DesignCon 2008

Performance Limitations of Backplane Links at 6 Gbps and Above

Jason Chan, Amphenol TCS
jason.chan@amphenol-tcs.com

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

Jose Paniagua, Amphenol TCS
jose.paniagua@amphenol-tcs.com

Abstract

As data rates exceed 6 Gigabits per seconds (Gbps), the size of many physical structures within backplane interconnects become significant. For example, the electrical length of a typical backplane connector produces standing wave resonances between 2 GHz and 3 GHz. This can result in dramatic changes in crosstalk, mode conversion, and insertion loss deviation at fundamental frequencies of interest. These passive interconnect issues can be further compounded by the common mode issues and other challenges silicon vendors and chip designers must confront. This paper focuses on the major performance limitations seen across numerous backplane designs and proposes methods to improve the overall performance of such systems.

Authors' Biographies

Jason Chan is a signal integrity engineer at Amphenol TCS. His activities include connector development, electrical analysis/synthesis of passive backplane systems, correlation of aggregate system measurements with models, and development of PCB launch structures. He received his MSEE degree from the University of New Hampshire.

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

Jose Paniagua is a signal integrity engineer at Amphenol TCS. His activities include connector modeling, link analysis of backplane channel systems and high frequency measurements. He has previously worked in PCB manufacturing and process control. He received his BSEE degree from the University of Massachusetts Lowell.

Introduction

As data rates exceed 6 Gigabits per seconds (Gbps), significant challenges arise for backplane interconnects as this corresponds to an inevitable increase in spectral content. This subsequently translates to an effective increase in the electrical size of backplane connectors that are susceptible to standing wave resonances at critical fundamental frequencies. Consequently, standing wave resonances can exacerbate crosstalk and insertion loss deviation at frequencies between 2.5 GHz and 5 GHz, which is a critical frequency range for the contemporary data rates of interest. Furthermore, these standing wave resonances severely diminish the signal-to-noise ratio. The continual push for lower loss interconnects can actually be counter-productive as no dissipative mechanism exists to mitigate the crosstalk and reflections. These issues are illustrated throughout the proceeding section.

Most SERDES design engineers are aware that decreasing bit times and subsequently faster rising edges of high data rate devices precipitate considerable challenges in terms of controlling the common mode voltage produced from transmitters. This is coupled to the ever present challenge of controlling skew in PCB structures. The interaction of high common mode voltages at the driver and the accumulation of skew throughout the interconnect produces significant noise levels through mode conversion crosstalk. These increased noise levels are produced by adjacent aggressors and mode conversion on the victim directly. The strongest resonances in conjunction with high mode conversion effects coincide with critical frequency ranges thereby leading to a compounding of the severe signal degradation. A qualitative explanation of the overall effect of mode-converted crosstalk on the backplane interconnect is investigated.

Performance Limitations at the Connector Level

The contemporary design paradigm that currently exists for backplane systems operating at 3.125 Gbps is predicated on the assumption that the electrical size of conventional connectors is sufficiently small. This lumped-element based presumption facilitates the usage of rule-of-thumb design techniques thereby expediting the design and analysis of such backplane systems. Many of the aforementioned rule-of-thumb techniques are essentially obsolete at higher data rates as designers must consider the full wave distributed nature of connectors as the transmission lines are electrically long relative to the spectral wavelengths of operation. Thus, a significant obstacle that system designers must surmount at and above 6.25 Gbps is the inherent susceptibility of connectors to sustain standing wave resonances. The following is a rudimentary quantitative illustration of such resonance effects. The electrical length of all backplane connectors is between 100 picoseconds for smaller connectors or shorter rows in larger connectors and 300+ picoseconds for the outer rows of larger connectors. The round trip delay of the connector determines the resonance frequencies within a connector structure. The fundamental resonance frequency in the longest row of a connector is approximately

$$\frac{1}{300\,ps * 2} = 1.67GHz$$

Manifested in shorter rows are higher harmonic standing wave resonances as

$$\frac{1}{100\,ps * 2} = 5GHz$$

It is important to note that all backplane connectors will have resonant structures between 6 and 10 Gbps. If not controlled, the deleterious effects of resonances manifest themselves with respect to three different electrical aspects of system design: significant insertion loss deviation, extremely high crosstalk, and mode conversion.

Insertion Loss Deviation

The most common issue with backplane interconnects is the poor impedance match between the connectors and the footprints. This is due to many different aspects. For conventional press-pin (often referred to as compliant pin) connector, the plated-through-holes (PTHs) or vias of the corresponding footprints must possess a sufficiently large barrel diameter to support the compliant pin. Furthermore, the aspect ratio of the backplane to the diameter of the PTH diameter is often confined from 10:1 to 12:1 in order to facilitate high volume manufacturing processes. For instance, the PTH diameter for a 250 mil thick backplane is typically constrained to approximately 20 mils. Overall, the size of the PTH, the signal density requirements, and associated ground via shielding (required to mitigate crosstalk in the corresponding footprint) yields a backplane footprint characteristic impedance that is usually between 70-80 ohms. This is one of the reasons

backplane interconnects, such as PCI Express, is evaluating 85 ohm characteristic impedance links.

An alternative that is also available in the interconnect landscape is the surface mount backplane connector. An intrinsic electrical benefit derived from this solution is the possible utility of smaller via diameters for improved impedance matching. However, two primary obstacles are in contention from an electrical and mechanical perspective. First, the corresponding surface mount via diameter is ultimately constrained by the maximum possible aspect ratio for a specific backplane thickness. In other words, thicker backplane stackups require larger via diameters. Thus, its impedance matching benefits are inversely proportional to the desired thickness. Secondly, new mechanical issues that come to fruition with this surface mount solution include soldering registration and co-planarity problems on large boards. In addition, finding substrate materials that concurrently conform to desired electrical parameters and the mandated lead-free soldering requirements can also be a significant challenge. Consequently, most backplanes still use press-pin connectors. [1]

The impedance of the daughter card footprint must not be ignored as it has its own set of challenges. The daughter card, which may consist of a line card, fabric card or CPU module, is not typically constrained by the aspect ratio limitations as the cards are normally between 63 and 125 mils thick. However, a backplane system usually has a single backplane that support 10 to 20 daughter card slots. Therefore, the cost of backdrilling every daughter card could be prohibitive due to the large number of boards. Given this sensitivity, system designers must trade off the cost savings for additional PTH stub capacitance in the daughter card footprints. Therefore, the connector possesses inherent impedance discontinuities at the daughter card and backplane footprints. These discontinues generate the standing wave modes within the connector thereby entailing insertion loss deviation.

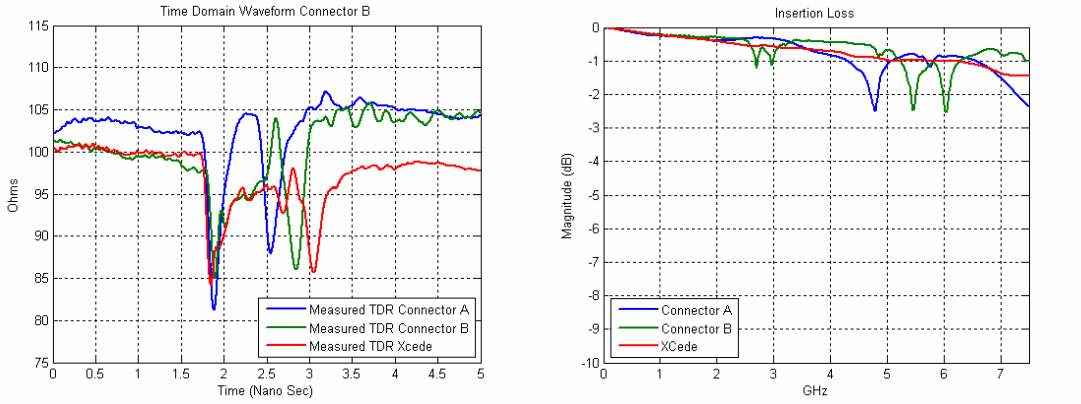

(a) (b)

Figure 1: (a) TDR, (b) Insertion Loss

Figure 1(a) depicts TDR measurements of a few typical connectors. Each connector exhibits similar impedance mismatches at the daughter card footprint at 1.9 nS (ns). The backplane footprint is the second impedance trough between 2.5 nS and 3nS. Note that all connectors were measured with daughtercard panels that all consist of identical stackups, substrate materials (FR4), and thickness (63 mils). The differences in measured

phase delay between the connectors are due to different rows lengths and dielectric composition encapsulating the respective leadframes.

