DesignCon 2019

Elimination of Highly Reflective Structures through Sliding Decision Feedback Equalization

Joseph Aday, The Raytheon Company
josephaday@raytheon.com

Hannah Lee, The Raytheon Company
hannah.lee@raytheon.com

Juli Olenick, The Raytheon Company
juli.olenick@raytheon.com

Hongtao Zhang, Xilinx Inc.
hongtao.zhang@xilinx.com

Geoff Zhang, Xilinx Inc.
geoff.zhang@xilinx.com

Abstract

Return-loss dominated high speed channels are not uncommon within the sectors of aerospace, defense, drilling, medical, and space, which require high-reliability products. In these environments, critical mechanical requirements can push electrical performance into a very challenging corner. Conventional SerDes design emphasizes channel ISI equalization and to some extent mitigating the impact of crosstalk, while leaving serious reflections to channel designers. This paper revisits the concept of the ST-DFE (Sliding Tap DFE), a.k.a. sliding tap DFE, and proves its applicability in high-reliability system-level solutions. We conclude with an example of a difficult impairment that is critical to product success and its resolution.

Author Biographies

Joseph Aday is a Principal Engineer at Raytheon, in Tucson Arizona, where he is focused on high-performance computing architecture and design. Prior to Raytheon, he did circuit design at the Boeing Company, in Anaheim, CA. Joseph holds a BS in Electrical Engineering, and a patent in antenna design. His current interests include system architecture, mixed-signal circuit design, and system level modeling.

Hannah Lee is a Senior Electrical Engineer at Raytheon. She has been with Raytheon for 10+ years working in high speed digital design and signal integrity analysis. She received a BS in Electrical Engineering from University of Texas at Arlington.

Juli Olenick has over 20 years of experience in Signal Integrity and five patents. A 1998 Penn State Graduate, she has been with Raytheon for over 5 years concentrating on signal and timing integrity. Prior experience was with a connector company, analyzing and creating high speed connector models to ensure performance. Currently, she is working with colleagues to optimize pinouts for various protocols and connectors.

Hongtao Zhang received his Ph.D. degree in Electrical and Computer Engineering from University of California, San Diego in 2006. He joined Xilinx in 2013 and is now a senior staff Design Engineer, working on SerDes architecture development and circuit design. From 2010 to 2013, he was with SerDes design team at Oracle Corporation, where he worked on circuit design and architecture modeling. Prior to that, he worked on SerDes characterization at Texas Instruments, Dallas. His current interests are SerDes architecture development and modeling, high speed mixed-signal circuit design and optimization, and system level modeling.

Geoff Zhang received his Ph.D. in 1997 in microwave engineering and signal processing from Iowa State University, Ames, Iowa. He joined Xilinx Inc. in June, 2013. Geoff is currently a Distinguished Engineer and Supervisor of Transceiver Architecture and Modeling Team, under SerDes Technology Group. Prior to joining Xilinx Geoff has employment experiences with HiSilicon, Huawei Technologies, LSI, Agere Systems, Lucent Technologies, and Texas Instruments. His

current interest is in transceiver architecture modeling and system level end-to-end simulation, both electrical and optical.

1. Introduction

High speed SerDes solutions have permeated the large majority of commercially available electronics, data centers, and test equipment, with data rates moving to 56Gbps and 112Gbps, and with the adoption of PAM4 modulation. This has been driven by the need for high-bandwidth video and data services. However, lower data rate NRZ still plays a vital role in SerDes solutions across the entire industry.

A clear example is with high-reliability products within the aerospace, defense, drilling, medical, and space sectors which have seen minimal penetration of reliable high speed solutions. This is not because vendors are not targeting these high-value customers, but rather because the solutions simply cannot be successfully implemented. Drivers in high-reliability systems center around extreme temperature, mechanically hardened interconnect solutions, and Class III PCB structures. Each of these critical design drivers introduces difficult channel impairments that are not often overcome by conventional SerDes designs, thus could result in failed product development cycles.

This paper is intended to share an essential method of resolving poor return-loss structures through the use of Sliding Tap Decision Feedback Equalization (ST-DFE). We will explain what a ST-DFE is and how it works in handling channel ISI and in mitigating channel reflections. We will provide theoretical analysis to show its advantages over the conventional DFE, by showing that the number of DFE taps can be reduced dramatically while the system SNR at the data slicers is largely maintained. We will also show different implementation algorithms that adaptively enable ST-DFE. Armed with the fundamentals of ST-DFE, we will show the performance difference using IBIS-AMI simulation with and without ST-DFE included in the model.

We will then explore on how an ST-DFE can be employed to implement a high-speed SerDes link in a return-loss dominated channel. We will then show a real example of a difficult impairment that is critical to product success, which was resolved satisfactorily with ST-DFE.

Although not a new topic in industry, its implementation has yet to proliferate because its use cases have not been well explored. This new application of ST-DFE will be of interest to users in the aforementioned sectors, and hopefully increase the level of attention from the design community.

The key takeaway is that return-loss dominated high speed channels are not uncommon with high-reliability products. SerDes solutions with ST-DFE are of great importance in providing a

system level optimal solution with excellent performance margin and great reliability. The working mechanism of the ST-DFE is discussed. A real product example with an ST-DFE is provided.

2. DFE Basics

Channel equalization to compensate for ISI (inter-symbol interference) is usually accomplished through TX de-emphasis and RX CTLE (Continuous Time Linear Equalization) and DFE (Decision Feedback Equalization). The reader is assumed to have a basic knowledge of these. The following serves as a review of DFE.

2.1 DFE for Post-Cursor ISI Removal

DFE is a standard technique to mitigate ISI without amplifying noise. It works by directly removing ISI from previous bits, allowing subsequent bits to be properly sampled. Figure 1 shows the block diagram of a simple DFE implementation. DFE starts with a "slicer for decision" to determine whether the current symbol is high or low. The resulting symbol goes through unit delays and gets multiplied by specific tap weights. The weighted delayed signals are then added to the input analog signals. If the tap weights are well selected, the result of this feedback loop is able to compensate for as many taps of post-cursor ISI as the DFE has.

