DesignCon 2016

New SI Techniques for Large System Performance Tuning

Donald Telian, SiGuys
telian@siguys.com

Michael Steinberger, SiSoft
msteinb@sisoft.com

Barry Katz, SiSoft
bkatz@sisoft.com

Abstract

Large systems with multiple configuration options and extended product lifecycles provide performance tuning opportunities such as SerDes setting optimizations and manufacturing improvements. This paper describes newly-developed techniques for equalization tuning and discontinuity reduction, offering additional design margin. Cost reductions are also achieved as new Signal Integrity (SI) techniques demonstrate performance parity removing non-essential re-timers and PCBs layers. This is the fourth in a series of DesignCon papers detailing the design and implementation of a system characterized by multiple thousands of interconnected serial links spanning dozens PCBs, operating at 3rd and 4th generation serial link data rates (6 to 12 Gbps).

Author's Biographies

Donald Telian is an independent Signal Integrity Consultant with SiGuys. Building on over 30 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.

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 of SiSoft, founded the company 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. Barry was the founding chairman of the IBIS Quality Committee. He received an MSECE degree from Carnegie Mellon and a BSEE degree from the University of Florida.

1. Introduction

Simulation advancements released over ten years ago [1,2,3] allowed examination of serial links in greater detail, identifying performance limiters [4, 5] which motivated further refinements in the same [6]. Once tuned, the technologies were scaled to rapidly scan thousands of serial links to identify failure modes in rogue channels [7] enabling their correction and confident transition into production using simulated Bit Error Ratios (BERs) as a qualifying metric [8]. This paper leverages and enhances these same technologies to provide large-system SerDes setting optimization and cost reduction, highlighting recent advancements in equalization optimization techniques and algorithms.

SerDes setting optimizations are applied to a wide range of channels across systems of various sizes. Optimization techniques are described, automated, applied to thousands of links and performance gains are quantified. SerDes tuning processes were simpler when one or two taps were available in the Tx and the Rx equalization options were few. However, with newer technologies the number of setting options grows exponentially as a larger number of Tx taps are available and tradeoffs must be made between Tx and Rx equalization. As such, new tuning techniques are necessary. Furthermore, using re-timers to handle excessively long channels (40+ inches) seemed essential in previous generations of serial links. However as PCB and SerDes technologies continue to improve – combined with the optimization techniques described – we find that re-timers may not necessarily warrant their associated cost, complexity, and real estate.

Manufacturing improvements that enhance performance are also described. One example of a short connection with seven discontinuities spread across multiple PCB layers that initially showed multiple impedance changes in the 25% range is demonstrated to become nearly transparent over time. Relentless measurements on bare PCB fabrications, good vendor communication, and manufacturing process improvements and controls are key. Breakout traces longer than ¼" should be compensated. However, intentional trace layout manipulation in the presence of impedance control and re-imaging during fabrication must be carefully managed. Dual-diameter via construction is shown to be a viable solution for reducing discontinuities when via lengths exceeds 200 mils, by using simulation confirmed by lab measurement to achieve 20% channel eye improvement in channels of various lengths. SI analysis also verifies acceptable performance in reduced layer-count PCBs to achieve lower cost.

This is the fourth in a series of DesignCon papers from the authors, detailing the signal integrity analyses associated with a very large system design. The system is characterized by 7,000+ square inches of circuit board, multiple thousands of interconnected serial links spanning dozens PCBs, operating at 3rd and 4th generation serial link data rates (6 to 12 Gbps).

2. Performance Tuning Using SerDes Setting Optimization

This section illustrates system-level SerDes setting optimization performance tuning examples across large, medium, and small systems. For each system, baseline "coded" SerDes settings will be contrasted with "optimized" settings, where "coded" settings represent the previously best-known configuration for the SerDes based on previous simulation and hardware testing [4,7,8]. In other words, "coded" represents the settings coded into the system before optimized settings were derived.

The system configuration examined required PCBs with newer SerDes technology to interoperate with PCBs using older SerDes technology, as shown in Figure 1 below. The system integration scenarios involved various connectors and modularity, not shown in Figure 1 for simplicity. As newer SerDes (red, at right) can compensate for additional loss, their PCBs typically include additional route length - marked in the Figure as an "Extension" to the "Original Channel" at left. This additional length raises the question of the need for Re-Timers (green, at center) when combined with older SerDes (blue, at left).

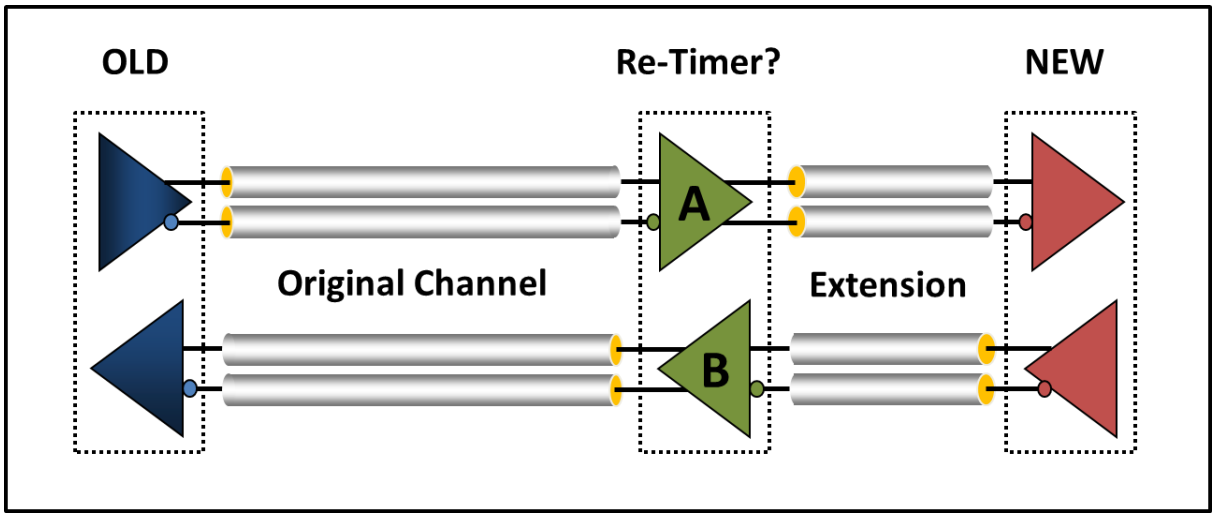

Figure 1: System Configuration Combining SerDes Generations

While the Tx in newer SerDes adds some new features, more significant changes have been made in Rx technology as CTLE (Continuous Time Linear Equalization) and DFE (Decision Feedback Equalization) implementations continue to improve. As a result, it can be shown that the upper signal path in Figure 1 (old Tx driving new Rx) functions acceptably with the additional route length rendering the Re-Timer unnecessary in position "A". More challenging is the lower signal path (new Tx driving old Rx), thus reducing the design problem to assessing the need for the Re-Timer in position "B".

Since older technology is involved, real-time auto-negotiation of Tx settings in the physical system is not an option and other forms of channel optimization must be found. A design-time methodology, more fully described in the next section of this paper, is utilized to derive optimal performance given the available configurations and setting options in the SerDes involved. Specifically, the design task is to determine if the FFE (Feed Forward Equalization) taps in the newer Tx can present an acceptable signal to the older Rx without the use of the Re-Timer in position "B". This analysis is performed using the system model previously described in [7,8].

The first, and most challenging scenario examined was the largest system which included thousands of serial links interconnected across dozens of PCBs. The plot at left in Figure 2 shows simulated eye heights and widths for ~1,000 channels implemented without a re-timer in position "B" versus design targets (horizontal and vertical markers). For this analysis, the channels were "coded" with best-known SerDes configuration settings. As shown, channels without the "extension" (red) perform acceptably against the targets, while some channels with the extension (green, with highlights in gold and black) are failing. Highlighted channels (gold

and black) below and near the eye height target are those with the lengths up to 25% longer than initial budgets. Failures of such a linear nature are typically related to additional length and loss.

To address the failing channels, an initial implementation of channel optimization algorithms was tested with the results shown at right in Figure 2 (same color schemes). As shown, the performance of all channels exceeds design targets with more than 60% improvement in eye height and width for the worst-case channels. Performance comparisons revealed improvement in more than 85% of channels using these initial optimizations, motivating further refinements in the algorithms. Most importantly, Tx settings had been derived that moved channel performance within design targets for the worst-case channels (gold and black). The axes in Figure 2 are aligned for easier comparison.

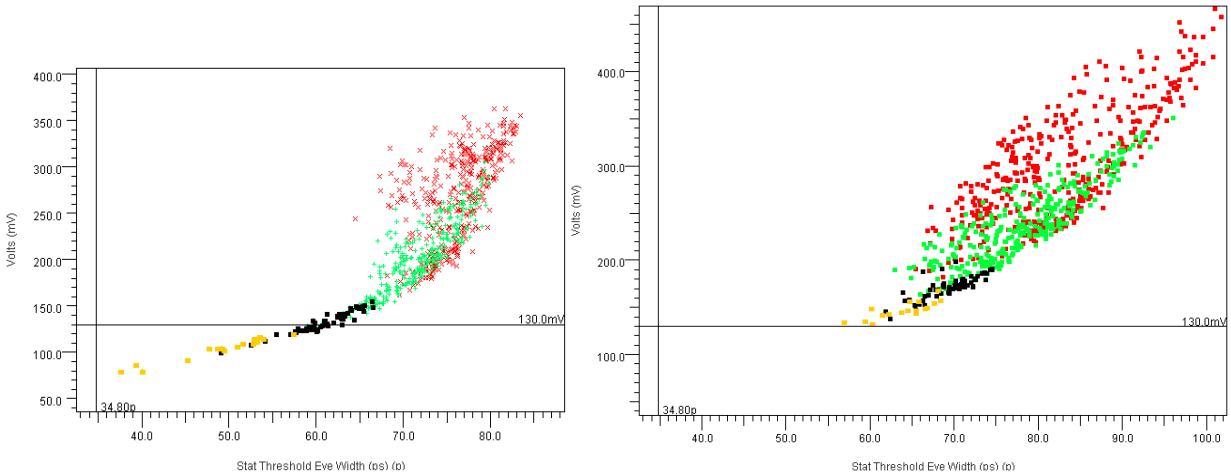

Figure 2: Simulated Results, Best-known Settings (left), Optimized Settings (right)

As similar PCB/SerDes combinations were deployed in systems of differing scale, performance was examined for all scenarios. The plots in Figure 3 compare the eye openings using default SerDes coding (blue) against optimized coding (red) for three system sizes: large (left), medium (center) and small (right). These simulations add additional system-level jitter and improved optimization algorithms, and illustrate how optimization moves worst case channels within design targets (lower left corner, all plots) while providing improved margin for all channels.

