DesignCon 2003

High Performance System Design Conference

Impact of Various Source Jitter Components on Equalized Link Performance

Ken Lazaris-Brunner
Gennum Corporation

John D'Ambrosia
Tyco Electronics Corporation

John Patrin
Craig Emmerich
Wavecrest Corporation

Abstract

The current state-of-the-art backplane design uses a 4-lane parallel interface running at a symbol rate of 3.125Gbps (XAUI). State-of-the-art backplanes today represent the legacy backplanes of tomorrow. There is considerable interest in increasing the bandwidth capacity of these backplanes, but it is not certain that they will be able to support the performance requirements necessary for 10Gbps. It is felt that these backplanes should be able to support 5 to 6.4Gbps, which would represent a 2x increase in the capacity of these backplanes, as well as extend their useable life.

This paper will focus on the ability of these backplanes to carry these data rates, and the full link characteristics required to achieve success. Various test backplanes will be used in conjunction with various available sources and an active analog receive end equalization technique to demonstrate the performance of the various links at 5Gbps. These backplanes will provide different test cases, where board material and thickness are varied, as well as trace conditions, such as trace length, width, and layer connection. Links will be composed of source, channel and receiver (active equalizer).

The XAUI standard refers to a total jitter, which includes random jitter, to achieve a BER of 10^{-12}, and it is likely a standard to upgrade the speed of these backplanes will have a similar requirement. A unique algorithm developed by Wavecrest is used to measure the random jitter accurately, without the length of time and judgments required to extrapolate a waterfall curve. Using this tool, the impact of the channel plus equalizer on total jitter (random plus deterministic) will be evaluated. Conclusions on the feasibility of upgrading current backplane designs will be reached.

Authors Biographies

Ken Lazaris-Brunner is a Product Definition Specialist with Gennum Corporation in their Datacom Products Division. Previous to joining Gennum, Ken spent 10 years as a communications systems engineer with various companies, most recently in the aerospace industry. Since joining Gennum he has been active in investigating new product ideas for Gennum. He holds a B.Sc(Honours, Physics) degree from Queen's University (Kingston, Canada) and a M.Eng from McMaster University (Hamilton, Ontario).

John D'Ambrosia is the Manager of Semiconductor Relations for Tyco Electronics in Harrisburg, PA. In this position he interacts with semiconductor vendors in the development of active interconnect solutions. In this capacity he helped organize the High Speed Backplane Initiative (HSBI), as well as lead the industry effort to develop interoperability testing for XAUI chipset solutions. Mr. D'Ambrosia received a B.S. in Electrical Engineering Technology from the Pennsylvania State University in 1989 and a Master's Degree in Engineering Management from the National Technology University in 1999.

John Patrin is currently the Director of Product Marketing at Wavecrest. He has more than eight years engineering and marketing experience in scientific instrumentation and semiconductor capital equipment. Prior to joining Wavecrest, John worked as a staff engineer and engineering manager for a semiconductor capital equipment company. He received a BS in physics from St. John's University in Collegeville, MN and a Ph.D. in Materials Science from the University of Minnesota in Minneapolis. He has published 15 papers in technical journals and has two patents and one pending.

Craig Emmerich is currently Product Marketing Engineer at Wavecrest. He has more than seven years engineering and marketing experience in high-speed design, applications, and marketing. Prior to joining Wavecrest, Craig worked as a design engineer at Honeywell. He received a BS in Electrical Engineering from the University of Minnesota in Minneapolis. He has written and presented several technical papers and given many seminars on jitter and signal integrity around the world.

Introduction

The frontier of high-speed backplane design is currently at a serial data rate of 3.125Gbps for XAUI backplanes. Given today's current economic climate, service providers are reluctant to undergo forklift upgrades to increase system capacity. An increase in capacity per port and port density per card has resulted in the need for backplane architectures that have capacity requirements where 640Gbps and greater are the norm. Thus, there is significant focus on enabling current backplane designs to go faster, enabling blade replacement rather than a complete new system.