$$y(t) = u(t) + \sum_{i=1}^{n} w[i] * y_d(t - i * UI)$$

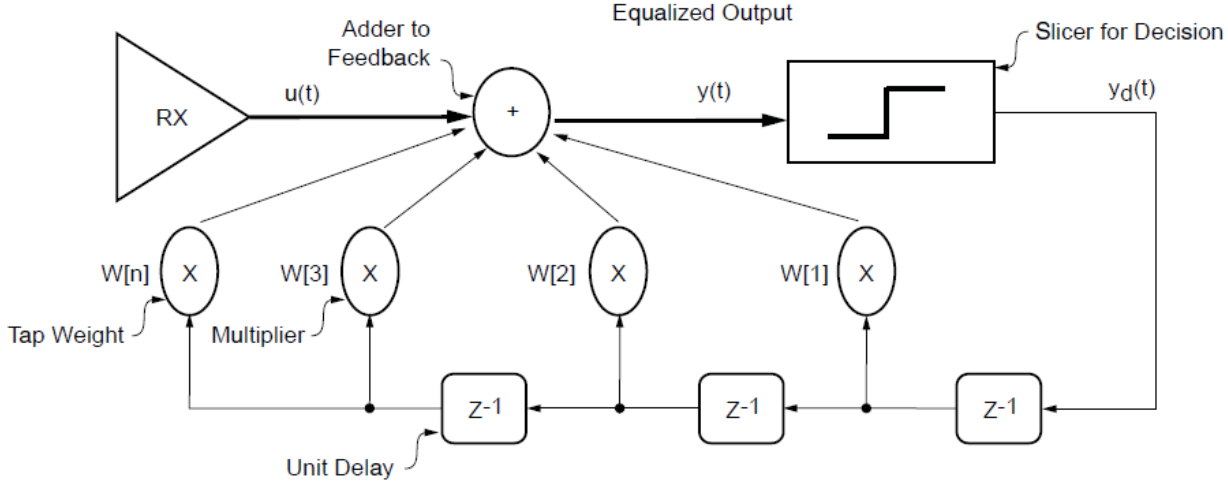

Figure 1. Simplified DFE block diagram.

Figure 2 and Figure 3 illustrate two data streams equalized by the DFE. The red curve in Figure 2 is the pre-equalized data stream of 010110000. Without any equalization, the second '0' would be sampled incorrectly. The green curve is the feedback stream added to the original data

stream with well-adapted DFE tap weights. The blue curve is the final data stream after applying the feedback stream. Compared to the red curve, the blue curve has much more margin and is likely to be sampled correctly. Similarly, Figure 3 shows the DFE equalizing a data stream of 011110000. The first '0' after the '1's and the first '1' after the '0' showed less margin before DFE. After DFE equalization, the decisions have more margin and will be sampled correctly.

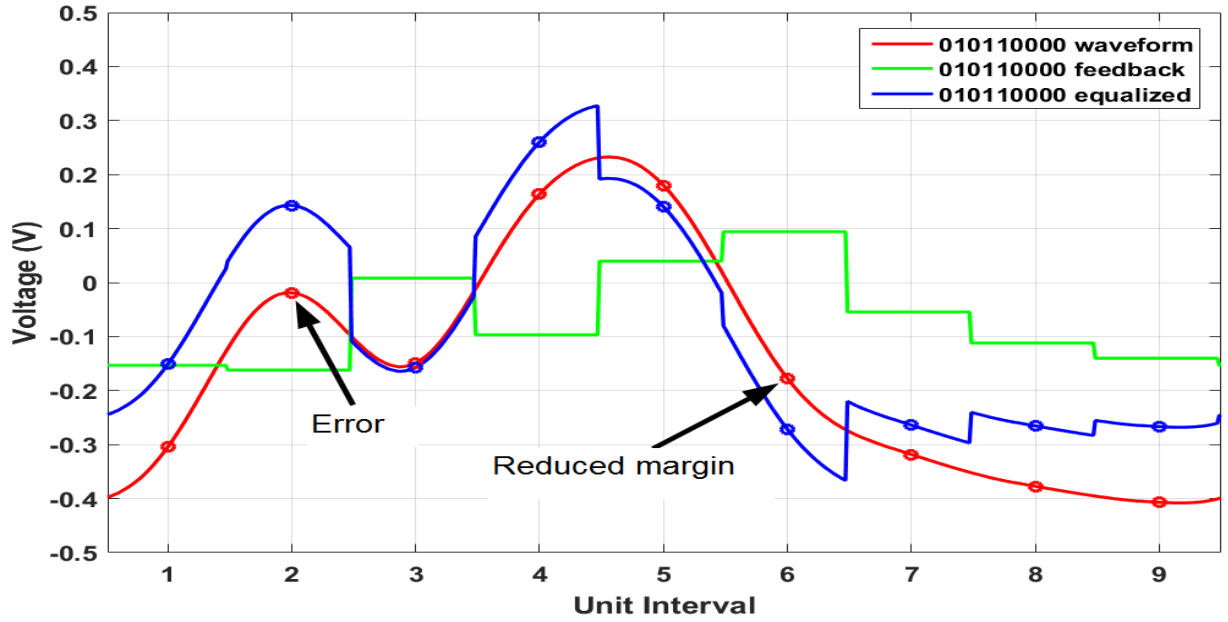

Figure 2. 0010110000 binary pattern with DFE equalization applied.

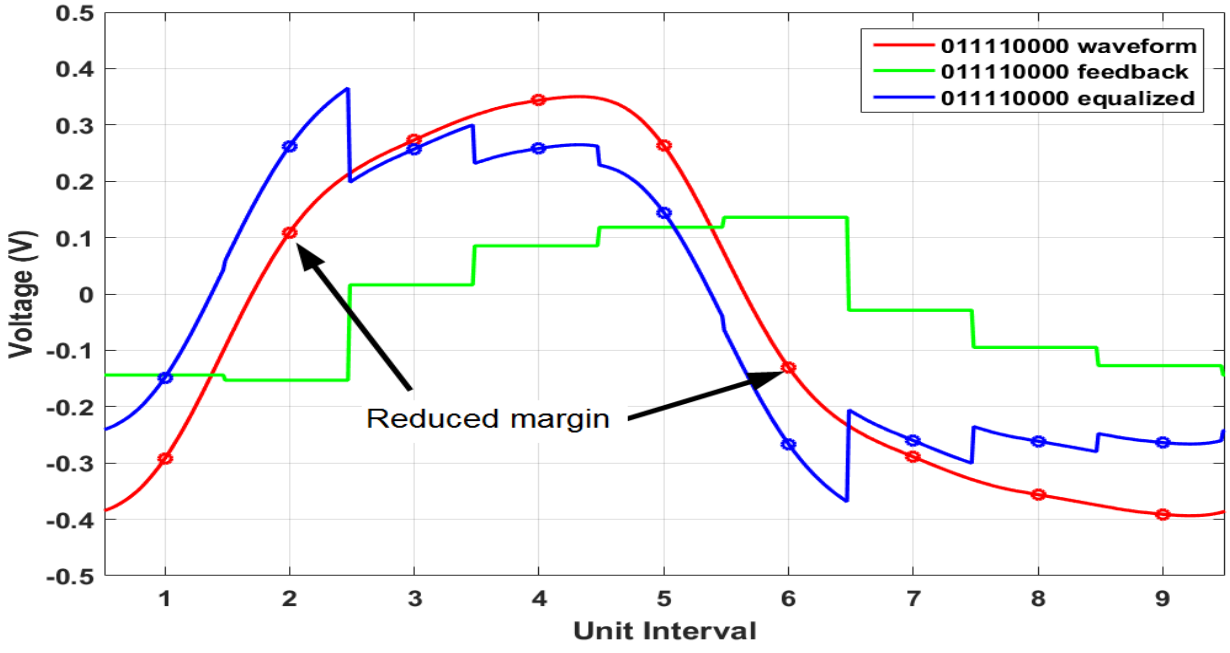

Figure 3. 0011110000 binary pattern with DFE equalization applied.