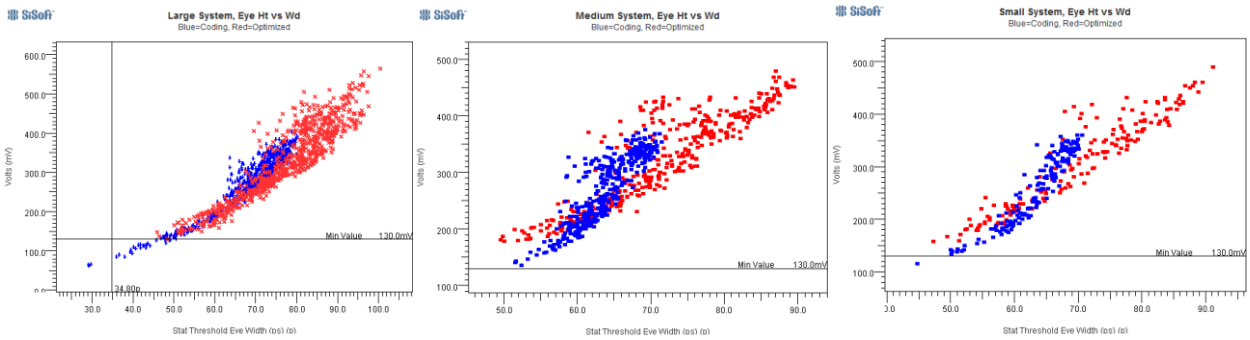

Figure 3: Default vs Optimal Eye Openings for Large, Medium, and Small Systems

For the largest system, Figure 4 shows optimization improvement by contrasting simulated eye heights using both the original algorithm (left) and the improved algorithm (right). Each point in the plots represents one channel. For each channel, the X axis plots simulated eye height using the original (baseline) SerDes coding and the Y axis plots the eye height derived from the optimized coding. As such, points on the diagonal black line represent channels that performed the same using either SerDes configuration. Channels above the line have better eye heights using optimization, while points below the line represent channels with better eye heights using the baseline coding. The plots show that the optimized settings exceed the baseline in 85% of the channels using the original algorithm (left) and 95% of channels using the improved algorithm (right). As desired, worst-case channels in the lower-left corners of both plots increasingly move away from the black line. The Y axes are aligned to show ~100mV gain in some channels.

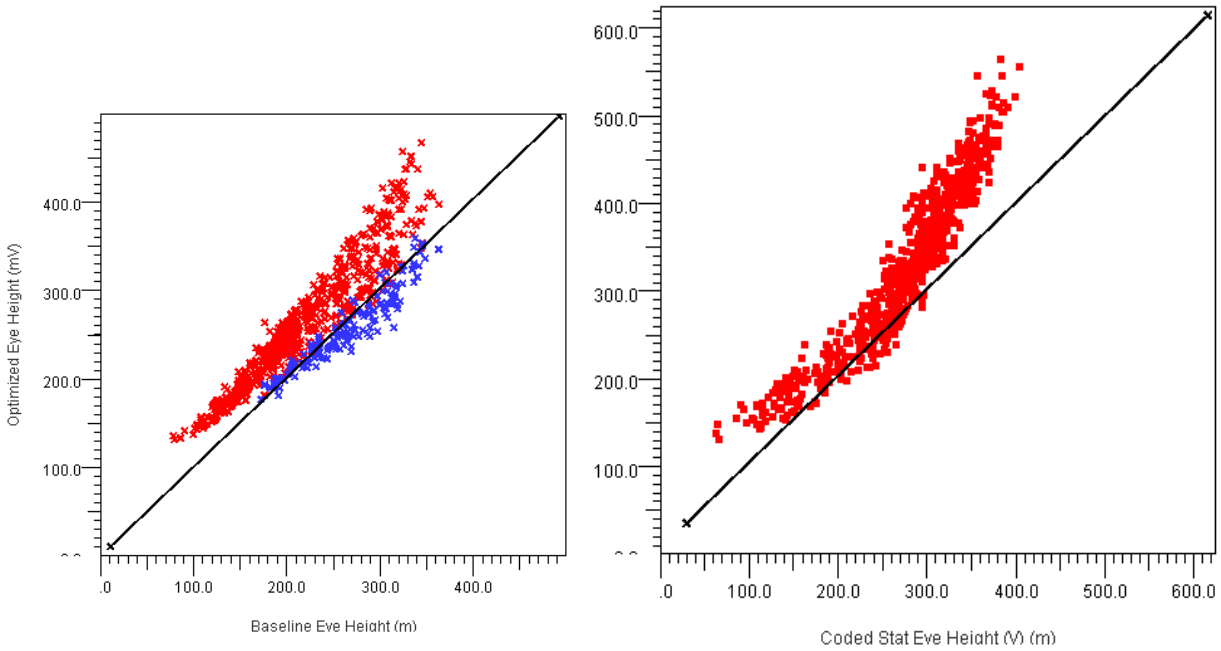

Figure 4: Eye Heights Contrasted, Original (left) and Improved (right) Optimization Algorithms

In both cases in Figure 4, the default (non-optimized) coding is better in some mid-range channels. The channels highlighted in blue (left) represent channels that did not trade tap one optimally between the Tx and Rx in the original algorithm, as this trade-off is not simple to derive. Although the optimization techniques described in the next section will detail the removal of intersymbol interference (ISI) by forcing the pulse response to zero in pre- and post-tap UIs, in some cases that method does not provide the optimal result. In Figure 5, the zero-forcing solution (blue) produces a smaller eye than the one derived by the refined automated algorithm (red). In some cases the complete elimination of ISI can remove too much amplitude from the main cursor (compare red and blue at ~8.25 nS). As such, optimization algorithms must intelligently trade-off amplitude, Tx/Rx taps, equalizable ISI, residual ISI in long tails, and other factors to arrive at the optimal solution. This is not a simple task – particularly since the "optimal" solution may vary depending on performance criterion.

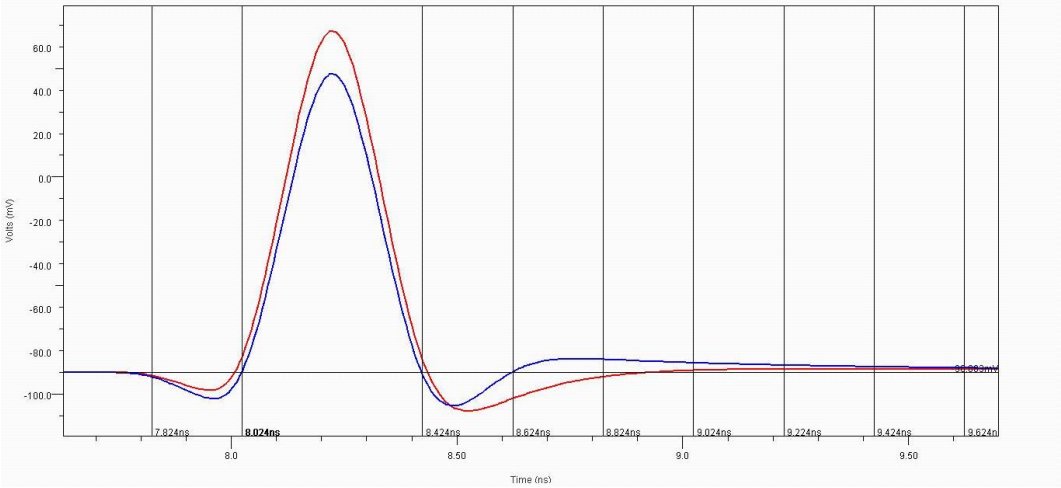

Figure 5: Pulse Responses Derived Using Conceptual and Automated Techniques

3. System-Level Channel Optimization Techniques

3.1 Overview

One of the tasks to be performed in the analyses shown in the preceding section was to optimize the combination of transmit and receive equalization separately for each channel in the system. Given the number of channels to be optimized, an automated procedure was required.

The optimization procedure we used is based on an analysis of the impulse response at the output of the receiver's IBIS-AMI model. The procedure is based on a pulse response derived from the impulse response and assumes that there are three equalization mechanisms applied to the channel:

1. Transmitter feed-forward equalization (FFE).
   - The FFE is assumed to be a synchronously spaced tapped delay line with linear taps, in which the equalizing taps can occur either before or after the main tap.
   - It is assumed that the maximum transmit voltage is constant, which is achieved by keeping the transmit swing constant and the sum of the absolute values of the taps equal to one.
2. Receiver continuous time linear equalization (CTLE).
   - The CTLE is assumed to have a finite number of configurations which can be chosen through the IBIS-AMI model's input parameter string.
   - No assumption is made concerning the relationship of one CTLE configuration to another.
3. Receiver decision feedback equalization (DFE).
   - The DFE is assumed to be a synchronously spaced tapped delay line with linear taps, driven by the detected data.
   - It is assumed that the detected data is almost always correct.

The overall procedure is

1. Define a starting configuration for the transmitter and receiver.
2. Solve for the impulse response of the passive interconnect network.
3. Apply the transmitter IBIS-AMI model to the channel impulse response, resulting in an impulse response at the input of the receiver model.

4. For each CTLE configuration:
   a. Set the receiver IBIS-AMI model to the selected CTLE configuration and apply the receiver model to the input impulse response.
   b. Using Hilbert space projection, compute the combination of FFE and DFE taps that will minimize the intersymbol interference.
   c. Using Banach space techniques, adjust the FFE and DFE taps to maximize the eye height.
   d. Record the best eye height, CTLE configuration, FFE and DFE tap values.
5. Set the FFE taps in the transmitter IBIS-AMI model to the selected tap values and set the receiver IBIS-AMI model to the selected CTLE configuration. It is assumed that the receiver's DFE will perform its own adaptation.
6. If the FFE tap values have changed significantly since the last time they were set, go back to step 3 and repeat the optimization algorithm. If the FFE tap values have not changed significantly, execute the statistical analysis in the normal manner.

