# DesignCon 2014

# Moving Higher Data Rate Serial Links into Production - Issues & Solutions

Donald Telian, SiGuys
telian@siguys.com

Sergio Camerlo, Ericsson
sergio.camerlo@ericsson.com

Kusuma Matta, Ericsson
kusuma.matta@ericsson.com

Michael Steinberger, SiSoft
msteinb@sisoft.com

Barry Katz, SiSoft
bkatz@sisoft.com

Walter Katz, SiSoft
wkatz@sisoft.com

## Abstract

Large systems moving thousands of higher data rate (12 to 25 Gbps) channels into production encounter a variety of issues not seen in the first three generations of serial links. This paper describes issues and solutions recently discovered in bringing 4th generation products into the mainstream. Solutions are substantiated using simulation, field solution, VNA/TDR/Scope measurement, and photomicrography of fabricated structures. Issues and solutions arise and are hence detailed in two broad categories: design and manufacturing.


## Author's Biographies

**Donald Telian** is an independent Signal Integrity Consultant with SiGuys. 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. His numerous published works on this and other topics are available at his website www.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 (ESV), which he joined through the Redback Networks acquisition. His responsibilities include the Chassis/Backplane infrastructure design, PCB Layout Design, System and Board Power Design, Signal and Power Integrity. He also serves on the company Patent Committee and is a member of the ESV Systems and Technologies HW Technical Council. 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. Before, Sergio spent close to a decade at Cisco Systems, where he served in different management capacities. Sergio has been awarded fourteen U.S. Patents on signal and power distribution, interconnects and packaging.

**Kusuma Matta** has been working with Ericsson Inc. as a hardware engineer in the field of signal integrity engineering since 2007. Prior to Ericsson, she worked at LSI also in the field of signal integrity engineering. Her research interests include signal integrity for SerDes and DDR interfaces, board and package level SI optimizations, and VNA and TDR measurements. Kusuma has M.S. in Electrical Engineering from University of South Carolina and B.S. in Electronics and Communications Engineering from JNTU (Jawaharlal Nehru Technological University), Kakinada, India.

**Michael Steinberger, Ph.D.**, Lead Architect for SiSoft, has over 30 years experience designing very high speed electronic circuits. Dr. Steinberger holds a Ph.D. from the University of Southern California and has been awarded 14 patents. He is currently responsible for the architecture of SiSoft's Quantum Channel Designer tool for high speed serial channel analysis. Before joining SiSoft, Dr. Steinberger led a group at Cray, Inc. performing SerDes design, high speed channel analysis, PCB design and custom RAM design.

**Barry Katz**, President and CTO for SiSoft, founded SiSoft in 1995. As CTO, Barry is responsible for leading the definition and development of SiSoft's products. He has devoted much of his efforts at SiSoft to delivering a comprehensive design methodology, software tools, and expert consulting to solve the problems faced by designers of leading edge high-speed systems. He was the founding chairman of the IBIS Quality committee. Barry received an MSEE degree from Carnegie Mellon and a BSEE degree from the University of Florida.

**Dr. Walter Katz**, Chief Scientist for SiSoft, is a pioneer in the development of constraint driven printed circuit board routers. He developed SciCards, the first commercially successful auto-router. Dr. Katz founded Layout Concepts and sold routers through Cadence, Zuken, Daisix, Intergraph and Accel. More than 20,000 copies of his tools have been used worldwide. Dr. Katz developed the first signal integrity tools for a 17 MHz 32-bit minicomputer in the seventies. In 1991, IBM used his software to design a 1 GHz computer. Dr. Katz holds a PhD from the University of Rochester, a BS from Polytechnic Institute of Brooklyn and has been awarded 5 U.S. Patents.

# 1. Introduction

Since the year 2000, the highest data rate interfaces on Printed Circuit Boards have increasingly been implemented as serial links. Beginning with 1st generation links at 1.5 Gbps, the industry has steadily doubled data rates every four to five years. Today we are bringing 4th generation 12 Gbps interfaces into production while we explore technologies aimed at taking us to 24 Gbps and beyond. (Admittedly, high-volume interfaces such as PCI Express are out-of-sync in both introduction date and data rate from the generations described here. There were also lower data rate links prior to year 2000 – this is simply a convenient industry-wide starting point for the purposes of this paper.)

Over time, new design techniques and features in each current generation become requirements in the next, as shown in Table 1 below. For example, Transmit (Tx) pre-emphasis was available in some 1st generation components yet generally required in 2nd generation devices. Similarly, Receive (Rx) Decision Feedback Equalizers (DFEs) became as commonplace in the 3rd generation as back-drilling in the 4th, though both technologies emerged in previous generations. As such, it is important to understand advances in design and manufacturing in each generation before they become imperative in the next.

| (nice to have)<br>(required) | 1st Generation<br>1.5 Gbps | 2nd Generation<br>3 Gbps | 3rd Generation<br>6 Gbps | 4th Generation<br>12 Gbps | 5th Generation<br>24 Gbps |
|---|---|---|---|---|---|
| P/N Symmetry | | | | | |
| Tx Pre-emphasis | | | | | |
| Rx CTLE | | | | | |
| Rx DFE | | | | | |
| Via Design | | | | | |
| Via Back-drilling | | | | | |
| Antipad Etch Tapers | | | | | |
| Via Cage Symmetry | | | | | |

**Table 1: Serial Link Generation and Feature Migration**

This paper highlights discoveries made by the authors in the process of putting 4th generation links into production that are likely to become requirements in 5th generation links and beyond. Solutions are substantiated using simulation, field solution, VNA/TDR/Scope measurement, and photomicrography of fabricated structures. Issues and solutions arise and are hence detailed in two broad categories: design and manufacturing.

The context of these discoveries is the design and manufacture of a family of large-scale systems implementing multiple thousands of serial links spanning numerous standard and proprietary protocols. Similar 3rd generation discoveries related to the same family of systems are chronicled by the authors in [1] and [2].

# 2. Design Improvements

This section details 4th generation design improvements in the areas of channel design, firmware settings, and layout techniques.

## 2.1 Removing Discontinuities

As existing SerDes equalization is primarily aimed at compensating for loss in longer channels, the worst-performing channels tend to be limited by impedance discontinuities and are often the shortest [2][3]. For shorter channels, performance is improved by removing these discontinuities that cause excessive ISI and hence eye closure.

Figure 1 plot at left below illustrates performance gains if discontinuities are removed from a short channel. The plot in red shows how this channel's eye height changes by more than 100% as the route length between the Rx and its closest discontinuity is lengthened (RXLEN on X axis, eye height on Y axis). Note how performance oscillates in relation to the round-trip time between the elements, and how the oscillation dampens as length is increased. In green we see performance improves nearly 400% if only **two** of the largest discontinuities are improved (which, in this case, are vias). Blue reveals only incremental improvement if **all** discontinuities (6 more, in this case) are improved to their theoretical best case, suggesting the importance of understanding which channel discontinuities are limiting performance and to what degree. At right the pulse responses (same color scheme, green channel's unequalized response is dark green) reveal that equalization (lighter shades, with flat signal at -225mV) improves main-cursor pulse spreading, yet is unable to correct discontinuity-induced noise in the original (red and magenta) channel seen at ~4nS.

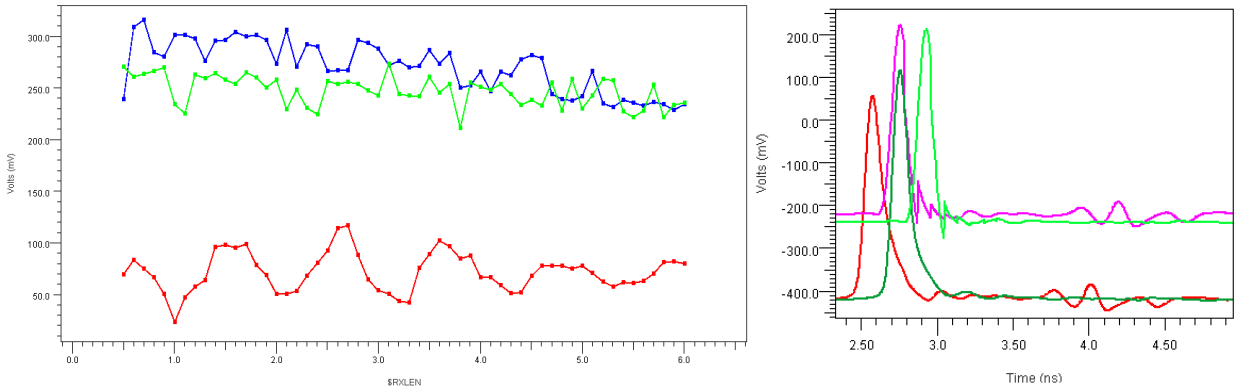


**Figure 1: Short Channel Performance Relationship to Discontinuities and Equalization**

Extending this same correction to hundreds of channels performing as shown at left in Figure 2 below (scatter plot of eye height on Y axis versus eye width on X axis), the plot at right reveals how dramatically performance can be improved if we implement this simple change on all channels. While the plots show that eye heights generally improve about 20%, the more significant change is the removal of interconnect/deterministic jitter from the system resulting in a very small distribution of eye widths. As shown, eye width variation is 4x wider at left. Furthermore, the eye width in the improved system is consistently as wide as the best widths seen in the plot at left. This performance improvement is impressive when you consider it

required only a slight adjustment to 1% of the channel's interconnect, demonstrating that the challenge is knowing which elements to adapt and in what way.