2.2 ST-DFE Concept

The DFE working principle illustrated above is for the so-called fixed taps (FT-DFE), in which the taps are delayed consecutively without holes. One can also have several fixed taps, and several sliding taps, the ST-DFE. This combination of fixed and sliding taps enables great compensatory capability while keeping the total number of taps small and therefore the power consumption low and cost in check. An example is shown in Figure 4. Sliding some of the DFE taps by a sufficient number of unit intervals (UI) can help cancel reflections and improve the system SNR.

The CTLE can remove a smooth long tail beyond the range of the DFE in a single-bit response, but it cannot correct for bumpy reflections that tend to alternate in a short period. This problem can be mostly remedied by introducing additional DFE taps as sliding taps and adjusting their position based on the actual location of major long-tail ISI components.

In order to avoid unwanted interactions between the sliding tap and CTLE adaptation loops, the CTLE adaptation logic is gated when the location counter is passing through a location also occupied by a sliding tap. In other words, the CTLE is adapted with a certain range of taps and any sliding tap landing in that range is skipped. Thus, those taps do not affect the CTLE adaptation.

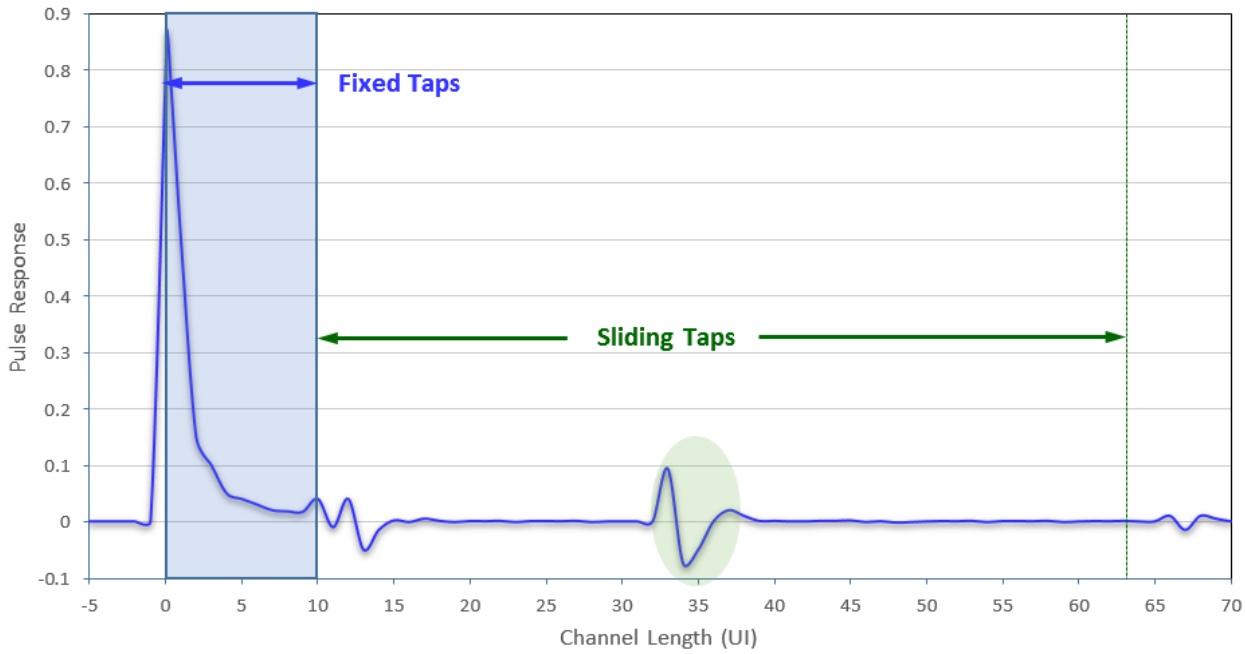

Figure 4. ST-DFE illustration.

3. Channel Description

Ruggedized stacking connectors are commonly used to mate together multiple circuit cards in a stack configuration. Aerospace-grade connectors that serve this purpose are available from various hi-rel connector vendors. These kinds of connectors are designed to be mechanically robust, at the understood expense of high-speed electrical performance. As such, they are often a limiting factor in channel performance because of their low impedance geometries and reflective stubs. A channel with hi-rel connector elements at each end and approximately 5" of low-loss PCB material was created as shown in Figure 5.

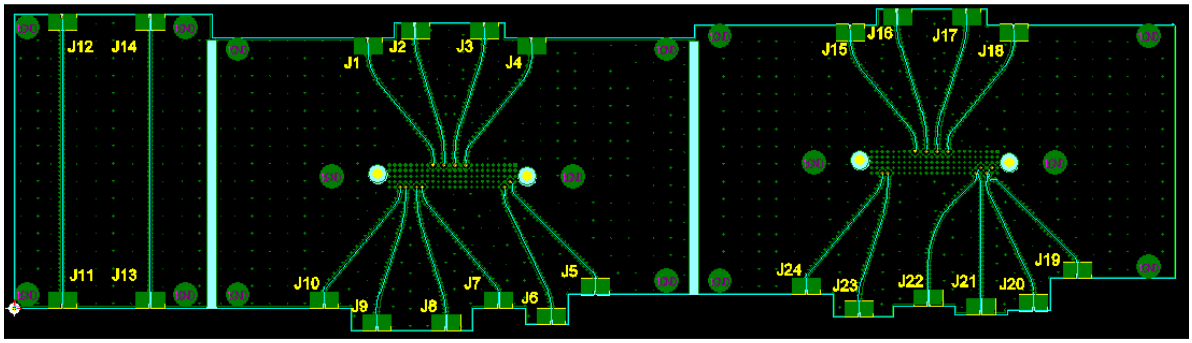

Figure 5. PCB Test Routing.

The PCB test structures need to support very large pins and pin pitches, which result in large via geometries and anti-pads. These large pins and wide pitches are the result of mechanical optimization for reliability. These large geometries limit pinout/performance optimization because any optimizations result in a larger pin field. For example, inserting blanks between signal and ground pins will increase the impedance, but would result in a connector that is too large to be feasible in a real design (Figure 6). This means that we have to find creative solutions to deal with largely reflective interconnect behavior caused by uncontrolled, low impedances and quarter-wave stub resonances.