All insertion loss measurements consist of a single connector with etch losses and dispersion effects de-embedded through TRL calibration techniques. Note, however that the aforementioned measurements account for the effect of PTH discontinuities as the TRL calibration methodology is unable to de-embed such inhomogeneities. The PTHs of the daughtercard test panels were intentionally designed to have a minimal impact by routing on the lowest layer of a thin board in order to alleviate PTH stub capacitance effects. Figure 1(b) illustrates the measured insertion losses that correspond with the TDR measurements. Note here that Connectors A and B utilize discrete ground shields within each wafer. These result in multiple different length return or ground paths in the connector. Connector A is a relatively short row with a propagation delay of approximately 105 pS. This induces the first resonance harmonic just below 5 GHz. The insertion loss of connector B, whose propagation delay is approximately 180 pS, exhibits dual resonant modes at centered at approximately 2.8 GHz. This particular connector features two ground paths per signal. Due to the topological nature of right-angle backplane connectors, the lengths of the accompanying ground shields are intrinsically mismatched. Therefore, the discrete ground shields resonate at two distinct frequencies centered around 2.8GHz.

The TDR plot for Amphenol-TCS XCede® connector, whose measured propagation delay is 180 pS, demonstrates a similar impedance profile with slightly lower connector impedance. The XCede® connector was optimally designed to produce the most linear insertion loss without adding excess attenuation. The insertion loss for XCede® depicts a linear response without any noticeable resonance troughs. The absence of any significant resonance harmonics is due to the effectiveness of the common ground path in conjunction with the hybrid shield construction, which is detailed in a previous paper. [2]

There are some common misconceptions about these resonances. First, it is expected that these resonances are damped by the etch losses in the complementary trace. Figure 2 demonstrates that this misconceived notion is implausible. One measurement was recorded with minimal PCB etch while the other was taken with 3 inches of etch on either side of the connector. The resonances appear in the same location with very similar magnitudes. Also observe that the resonance is actually greater in the measurement with etch included, although this maybe due to the resolution of the measurement. Overall, though, figure 2 clearly delineates the coincidental resonance harmonics in both measurements. It is also important to note that the high “Q” resonance in figure 2(b) is actually outside the IEEE 802.3ap insertion loss deviation guideline despite the fact that this is essentially a *best case* scenario with minimal via parasitics.

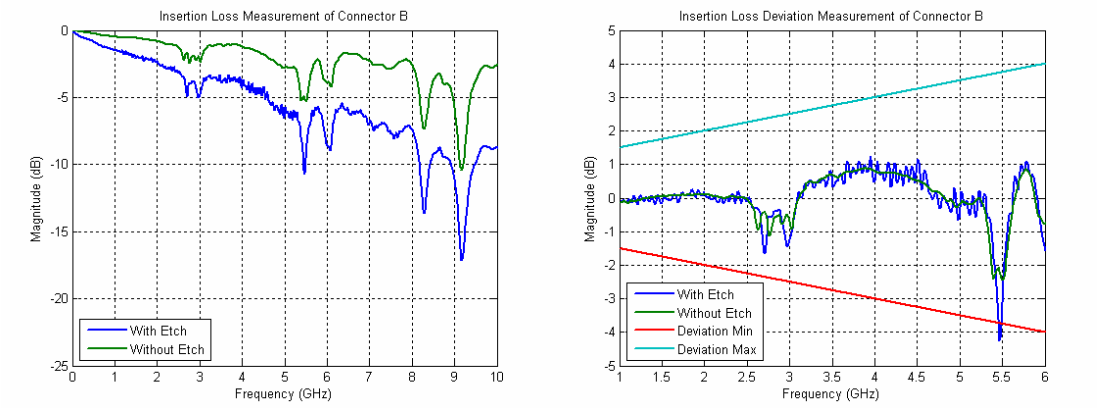

Figure 2: (a) Insertion Loss, (b) Insertion Loss Deviation

Additional resonances commonly occur in backplanes where partial standing wave modes are manifested between the footprints of closely spaced connectors. These standing wave modes can be sufficiently attenuated by exploiting the trace losses between the connectors. This implies that it is better to have some loss between impedance mismatches to dampen the reflections therefore entailing a subsequent improvement in the insertion loss deviation.

The two largest impedance mismatches on the backplane interconnect are commonly located in the daughter card and backplane footprints. As there is already very little loss within a connector, connector OEM vendors have attempted to mitigate the aggregate losses even further with the use of air and other very low-loss dielectrics. The resulting ultra low-loss cavity produces significantly high-Q resonance harmonics within the transmission line itself. Contemporary connector designs that integrate discrete ground structures exacerbate these resonance harmonics due to the concurrent excitation of additional resonance modes within the accompanying ground shields themselves. These produce significant issues for interconnects designed to run above 6 Gbps.

Figure 2(b) shows the insertion loss with optimized via structures, which are not typical to actual board layouts. For a more realistic example, interconnects with 2 backplane connectors are investigated with varying daughter card PTH effects. Figure 3 depicts a measurement of an orthogonal connector system where Connector A was the connector platform of choice. In such a system, the connector on each side of the midplane plugs into the same PTH. The midplane via was optimized for electrical transparency, i.e. impedance matched to nearly 100 ohms, and inherently contains no via stub as it serves as an intermediate pass-through structure. Each daughter card has 3 inches of FR4 etch; the corresponding PTH stub was varied from 20 mils to 80 mils to represent the best and worst-case break out layers. Figure 3 shows the insertion loss and insertion loss deviation for these 2 cases.

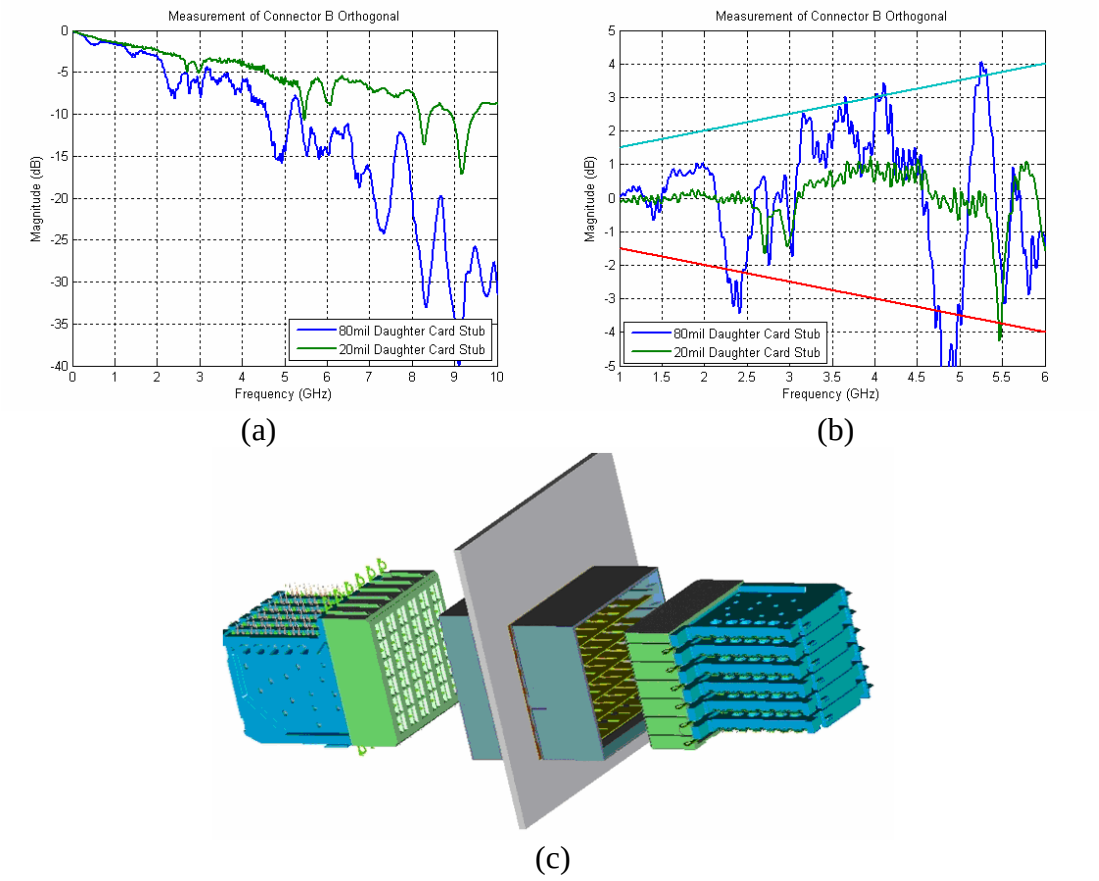

Figure 3: (a) Insertion Loss, (b) Insertion Loss Deviation, (c) Generic Orthogonal Picture