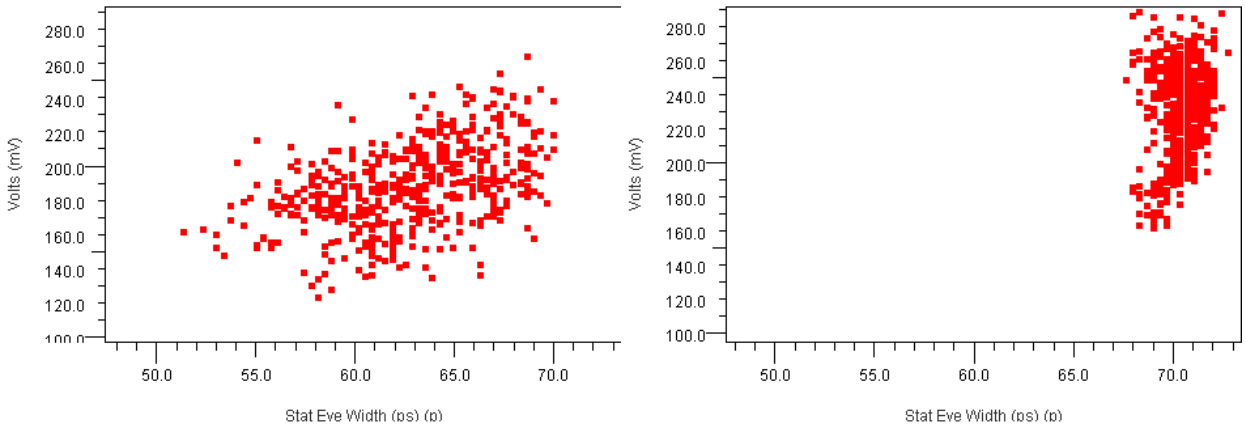


**Figure 2: Performance of Hundreds of Channels with (left) and without (right) Discontinuities**

The TDR plots in Figure 3 below illustrate how one of the channels has changed, as the large discontinuities seen at ~3 and 8 nS have been substantially removed. As noted above, the other discontinuities seen in the plots are not significantly limiting performance at this data rate.

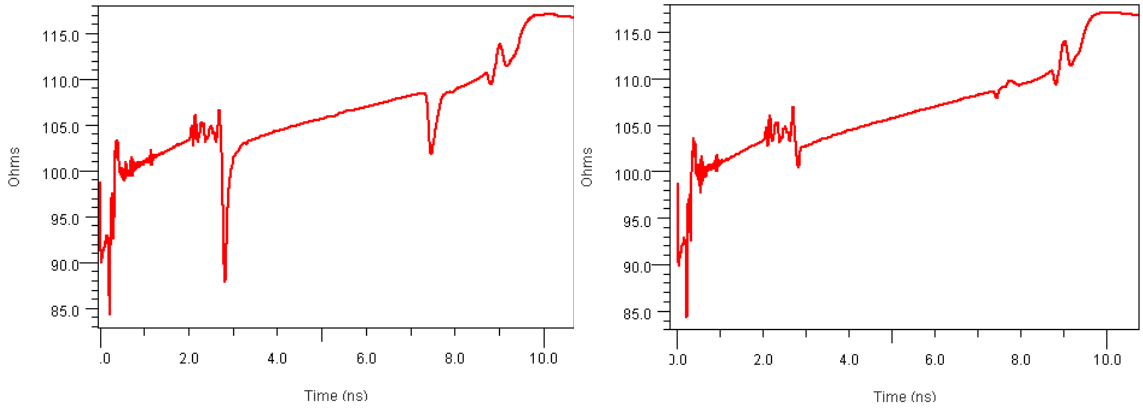


**Figure 3: TDR Plots Showing Sample Channel with (left) and without (right) Discontinuities**

## 2.2 SerDes Settings

As data rates increase, so do SerDes equalization features and setting options. Interestingly, power-up "default settings" are – in many cases – not well-optimized for the majority of channels. This occurs because defaults are targeted at worst-case channels, causing them to over-compensate with unnecessary noise and power. Correcting non-optimal settings requires interaction between hardware and firmware teams that can be difficult to achieve. This section examines how adapting settings improves performance for two systems: 12 Gbps native channels, and 12 Gbps components interoperating with 6 Gbps devices.

### 2.2.1 12 Gbps Channels

Driving default settings into over 1,000 channels ranging from one to three feet in length yields the performance distribution shown in Figure 4 at left (as in the previous section, eye height on Y axis versus eye width on X axis). In addition to a rather wide distribution of eye heights, two

performance bands are seen. At right we plot eye height versus channel length and observe it is fairly flat to the vertical line, after which eye height begins to slope downward.

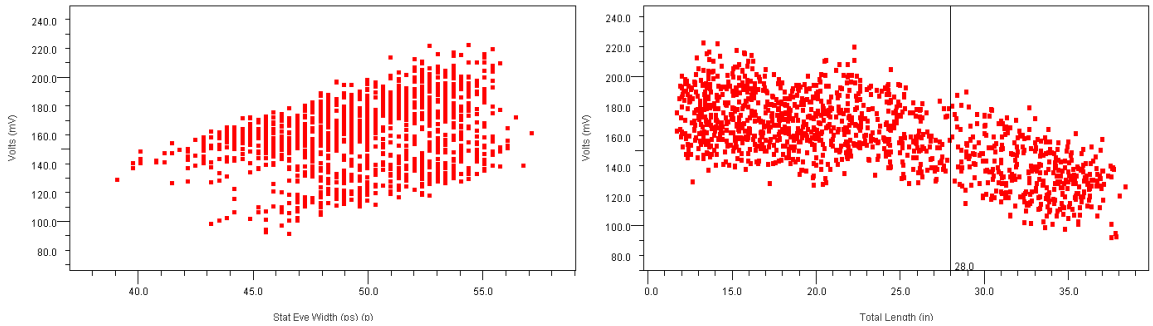


**Figure 4: Channel Performance with Default Settings**

Recognizing that eye heights are fairly flat until 28" channel lengths (vertical line, above right), we color channels with lengths greater than 28" blue in Figure 5 at left below revealing the bimodal distribution is caused by excess loss in the longer channels. This is further confirmed by calculating each channel's inherent "unequalizable energy" of the passive interconnect [3, page 19] as shown at right below. Pulse responses in longer networks – though they decay evenly – possess the largest amounts of unequalizable energy, as shown in green below right highlighting values greater than 3.0.

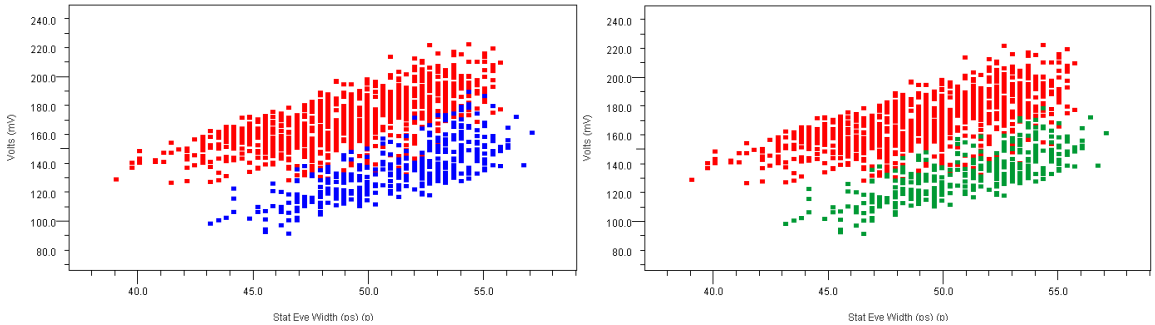


**Figure 5: Channel Performance Highlighting Long Lengths (left) and High UnEqNrg (right)**

As more amplitude/energy can be added to the Tx defaults we must determine if it should be added to the main cursor or post cursors. Testing these choices at the system-level we find adding Tx post-cursor decreases eye performance by more than 20% and causes Rx DFE post-cursor to decrease by the same amount at the Tx was increased. Indeed the two are typically interchangeable, as was also demonstrated in [2]. Adding the same amount of amplitude to the Tx main cursor we realize a 40% improvement in eye performance, as the energy is applied to signal transmission instead of equalization.

Adding additional amplitude to the longest channels we find the distribution is now better clustered with the amplitude on long channels now improved in both height and width (Figure 6 below, long channels again highlighted in blue). With bimodal behavior removed, we can now focus on correcting outliers in light of performance targets.