The procedure described was first implemented in a simplified form in 2002. While the initial approach was reasonably successful, the subsequent evolution has made the procedure much more complex and introduced a number of subtle details. These details, and the optimal automation of the same, are beyond the scope of this paper.

This paper will instead present techniques that system developers can apply manually to a pulse response in a waveform viewer either to obtain an equalization configuration that is close to optimum or to gain insight into the tradeoffs that the optimization must address. The emphasis is on insight rather than automation.

The manual procedures to be described in subsequent sections are

- Clock Recovery: The optimization procedure is critically dependent on the position of the recovered clock. An algorithm we call the "hula-hoop" algorithm quickly and accurately determines the average clock timing to be expected from a bang-bang clock recovery loop.
- Minimize Intersymbol Interference: Choose FFE and DFE tap weights that will minimize intersymbol interference, and understand how these choices will affect the eye height.
- Maximize Eye Height: Understand how minimizing the intersymbol interference improves eye height but does not obtain exactly the maximum eye height.
- CTLE vs. FFE Tradeoffs: The effect of CTLE can be very similar to the effect of FFE, often making the two mechanisms almost interchangeable. The low frequency gain of a given CTLE configuration determines how desirable that configuration is compared to a similar result obtained using FFE.
- FFE vs. DFE Tradeoffs: In some ways, the effect of DFE can be very similar to the effect of FFE; however, there are also some important differences that affect the way the two mechanisms should be used together.

3.2 Clock Recovery

An examination of the threshold crossing time distribution due to intersymbol interference for a number of different examples led to the conclusion that these threshold crossing time distributions are always symmetrical, or nearly so. This makes sense in that the effect of an intersymbol interference contributor on the transition time of a data edge is reasonably close to linear. That is, if a given intersymbol interference contributor delays the data transition by a

specific amount when the interfering bit is a "one", then the data transition will be advanced by almost exactly the same amount when the interfering bit is a "zero".

Therefore the effect of all the interfering bits balances out (or nearly so) and both the median and mean transition times are determined by the primary transition from a "zero" in the preceding bit to a "one" in the bit to be detected, and back to a "zero" in the bit following the bit to be detected. In other words, the recovered clock timing will be half way between the transition times for a 010 data pattern in isolation.

Equivalently, we can examine the pulse response, looking for two non-zero values which are exactly one bit time apart and equal to each other. The recovered clock will be half way between those two samples. We call this algorithm the “hula hoop” algorithm because if one imagines the pulse response to be a solid object (such as a length of bent wire) and one were to drop a ring with one UI diameter (the “hula hoop”) on that object and level it, the center of the ring would be at the recovered clock time.

Figure 6 illustrates the procedure as applied in a waveform viewer. This procedure only takes a minute or two, and is quite precise.

1. The user starts by placing two vertical markers exactly one UI apart in a position that straddles the main pulse.
2. The user places a horizontal marker that is centered between the points where the vertical markers intersect the pulse response.
3. The user shifts both vertical markers to approximately the intersection of the horizontal marker with the pulse response, while keeping them exactly one UI apart.
4. Steps 2 and 3 are repeated until both the vertical markers and horizontal marker intersect the pulse response while the vertical markers have remained one UI apart.
5. The recovered clock time is half way between the two vertical markers. This recovered clock time and times before and after that are an integer number of UI away are the times at which the intersymbol interference is to be evaluated.

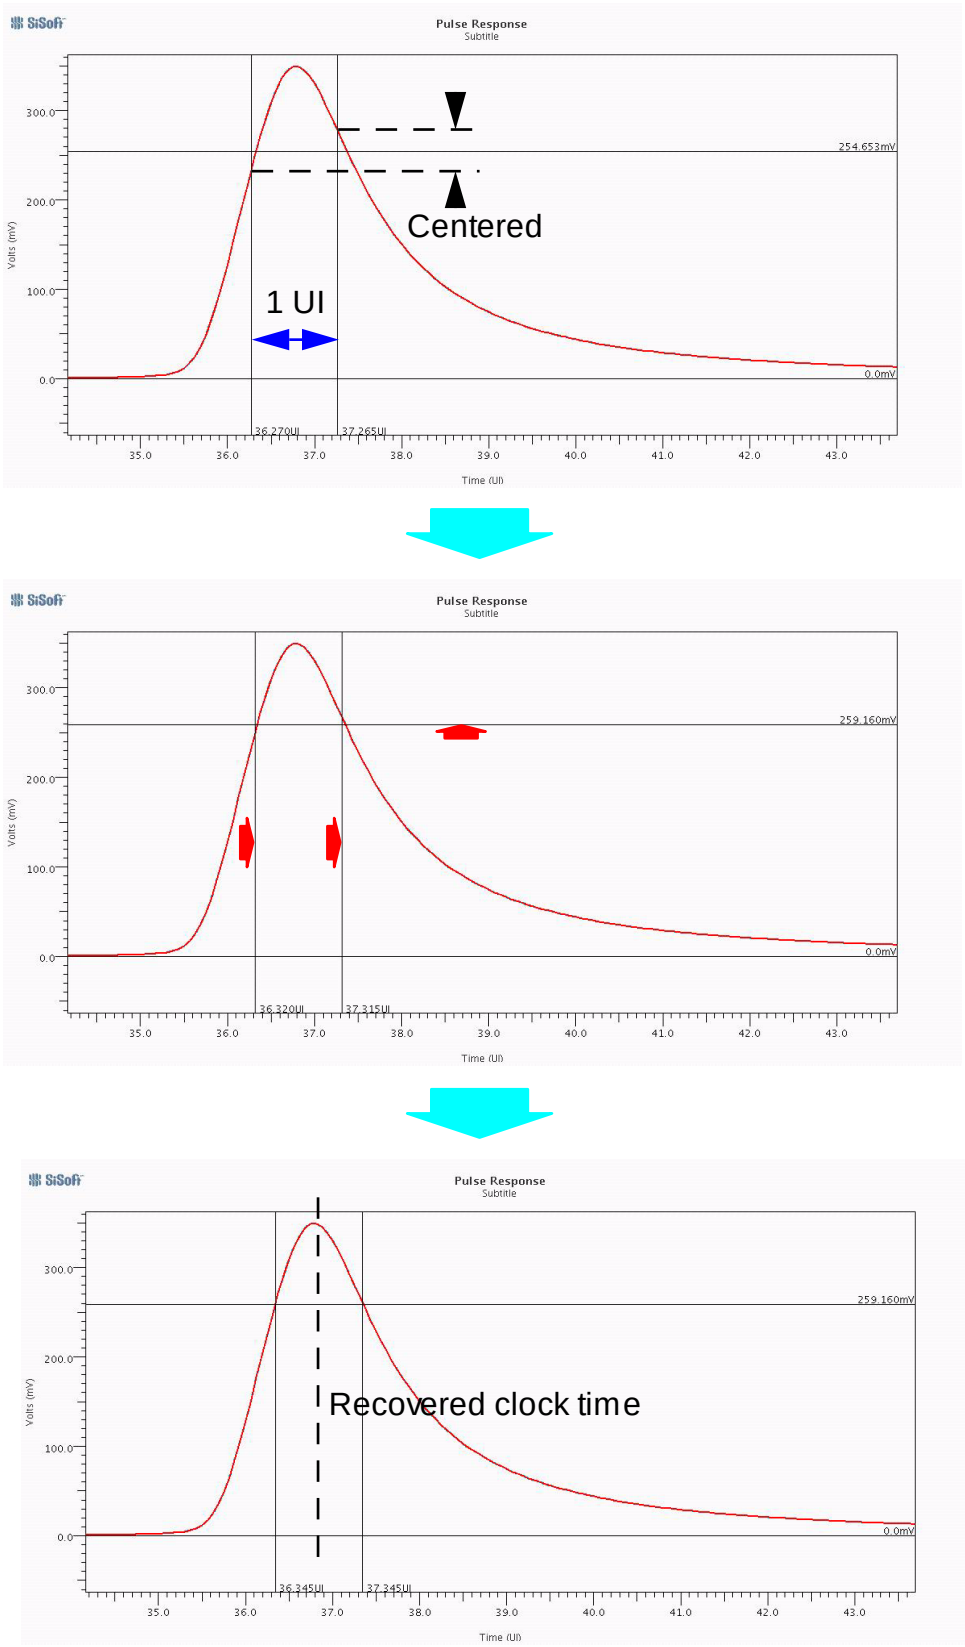

Figure 6: Hula Hoop Algorithm as Implemented in a Waveform Veiwer

3.3 Minimizing Intersymbol Interference

Given the recovered clock time, the next step is to choose FFE tap weights that minimize the intersymbol interference. Note that changing the FFE tap weights can affect the recovered clock time; so given the FFE tap weights that minimize intersymbol interference, it may be desirable to recover the clock time again and then fine tune the FFE tap weights. Figure 7 shows the same pulse response as in Figure 6, indicating voltage samples that are spaced an integer number of UI apart. The desired sample is *V*[0] and the remaining samples are intersymbol interference.