The measurements demonstrate the oft-neglected effects the daughter card via imposed on the overall system level performance. For instance, the 80mil stub case on the daughter card, even relatively transparent pass-through backplane vias, severely fails the IEEE 802.3ap mandated insertion loss deviation specification at spectral regions of critical importance including the 5 GHz fundamental frequency for 10 Gigabit Ethernet. The PTH parastics in the footprint reflect significant amounts of energy which resonates within the connector. Since the dissipation mechanisms have been designed out of the connector, a significant portion of the energy is converted to crosstalk. For the connector B, it is low-loss connector within an inherently resonant structure and its performance doesn't support the data rates above 6 Gbps.

Crosstalk Considerations

The previous section describes a backplane connector as a low-loss structure with an impedance discontinuity at each footprint. The reflections from these impedance discontinuities form a resonant cavity within the connector that has a very high "Q" as the structure has no dissipative mechanism. The reflected energy within the cavity produces insertion loss deviation as the reflections add in a constructive or destructive manner. This reflected energy reverberates within the connector thereby exacerbating

crosstalk at the resonant frequencies. The resonance harmonics are immediately identified as notches in the associated insertion loss measurements.

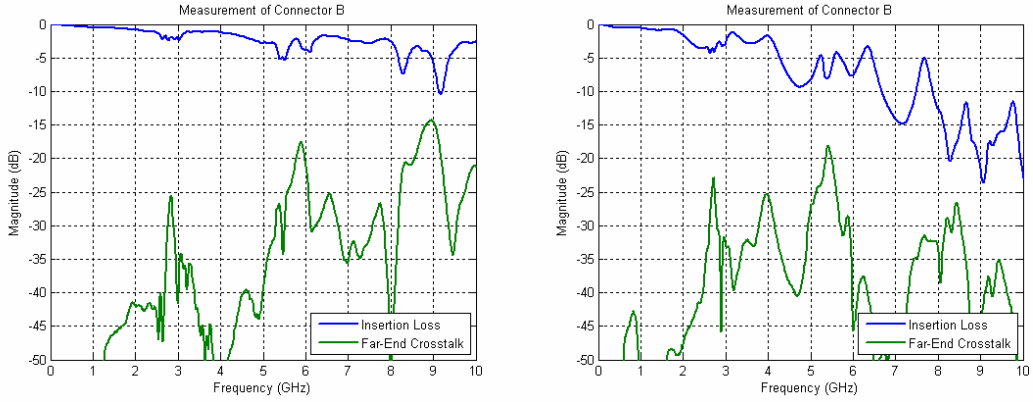

(a) (b)

Figure 4: (a) Optimized Daughter Card, (b) Daughter Card with PTH Stub

Using the example from the previous section, figure 4 illustrates a spectral overlay of the crosstalk of a single aggressor with the insertion loss of the victim pair as shown in figure 4. Figure 4(a) illustrates the insertion loss and far-end crosstalk (FEXT) for the interconnect with optimized PTHs incorporated into the daughtercard footprint, while figure 4(b) displays the same data for the daughter card with a non-ideal 80mil PTH stub. Both data sets assume an optimized orthogonal midplane pass-through PTH design. Each of the insertion loss resonances in figure 4(b) corresponds with a significant increase in crosstalk of approximately 15 dB. The measurements for the longer PTH stub demonstrate a significant increase in the overall crosstalk, especially between 3 and 6 GHz. This particular scenario also exemplifies that the connector is an inherently resonant structure with no mechanism to dissipate undesirable energy emanating from resonant sources. The internal reflections sustained by this structure can also exacerbate the aggregate crosstalk as the reflected modes are partially converted to crosstalk.

The above example examined the crosstalk effects of a single aggressor. The issue is highly accentuated when accounting the aggregate crosstalk due to multiple aggressor sources. The Insertion-loss-to-Crosstalk Ratio (ICR), defined in the IEEE 802.3ap specification, is used as a performance reference in figure 5. The ICR is analogous to the more commonly used signal-to-noise ratio. The insertion loss and total power summation of all FEXT crosstalk aggressors are shown in figure 5(a). When compared to the IEEE 802.3ap recommendations, the ICR fails by 10 dB at approximately 3 and 5 GHz. It should also be noted that the channel with the larger PTH stub exhibits additional degradation, further compounding the failures already seen on the insertion loss deviation.

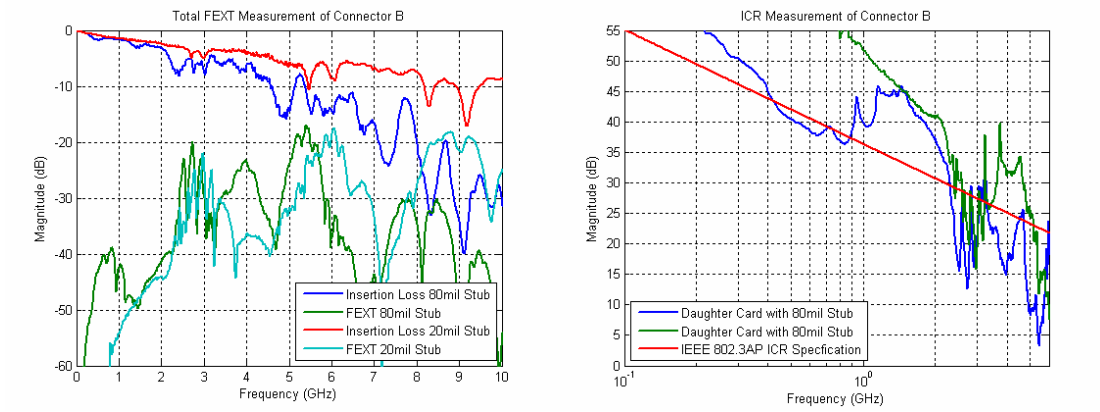

(a) (b)

Figure 5: (a) Total Crosstalk, (b) IEEE ICR Plot

Let us compare the aforementioned orthogonal topology with an XCede®-based traditional backplane topology as shown in figure 6. The XCede® 4-Pair daughtercards with identical routing and stackup topologies, and general PCB characteristics. Furthermore, all of the etch lines escape in the bottom-most layer thereby yielding 20 mil stubs. In this particular scenario, the backplane has 2 inches of etch. Judiciously chosen was Nelco 4000-13 SI substrate material in order to maximize the transparency of the backplane PTHs. Despite the utility of high-grade FR4 in the XCede® backplane, the impedances of these backplane vias are still not as optimal (i.e. close to 100 ohms nominal differential impedance) as it's orthogonal counterpart. Figure 7(a) depicts the insertion loss and crosstalk aggressors for the XCede® backplane link. The insertion loss exhibits no significant insertion loss deviation or resonances, even though the backplane is not as optimal as the shared via midplane in the previous example. Additionally, the corresponding crosstalk aggressors do not exhibit any additional resonance-based deterioration. It should also be noted that scale on the graph was changed to display the noise aggressors. The ICR plot illustrates about 10dB of broadband margin with 25 to 30 dB of additional noise margin over the previously discussed interconnect solutions. The mechanisms for crosstalk reduction have been covered in a previous publication. [2]