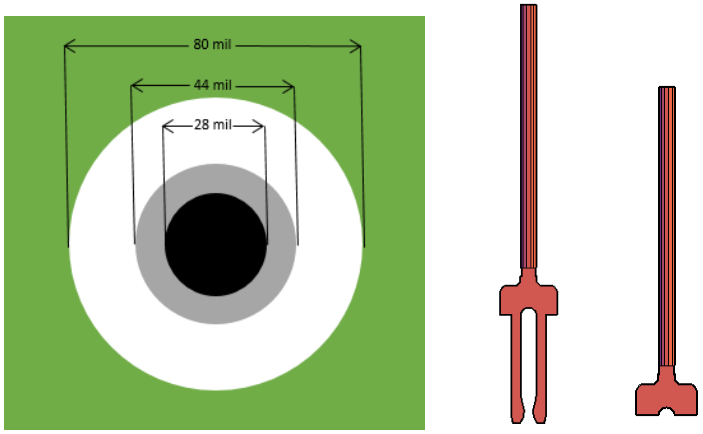

Figure 6. Press-fit Pin Geometries and Pin/socket Examples

In Figure 7 several compact pinouts were evaluated. These included GSGSG, GSSG, and a diagonal GSSG. High-reliability pins typically have multiple points of contact in their mate to the socket, which drives additional area and limitations to electrical performance. Although

improvement can be had in pinout optimization, its overall success is still limited by board real estate. This is because the pins will have a pitch of at least .062”- any additional spacing to optimize pinout performance quickly grows the connector and/or takes away from needed pin resources. Additionally, high-reliability requirements ensure multiple contacts for each pin.

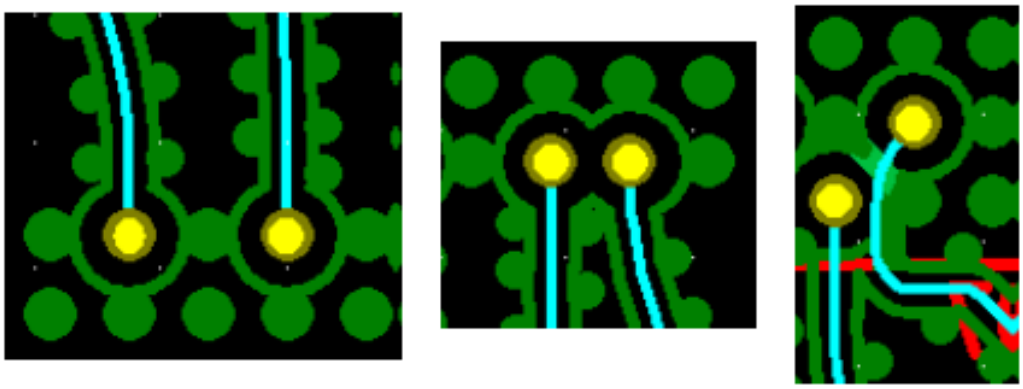

Figure 7. Test Breakouts.

A test board representative of realistic interconnect demands was created to evaluate the performance of the connector in-system. This is shown in Figure 8.

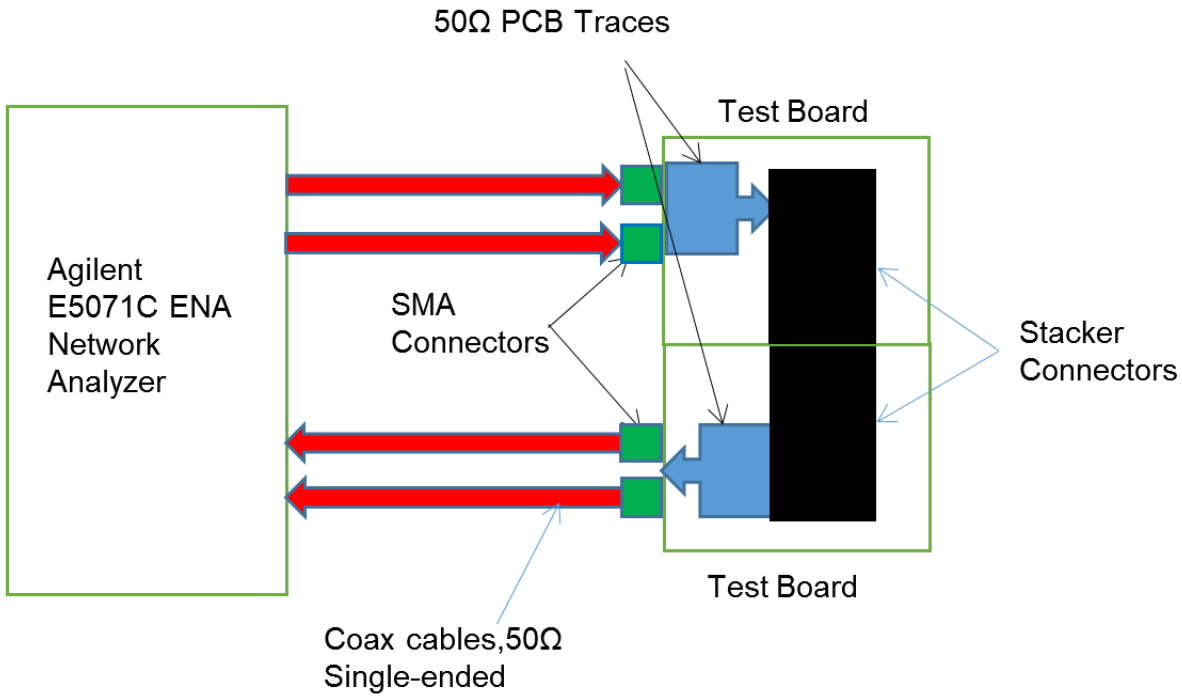

Figure 8. Test Measurement Setup

Test boards were created to measure the connector and evaluate its performance in a stack of PCBs, as illustrated in Figure 9. In the below example, an HD Stacker was evaluated, after passing testing per MIL-DTL-55302G, including mechanical shock & vibration. The connector

socket stub extending from the top side of the board is approximately 150mils. Since the connector pins are press-fit, they also extend from the bottom side of the board by approximately 50mils. The multiple mates of the pin to socket also result in some smaller stubs which are considered out-of-band for the purposes of this paper.

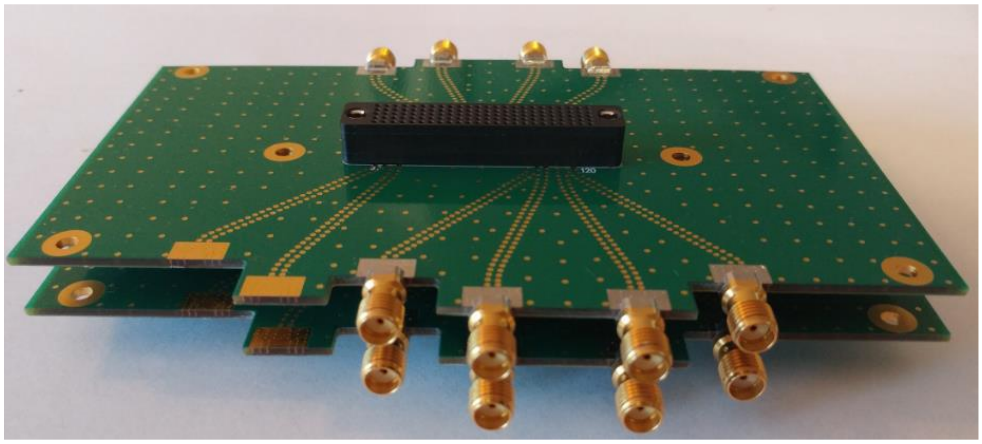

Figure 9. An Example High-Reliability Stacking Connector.