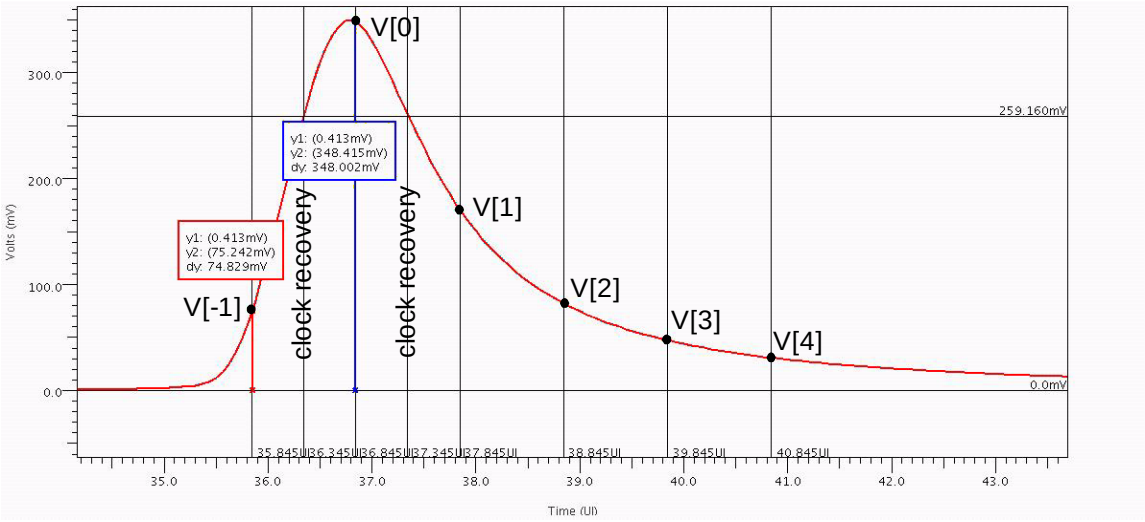

Figure 7: Intersymbol Interference Samples

Since in a real data signal the adjacent bits can be either "zero" or "one", it's the absolute values of the intersymbol interference samples that matters. Assuming that the data signal is a balanced NRZ signal, the corresponding signal voltages are -0.5 and +0.5. The minimum eye height is then

$$V_{min} = V[0] - \sum_{i \neq 0} |V[i]| \qquad \text{(EQ1)}$$

One of the most important characteristics of FFE is that the effect of an individual tap persists much longer than the time at which the tap was applied. Thus, when applying the manual procedure, it's important to apply the earlier FFE taps before the later ones because each tap affects the impact of subsequent taps.

In the current example, the first tap is a single precursor tap. We will therefore determine its tap weight before addressing the postcursor taps. Suppose that the tap weight for the precursor tap is *W*[-1] and the tap weight for the main tap is *W*[0]. We have assumed that the sum of the absolute values of all the taps is one. Furthermore, we are assuming that the main tap is positive, and we can observe that the precursor tap is going to be negative. Therefore the tap weight constraint is

$$-W[-1] + W[0] = 1 \qquad \text{(EQ2)}$$

We also desire that the precursor tap should zero out the intersymbol interference at the precursor position. In other words, the signal contribution from the precursor tap at the precursor tap time ($W[-1]V[0]$) should cancel out the signal contribution from the main tap at the precursor tap time ($W[0]V[-1]$). Still keeping in mind that the main tap will be adjusting the amplitude of the main pulse, the desired condition is

$$W[0]V[-1] + W[-1]V[0] = 0 \tag{EQ3}$$

Solving these two equations,

$$W[-1] = -\frac{V[-1]}{V[-1] + V[0]} \tag{EQ4}$$

$$W[0] = \frac{V[0]}{V[-1] + V[0]} \tag{EQ5}$$

Applying these equations to the values shown in Figure 7 produces the result shown in Figure 8. Note in Figure 8 that when the precursor tap response (in green) modifies the original pulse response (in red), the resulting pulse response (in blue) goes through zero at the precursor tap position. Note also that the amplitude of the main pulse has been reduced significantly, and that there was some equalization at the postcursor tap positions.

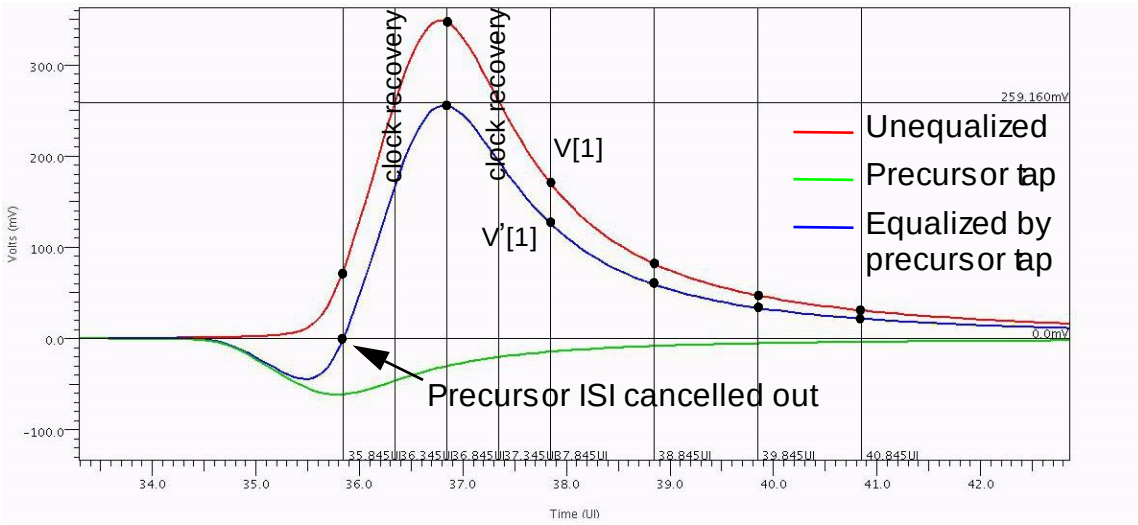

Figure 8: Pulse Response with Precursor Tap Applied

The calculation of subsequent tap weights follows the same process; however, the equations become much more difficult to solve algebraically. Note the voltage $V'[1]$ measured at the first postcursor bit position in Figure 8. For the first postcursor tap, the equations are

$$W[-1] \approx -\frac{V[-1]}{V[-1] + V[0] + V'[1]} \tag{EQ6}$$

$$W[0] \approx \frac{V[0]}{V[-1] + V[0] + V'[1]} \quad \text{(EQ7)}$$

$$W[1] \approx -\frac{V'[1]}{V[-1] + V[0] + V'[1]} \quad \text{(EQ8)}$$

and the resulting pulse response is shown in Figure 9.

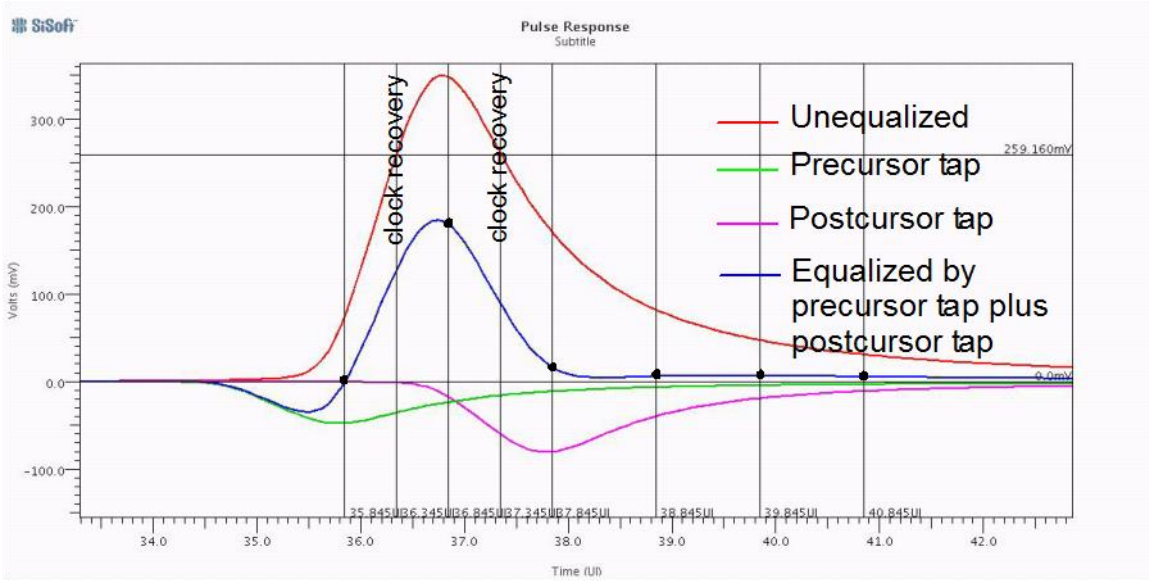

Figure 9: Pulse Response Equalized by Precursor Tap plus Postcursor Tap

Several observations can be made from Figure 9:

1. One could improve the equalization solution by increasing the magnitude of the postcursor tap by a little bit. Note that when doing so, the other two taps would have to be adjusted as well, both to maintain tap weight normalization and to maintain good equalization at the precursor tap position.
2. The amplitude of the main pulse has been reduced by approximately a factor of two. This is one of the factors to consider when choosing FFE taps: equalization costs amplitude, reducing the energy delivered to the receiver.
3. The intersymbol interference at later symbols has been all but eliminated. This is especially true for a well behaved pulse response as shown in this example; however, the same phenomenon occurs for less well behaved pulse responses. In general, one FFE tap affects the intersymbol interference at multiple bit positions.

Equation 1 states that the minimum eye height is the main pulse response amplitude minus the sum of all the intersymbol interference amplitudes. FFE derives its effectiveness from the fact that even though it reduces the main pulse amplitude, it also reduces the intersymbol interference amplitude at many positions, thus gaining leverage from the amplitude that was invested in equalization.

The general form of Equation 6 through Equation 8 can be applied to estimate additional tap weights. However, for best results those tap weights should be adjusted manually after the equations have been used to provide the initial estimate. The overall procedure can be completed in less than an hour - practical for a few test cases but not for equalizing thousands of channels in a system.

3.4 Intersymbol Interference vs. Eye Height

While minimizing intersymbol interference goes a long way toward improving channel performance, intersymbol interference as a performance measure is not the same as eye height or eye width. Furthermore, the optimum configuration for one performance measure will probably not be exactly the optimum for another performance measure. Rather, the optimum configuration will be a function of the performance measure that was chosen. This section will concentrate on eye height because it's almost always harder to maximize than eye width.

In the case of the FFE tap weight calculations described above, the weight given to the equalizing taps subtracts from the main pulse, and therefore subtracts from the eye height. Therefore backing off slightly from the tap weights for minimum intersymbol interference typically yields a slightly greater eye height.

Figure 10 is an example of the tradeoff that can be expected for the case of a single tap FFE and a channel that is relatively easy to equalize. The green line is the average amplitude in the eye diagram at the receiver decision point, which is proportional to the main tap weight. The blue curve is the maximum intersymbol interference at the recovered clock time and the red curve is the minimum eye height at the recovered clock time. The red curve is equal to the green curve minus the blue curve (i.e., the average amplitude minus the intersymbol interference). In this case, reducing Tap 1 by a little over 4% increases the eye height by a little over 1%.