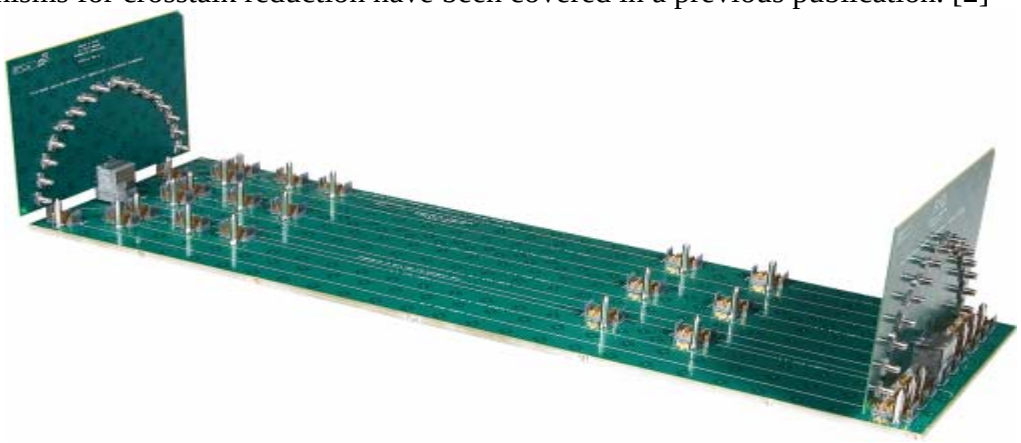

Figure 6: XCede 4-Pair Evaluation Backplane Configured utilizing a Traditional Backplane Topology

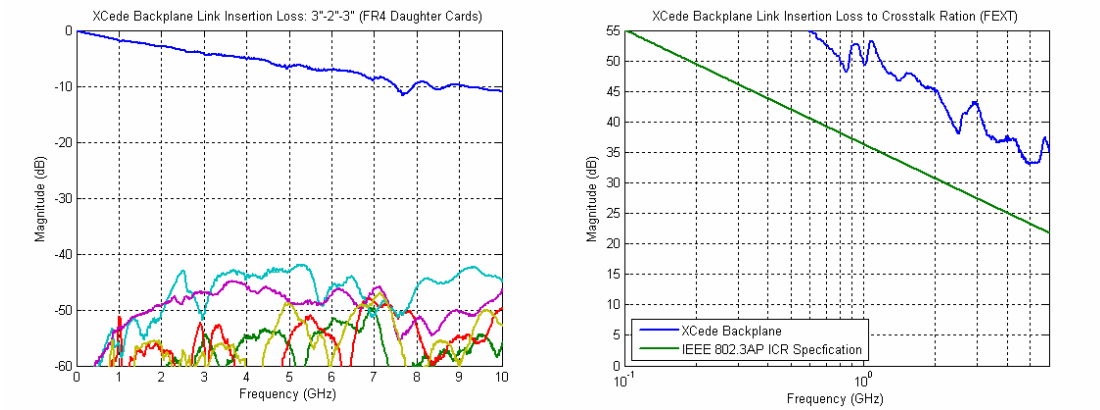

(a) (b)

Figure 7: (a) Crosstalk Aggressors, (b) IEEE ICR Plot

Mode Conversion Issues

Mode conversion is a general problem and its effect on the overall link performance is detailed in the following section. However, it is important to note that as data rates increase, it becomes increasingly difficult to control common mode transmission from the transceivers. Furthermore, skew throughout the system becomes a larger consideration of the overall performance budget. It is not uncommon for common mode to differential coupling to be of equal or greater magnitude than the pure differential coupling. Therefore, the mode conversion properties can be a significant limitation in performance. Figure 8 displays the mode conversion properties associated with connectors discussed in the previous sections. Once again, the cavity resonances, due to the physical structure of the connector, produce significant levels of crosstalk and possible EMI concerns at the same frequencies that entail the greatest amount of insertion loss deviation and differential crosstalk. Consequently, the aforementioned effects compound the degradation of the overall performance at the important frequencies of 3 and 6 GHz. The XCede® connector system with the controlled resonances shows a 20dB reduction in both common mode coupling and mode conversion.

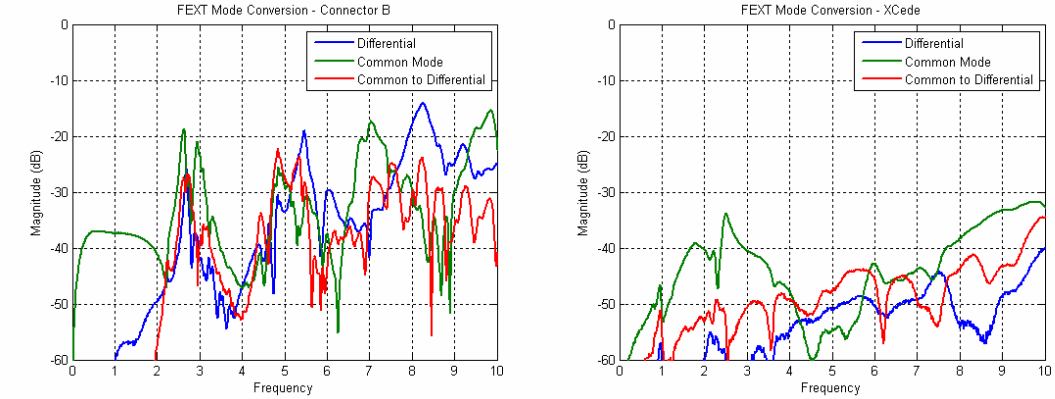

(a) (b)

Figure 8: (a) Connector B, (b) XCede®

Mode Conversion Effects at the Link Level

Conventional analysis techniques, including those cited in the IEEE 802.3ap standard, emphasize on the evaluation of signal-to-noise with respect to pure differential mode propagation throughout the entire channel. However, silicon devices operating at higher data rates are imperfect as the embedded drivers are inherently imbalanced due to process limitations. Consequently, the aggregate signal that SERDES drivers inevitably transmit is comprised of a linear combination of the intended differential signals and unwanted common mode voltage spikes. One such example is illustrated in figure 9 where an imbalanced 5 Gbps differential silicon driver yields the intended 5 Gbps differential signal in conjunction with a spurious 50 mV common mode voltage spike. These 50 mV common spikes are a manifestation of the inherent imbalance of the differential drivers.

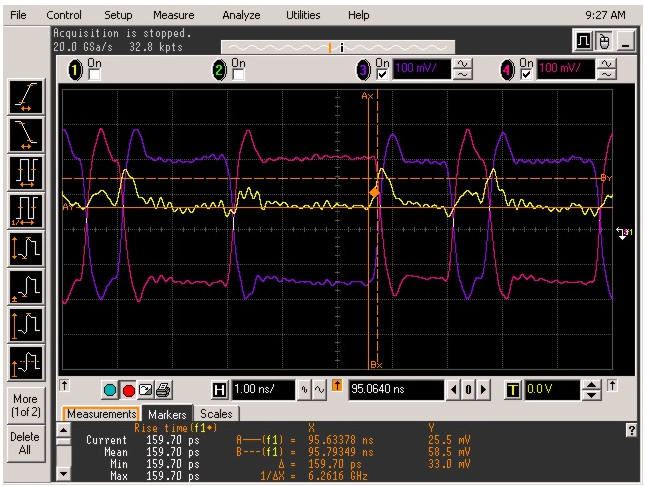

Figure 9: Common Mode Voltage Spikes originating from a 5 Gbps SERDES Driver (Yellow)

The yellow common mode voltage spikes shown in figure 9 illustrate that the width of each spike spans across a single bit time. One can reasonably extrapolate that common voltage spikes can potentially span at or greater than a single bit period at data rates beyond 6 Gbps. For instance, common mode spike widths could potentially span less than 100 pS that is equivalent to 10 Gbps signaling. As data rates continue to increase beyond 6 Gbps, the deleterious effects of such spectrally rich common mode energy propagation through an interconnect could become another significant issue that system designers must confront. As fully documented in known literature, common mode propagation can induce an undesired amount of electromagnetic radiation in leaky sources. Among such leaky electrical elements in the interconnect are poorly shielded and highly resonant connectors that posses a high degree of mode conversion. However, another seemingly inconspicuous facet of common mode propagation that system designers can ill-afford to neglect is the degrading effects of mode-converted crosstalk.

A qualitative explanation of the overall effect of mode-converted crosstalk is as follows. Figure 10 depicts all significant sources of differential crosstalk that is captured by a receiver channel, denoted as “Rx Victim Channel”. The crosstalk driver, labeled as “XTLK Source”, is comprised of differential and common mode components.