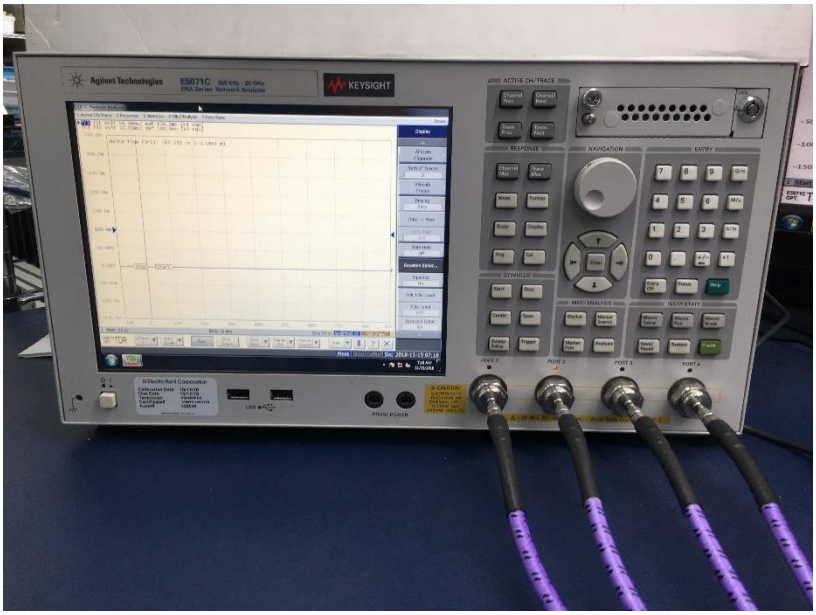

Figure 10. Vector Network Analyzer Used for Measurement

Note that the test fixture PCB and test cabling losses were de-embedded to show the performance of the stacked connector pair only. Here is an example of the GSGSG pinout. The primary differential insertion loss is caused by the connector pin stubs that are a result of the stacking nature of the connectors. Additional reflections are caused by the low-impedance pins at each end of the channel.

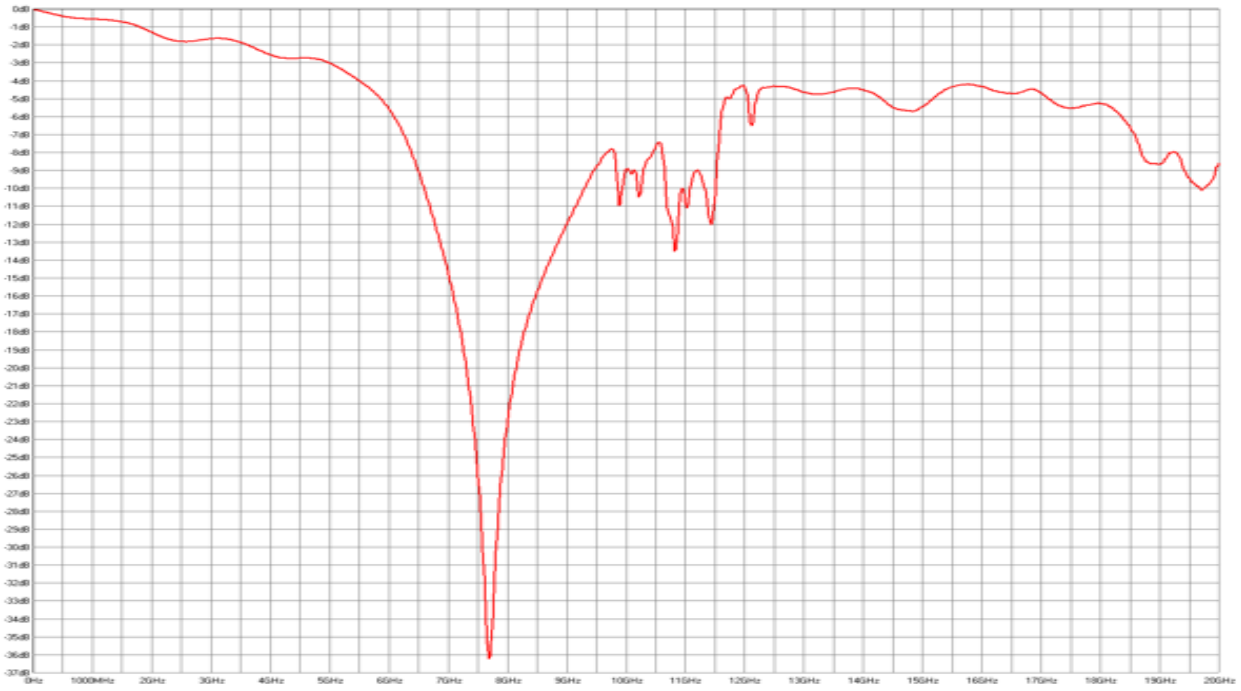

Figure 11. Example Mated Connector SDD21.

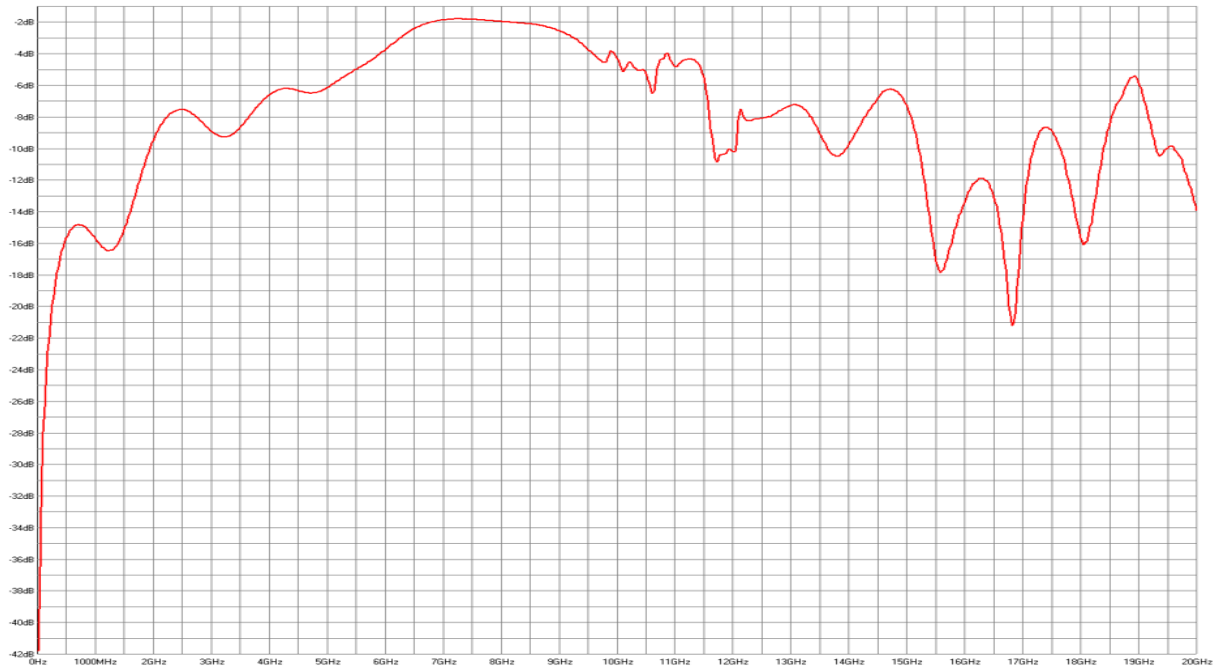

Figure 12. Example Mated Connector SDD11.