By doubling the data rate to 5Gbps (for SONET) or 6.25Gbps (for Ethernet) the backplane designer will widen the range of problems encountered or alternately increase the impact of current problems. Transmit/receive equalization has been introduced to enable higher speeds of transmission, but is causing the system interconnect to evolve from a passive to a more sophisticated active nature. This evolution requires the system designer to understand the synergistic roles between the transmitter, receiver, and channel.

Looking at the problem in both time and frequency domain will provide a more in-depth understanding of the task at hand. The time domain signal is made up of various frequency components, so it is imperative to understand the ability of the channel to pass those frequency components. This becomes even more problematic as frequency increases, since the amount of signal getting through will be reduced while the amount of noise present will increase. The system designer is faced with the daunting task of dealing with these issues while optimizing overall system performance to meet the intended BER. Designing system interconnects that exceed BER of 10^{-15} or better is becoming the norm. Thus jitter analysis in the time domain provides the system designer with the analysis technique needed to meet this requirement.

As system speeds increase, however, the variation in channel-to-channel performance within the backplane environment becomes a significant issue. The use of transmit/receive equalization then interacts with the crosstalk performance of the system that must be addressed. Thus, the passive and active components within the system interconnect combine to impact the overall jitter performance of the link. And with increasing system speeds, shorter symbol durations further reduce the amount of allowable jitter.

In light of the above, it is critical that the backplane designer have some guidelines to evaluate whether the current backplane design will support higher data rates, or to help him design a backplane that can be easily upgraded. Although a brief note of the impact of channel ripple on jitter out of an equalizer was made in a previous paper by two of the authors,[1] it needs to be verified with experimental measurements to provide the designer with a greater level of confidence in using the formula.

Real channels provided by Tyco Electronics were used in conjunction with an active receive equalization device, the GN1404, supplied by Gennum to examine the impact of various passive characteristics on the performance of the channel/equalizer combination. S-parameter measurements of the channels under consideration are presented to provide a qualitative understanding of channel variance. Quantitative analysis using tools from Wavecrest to accurately predict channel performance out to a BER of 10^{-12} was performed on channels running at 5Gbps. Analysis of the data will permit conclusions to be drawn, which will provide the system designer both qualitative and quantitative understanding of how the passive channel and various sources of jitter can impact analog receive end equalization. The conclusions will be of aid to designers in evaluating legacy backplanes for higher data rates, and similarly in designing new backplanes to be compatible with higher data rates.

Thinking in Frequency Domain

Signal Throughput

The connector/board interface, specifically the plated through-holes into which the connectors are inserted, provide the connection path between the daughtercard PWB, the connector, and the backplane PWB. These through-holes also provide the transition path to the various layers on both PWB's. When the intended signal layer for routing is the top layer, the signal transitions from the signal pin of the connector to the signal layer through the minimum amount of the plated through hole. The remaining unused portion of the plated through hole provides excess capacitance. In the case of the bottom layer, the signal transitions from the signal pin of the connector to the intended signal layer for routing through the maximum amount of the plated through hole. This minimizes excess capacitance due to the plated through hole. Thus, the impedance discontinuity due to signal layer connection will increase as the intended layer for routing moves towards the top layer.