Differential crosstalk is superimposed onto the victim signal through two avenues. First, there is pure differential crosstalk, highlighted by the red path in figure 10, which is primarily due to pure differential coupling in PCB footprints and connectors. The second, albeit inconspicuous, method of inducing additional crosstalk onto the victim channel is through the introduction of skew into the system. As delineated by the green path in figure 10, skew in the interconnect is manifested through intrinsic imbalances in PCB etch, connectors, and the SERDES package itself. Note that aside from the Tx SERDES drivers being imbalanced, additional imbalance can be induced via additional preprocessing schemes such as the presence of Tx pre-emphasis feed-forward equalizers. The aforementioned elements induce the excitation of common mode. This falsely innocuous excitation, however, can mode-convert back into the differential mode in some of the electrical elements including poorly designed connectors. This double conversion process thereby produces an extraneous amount of additional crosstalk. Lastly, additional differential crosstalk is superficially induced through a direct mode conversion process from common mode to differential mode as highlighted by the blue path.

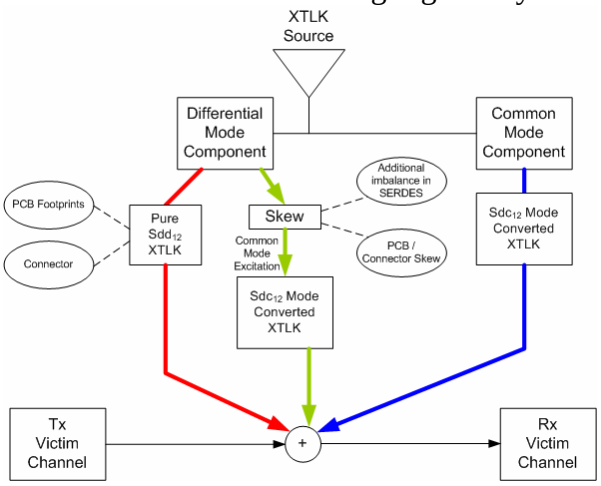

Figure 10: Crosstalk Sources

The following hypothetical scenarios exemplify the aforementioned phenomena in the frequency domain. Let us consider a hypothetical interconnect model, illustrated by the high-level block diagram in figure 11, where it consists of a concatenation of line and fabric daughtercard footprints, backplane footprints, connectors, and PCB etch. The resultant model is a mixed composite of RLGC etch models and s-parameter matrices associated with footprints and connectors. Not integrated into the aggregate model are SERDES transceiver models and accompanying BGA footprints due to an absence of accessible models. The corresponding connector system of choice is the Amphenol-TCS XCede® 4-Pair platform. Note that all resultant footprint s-parameters were derived from Ansoft HFSS V.10.1.1 while the XCede® 4-Pair connector model was generated with CST Microwave Studio. Lastly, the RLGC models were determined with the Ansoft Maxwell 2-D V.7.1 quasi-static field solver.

Outlined in table I are the accompanying modeling parameters for the daughtercard and backplane footprint models. The 300 mil thick backplane consists of traces breaking out in the bottom-most layers of the simulated stackup; the deep depth z-axis topology represents a worst-case scenario whereby crosstalk is maximized among adjacent vias. Illustrations of the aforementioned footprints are shown in figures 12(a)

and 12(b), respectively. Additionally, imposed in both models is a mid-grade type of FR4, Nelco N4000-13 in this particular case, to mitigate substrate-related losses. The corresponding cross-sectional etch dimensions are also listed in table I.

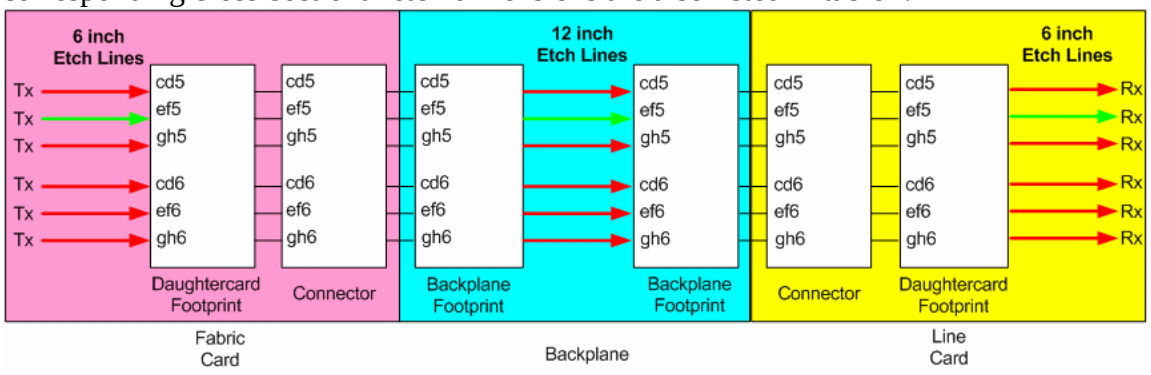

Figure 11: High-level schematic of a conventional backplane topology

	Daughtercard Footprint	Backplane Footprint
Thickness	120 mils	300 mils
Via Hole Diameter	22.5 mils	21.7 mils
Surface Pad Diameter	32.5 mils	31.7 mils
Signal Pad Diameter	22.5 mils	31.7 mils
Dielectric Constant: Dk	3.70	3.70
Dielectric Loss Tangent: Df	0.012	0.012
Etch Dimensions (Width/Space/Width)	4.25/8.00/4.25 mils	7.00/8.00/7.00 mils

Table I: Daughtercard and Backplane footprint model Parameters

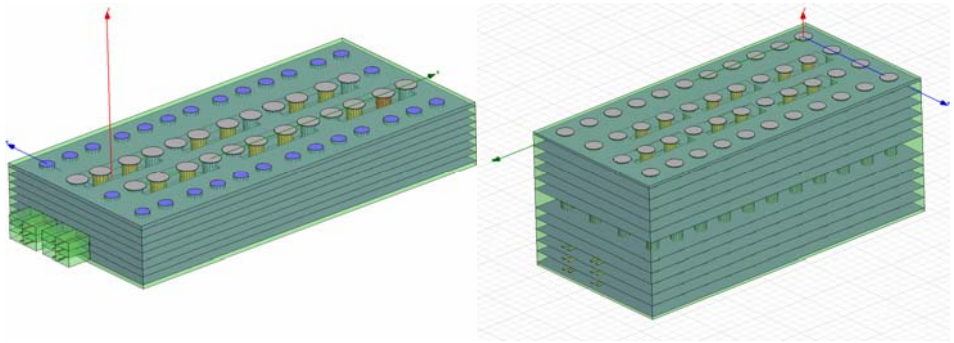

(a) (b)

Figure 12: (a) Daughtercard Footprint Example, (b) Backplane footprint example

In order to facilitate the illustration of the mode-converted crosstalk phenomena, let us consider the following modeling scenarios. Each scenario hypothetically examines the severity of the overall crosstalk performance with respect to increasing skew and the introduction of common mode signaling components. The aggregate skew is a compilation of additional skew in the silicon package and extraneous etch-related skew in the corresponding daughtercard. The first modeling scenario, designated as Case I, evaluates the crosstalk suppression performance of a backplane system where a typical connector platform without any resonance dampening technology is utilized in the

topology. Incorporated into the system model is an altered version of Amphenol-TCS XCede® 4-Pair connector. The intent of this alteration is to emulate the properties of other typical connectors where the resonance dampening mechanism is absent. Note, however, that the corresponding footprints utilize a generic 1.4mm square grid to emulate the generic topology of other competitive connector platforms. Likewise, Case II evaluates an identical backplane system topology where it considers the implementation of the actual Amphenol-TCS XCede® 4-Pair connector system with its unique footprint topology integrated into the system design. The fundamental objective of this particular case is to demonstrate the effectiveness of the resonance dampening mechanism pertaining to the suppression of mode-converted crosstalk and harmonic resonances.

Case I: Conventional Backplane Topology complemented with a Typical Connector Platform

Let us perturb the model by intentionally injecting an artificial amount of skewed etch into the fabric card slot of the model as illustrated in figure 13. The inclusion of skewed etch permits us to emulate the accumulated skew in the daughtercard footprint due to imperfectly balanced PCB etch and SERDES drivers. As shown in figure 10, an ideal differential input stimulus is imposed, i.e. an Sdd_{12} type s-parameter analysis is to be performed. Figures 15(a) and 15(b) show an overlay of the total insertion loss of the victim line and total crosstalk, respectively, versus skewed daughtercard etch whereby a power summation was utilized assuming zero statistical crosstalk correlation. The 0 pS case ideally delineates the scenario where the etch lines possess zero skew.