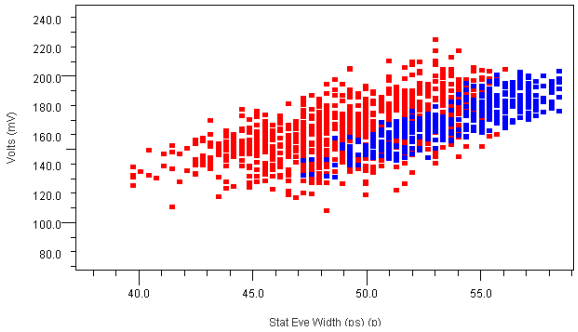


**Figure 6: Performance Improvement Adding Tx Amplitude on Longest Nets (blue)**

While eye heights and widths are good performance indicators, the system metric of primary interest is Bit Error Rate or BER. In Figure 7 below we highlight BERs below our target (with a bit of margin added in) in blue. This shows us graphically that it is the *combination* of eye height and width we care about, and allows us to establish a performance line in black. Note that near this line there are some passing and failing channels, again reflecting our margin. This line establishes a baseline performance target, allowing us to examine failure modes and design enhancements aimed at moving all channels above the line. While it is possible to tightly cluster performance as shown in the previous section, doing so would be over-design imposing unnecessary effort and expense as – in this case – only 12% of the channels need BER improvement.

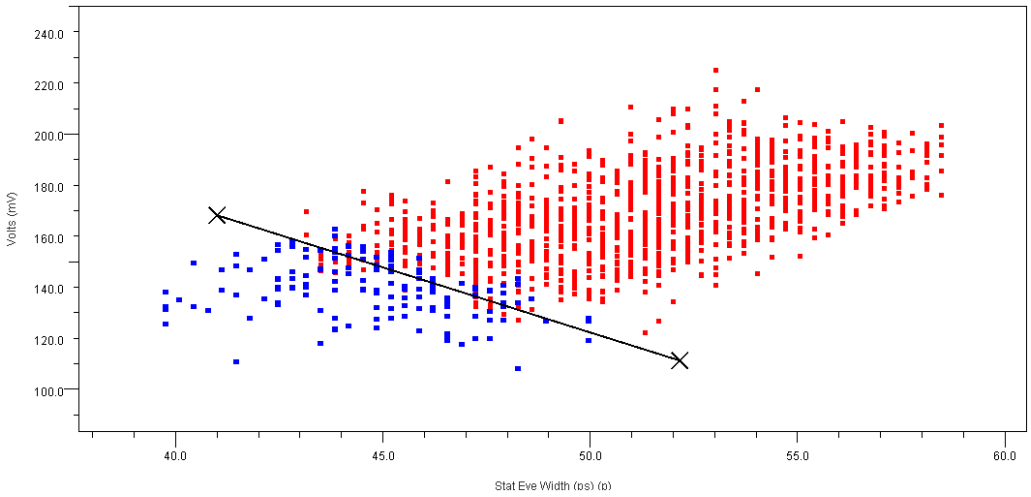


**Figure 7: Channel Performance, Failing BERs in blue**

Due to our large amount of channels, poor performance is related to both long and short channel problems, or loss and ISI respectively. As such, we next apply the two corrections that have been previously demonstrated to improve performance:

1. Moderately increasing Tx amplitude on medium-length channels, and
2. Improving discontinuities on shorter channels

Figure 8 highlight nets that can be improved by these two changes. At left in green we superimpose the 12% of channels with the highest unequalizable energy with equalization (a measure of non-cursor energy that cannot be removed by available Tx/Rx equalization [3]). This

metric isolates channels whose pulse response has signal amplitude beyond the time targeted by the available Rx DFE taps, which may be due to either a long decay (longer, high-loss channels) or excessive ringing (shorter, high-ISI channels). Note that this metric identifies many, but not all, of the signals below the black line. At center, in black, are channels with routes that are both short and have significant discontinuities (longer vias, in this case). Note the similarities between the green signals determined by a system-level electrical metric and the black signals sorted entirely by physical parameters. These signals would likely be shifted up and to the right if the channels are lengthened or the discontinuities improved. At right, in light-blue, are medium-length channels a bit shorter than the bimodal range corrected above. Applying more amplitude to these signals should move all light-blue points up and to the right. Note that the combination of light-blue and black signals selects all signals that had failing BERs below the line, in addition to various signals above the line.

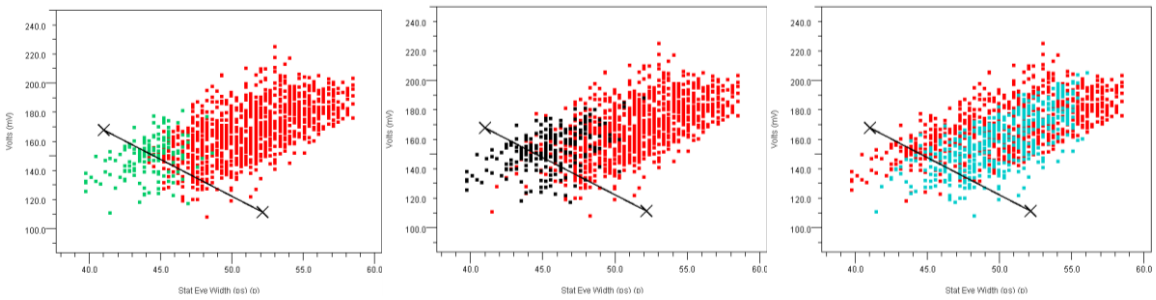

**Figure 8: Channels with Discontinuities and Medium Lengths Highlighted**

Combining the changes described above, we elect to reduce the discontinuities on the shortest channels and slightly increase the amplitude on the medium length channels. In Figure 9 at left below we plot our original baseline with problematic BERs highlighted in blue. At center we introduce only the amplitude increase and see it has removed 50% of the signals below the line, highlighting in green the remaining channels with problematic BERs. Adding the additional discontinuity improvement at right we see we have achieved our goal of moving the signals with the worst BERs above the line. The two signals close to the line can be examined for additional corrections, as desired.

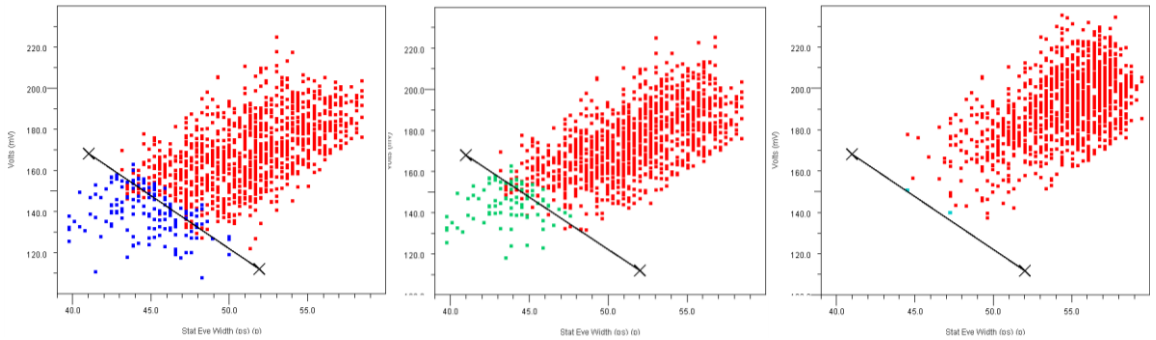

**Figure 9: Incremental Improvements to Remove Failing BERs**

This section has demonstrated a 4$^{th}$ generation design approach in which both the inherent channel construction and SerDes settings can be adapted to achieve performance targets. The next section examines the use of 4$^{th}$ generation cards in legacy 3$^{rd}$ generation systems – situations in which corrections can only be implemented through changes in firmware settings.

### 2.2.2 Interoperability of SerDes Generations

As systems and board designs transition to higher data rates newer SerDes typically need to interoperate with SerDes from the previous generation. While the newer generation brings an expanded set of equalization features, these features are typically targeted at their higher/native frequencies. In addition, the newer generation may operate at lower voltages and hence require smaller amplitudes from the previous generation's components. As the components interoperate they're required to use the data rate of the previous generation. This becomes an interesting design problem that is typically not handled well by the default settings of either generation.

In Figure 10 at left below we see initial system performance for over 1,000 channels with default settings applied, revealing interesting clustering of eye openings (eye width on Y axis, eye height on X axis). As both generations specify defaults for long and short channels, at right below we highlight the four resulting setting groups. This plotting method allows us to quickly see that the red and green setting groups are performing well, while blue and gold ("long" channel settings) are not. As shown, the blue signals are amplitude constrained and gold signals are not performing well in both height and width.

