




New SI Techniques for Large System Performance Tuning (Part 2)

by Donald Telian, SIGUYS,
and Michael Steinberger, Barry Katz,
SISOFT

This paper was originally published in the proceedings of DesignCon 2016. Part 1 of this paper was published in the July 2016 issue of [The PCB Design Magazine](#).

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.




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

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.




Figure 13: Eye Diagrams Produced by FFE, Typical 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.

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.

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: TDR of 7 Discontinuities Across Design Iterations

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.

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 [4] 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 [4] 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: System-level TDR Contrasting Normal and Dual-Diameter Vias

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.

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.




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 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.

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 deAlbuquerque 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] Donald Telian, 2004
- [2] Anthony Sanders, EETimes 2007
- [3] Steinberger, Westerhoff, SNUG Boston 2007
- [4] [“Simulation Techniques for 6+ Gbps Serial Links”](#) Telian, Camerlo, Kirk, DesignCon 2010
- [5] [“When Shorter Isn’t Better”](#) Steinberger, Wildes, Higgins, Brock and Katz, DesignCon2010
- [6] Steinberger, Brock, Telian, DesignCon 2013 paper 8-TA1
- [7] [“Simulating Large Systems with Thousands of Serial Links”](#) Telian, Camerlo, Steinberger, Katz, Katz, DesignCon 2012
- [8] [“Moving Higher Data Rate Serial Links into Production – Issues & Solutions”](#) Telian, Camerlo, Matta, Steinberger, Katz, Katz, DesignCon 2014 Best Paper

Rev 1.0