The swept perturbation shown in figures 15(a) and 15(b) illustrates the degradation of signal-to-noise ratio as crosstalk suppression performance deteriorates with respect to increasing skew in the fabric daughtercard. Another perspective of this general trend, as depicted in figure 16, illustrates the deterioration of the least squares fitted ICR performance with increasing skew. The degradation of the overall ICR performance with increasing skew further decreases the viability confidence of this potential 10 Gbps system design with respect to the informative mask defined by IEEE 802.3ap. In summary, the aforementioned observations indirectly confirm the initial hypothesis as schematically depicted by the double-conversion process in figure 14.

Further scrutiny of the frequency domain plots in figure 15 reveals that the introduction of the aggregate skew exacerbates the deleterious effects of the crosstalk resonance peaks. These resonance peaks are a manifestation of the formation of internal reflections within the connector structure. For instance, a 16 pS aggregate skew in the fabric slot entails the exacerbation of the secondary harmonic peak at 5 GHz by an additional 6 dB. This could have serious implications on the overall robustness of the system as this resonance peak coincides with the fundamental frequency of 10 Gigabit Ethernet signaling. This observation exemplifies the compounded two-fold penalty designers could confront. First, the absence of any resonance dampening technology leads to unwanted crosstalk resonances in the overall topology. Subsequently, the failure to mitigate mode-converted crosstalk via common-mode attenuation leads to the accentuation of the undesired crosstalk resonances.

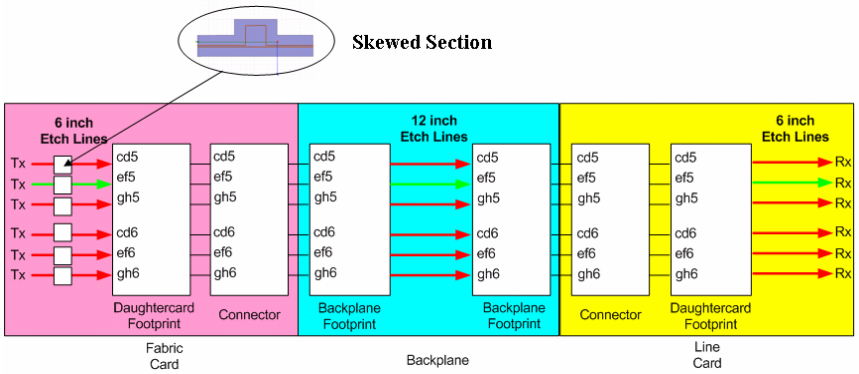

Figure 13: Conventional backplane topology with specified pin mapping and location of skewed elements

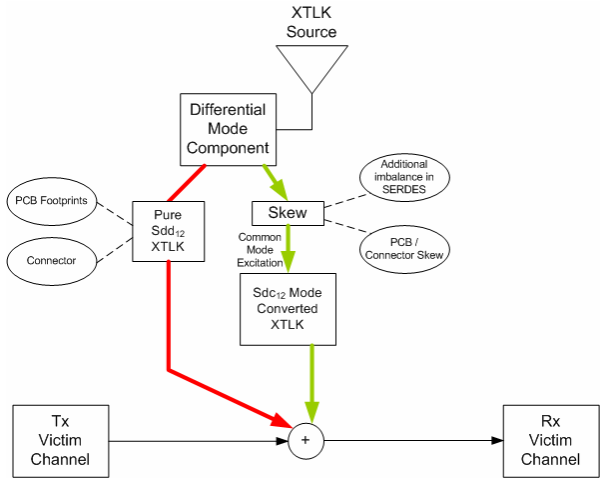

Figure 14: High-level schematic of figure 13 illustrating injection of pure differential stimuli into crosstalk sources

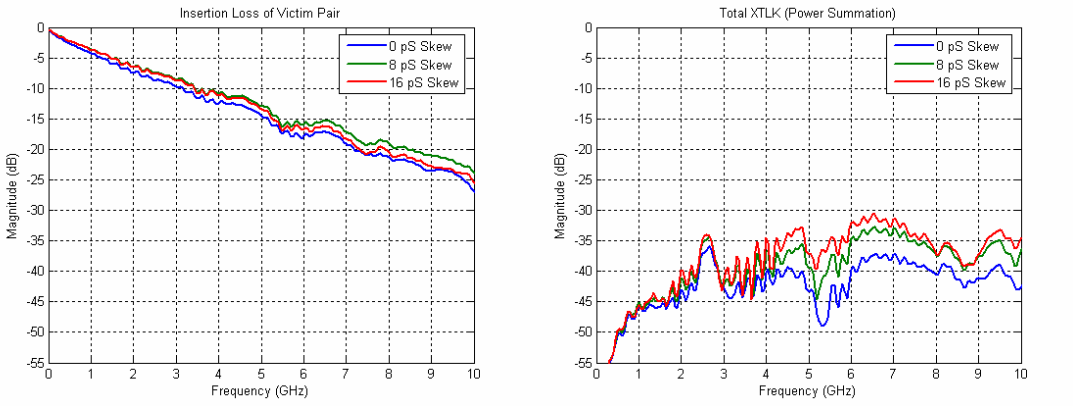

(a) (b)

Figure 15: (a) Case I insertion loss of ef5 victim pair, (b) Case I total crosstalk assuming pure differential stimuli injection into all crosstalk sources

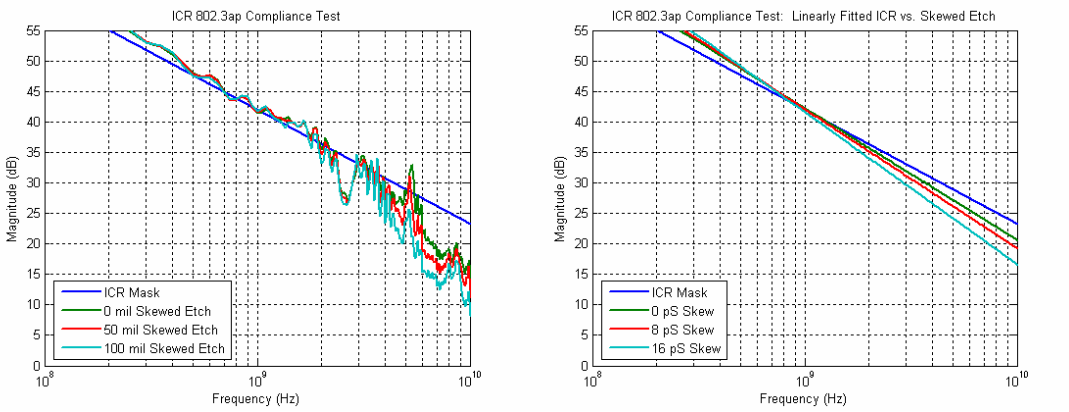

(a) (b)

Figure 16: Case I ICR versus skew assuming pure differential stimuli injection into all crosstalk sources; (a) Raw ICR, (b) Least squares fitted ICR

Now consider the scenario where a mixed composite stimulus, comprised of differential and common modes, is injected into each crosstalk source. The blue path in figure 10 illustrates that the effect of mode-converted crosstalk, Sdc_{12}, must be superimposed with the other crosstalk components. The direct power summation of the Sdc_{12} crosstalk magnitudes with the other Sdd_{12}-based crosstalk components is implausible as it implies that the common mode input amplitude is comparable to its differential counterpart. As exemplified in figure 9, this is certainly not the case as the common mode amplitude is a fractionally scaled version of the differential amplitude. Thus, we can compensate for this amplitude scaling by multiplying each Sdc_{12} mode-converted crosstalk magnitude by the ratio of the common amplitude to the differential amplitude. This simplified scaling technique approximates the spectral power content of the common mode spike that is scaled down with respect to broadband differential spectrum. The broadband assumption is well justified by the following qualitative argument. The common mode voltage spikes in figure 9 are temporally narrow thereby indicating rich spectral content in the frequency domain. We can approximate each narrow spike as a sinc[†] function. The Fourier Transform of a sinc function is a rectangular spectral pulse whose bandwidth is inversely proportional to the width of the sinc pulse. As previously mentioned, spike widths at 10 Gbps could be less than or equal to 100 pS. This corresponds to an approximate spectral width in excess of 10 GHz, which is essentially broadband with respect to the band-limited analyses performed here.