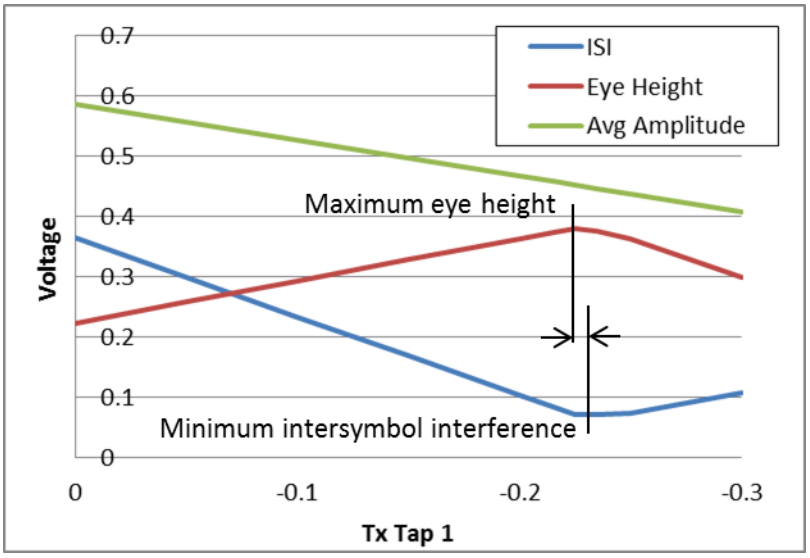

Figure 10: Tradeoff Between Intersymbol Interference and Eye Height for a Single Tap FFE

While more sophisticated algorithms are possible, for manual optimization it should be sufficient to reduce all of the equalizing tap weights in the FFE by a fixed percentage and add that

amplitude to the main tap. It would be practical to evaluate the eye height for several different values of the fixed percentage; however, in most cases a value from 0% for channels that are difficult to equalize to about 5% for channels that are easy to equalize should be about right.

3.5 FFE vs. CTLE vs. DFE

While FFE, CTLE and DFE are all effective forms of equalization, they have different characteristics, both in the time domain and the frequency domain. Often, several of these forms of equalization are combined in a single channel; and so the characteristics of each form of equalization must be considered when searching for an optimum solution.

Figure 11 presents the example of a single pulse response equalized by FFE, CTLE, and DFE, thus illustrating the differences in characteristics in the time domain. The unequalized pulse responses are shown in red while the equalized pulse responses are shown in blue. Note that the equalized pulse responses from FFE and CTLE look remarkably similar, although the pulse response from CTLE has lower amplitude than that of the pulse response from FFE. This difference will be explained below in the context of the frequency domain differences.

The pulse response due to DFE is quite different from those due to FFE or CTLE. Whereas the FFE and CTLE pulse responses are relatively smooth, the DFE pulse response contains discontinuities spaced one UI apart. These discontinuities are due to the rising and falling edges of the recovered data driving the DFE taps. Note also that DFE does not affect the amplitude of the main pulse.

One way to look at the time domain difference between FFE and DFE is that whereas a single FFE tap affects the intersymbol interference at multiple bit positions, a single DFE tap only affects a single bit position. Thus, DFE is a more flexible form equalization that is well suited for cleaning up any intersymbol interference left over from the FFE/CTLE combination provided it has a sufficient number of taps.

Another way to look at the time domain difference between FFE/CTLE and DFE is that FFE and CTLE tend to equalize across the entire bit time whereas, because of the discontinuities, DFE is only trying to equalize in the middle of the eye. Thus, FFE and CTLE tend to be better suited for removing the bulk of the intersymbol interference.

It’s also important to recognize that in most cases, only FFE can perform equalization at the precursor bit position. That is, most CTLE and DFE designs cannot equalize the effects of a bit that hasn’t been received yet. Thus, when working with a DFE that has a large number of taps, the optimal configuration will usually use the FFE for precursor equalization only and let the DFE do the postcursor equalization [7]. (In other words, set the FFE postcursor equalization taps to zero.)

The tradeoff between FFE and CTLE is addressed later in this section.

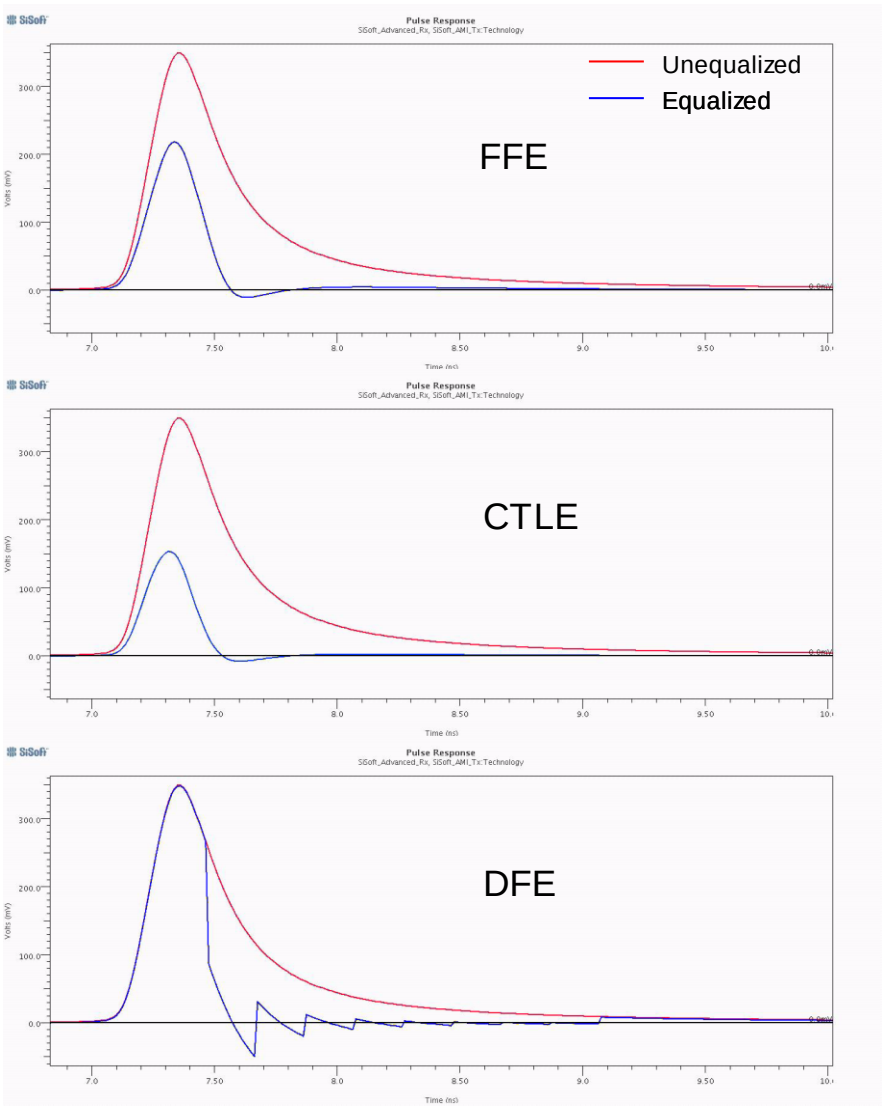

Figure 11: Comparison of Equalized Pulse Responses from FFE, CTLE and DFE

Figure 12 presents the transfer functions of an example channel equalized by FFE, CTLE, and DFE. The unequalized transfer function is shown in red and the equalized transfer functions are shown in blue.

For the CTLE there is an additional transfer function (shown in green) which would be produced by a CTLE with a better design. The CTLE transfer function shown in blue is typical of many CTLE designs - the gain at low frequencies is reduced so as to create an increase in gain at higher frequencies. This is a comparatively simple circuit to design, for example by inserting degenerative feedback in the source circuit of a differential amplifier. It's more difficult to design a circuit that has unity gain at low frequencies and then produces a gain peak at higher frequencies.

The additional information is in the comparison between the FFE and CTLE responses. At the frequencies that matter the most (the lower frequencies), the shapes of the transfer functions are

nearly identical. The most important difference is in the overall gain. At a frequency equal to one half the symbol rate, the FFE has exactly unity gain. In comparison, the lower performance CTLE design has a small amount of loss at that frequency and the higher performance CTLE design has a significant amount of gain. Thus, for equivalent equalization, the lower performance CTLE produces a lower eye height than the FFE while the higher performance CTLE produces a greater eye height.

The net result is that if both FFE and CTLE are present in the channel and have similar equalization capabilities (usually the case), the choice between FFE and CTLE will depend on the net gain of the CTLE. In the case of the lower performance CTLE, one would depend primarily on the FFE and, if anything, disable the CTLE; whereas in the case of the higher performance CTLE, one would definitely choose the CTLE and set the FFE to unity gain.

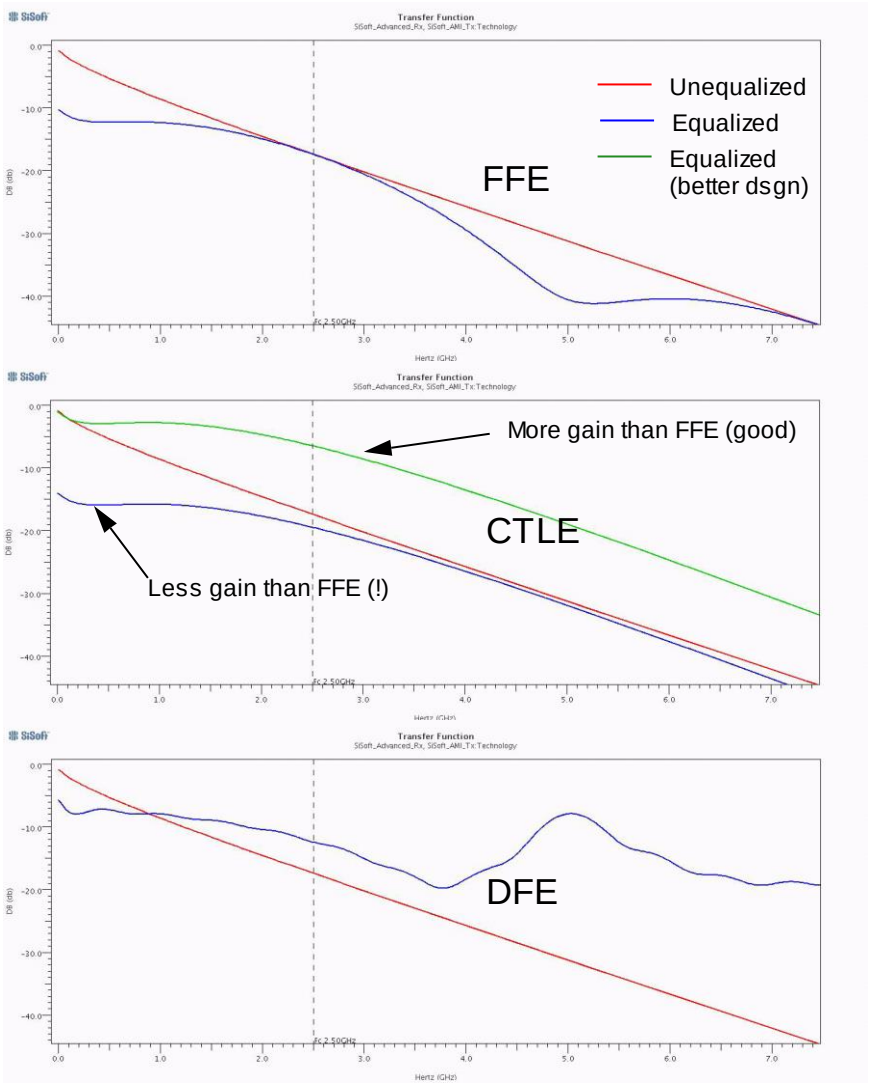

Figure 12: Comparison of Equalized Transfer Functions from FFE, CTLE and DFE

For the sake of completeness, Figure 13 is a comparison of the eye diagrams produced by the three different types of equalization.