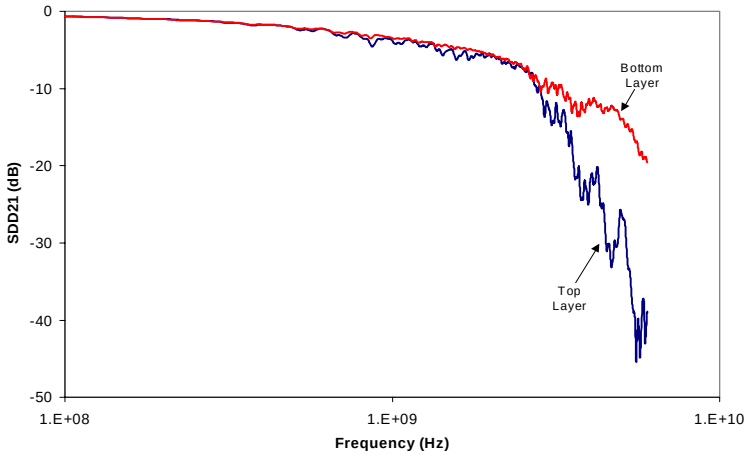

Figure 1 – Impact of Layer Connection on S_{DD21}

Figure 1 illustrates the impact of the layer connection in a backplane environment on the performance of the channel in the frequency domain. With a top layer connection, there is a significant roll-off in throughput in comparison to the bottom layer. The roll-off will translate to a reduced eye-opening and increased jitter. Therefore, as the thickness of the backplane increases, greater variation in performance between the signal layer extremes should be expected. Figure 1 also demonstrates the low-pass filter nature of a channel. Attenuation of the signal reduces the strength of the signal that will be seen at the receiver, and as high-frequency content is limited, jitter on the edges of the signal is induced. Furthermore, the channel permits greater low frequency content through than higher frequency content. This is a significant contributing factor to ISI. Therefore, as the frequency of operation increases, a top layer connection channel will exhibit an eye that is reduced in both amplitude and time as compared to a bottom layer connection.

Figure 1 demonstrates that in a backplane system the top layer connection will have an unpredictable S_{DD21} profile, while the bottom layer connection is more predictable. Thus, the channel consists of two types of losses – predictable and unpredictable. The predictable losses are those losses associated with the transmission media itself. These are mostly skin effect losses that increase with the square root of the frequency, and dielectric losses that increase linearly with frequency. On the other hand, unpredictable losses are caused by impedance discontinuities within the channel, which typically result from transitions between transmission media, such as in the connector/PWB interface. Furthermore these unpredictable losses cannot be dealt with by merely dealing with the predictable loss aspect of the system interconnect. Instead the impedance discontinuity itself must be dealt with in order to minimize the unpredictable nature of the channel.

Channel Crosstalk

With the introduction of various transmit and receive equalization approaches the life of copper has been extended. During the development of a backplane the engineer is challenged to extend the reach and frequency of operation of all the channels within the backplane environment. This has led to a fascination within the industry of a given device's ability to run greater distances or higher speeds over a given channel. This focus on the "given" channel has led to intense analysis of the loss characteristics of a channel and identifying all of the various components that influence its behavior. Engineers have characterized the channel using S-Parameter data, specifically S_{DD21}. This has resulted in a basic understanding of channel performance and its throughput performance in frequency domain.

Transmit and receive equalization forces the engineer to be more concerned about the inherent crosstalk behavior over frequency between a given victim and all of the neighboring aggressors. The reasons for this are specific to the form of equalization employed, be it either transmit or receive equalization, or a combination of the two. Any waveform is a function of the transmit characteristics and its channel. In applying this to crosstalk, the amount of noise generated will depend on the frequency content of each aggressor signal and the channel or coupling between it and the victim. In transmit equalization the frequency content of the launch signal is altered to compensate for the behavior of the channel. This results in a smaller eye at the receiver, which effectively increases the receiver sensitivity to noise. Forms of receive equalization that are active employ gain functions, which will amplify both signal and noise. This amplification is frequency dependent. Thus, the input to a receive equalizer that employs gain will be sensitive to the frequency throughput of the forward channel, as well as all of the crosstalk channels. This problem is further complicated with the use of both transmit and active receive equalization.

Microwave engineers have understood this concept for years. It is simply the signal to noise ratio. In a system environment this concept is fairly complex given all of the system variables that need to be considered. Signals of varying length and direction can be located next to each other. Thus, both the signal and noise characteristics that must be considered are extremely dynamic. This is demonstrated in Figure 2. The impact of system throughput and crosstalk ultimately impact the observed jitter performance on a given link. This jitter may impact the effectiveness of equalization.

