602-605, http://www.bethesignal.net/bogatin/bookinfo.php

[15] IEEE Std 802.3ap™-2007, 22 May 2007, "Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications"

[16] "Advanced Techniques for Channel Analysis", Donald Telian/Siguys webinar, http://events.unisfair.com/index.jsp?eid=497&seid=25

[17] "PCB Transmission Line Modeling for Multi-Gb/s Link Analysis", Liang, Hall, Heck, Brist, DesignCon 2005

[18] "Performance Limitations of Backplane Links at 6 Gbps and Above", Chan, Kirk, Paniagua, DesignCon 2008

[19] "Advanced Design Techniques to Support Next Generation Backplane Links Beyond 10 Gbps", Cartier, Chan, Cohen, Kirk, DesignCon 2007

[20] "Practical Design Considerations for 10 to 25 Gbps Copper Backplane Serial Links", Kollipara, Chia, Lambrecht, Yuan, Zerbe, Patel, Cohen, Kirk DesignCon 2006

[21] "Impact of Manufacturing Parametric Variations on Backplane System Performance", Kollipara, Chia, Lin, Zerbe, DesignCon 2005

Rev 1.1