Figures 17(a)-(c) display the system level crosstalk plots with respect to increasing aggregate skew in the fabric cards. Additionally, each of the subplot figures delineates a parametric overlay of crosstalk spectra with increasing broadband common mode amplitude. As definite common mode amplitude values are defined, an 800 mV differential reference amplitude is assumed. Observe in the hypothetical 0 pS skew case, represented by the spectral overlay in figure 17(a), where it underscores the degrading effects of direct common mode injection into the system. A common mode injection amplitude of 250 mV (31.25% of the reference amplitude) entails an exacerbation of mode-converted crosstalk where harmonic peak at 5 GHz increases by 6 dB.

An interesting trend to note, though, is that the overall impact of common mode injection seems to be less accentuated with increasing aggregate skew in the fabric card.

For instance, unlike the 6 dB deterioration in the secondary resonance peak at 5 GHz for 0 pS skew, figure 17(c) demonstrates a marginal increase of about 2 dB. One must note the existence of elevated common mode amplitude as induced by the intrinsic aggregate skew in the 16 pS scenario. As such, the relative power contribution of a pure common mode stimulus with respect to the overall common mode power is less significant, thereby implying the relatively marginal increase in the mode converted crosstalk. Nonetheless, this does not discount the fact that the injection of additional common mode power contributes to additional degradation in the system level signal-to-noise ratio. In summary, the aforementioned example exemplifies the fact that system designers must be cognizant of the compounding effect of mode-converted crosstalk as highlighted by the blue and green paths in figure 10.

† Given an independent variable "x", sinc(x) = sin(x)/x

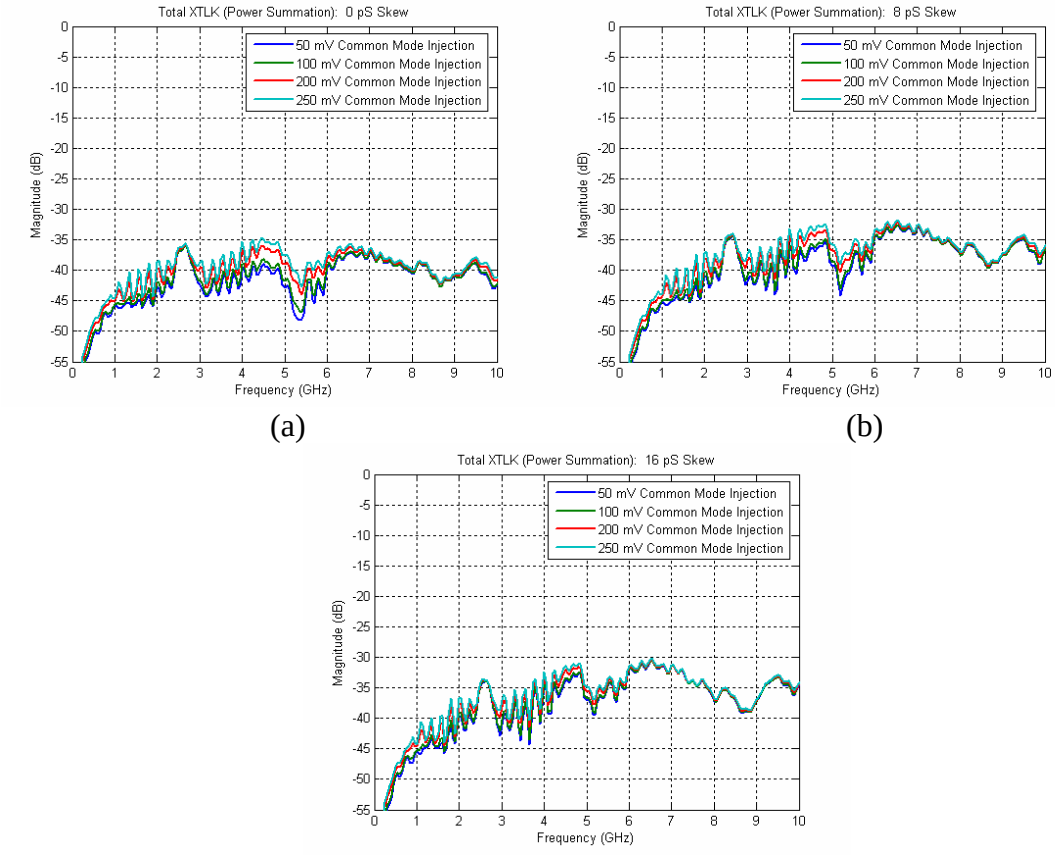

Figure 17: Mixed mode signal injection into the Case I system with respect to (a) 0 pS Skew, (b) 8 pS skew, and (c) 16 pS skew

Case II: Conventional Backplane Topology complemented with the Amphenol-TCS XCede® 4-Pair Connector Platform

Once more, let us perturb the model by intentionally injecting an artificial amount of skewed etch into the fabric card slot of the model as illustrated in figure 13. The inclusion of skewed etch permits us to emulate the accumulated skew in the daughtercard footprint due to imperfectly balanced PCB etch and SERDES drivers. As shown in figure 13, an ideal differential input stimulus is imposed, i.e. an Sdd_{12} type s-parameter analysis is to be performed. Figures 18(a) and 18(b) show an overlay of the total insertion loss of the victim line and total crosstalk, respectively, versus skewed daughtercard etch whereby a power summation was utilized assuming zero statistical crosstalk correlation.

The swept perturbation shown in figures 18(a) and (b) illustrates the degradation of signal-to-noise ratio as crosstalk suppression performance deteriorates with respect to increasing skew in the fabric daughtercard. However, unlike Case I, the deteriorating effect of increasing aggregate skew is not as accentuated with the XCede® 4-pair connector platform. A comparison of figures 15(b) and 18(b) depicts two notable observations. First, the crosstalk resonance harmonics are suppressed in this XCede®-based backplane system. Thus, the alleviation of unwanted crosstalk resonances assures designers sufficient noise margins without possibly compromising system performance at especially sensitive spectral ranges. Secondly, the broadband differential crosstalk suppression performance of the XCede®-based system is considerably improved over the typical connector scenario. The skew perturbation plots depict that, unlike the Case I system, this system is not as susceptible to a significant deterioration in the aggregate crosstalk with increasing skew. This example implies the effectiveness of the resonance dampening technology as it functions as a means of limiting the additive increase in mode-converted crosstalk. The signal-to-noise ratio perspective illustrated in figure 19 serves to illustrate the effectiveness of the optimized footprint topology in conjunction with the resonance dampening mechanism. The influence of mode-converted crosstalk on the overall ICR performance poses a minimal effect on the viability confidence of this potential 10 Gbps system design with respect to the informative mask defined by IEEE 802.3ap. Overall, these differentiating attributes are consequences of the footprint topology providing optimal suppression of differential crosstalk, and an integrated resonance dampening mechanism entailing both mode-converted crosstalk and resonance mitigation. In addition, **the electrical noise margins yielded by the XCede® based solution are also indicative of the fact that it is far more tolerant to undesirable increases in skew and undesired common mode propagation**.

It has been shown that the resonance dampening technology provides a sufficient means of mitigating mode converted crosstalk. This is due to its inherent common mode attenuation properties as qualitatively justified by the following. Suppose that no skew is intentionally introduced into the PCB, however an application of common mode power is supplied to the system. Observe in figure 20(a) that perturbing the system model with common mode signaling from 50 mV to 250 mV yields a maximum increase in crosstalk of 5 dB at certain regions of the spectrum. Note that this is not a substantial increase from an absolute magnitude perspective as the initial absolute magnitude is already low. For instance, the relative increase at 5 GHz increases from -48 to -44 dB. Though it may

appear to be a serious increase, note that the corresponding linear magnitude increase (relative to 800 mV differential signaling) is only from 3.2mV to 5 mV. Likewise, figures 20(b) and (c) both exemplify common mode suppression effects as manifested by the mitigated increases in mode-converted crosstalk. Conversely, the crosstalk increase for Case I at 5 GHz is from -42 to -37 dB as illustrated in figure 17(a). The corresponding linear magnitude increase is from 7.9 mV to 14.1 mV. In summary, the Case I linear magnitude analysis yields a 6.2 mV increase whereas Case II yields a more palatable 1.8 mV. This rudimentary analysis, therefore, heuristically illustrates the common mode suppression phenomenon.