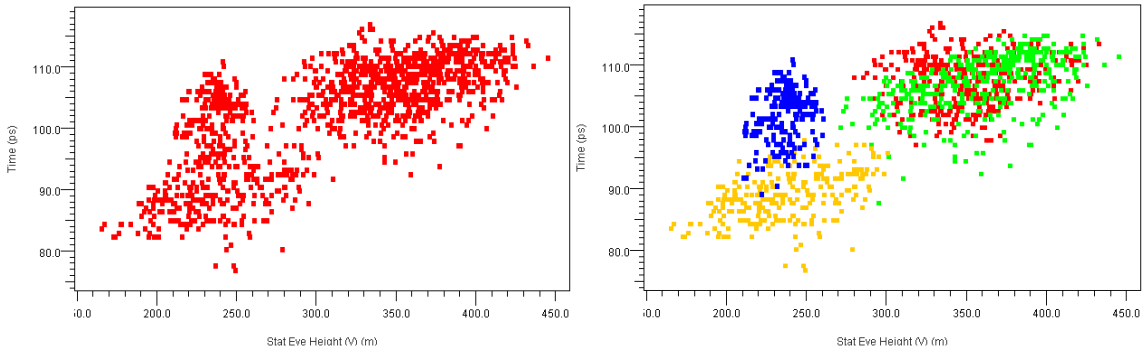


**Figure 10: Interoperability Performance with Default Settings**

Studying the problem in detail, and in this case *adapting only SerDes settings* (newer devices must typically operate leaving legacy system hardware as is) we achieve the performance improvements shown in Figure 11 at right below contrasted with default performance at left. Note that – using only firmware setting changes – we have improved poorly performing channel eye heights and widths by more than 20%. While eye widths are clustered well (Y axis), some amplitude constraints remain.

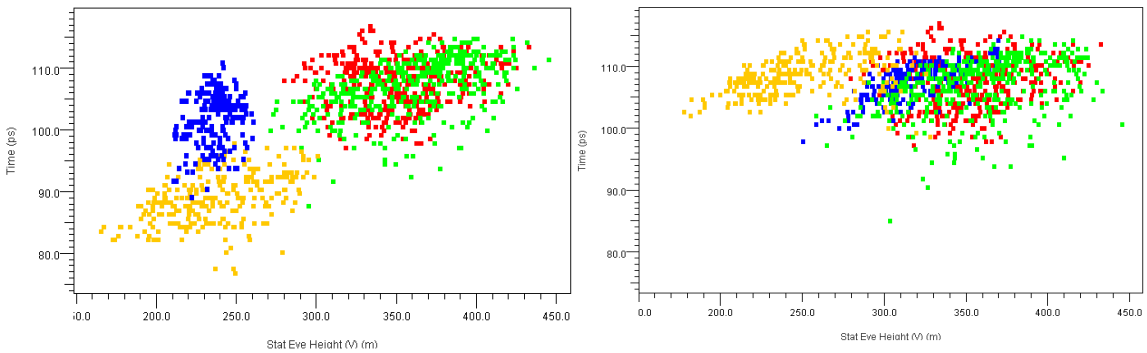


**Figure 11: Channel Performance Before (left) and After (right) Adapting Default Settings**

Plotting eye height vs channel length in Figure 12 below with default (left) and improved settings (right) we see that previous driving new generation signals (green/blue) greatly benefit from the improved equalization in the newer Rx, while the simpler Rx equalization seen when the new generation drives the previous (red/gold) decreases fairly linearly with length.

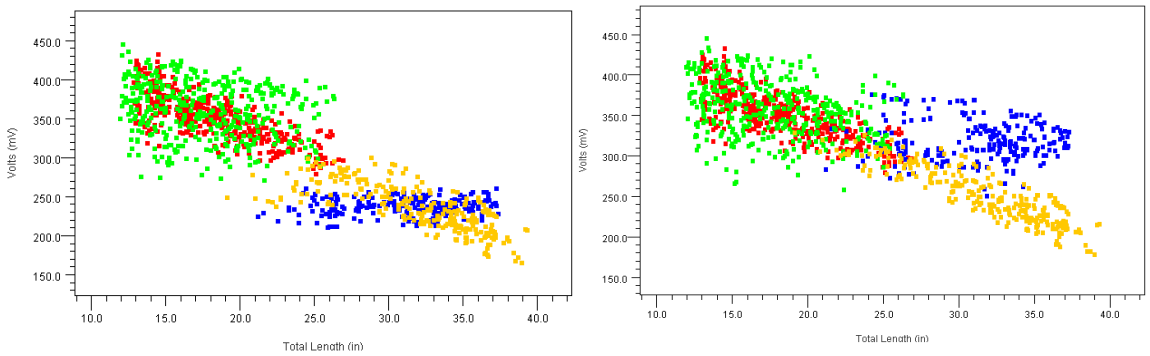


**Figure 12: Eye Height versus Channel Length Before (left) and After (right) Adapting Default Settings**

Plotting eye width versus length in Figure 13 below with default (left) and improved (right) settings we see eye width is now very consistent in both directions regardless of channel length. Some outliers on the green channels (green) are noted, again related to ISI in short channels that equalization – even in newer components – cannot remove.

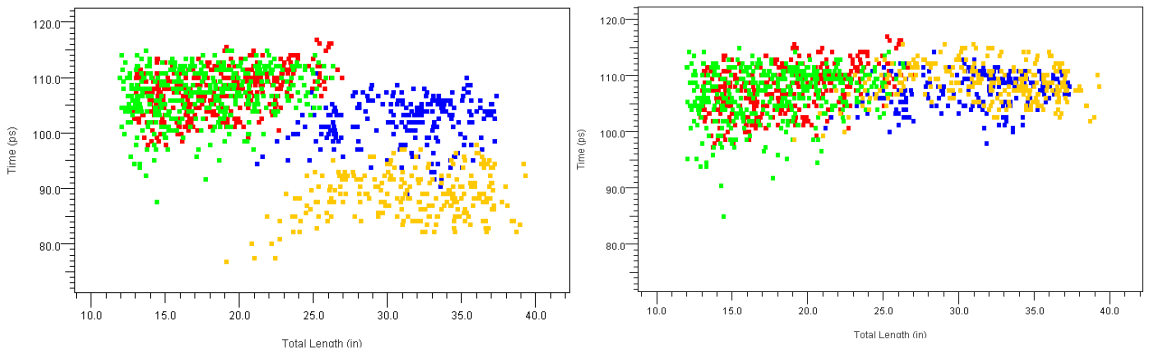


**Figure 13: Eye Width versus Channel Length Before (left) and After (right) Adapting Default Settings**

The previous sections demonstrate the value of observing and measuring eye performance at the *output* of Rx equalization, allowing us to balance equalization at the system level rather than tuning each end separately. This is not only necessary because there is no observable eye at the Rx input [6] but also allows us to trade-off amplitude and compensation, as previously demonstrated in [2].

## 2.3 Layout Adjustments

The performance of 4$^{th}$ generation links is surprisingly influenced by structures that are typically less than 30 mils, such as via stubs, antipad traces, and serpentine route corrections. While the effect of via stubs is well-documented elsewhere, this section examines lesser-known structures that also have a measurable effect on 12 Gbps performance.

### 2.3.1 Tapered Antipad Traces

Though very short, the floating nature of traces in antipad regions cause a high impedance that is easily seen in 3rd generation TDR measurements, as shown at left in Figure 14 after the initial via dip (measured via in blue, simulated in gold). Lowering the antipad trace impedance by using wider traces using either of the methods shown at right has been shown to improve eye performance by 7%.

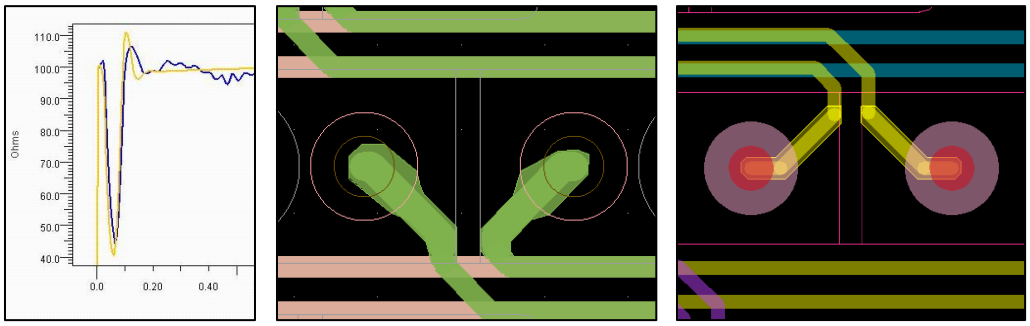


**Figure 14: Trace Corrections in Antipad Regions**

### 2.3.2 Irregular and Discontinuous Distributed Impedances

Measured data consistently reveals unanticipated irregular impedances in manufactured traces. This section explores the sources of these discontinuous distributed impedances. While some of the causes can be extracted from a layout database and others cannot, all sources should be understood and either designed out or compensated for at various data rates.

A sample TDR measurement showing irregular impedance is shown in Figure 15 below. In gold is a simplistic simulation of the structure, which consists of a short trace with vias on each end. In red and blue are TDR plots from both ends of the structure. These plots were derived from 4-port S-parameter VNA measurements obtained by landing high-frequency probes directly on the vias and their nearby grounds. Note that there are various irregularities in the measurements that span a ~5 Ohm range.