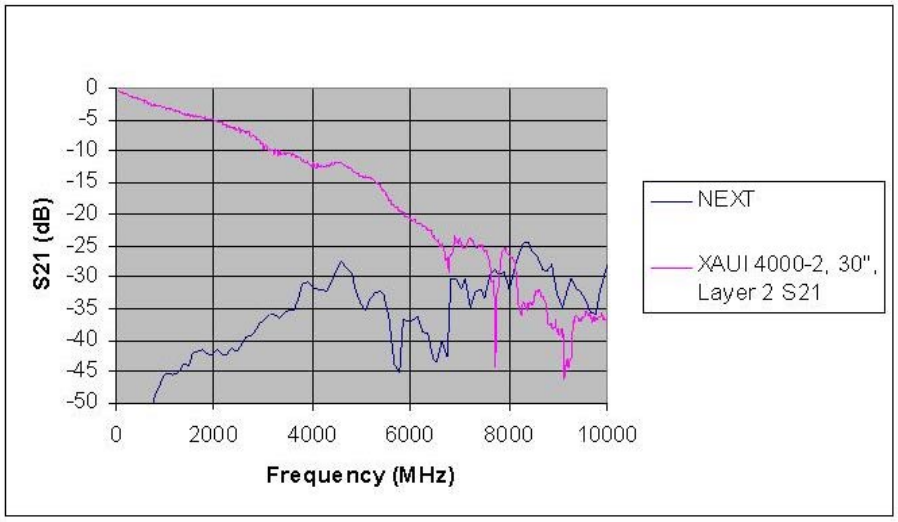

Figure 2 – Dynamic Signal to Noise Characteristics

Thinking in Time Domain

Understanding Jitter

Traditionally, measuring jitter has been critical to determining the performance of high-speed digital communications systems. Recently, as internal and external data rates of computers and networks have increased to unprecedented levels, reducing jitter has become an even higher priority for ensuring high reliability in high-speed databuses and integrated circuits. Jitter is the deviation of a timing event of a signal from its ideal position, as shown in Figure 3. Jitter affects a system as a whole, and can be introduced by every circuit element used to generate, convey and receive signals. As a result, understanding the amount of jitter introduced by each element of a system is imperative for predicting overall system performance.

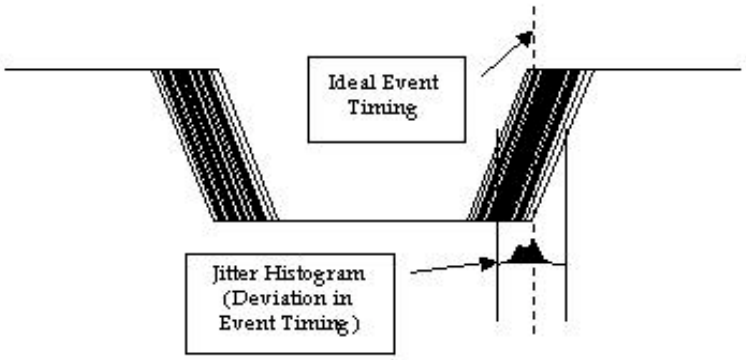

Figure 3 - Timing jitter

Total jitter (TJ) is the convolution of all independent jitter component Probability Density Functions (PDF). A PDF describes the likelihood of a given measurement relative to all other possible measurements, and is typically represented by a normalized histogram. TJ includes contributions from all deterministic and random components, and is a pk-pk value specified for a given sample size or Bit Error Rate (BER). It is common to plot TJ as a function of BER and this plot is called a bathtub curve as shown in Figure 4. For many high speed serial standards TJ is specified at a BER of $10^{-12.}$