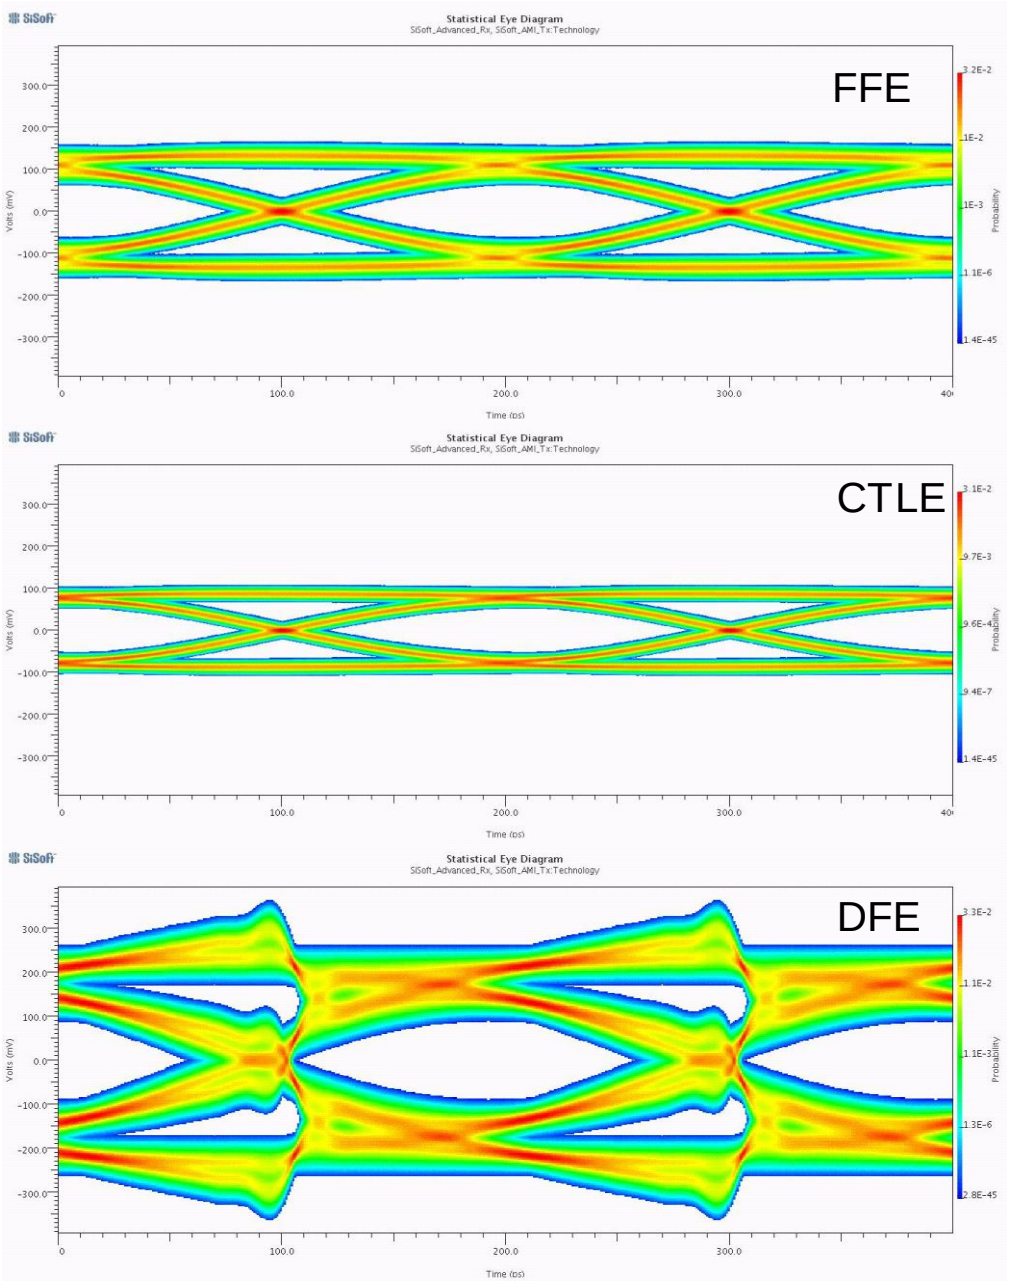

Figure 13: Eye Diagrams Produced by FFE, Typical CTLE and DFE

3.6 Manual Optimization Summary

1. The procedure is based on an analysis of the pulse response.
2. Recover the clock from the pulse response using the hula hoop algorithm.
3. If the FFE has a precursor tap, determine that tap value using the procedure in Section 3.3.
4. If the CTLE has sufficient gain, choose the CTLE configuration which minimizes intersymbol interference. Otherwise, if the DFE has enough taps, depend on the DFE for the bulk of the equalization. Otherwise, use the procedures in Section 3.3 and Section 3.4 to choose the FFE tap weights.

4. Cost/Performance Tuning with Manufacturing Techniques

Manufacturing improvements that enhance performance and/or reduce cost are described in this section.

4.1 Removing Discontinuities Using Design and Process Control

Serial link performance is directly related to the existence, placement and magnitude of impedance discontinuities in the signal path. While some discontinuities are unavoidable, by coordinating multiple disciplines over time it is possible to significantly reduce the magnitude and impact of discontinuities.

Figure 14 shows the impedance of similar signal paths across three design iterations, as measured on three different PCBs using Time Domain Reflectometry (TDR). The physical requirements of this interconnect required seven discontinuities in less than four inches across up to seven different PCB layers, both microstrip and stripline. The plot illustrates how the magnitude of the discontinuities were reduced over time in relation to our target impedance (black line); the first iteration (red) showing variations up to 20%, the second iteration (blue) becoming more consistent yet still varying up to 15%, and the third iteration (green) looking consistent with variations now within typical tolerances of 8% and mostly related to the external component over which we have less control. The second iteration (blue) shows good progress in the discontinuities under design control, yet highlights the challenge of achieving consistent trace impedance when using new PCB materials. The third iteration (green) shows excellent progress in reducing discontinuities by using both design and process control across seven signals spread across seven PCB layers.

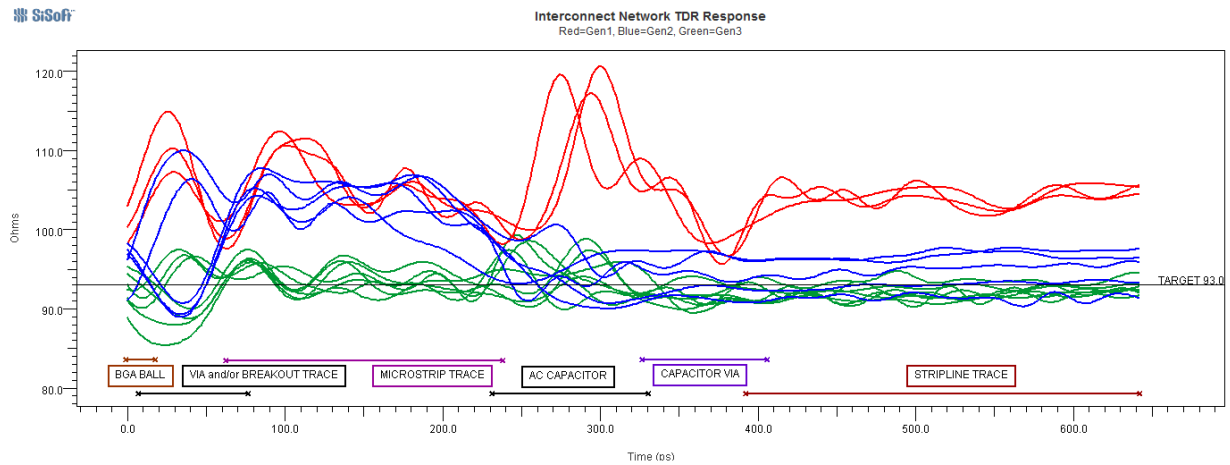

Figure 14: TDR of 7 Discontinuities Across Design Iterations

While TDR plots reveal the magnitude and location of the discontinuities of concern, Figure 15 illustrates the impact of these discontinuities in terms of more familiar eye openings. The first through third design iterations are shown from left to right. The top row shows the variation due to only these traces, revealing their associated ISI and impact on an eye opening due to the discontinuities. As the discontinuities shown represent only one section of a larger channel, the bottom row adds 12" of PCB trace to examine their system level impact. Both rows are

simulated at 11.5 Gbps and utilize the trace's measured S-parameters, from which the TDRs above were derived.