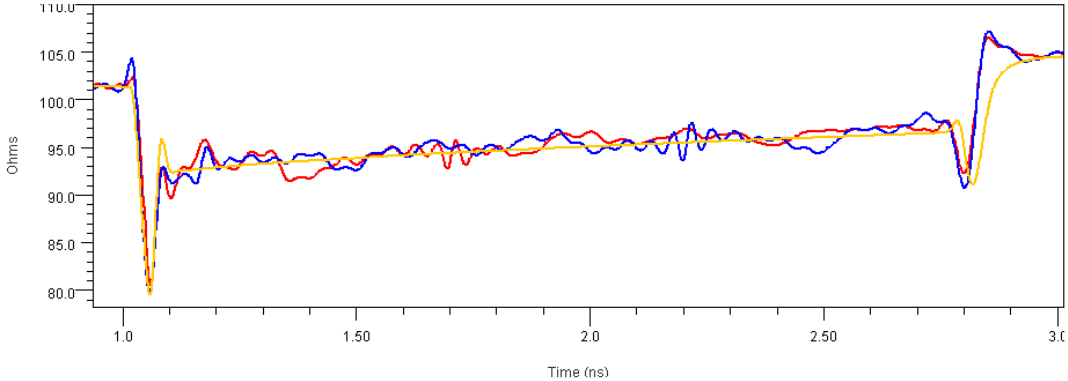


**Figure 15: Irregular and Discontinuous Impedance**

Four sources of irregular impedances are shown in Table 2 below. Based on their relative impedance discontinuities, it is important to understand and mitigate these four sources when

implementing 4th generation serial links. Note that most of them can be comprehended by simulation of extracted structures, yet fiberglass weave cannot.

| Rank | Cause | Magnitude (Ohms) | Extractable? |
|---|---|---|---|
| 1 | Serpentine Routes | 15 | Yes |
| 2 | Reference Plane Voids | 5 | Complex |
| 3 | Fiberglass Weave | 5 | No |
| 4 | Reflections | 3 | Yes |

**Table 2: Features Causing Irregular Impedances, and their Magnitudes**

**Serpentine Routes.** Serpentines are often used on one side of a differential trace to match signal lengths – particularly in breakout regions – and show variations up to 15 Ohms on microstrip layers. Differential (purple) and single-ended (blue/gold) TDR measurements of the three highlighted trace routes are shown in Figure 16 below. In each measurement the single-ended TDR of the net with length-compensating serpentine is shown in gold, revealing greater than 15 Ohm impedance discontinuities. These disturbances show a similar differential magnitude as the impedance of the non-serpentine routes shown in blue are more consistent.

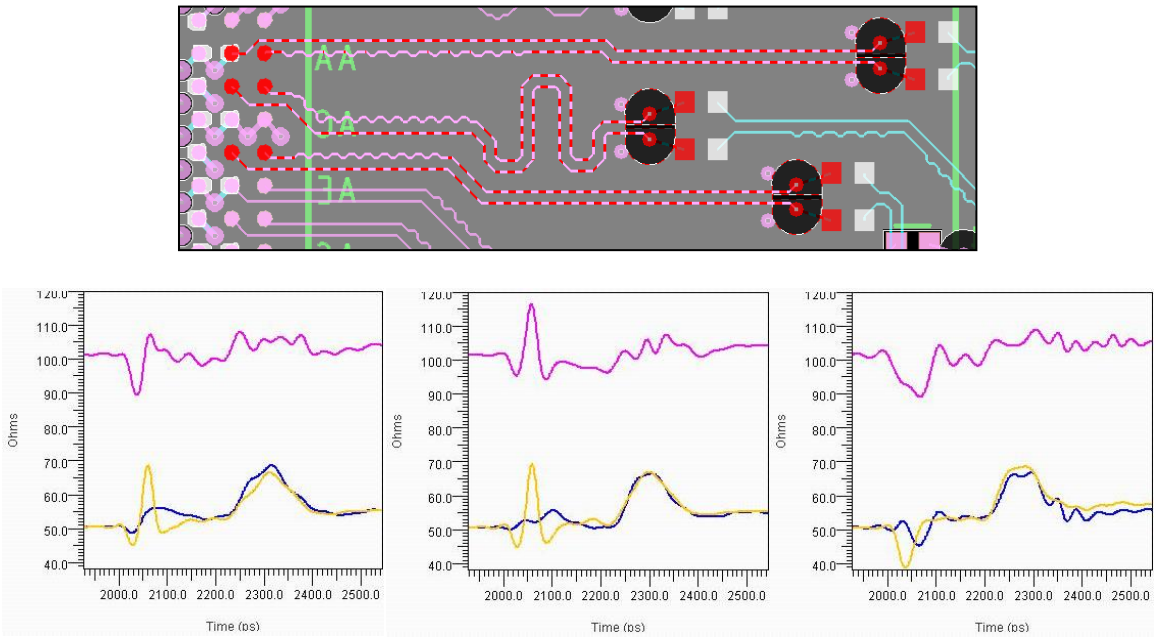


**Figure 16: Serpentine Route, Physical and Electrical Views**

**Reference Plane Voids.** Signals traveling through dense connector regions (or other structures) often encounter voids in their reference plane less than half a trace width away. When voids are longer than ~100 mils impedance variations on the order of 5 Ohms are seen. The example in Figure 17 below shows three distinct impedance bumps in the single-ended TDR (left) closely match the three connector pin fields the signal routes through on the horizontal segment (right).

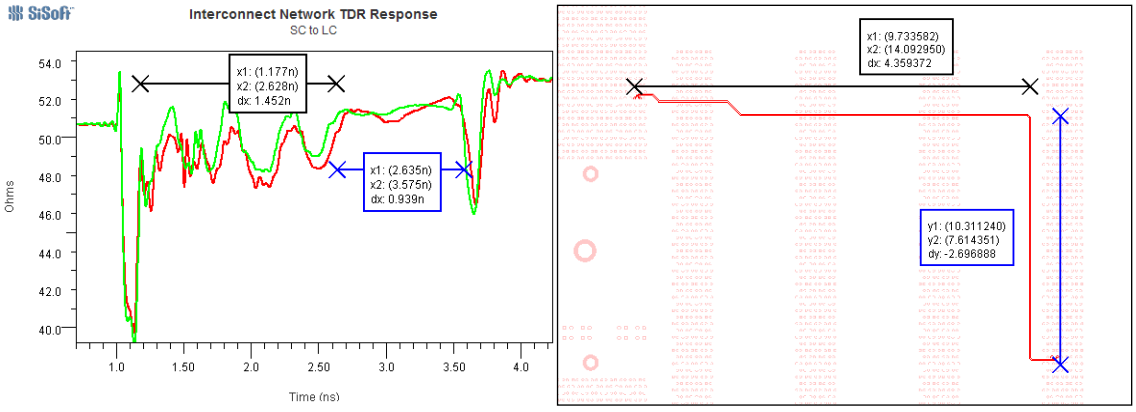


**Figure 17: Routing Near Voids, Electrical and Physical Views**

**Fiberglass Weave.** Dielectric weave effects are believed to be another explanation for many of the irregular impedances observed in PCB measurements, though this phenomenon is difficult to confirm. Impedance variations of ~5 Ohms are seen, as well as p/n time variations of 1%.

Figure 18 at left overlays a large group of measured signal's SDC21 mode conversion plots. While nearly all signals stay below 20dB, one signal rises to nearly 10dB suggesting that 25% of a 6 GHz differential signal on this route might be converted to common mode. At right, TDT plots of the p and n sides of this differential pair reveals a time difference of 18pS. Though the signals are matched to less than 1 mil, this is the type of difference that would be seen if they were miss-matched by one-hundred times that amount. Fiberglass weave is believed to be the only explanation for the variation. This type of variation was shown to cause a 10% reduction in eye height.

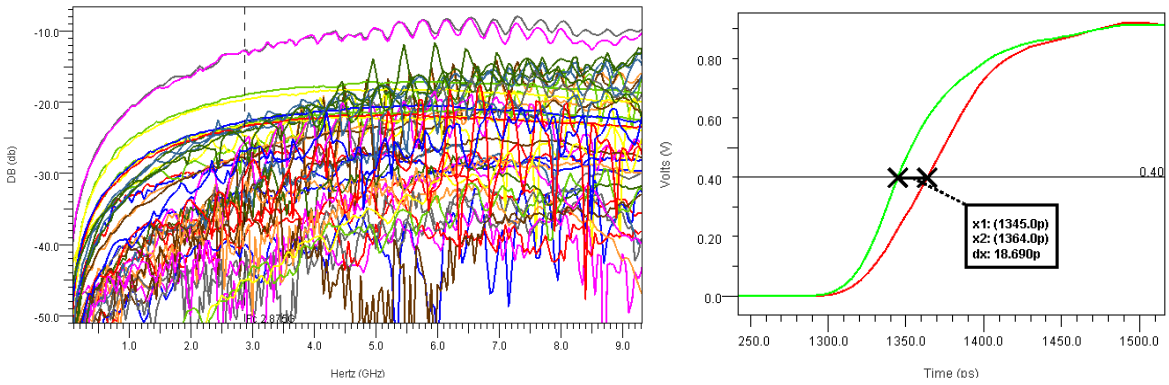


**Figure 18: Channel with High Mode Conversion**

Figure 19 shows how the three above phenomena can be identified on a signal that would have been expected to show more consistent impedance. The TDR measurement of this 6.6" differential route is shown between 2 and 4 nS. While the only anticipated discontinuity is an AC capacitor structure within the first 1" of trace, the measured impedance of the entire trace is far from continuous due to the route structures noted. Three sections of diagonal/horizontal route transitions in the center of the trace suggest fiber weave variations on the order of 5 Ohms.