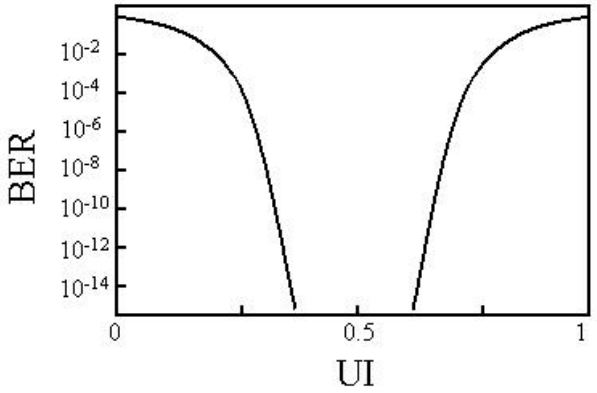

Figure 4 - TJ as a function of BER

Figure 5 shows how total jitter (TJ) includes deterministic jitter (DJ) and random jitter (RJ). DJ can be further separated into periodic jitter (PJ), data dependent jitter (DDJ), duty cycle distortion (DCD), and intersymbol interference (ISI). Random jitter (RJ), one of the main components of TJ, is characterized by a Gaussian distribution and assumed to be unbounded. This section will only review deterministic jitter, the main jitter source observed in backplanes. A more in comprehensive jitter review can be found elsewhere.[2]

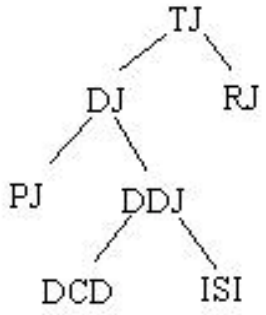

Figure 5 - The Components of Jitter

Deterministic Jitter

Deterministic jitter (DJ) has a non-Gaussian PDF and is characterized by its bounded pk-pk amplitude as shown in Figure 6. There are several types of DJ, including periodic jitter (PJ), duty cycle distortion (DCD) and intersymbol interference (ISI). DCD and ISI are types of data dependent jitter (DDJ). (Other types of DDJ are still being investigated.) PJ, also referred to as sinusoidal jitter, has a signature that repeats at a fixed frequency. For example, PJ could be the result of unwanted modulation, such as electromagnetic interference (EMI). PJ is quantified as a pk-pk number, specified with a frequency and magnitude. DCD is the result of any difference in the mean time allocated for the logic states in an alternating bit sequence (e.g., 0, 1, 0, 1). Different rise and fall times and threshold variations of a device could cause DCD. DCD and ISI are functions of the data history that occur when the transition density changes. For example, Fibre Channel systems and devices are commonly tested with a Compliant Jitter Tolerance Pattern (CJTPAT) that stresses DCD and ISI by alternating long strings of zeros or ones with short strings of zeros or ones within the pattern. It is the DCD and ISI caused by the time difference that is required for the signal to arrive at the receiver threshold when starting from different places within the bit sequence (symbol). ISI occurs when the transmission medium propagates the frequency components of data (symbols) at different rates. One example of DCD and ISI is when jitter changes as a function of edge density.

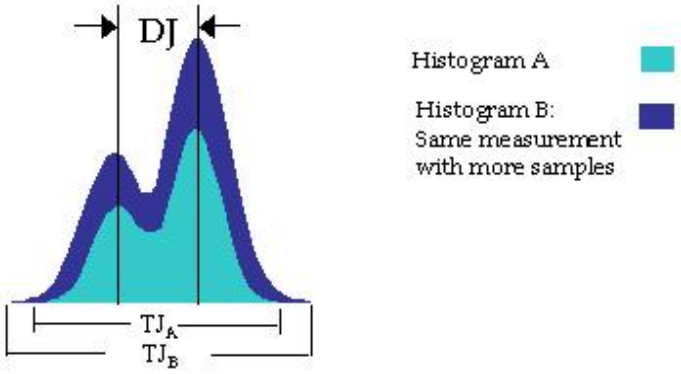

Figure 6 - Deterministic Jitter Histogram

Measuring DJ

Quantifying jitter components from measured data is the foundation of true signal integrity analysis. It involves statistics, DSP, algorithms and basic assumptions about the data histograms. In the time domain, jitter data are typically collected from one particular edge to another edge. For example, a period measurement is taken between a rising edge and the next rising edge. The histogram of period measurements contains a mixture of DJ and RJ processes. Traditionally, the TJ histograms included DJ and RJ components, and were quantified by a pk-pk value and a 1σ. However, given the Gaussian nature of the random component, it is incorrect to quantify a jitter histogram with a pk-pk number without specifying the number of samples. Therefore, for a given jitter histogram containing RJ, the pk-pk value will increase with more samples. Furthermore, in the presence of DJ, the 1σ of the total distribution does not depict the Gaussian component RJ. The TJ histogram represents the TJ PDF. However, if the DJ and RJ processes are independent, then the total PDF is the