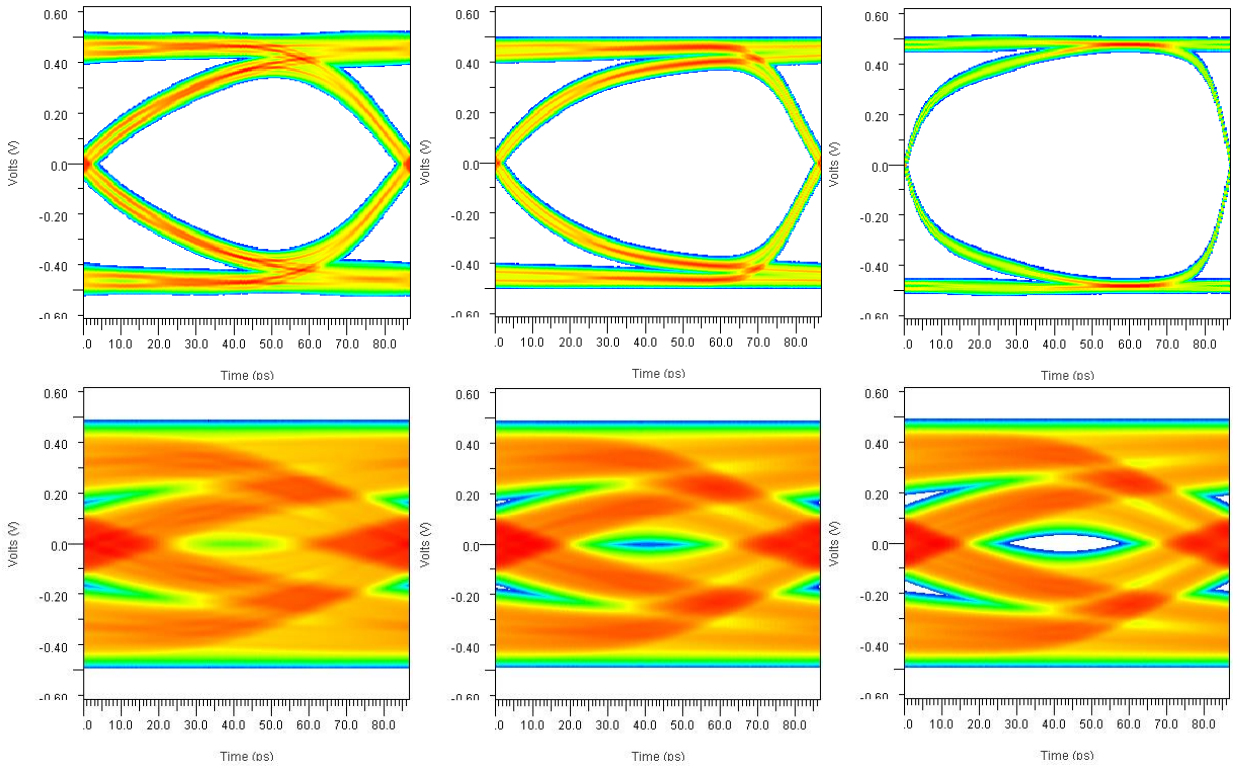

Figure 15: Eye Opening Iterations, Discontinuities Only (top) or at System Level (bottom)

The system-level plots above illustrate the importance of simulating and measuring a sufficient number of bits and/or failing bit patterns, without which the three eyes might look the same. In other words, ISI caused by discontinuities may not appear to affect eye openings in all situations. One challenge in developing serial links is they give the illusion of working when they are not working well.

When working to reduce discontinuities, the following items are helpful:

1. Use 2D and 3D field solvers derive dimensions that produce desired impedances for physical structures such as differential traces, vias, BGA pads and breakouts, capacitor plane cutouts, etc.
2. Work closely with PCB fabrication vendors to achieve and demonstrate predictable and consistent impedances – particularly when working with new materials, processes, and fabrication facilities.
3. Measure, measure, measure. Always measure actual hardware whenever possible. If you do not have the equipment or capability to produce reliable measurements, find a third party that can do so. The cost of performing measurements is much lower than the cost of debugging products in the field.
4. Simulate your design before and after fabrication, comparing and improving the results using both extracted and measured structures. Surprises happen.

5. Learn how to read TDR information from both simulation and measurement. This helps you pinpoint the cause, location, and magnitude of each discontinuity enabling you to determine which discontinuities are of concern and what to do about them.
6. Develop an intuitive sense of which structures are capacitive and inductive, and how that relates to impedance. This enables you to make changes to physical structures in layout, field solvers, and simulators to reduce discontinuities. Capacitive structures are fat and close to ground, while inductive structures are skinny and further away. Z=sqrt(L/C).
7. Determine what level of tolerance is sufficient for the technology and data rate at hand. This enables you to determine when the magnitudes of your discontinuities are "good enough".

4.2 Reducing Discontinuities Using Dual-Diameter Vias

In the authors' DesignCon 2014 paper [8] we demonstrated performance improvements up to 400% by improving discontinuities in less than 1% of a channel's interconnect, or more specifically two of the vias in the channel. In practice, one way to improve a via's impedance is to use a "dual-diameter" via structure that uses two drills to allow as much narrow hole as possible. This section augments the analysis shown in [8] by providing measured confirmation of the improvements offered by increasing via impedance.

Measured data comparing normal and dual-diameter versions of eleven via's differential impedance is plotted at left in Figure 16, organized with deeper PCB layers from left to right. At right is a sample plot comparing one of the layer 23 vias (red=normal, green=dual-diameter) overlaid with a layer 8 reference via. Note that dual-diameter via impedance on deeper layers is difficult to determine because it is not "flat", as shown in green. This is due to the various impedances seen in the dual-diameter structure such as the large diameter, small diameter, via stub, and pads. As such an average must be used, as shown by the marker at 84.5 Ohms. This irregular impedance is in contrast to normal vias that show more consistent impedance, as shown in red.

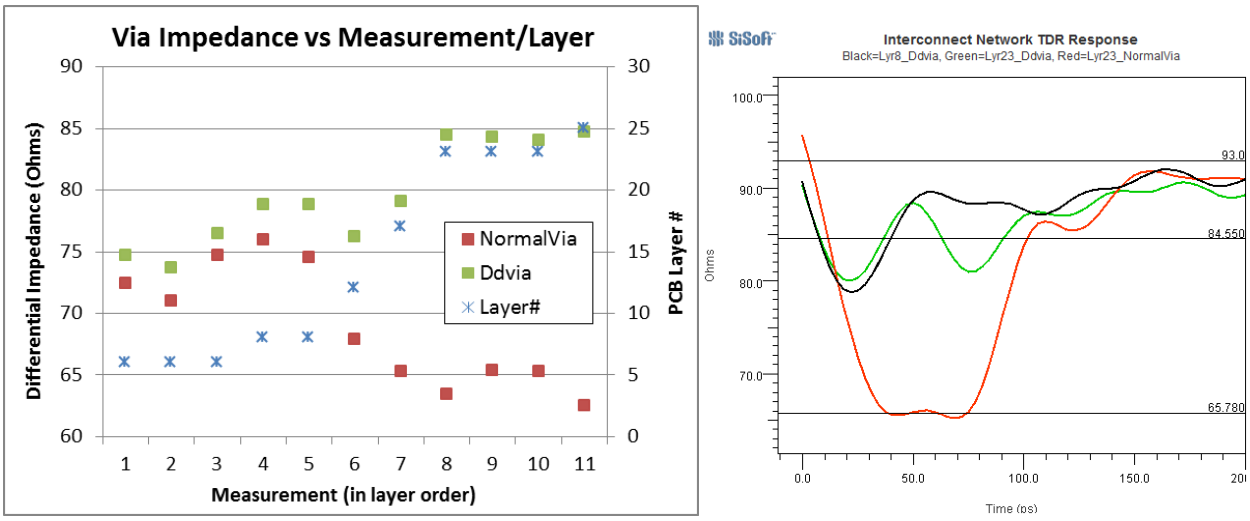

Figure 16: Measured Impedances, Normal and Dual-Diameter Vias

Dimensionally, signals on upper layers never see the smaller diameter as the transition occurs near those layers. As expected, upper layers do not see an impedance increase. In general, it is these deeper layers that are of concern as they present a more significant discontinuity. These measurements confirm that the dual-diameter structure realizes a ~20 Ohm improvement in differential impedance on deeper layers, as desired.

Figure 17 shows a TDR measurement of the same vias in an end-to-end channel in which two of the normal vias (red) are replaced with dual-diameter vias (green), as seen at ~2.6 nS and 4.2 nS. In this plot, the full 20 Ohm improvement is not evident because the probes were placed further away from the vias. Note that identical structures are assembled to the left and right of the altered vias.

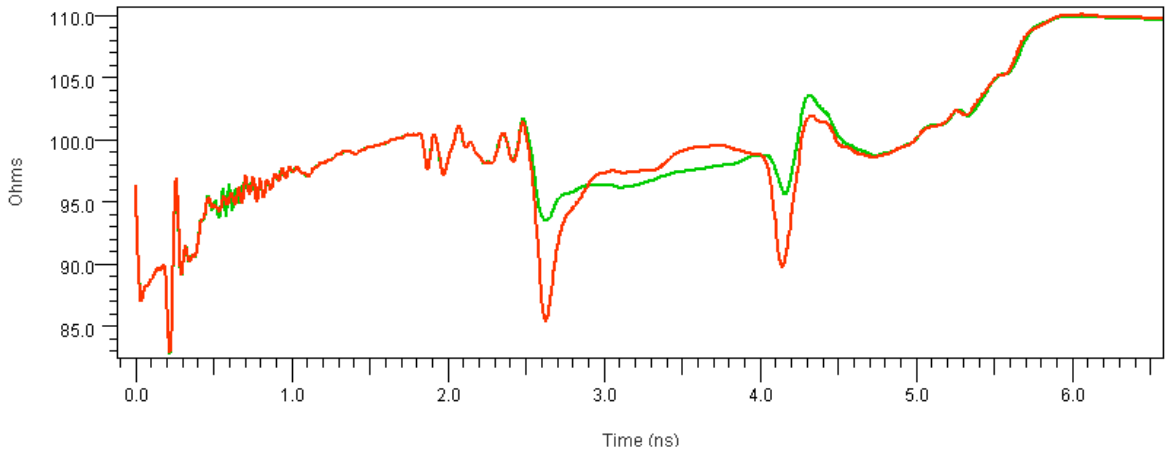

Figure 17: System-level TDR Contrasting Normal and Dual-Diameter Vias

Simulating the signals using the end-to-end measurements as channel models confirms consistent up to 30% improvement in eye height and width when comparing the same channel with normal vias (red) or dual-diameter vias (green), as shown in Figure 18. In addition, there is a much more consistent clustering of performance (compare green against red in plot at left below) across channels varying in total length from 10 to 20 inches.