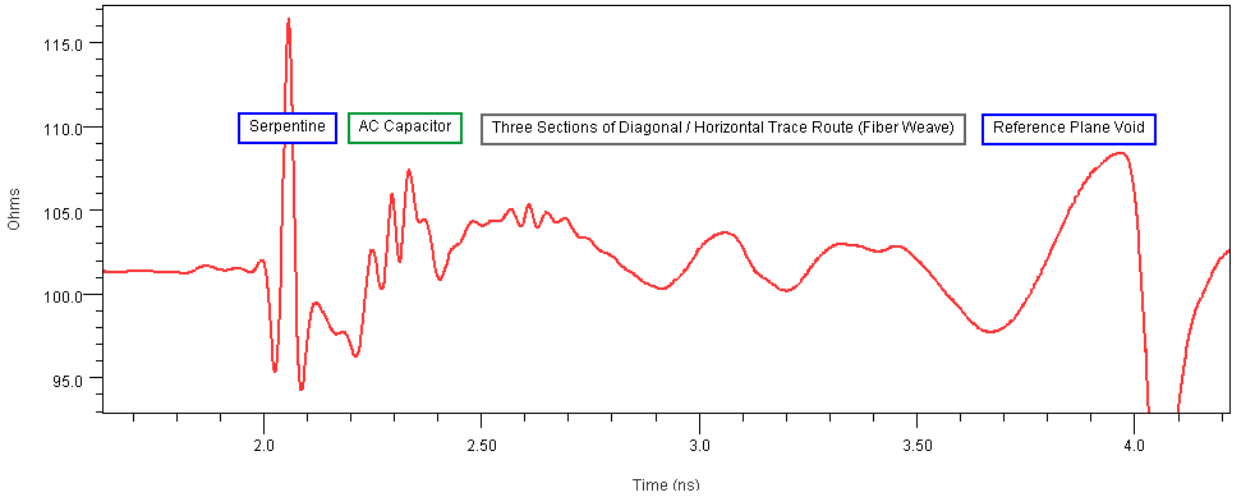


**Figure 19: TDR Measurement with Impedance Variations Highlighted**

Once familiar with the variations illustrated above, it becomes possible to look at a physical route and guess at what might be seen by a TDR. Layout designers who acquire this skill become adept at avoiding these problems, and hence create routes with fewer discontinuities.

A final example shown in Figure 20 below illustrates how some of the disturbances seen in TDR plots are not physical discontinuities, but rather reflections off existing discontinuities. These types of disturbances can also be seen in simulated TDRs (green below, measured data of the same structure in red), again even though there are no physical structures at that time position in the model. Specifically the disturbance near vertical marker ONE is found at 2x the time delay to the discontinuity at ~1.2 nS, hence it is due to the round-trip time to the source and back. Similarly TWO is the round-trip time from the discontinuity at ~1.7 nS back to the discontinuity at ~1.2 nS, and THREE is the due to the round-trip from the discontinuity at ~1.7 nS to the source and back. These types of irregularities are seen in both simulation (green) and measurement (red) and are on the order of 3 Ohms. Simulation tools account for these types of discontinuities, and hence they are comprehended during analysis.

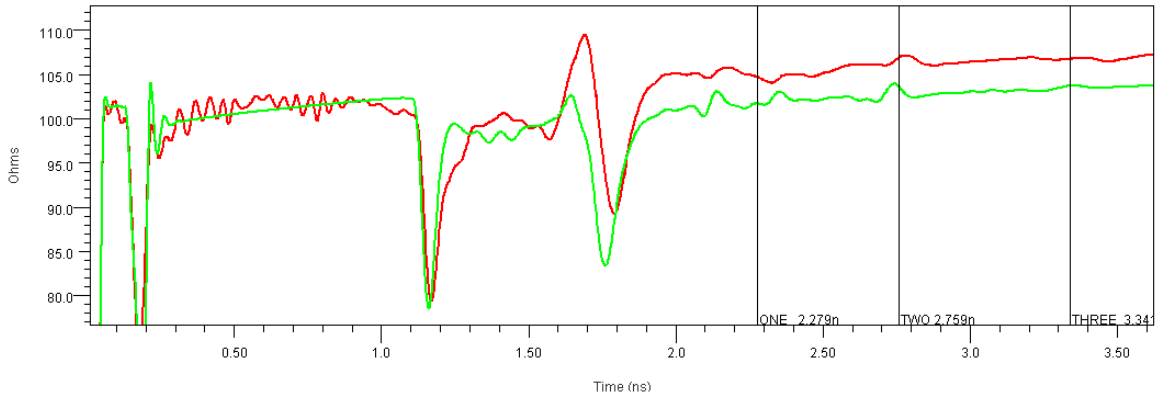


**Figure 20: TDR of Measured (red) and Simulated (green) Channel Showing Reflections**

### 2.3.3 Via Cage Asymmetry

Asymmetries in the placement of ground vias surrounding signal vias is consistently shown to cause single-ended loss notch frequencies proportional to via length. Figure 21 shows differential (black) and single-ended (red/green) insertion loss for signals with the same via length and structure, yet fabricated on different PCBs with different materials and embedded in signals of different lengths. Note that, regardless of fabrication, materials, and signal length, each plot shows a distinct notch at ~13 GHz of similar depth and perhaps smaller notches ~5 GHz both above and below 13 GHz. The fact that the single-ended plots (red and green) differ is believed to be due to the only difference between the signals: ground via placement around signal vias.

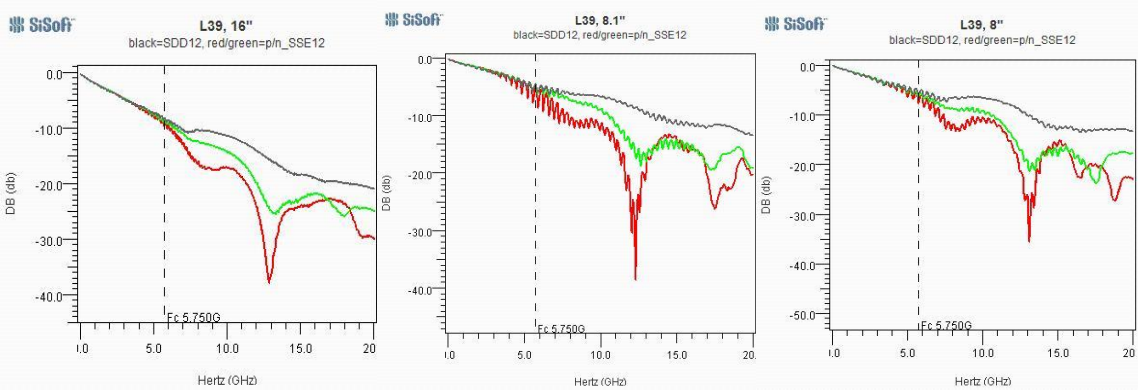


**Figure 21: Single-ended and Differential Loss Measurements, 3 PCBs**

Compiling a larger set of measurements we develop the graph of notch frequency versus via depth shown in Figure 22 (each mark represents a measurement). The linear trendline slopes reveal that the notch decreases roughly 1 GHz for every 30 mils of via depth. While these notches show a 5% impact on 12 Gbps signal performance, over 20 Gbps they must be either comprehended and corrected or designed out.

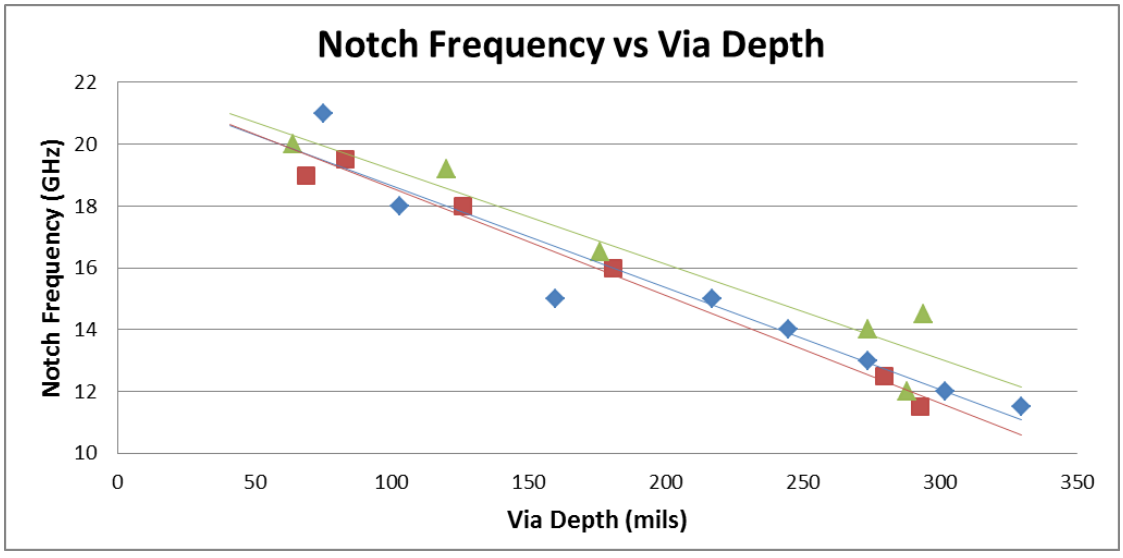


**Figure 22: Measured Notch Frequencies, 3 PCBs**

# 3. Manufacturing Improvements

This section examines important performance variations that should be comprehended when manufacturing 4th generation serial links.