The next step is to build up a sample channel using these structures. For the signal to travel from mated connector pair, across the PCB, and through the second mated connector pair, the following structure was built:

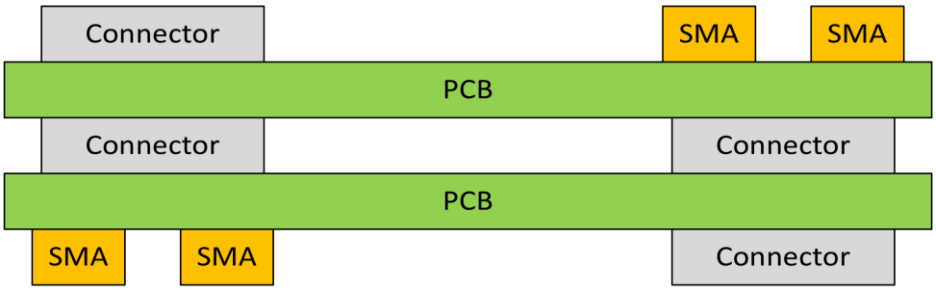

Figure 13. Example of Test Circuit for Measurement.

The TDR of the Test Circuit shows large impedance dips at the locations of the connectors (Figure 14). The lowest dip is approximately 22 ohms, which is about half of the desired 50 ohms. This is an additional artifact of common high-reliability connectors that is not dealt with in typical high-volume consumer applications. This is, again, a result of the drive for mechanical reliability and safety that pushes electrical performance into a lower level of criticality. For this reason, typical high-speed, impedance-controlled connectors cannot be used. As discussed previously, the structures needed to meet high-reliability performance goals typically have poor impedance and stub length control.

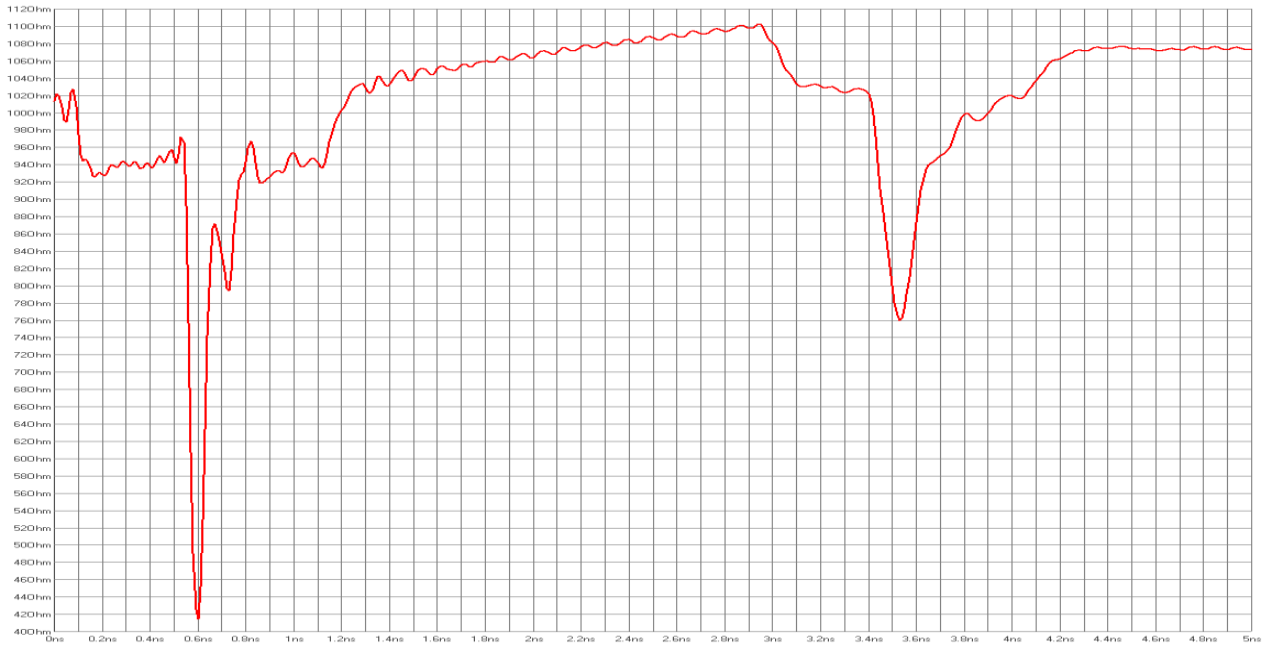

Figure 14. TDR Results of Test Circuit.

The Single-Bit Response of the channel is shown in Figure 15. It is seen that additional ISI is created approximately 3nS after the main cursor. This is caused by a delayed reflection bouncing back and forth between the low-impedance connectors. The circuit between the two reflective end points becomes a “noise trap” that produces a smaller but delayed ISI. The distance to the secondary ISI is dependent upon the length of the channel between the two resonant structures. Using lower loss material only makes the problem worse because the reflection would be even less attenuated. In this instance, it would require over 35 fixed DFE taps at a data rate of 10GT/s to adequately resolve the entire channel. This is clearly not feasible due to cost.