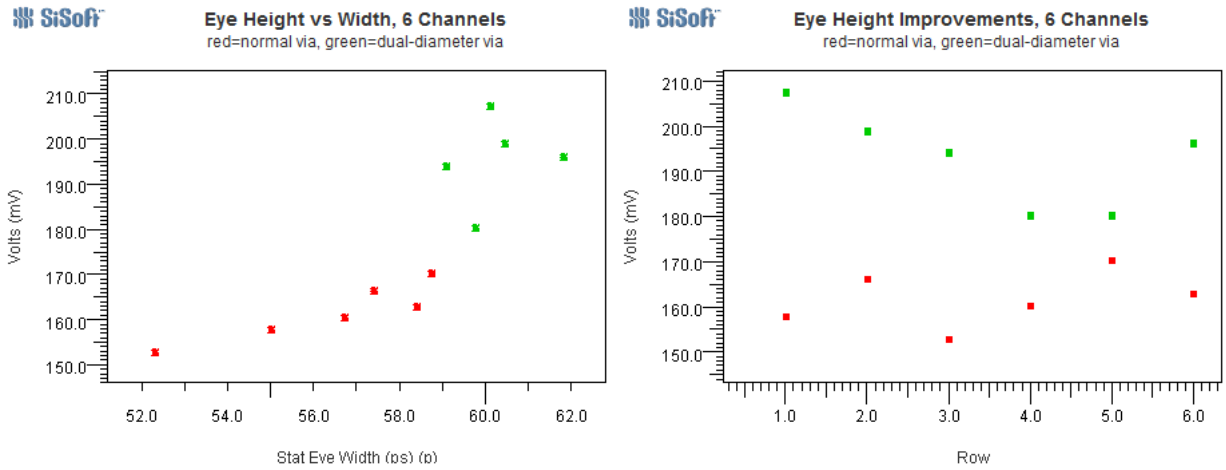

Figure 18: Eye Opening Metrics, Channels with Normal and Dual-Diameter Vias

4.3 Trace Compensation, Improvements and Challenges

It is common knowledge that the differential impedance of differential traces increases when they become uncoupled, as often occurs when routing into a BGA pin field as shown at left in Figure 19. Below the route are field solved impedances, predicting an 8 Ohm increase for this trace's construction. Measured impedance in the TDR at right confirms an ~8 Ohm increase on two revisions of this PCB, for this trace and a shorter trace.

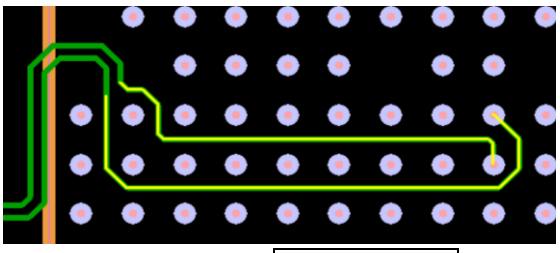

	Diff Imp
Uncoupled	102 Ohms
Coupled	94 Ohms
Difference	8 Ohms

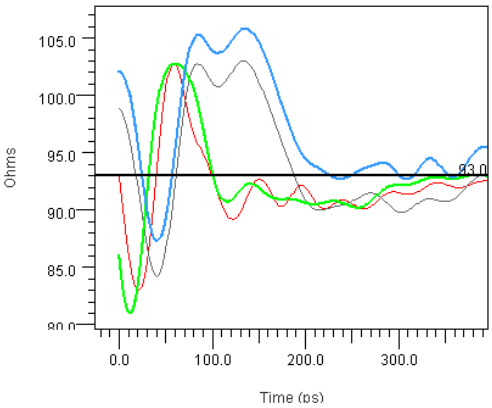

Figure 19: Uncoupled vs Coupled Impedances

Further investigation reveals that these predictable discontinuities cause, on average, an 8% impact on eye openings when the breakout traces exceed ¼" in length. As such, it becomes desirable to compensate the impedance by simply widening the trace in the uncoupled region highlighted in yellow.

While this adaptation is simple enough to comprehend in theory, it can be more difficult to achieve in practice. This is because on "controlled impedance" PCBs *fabrication vendors typically alter line widths on a given layer according to mapping tables tuned for their materials and process.* As such, if only the impedance of the differential trace is specified, it's possible the trace in the coupled region will be fabricated wider than the trace in the uncoupled region – making the problem even worse. The way to correct this problem is to also specify the desired impedance of the uncoupled single-ended portions of the trace.

4.4 Reducing Cost by Removing PCB Layers

As higher volume PCBs are revised to reduce cost and layers, it is imperative to confirm performance parity. The plots below compare two versions of a PCB before and after layer count reduction by examining their simulated performance within the larger system model. Both plots compare the original PCB's performance on the Y axis with the reduced-layer count PCB's performance on the X axis, showing eye widths at left (blue) and eye heights at right (red). There are over 2,500 dots in each diagram, with each dot representing the same channel in each system model. As such, dots on the black lines represent channels that perform the same on both PCBs. Dots above the line represent channels that perform better on the original PCB, while dots below the line represent channels that perform better on the reduced-layer PCB.

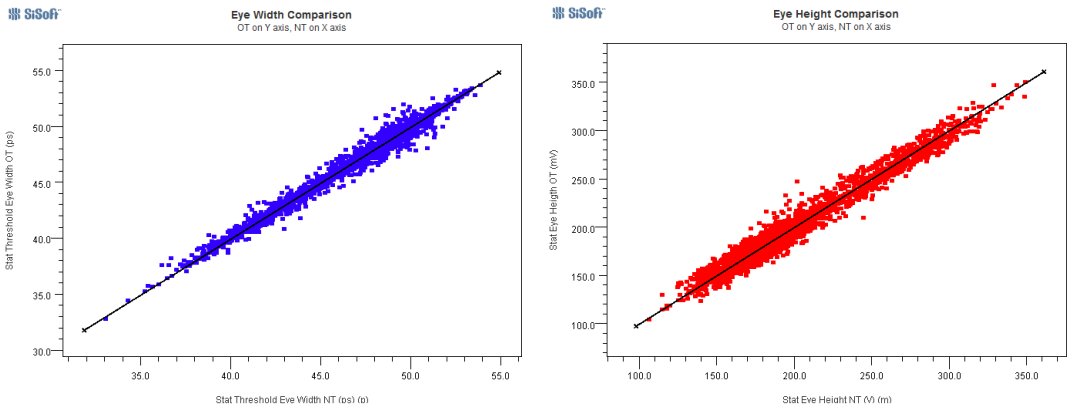

Figure 20: Eye Width (left) and Height (right) Contrasted Across PCB Revisions

The plots above demonstrate that eye widths (above left, in blue) generally vary only up to ~2% and the greatest variations are seen with channels with lots of margin (i.e., the variation from the black line gets wider as the plot moves to the right). As such, eye width variation is not relevant and below the anticipated accuracy level of the analysis. Eye heights (above right, in red) generally vary up to ~5% and variation is wider in channels with less margin – however the worst-case channels perform the same (on the black line at the lower left).

These plots confirm that the reduced-layer count PCB performs on par with the original, with neither version significantly out-performing the other. As such, system-level analysis is used to confirm adequate performance of the reduced-layer implementation allowing us to realize cost savings associated with laminating fewer PCB layers.

5. Summary and Conclusions

This paper has demonstrated and described new techniques for optimizing performance in high-speed serial links through the system-level manipulation of SerDes equalization settings. The manual optimization approach described minimizes intersymbol interference (ISI) by deriving Tx tap weights from a channel's pulse response. This technique improves performance, increases the system developer's understanding of relevant tradeoffs, and has been automated and scaled to be applicable to thousands of channels. For the systems shown, automated optimization improves simulated performance in 95% of channels across a 4x range of lengths. These improvements are achieved by managing amplitude/ISI tradeoffs resulting from Tx/Rx equalization trading to achieve required and optimal eye heights and widths. Performance of worst-case channels routed 25% longer than anticipated is shown to improve by more than 60%.

This paper also detailed methods for tuning performance using manufacturing process improvements. Multiple discontinuities spread across various PCB layers were demonstrated to become nearly transparent over time. Dual-diameter via construction and breakout trace compensation were also detailed as ways to reduce the impact of discontinuities. SI analysis also verified acceptable performance in reduced layer-count PCBs to achieve lower cost.

Acknowledgements

The authors wish to thank and acknowledge Sergio Camerlo at Ericsson as the visionary and motivator behind this series of papers, and for his relentless pursuit of engineering excellence. The authors also wish to thank Wheling Cheng, Kusuma Matta, Robert Wu, and Radu Talkad at Ericsson and Walter Katz, Frank deAlbequerque and Todd Westerhoff at SiSoft for their support. Additional thanks to Orlando Bell at GigaTest Labs for consistently delivering high quality measured data. Without the efforts of these and others this work would not have been possible.

References

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

[2] "Statistical Signal Analysis (SSA) Demystified" Anthony Sanders, EETimes 2007
http://www.eetimes.com/document.asp?doc_id=1275070

[3] "Multi-Gigabit Serial Link Analysis – Piecing Together a Design and Verification Strategy" Steinberger, Westerhoff, SNUG Boston 2007
http://www.sisoft.com/elearning/secure/files/Sept07_SNUG_Multi_Gigabit_paper.pdf

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

[5] "When Shorter Isn't Better" Steinberger, Wildes, Higgins, Brock and Katz, DesignCon2010
http://www.sisoft.com/elearning/secure/files/DesignCon_2010_When_Shorter_Isnt_Better.pdf

[6] "Fast, efficient and accurate: via models that correlate to 20 GHz" Steinberger, Brock, Telian, DesignCon 2013 paper 8-TA1
http://www.siguys.com/resources/2013_DesignCon_FastEfficientAccurateVias_paper.pdf

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

[8] "Moving Higher Data Rate Serial Links into Production – Issues & Solutions" Telian, Camerlo, Matta, Steinberger, Katz, Katz, DesignCon 2014 Best Paper
http://www.siguys.com/resources/2014_DesignCon_MovingToHigherDataRates_paper.pdf

Rev 1.2