## 3.1 Feature Predictability

After performing measurement to simulation correlation on numerous PCBs the following trend is consistently observed: *there is more precision in our ability to predict manufactured via performance than trace performance.* This is an important statement in light of the relationship between vias and channel performance shown in the previous section, and in [1][2].

The manufacturing reason for this difference in predictability is illustrated in Table 3 below. The Table lists the trace and via parameters that have the most significant impact on manufactured impedance, along with typical 4th generation channel values for each variable. The next column shows each parameter's observed manufactured variations, determined using photomicrography and measured TDR data. The final column reveals that, with present technology dimensions and observed variations, we can more reliably manufacture correct via impedance than trace impedance. This statement is particularly true for microstrip trace construction, which – due to its construction – is more sensitive to dielectric height and etching than stripline. Due to this large variation in impedance, microstrip trace should not be used in 4th generation channels. If it must be used, either trace lengths should be restricted or manufacturing variables more carefully controlled. Note that via variable "Stub Length" is the distance of via metal left below a signal layer after back-drilling.

| Trace Variable | Typical 4th Gen Value | Observed Variation | Unit | Impact on Impedance |
|---|---|---|---|---|
| **microstrip** | | | | |
| Pre-preg height | 3 | 1 | mil | 20% |
| Trace Width | 4 | 0.1 | mil | 1% |
| Trace Angle | 70 | 20 | degree | 10% |
| Dielectric Constant | 3.5 | 5% | | 2% |
| **stripline** | | | | |
| Pre-preg height | 3.5 | 1 | mil | 10% |
| Trace Width | 4 | 0.1 | mil | 1% |
| Trace Angle | 80 | 10 | degree | 2% |
| Dielectric Constant | 3.5 | 5% | | 2% |
| | | | | |
| **Via Variable** | | | | |
| Drill Size | 10 | 0.1 | mil | 1% |
| Stub Length | 10 | 8/4 | mil | 5% |
| Pad/Antipad Gap | 10 | 0.1 | mil | 2% |
| Dielectric Constant | 3.5 | 5% | | 2% |

**Table 3: Impedance Variation versus Manufacturing Variable**

Table 4 below illustrates this predictability by comparing measured versus designed/simulated impedance values extracted from seven signals on a typical PCB. The "measured minus simulated" comparisons represent differential impedances in Ohms, as measured from TDR data. While the average variation is similar (trace impedance measures 0.6 Ohms high while via

impedance measures 0.6 Ohms low, on average), the standard deviation for via variation is much lower than traces. Two types of vias are compared: one that is "complex" in size, antipad shape, and nearby grounds, and one that is "simple" using typical construction. Note that the complex via is either more difficult to build or model with the same precision than the simple via, as might be expected, yet both show more predictable variations than traces.

| Signal | Trace Impedance (meas-sim) | Complex Via Impedance (meas-sim) | Simple Via Impedance (meas-sim) |
|---|---|---|---|
| 1 | -6 | -3 | 1 |
| 2 | 4 | -2 | -2 |
| 3 | 4 | 3 | 0 |
| 4 | -8 | 0 | 4 |
| 5 | 5 | -5 | 1 |
| 6 | 0 | 6 | 3 |
| 7 | 5 | -4 | -2 |
| **Average** | 0.6 | -0.7 | 0.7 |
| **Std Dev** | 5.1 | 3.7 | 2.1 |

**Table 4: Measured and Simulated Impedance Variations, Traces and Vias**

Figure 23 shows a sample overlay of measured (red/blue) and simulated (green/gold) TDR data for a sample signal. The waveforms show via and breakout traces on each end and a 3.2" trace. Note that the trace impedance has been adjusted in this plot so via impedance variation can be measured. In the plot, correlation is achieved when red matches green and blue matches gold.

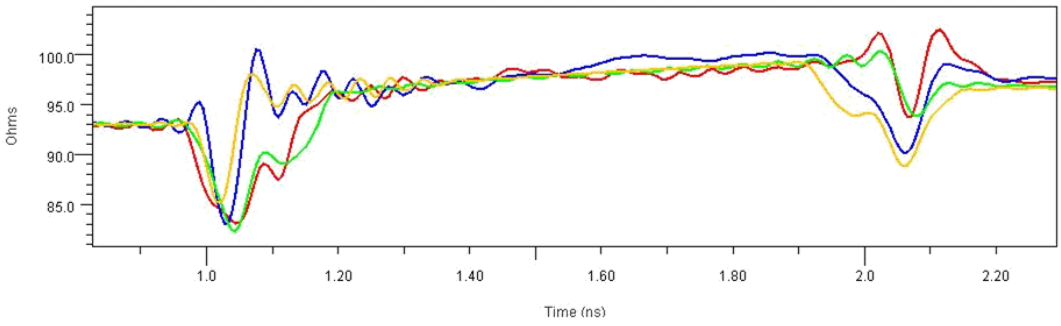


**Figure 23: Measured/Simulated TDR Overlay, Sample Signal**

The next sections examine the most significant variables found to alter manufactured impedance of traces and vias: pre-preg thickness and back-drill stub length, respectively.

### 3.1.1 Pre-preg Thickness Variation

Pre-preg thickness variation was shown in the previous section to be the most significant manufacturing variable affecting trace impedance. The photomicrograph at left in Figure 24 illustrates the non-homogeneous nature of dielectrics found in the cross-section of a PCB – both core and pre-preg. While variation in core thickness was found to be minimal, the variation in pre-preg thickness is more substantial. Figure 24 at right compares 24 manufactured pre-preg thicknesses across two PCBs with their design spec in red. Average variation is 0.53 mils with a

standard deviation of 0.46 mils, suggesting the 1 mil delta used in the previous section's calculations is likely to occur in practice. Due to this variation, use of microstrip traces should be either avoided or managed carefully. Furthermore, manufactured pre-preg thickness variation should be taken into account when designing a stackup for stripline traces. For dimensions defining the impedance of stripline traces, critical thicknesses should be implemented with core rather than pre-preg lamination layers.

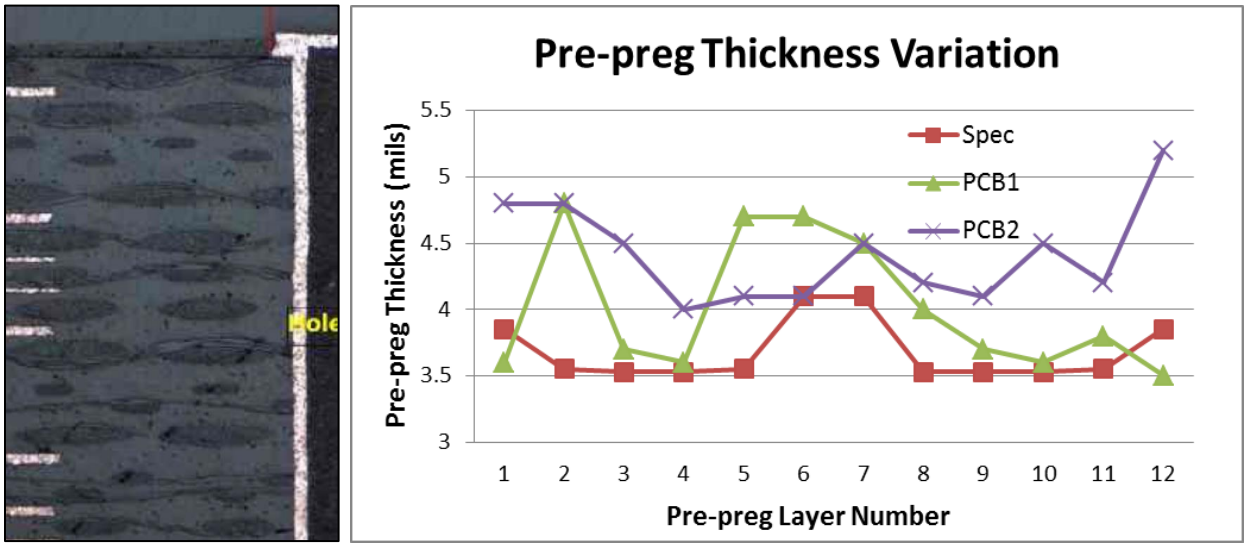


**Figure 24: Design vs Manufactured Pre-preg Thickness Variation**

### 3.1.2 Back-drill Stub Length

Back-drill stub length variation was shown previously to be the most significant manufacturing variable affecting via impedance. Figure 25 shows the measured variation in back-drill stub length for 26 vias. While typical manufactured tolerance is 10 mils +/- 8 mils, these measurements found lengths from 2 to 13 mils. Average stub length is 8.8 mils with a standard deviation of 2.2 mils. A photomicrograph of via #2's 2 mil stub is shown at left.