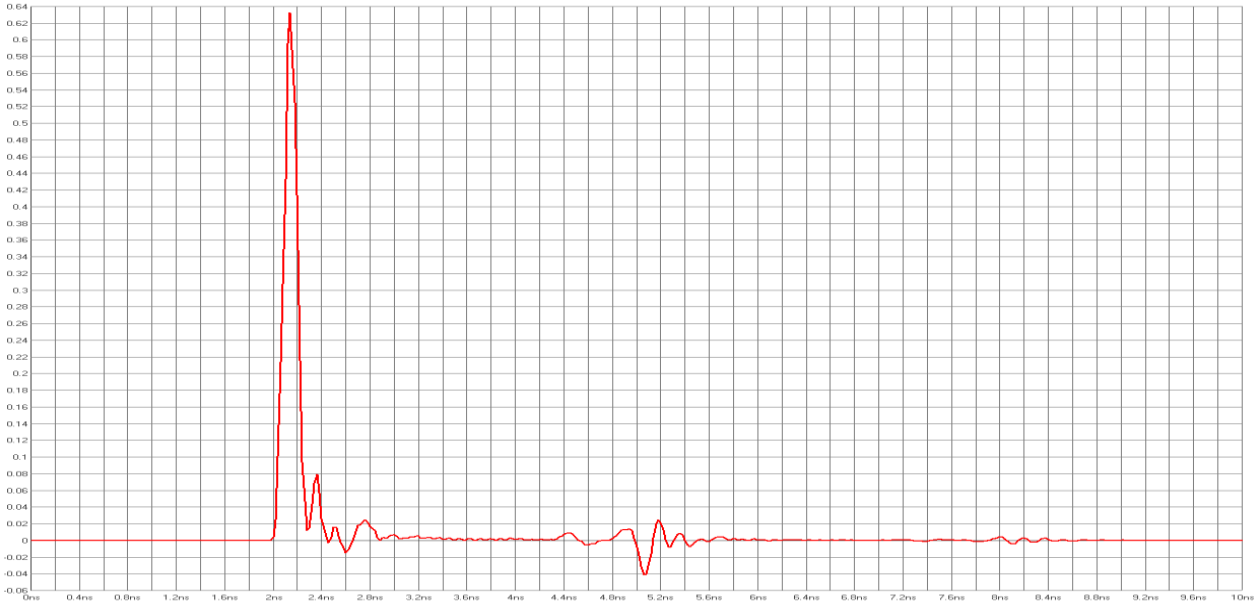

Figure 15. Single-bit Response of Test Circuit.

4. IBIS-AMI Simulations

4.1 Simulation setup

The SerDes architecture used for this study has TX FFE, RX CTLE and DFE. The TX FFE is not enabled in the simulation. The CTLE has multiple stages to handle high, medium, and low frequency bands of equalization. CTLE settings and AGC values are all adaptively tuned.

The DFE in this specific design has a total of 15 taps. There are two modes for configuring the taps, besides another register that can disable the DFE completely.

- In DFE_mode0: this is for FT-DFE. The tap locations are from the first post-cursor to the fifteenth post-cursor. The adaptation engine only computes tap strengths.
- In DFE_mode1: this is for ST-DFE. The last 5 taps can be anywhere from 11th to 64th post-cursor locations. The adaptation engine not only computes tap strengths, but more importantly, decides the tap location on the fly.

The above system is simulated with an IBIS-AMI model. The data rate is set to 10GT/s using a PRBS23 data pattern. The simulation test bench is shown in Figure 16. Impairment parameters such as jitter and noise are included in the simulation. However, they are less important as the purpose is to compare the link performance with the two different DFE schemes.

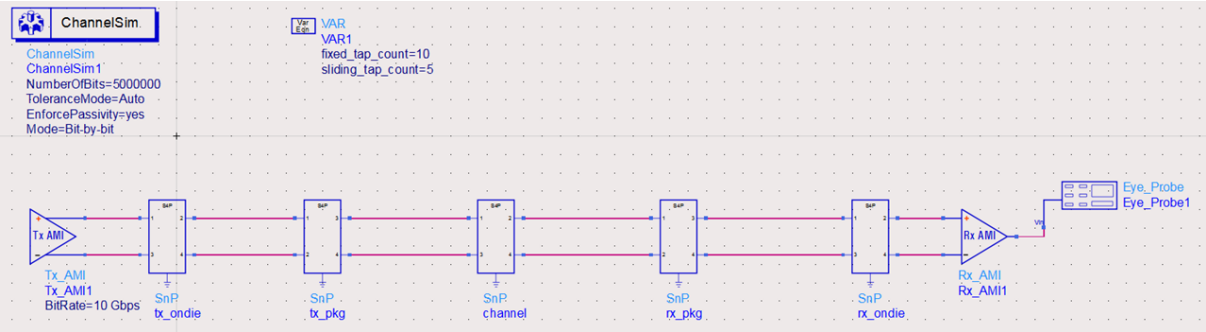

Figure 16. IBIS-AMI Model Simulation Setup

4.2 Simulation results

Table 1 summarizes the results using the above setup. Eye diagrams are shown for a first order conclusion as to whether ST-DFE has advantages or not. Then the eye height and eye width at BER equal to 1e-12 and 1e-15, i.e., EH12/EW12 and EH15/EW15, are computed from the simulated data. It is now obvious that the ST-DFE scheme has advantages over the fixed-DFE architecture in terms of performance.

DFE Mode	DFE_mode0	DFE_mode1
Eye diagrams	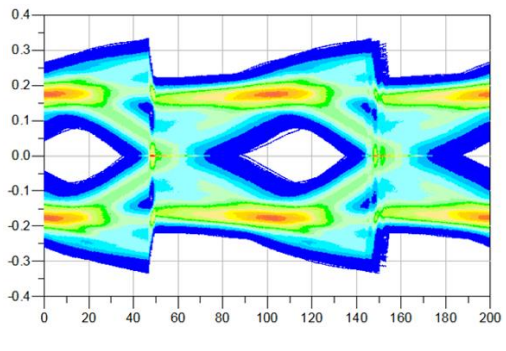	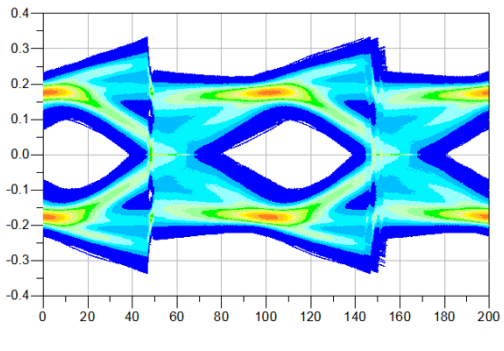

EH12 (mV)	149	195
EW12 (UI)	0.465	0.590
EH15 (mV)	137	183
EW15 (UI)	0.425	0.560

Table 1. Simulation Eye Diagrams and Eye Metrics at BER = 1e-12 and 1e-15.

Figure 17 shows DFE tap convergence in which the sliding taps are enabled. The x axis is the time and the y axis is the converged DFE tap value for each corresponding tap. The color coding maps to the tap location. h1 to h10 are the fixed taps. To better show the convergence details for the higher order DFE taps, the largest DFE tap h1 is not shown in the plot. In addition to the 10 fixed taps, there are 5 sliding taps whose positions are adaptive and do not need to be continuous. It can be seen that the 5 sliding taps converged to locations 33, 34, 13, 12, and 15, in the order of DFE tap magnitudes. This implies that the ISI at cursors 11 and 14 are less dominant than those at 33 and 34 (as shown in Figure 18.)