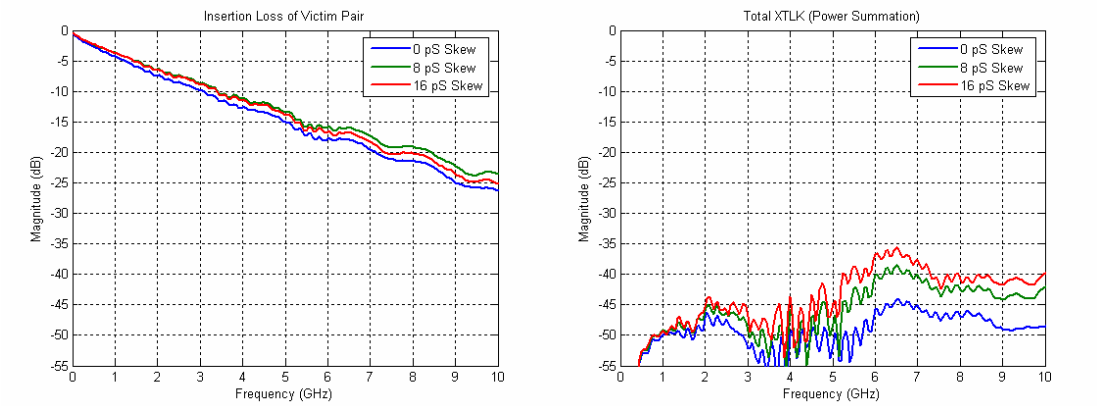

(a) (b)

Figure 18: (a) Case II insertion loss of ef5 victim pair, (b) Case II total crosstalk assuming pure differential stimuli injection into all crosstalk sources

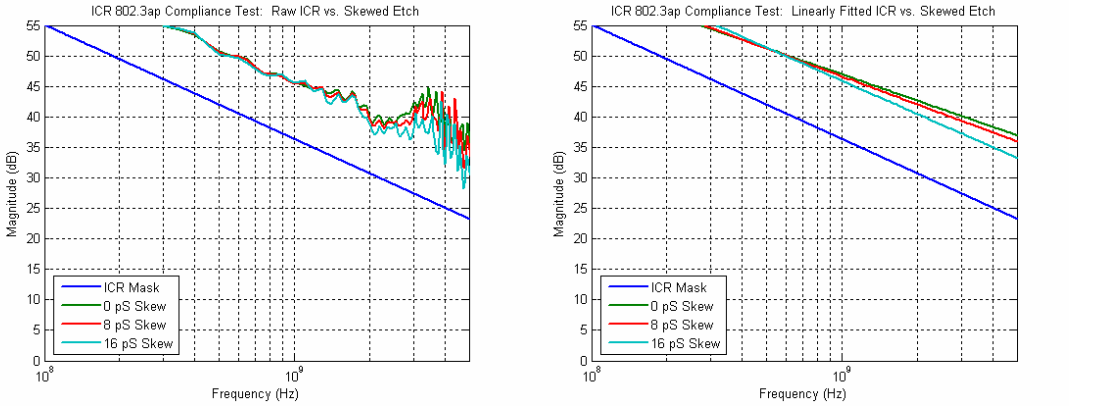

(a) (b)

Figure 19: Case II ICR versus skew assuming pure differential stimuli injection into all crosstalk sources; (a) Raw ICR, (b) Least squares fitted ICR

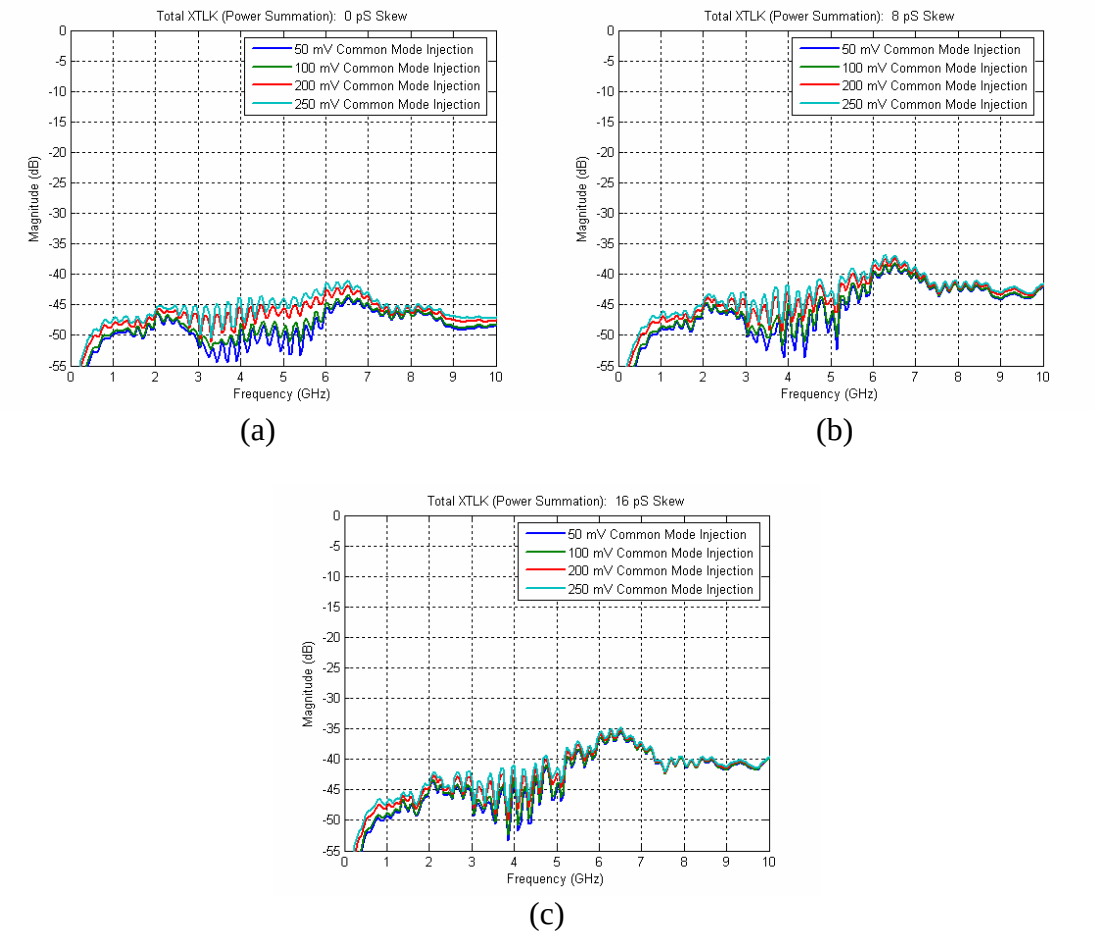

Figure 20: Mixed mode signal injection into the Case I system with respect to (a) 0 pS Skew, (b) 8 pS skew, and (c) 16 pS skew

Conclusion

This paper investigates the technical obstacles that arise for backplane interconnects as data rates exceed 6 Gigabits per seconds (Gbps). The physical size of backplane connectors can produce standing waves at the fundamental frequency of interest thereby exacerbating insertion loss deviation, crosstalk, and mode conversion problems. It is therefore critical to control the magnitude of these detrimental effects within the connector. Furthermore, the recent trend of designing the lowest loss connector with an inherently resonant structure can lead to severe system degradation. Finally, the effects of mode conversion on the entire link are explored in depth to understand the interaction between the silicon and passive components within the link. Improved control of the common mode energy from the silicon coupled with skew and mode conversion management in the passive components can yield 30 dB reductions in crosstalk at data rates above 6 Gbps.

References

[1] Phil Stokoe, Burke Hunsaker and Douglas Cannon, “Advances in SMT Backplane Connectors,” DesignCon2006, Santa Clara, CA, USA.

[2] Jason Chan, Marc Cartier, Tom Cohen and Brian Kirk, “Advanced Design Techniques to Support Next Generation Backplane Links Beyond 10Gbps,” DesignCon2007, Santa Clara, CA, USA.