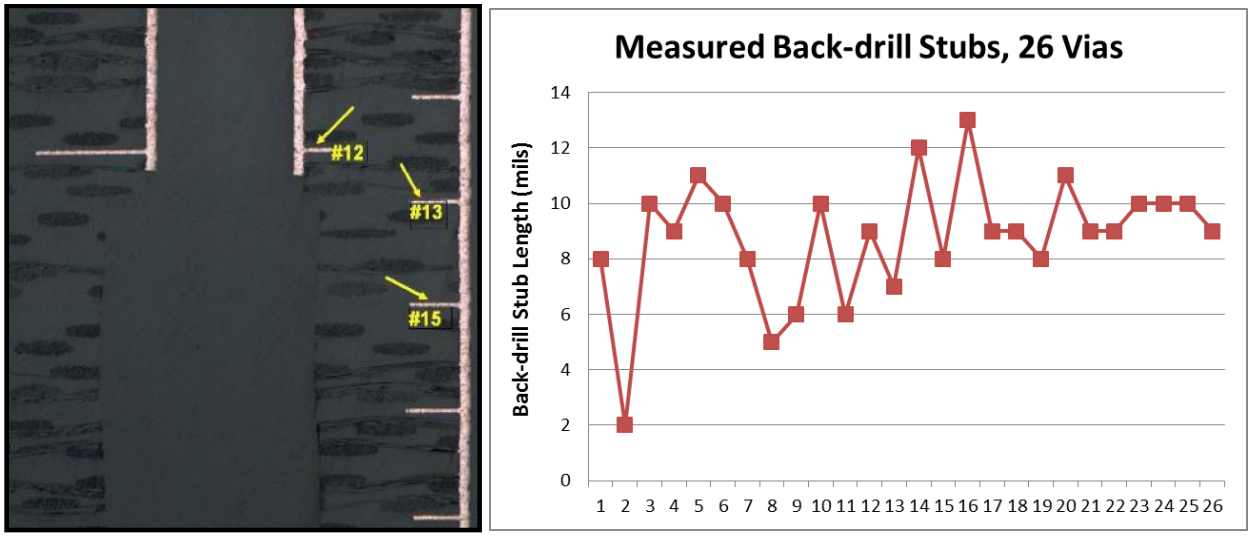


**Figure 25: Manufactured Back-drill Stub Lengths**

Figure 26 shows that impedance varies +/- 2.5 Ohms with stub length for both the simple via (left) and complex via (right). This amount of variation in stub length alone suggests it may be difficult to achieve correlation with measurement to better than 3 Ohms, all manufacturing variables considered.

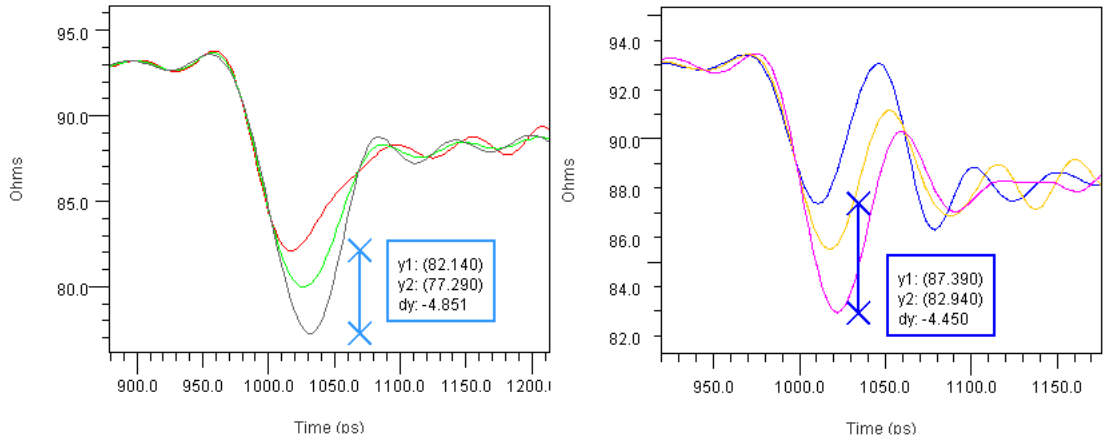


**Figure 26: Via Impedance Variation versus Stub Length, Simple and Complex Vias**

### 3.1.3 Insertion Loss

Another important parameter that varies in manufacturing – and more importantly varies between simulation and manufacturing – is insertion loss. [4] details important higher-frequency loss calculations derived from measured data that must be comprehended during the design of 4th generation channels. While traditional calculations show good correlation up to ~5 GHz, these methods consistently under-predict measured loss. Figure 27 at left compares measured differential insertion loss (blue) with a traditional loss calculation for the same structure (gold). A difference of 1.4 dB or 15% is observed at 6 GHz, and becomes greater with increasing frequency. This difference translates directly to a 15% decrease in eye height (amplitude) between simulated (gold) and measured (blue) shown at right. As illustrated, loss variation between design and manufacturing is an important item to resolve as frequency of operation continues to increase.

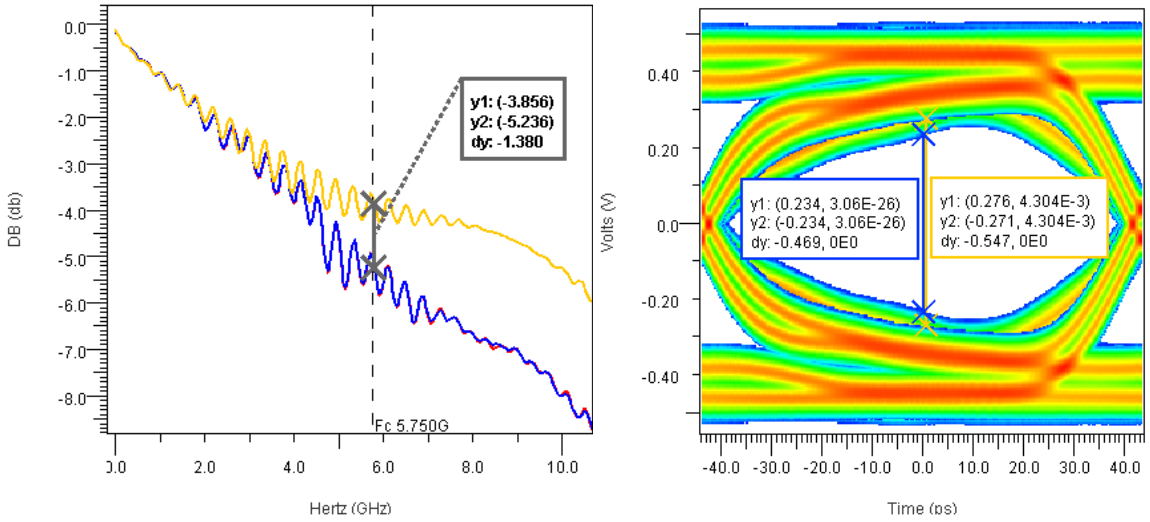


**Figure 27: Loss Variation Impact on Eye Height**

## 5. Summary

Fourth generation serial links operating at 12 Gbps and beyond require careful attention to new details in both design and manufacturing. Discontinuities ignored in previous generations introduce performance limitations and need to be removed or reduced. SerDes setting defaults are typically not well-optimized and should be adapted. Channel routing should also be adapted to manage impedance changes caused by layout structures in the 50 to 100 mil range.

In manufacturing, pre-preg thickness, via stub length, and insertion loss are important variables impacting 4th generation link performance. Typical dimensions and tolerances introduce greater variation in manufactured trace impedance than via impedance.

## Acknowledgements

The authors wish to thank Shashi Aluru, Minh Nguyen and Radu Talkad at Ericsson and Todd Westerhoff at SiSoft for their support. Many thanks to John Lehman, Jose Paniagua and Brian Kirk at Amphenol-TCS for their continued help in advancing channel design and improving the connector footprint. Additional thanks to Orlando Bell at GigaTest Labs for consistently delivering high quality measured data. Without the efforts of these and other people this work would not have been possible.

## References

[1] "Simulation Techniques for 6+ Gbps Serial Links" Telian, Camerlo, Kirk, DesignCon 2010 http://www.siguys.com/resources/2010_DesignCon_6GbpsSimTechniques_Paper.pdf

[2] "Simulating Large Systems with Thousands of Serial Links" Telian, Camerlo, et. al. DesignCon 2012 http://www.siguys.com/resources/2012_DesignCon_SimulatingThousandsofLinks_paper.pdf

[3] "When Shorter Isn't Better" Steinberger, Wildes, Higgins, Brock and Katz, DesignCon2010, Feb 2, 2010

[4] "Fast, efficient and accurate: via models that correlate to 20 GHz" Steinberger, Brock, Telian, DesignCon 2013 paper 8-TA1

[5] "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

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

Rev 1.0