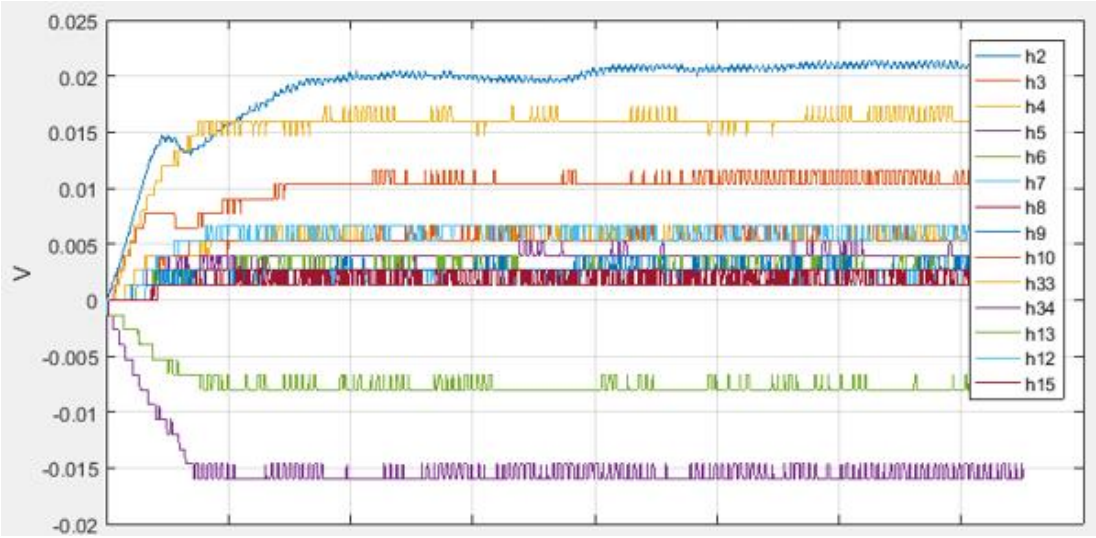

Figure 17. DFE Tap Convergence Profiles: 10 Fixed and 5 Sliding.

Figure 18 shows the single-bit response (SBR) of the channels cascaded with SerDes TX and RX packages. The real SBR should be modified by the TX and RX front-end models, as well as the converged CTLE settings, AGC setting, and the CDR sampling phase. So the SBR we will discuss in Figure 18 is only an approximation, but clearly highlights the use case. From the main cursor, it is seen that cursors 33 and 34 are much larger than those beyond tap 10. This shows that the design of the hardware and the adaptation algorithm are both working. Since the goal of the algorithm must be to find the max ISI location, we have verified that it converged correctly to the expected locations of main interest. Although power is not part of an AMI simulation, it

should be expected that floating DFE taps further out in time would result in an increase of power utilization.

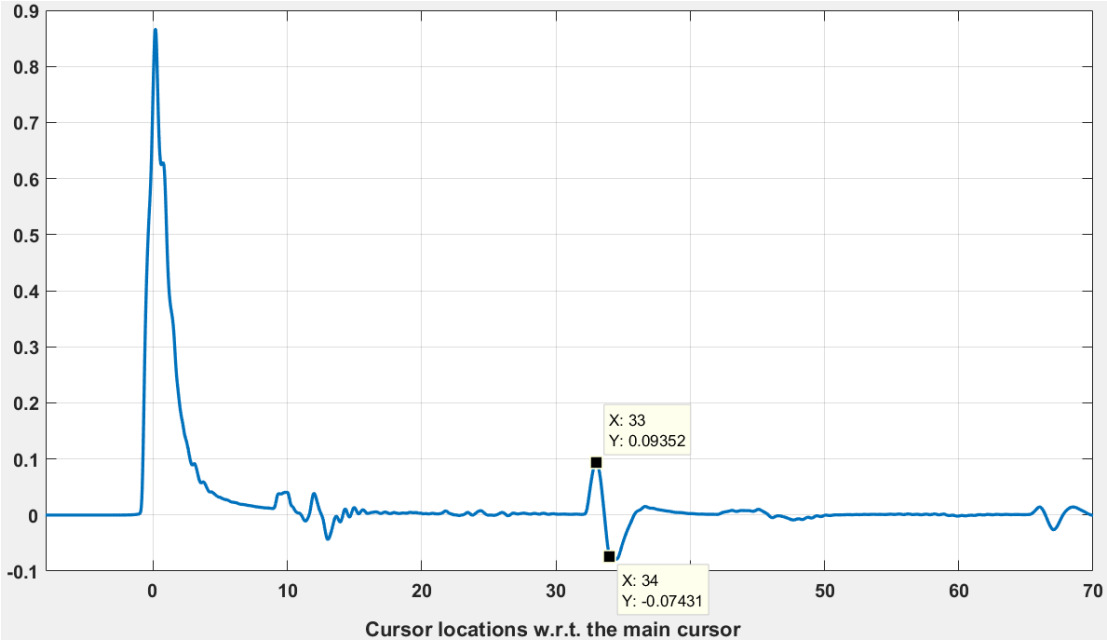

Figure 18. Single-bit Response of TX-package + Channel + RX-package.

5. Summary

In this paper, we presented the case for using mechanically robust, high-reliability connectors. We reviewed why these kinds of connectors have such poor Signal Integrity, and then shared a circuit employing them. We discussed how such a link can demonstrate large reflective behaviors, and identified when a work-around is needed. We discussed how a cascade of two large reflective structures creates a "noise trap." We proved that a channel with these multiple, significant reflections can produce very large ISI far away from the main cursor. We introduced the concept of a sliding-tap DFE to reach the additional ISI and resolve it.

Through IBIS-AMI simulations, we showed the advantage of the ST-DFE over the conventional FT-DFE. Comparison data, from a real test case, showed that eye margins were improved by approximately 30%. We also noted potential system trade-offs for the user should include consideration of additional power needed as well as implementation complexity. In conclusion, we have demonstrated a valid technique to enable high-reliability systems to meet eye mask requirements with needed margin at higher data rates.

References

[1] Harry Fu, “Equalization for High-Speed Serial Interfaces in Xilinx 7 Series FPGA Transceivers”, White Paper: 7 Series FPGAs, WP419 (v1.0) March 27, 2012

[2] Jafar Savoj, et al, “Design of High-Speed Wireline Transceivers for Backplane Communications in 28nm CMOS”, Proceedings of the Custom Integrated Circuits Conference, September 2012

[3] Glenair: G. Hunziker and L. Blackwell, “High Speed Characterization Report for Differential Applications”, May 21, 2018

[4] Glenair, “HD Stacker™ Rugged Board-to-Board Stackable Connectors”, https://www.glenair.com/hd-stacker/index.htm

[5] Glenair, “HD Stacker™ GSTBL single-sided mating connector”, https://cdn.glenair.com/hd-stacker/pdf/gstbl-single-sided-mating-connector.pdf

[6] Glenair, “HD Stacker™ GSTB (.095) single-sided mating connector

[7] Bottom-of-stack”, https://cdn.glenair.com/hd-stacker/pdf/gstb-single-sided-mating-connector.pdf

[8] Airborn, “RC Series”, https://www.airborn.com/products/connectors/rc-series

[9] Airborn, “R Series”, https://www.airborn.com/products/connectors/r-series-home-page