convolution of the RJ PDF and DJ PDF. If DJ was absent from the jitter histogram, then the distribution would be Gaussian. Adding DJ to the histogram effectively broadens the distribution while maintaining Gaussian tails. Adding DJ to the distribution effectively separates the mean of the right and left Gaussian distribution. The difference between the two means, μ_L and μ_R, is the DJ (see Figure 7). The tail portions of the histogram represent the RJ component of the TJ histogram. Other methods can also be used to determine RJ and DJ.[3]

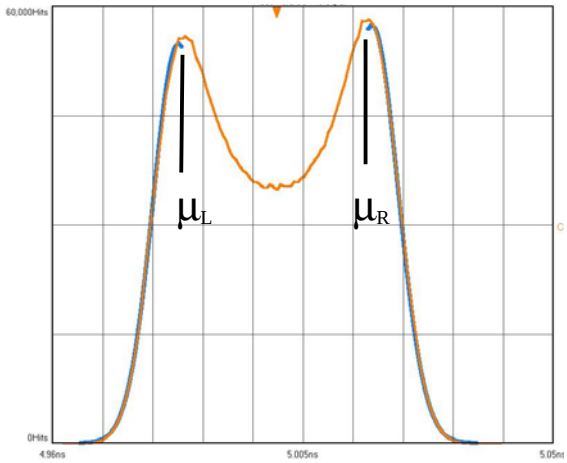

Figure 7 - Bimodal Distribution Containing RJ and DJ

Sources of Deterministic Jitter

Understanding the underlying cause of jitter is crucial to signal integrity analysis. Determining the source of jitter allows you to characterize and eliminate the potential problem. Here, we examine the most frequent causes of DJ and RJ. Some common sources of DJ include EMI, crosstalk and reflections. EMI is the result of unwanted radiated or conducted emissions from a local device or system. Switching-type power supplies are common sources of EMI. These devices can radiate strong, high-frequency electric and magnetic fields, and they can conduct a large amount of electrical noise into a system if they lack adequate shielding and output filtering. EMI can couple or induce noise currents in a signal conductor and corrupt the signal by altering its bias. Because the interfering signal is deterministic, the resulting jitter is also deterministic. EMI may also corrupt a ground reference plane or a supply voltage plane by introducing transient noise currents. Noise currents can sporadically alter the effective input thresholds of signal receivers. Given that logic signals require a finite time to change states, a sporadic change in receiver threshold results in signal jitter.

Crosstalk occurs when the magnetic or electric fields of a signal on a conductor are inadvertently coupled to an adjacent signal-carrying conductor. The coupled signal components algebraically add to the desired signal, and can slightly alter its bias depending on the amount of coupling and the frequency content of the interfering signal. The altered bias translates into jitter as the signal transitions the receiver's threshold.

Reflections in a data signal channel create DJ due to the signal interfering with itself. Signal reflections occur when impedance mismatches are present in the channel. With copper technology, optimum signal power transfer occurs when the transmitter and receiver have the same characteristic impedance as the medium. If an impedance mismatch is present at the receiver, a portion of the energy is reflected back through the medium to the transmitter. Reflections typically come from uncontrolled stubbing and incorrect terminations. Reflected energy, or energy not available to the receiver, reduces the signal-to-noise ratio at the receiver and increases jitter. If the transmitter is also mismatched, the transmitter absorbs a portion of the reflected signal energy while the remainder is reflected toward the receiver (again). Eventually, the delayed signal energy arrives at the receiver, out of phase with the original signal. The portion that is absorbed is algebraically summed with first-time arriving signal energy, resulting in DJ (specifically, ISI) from the receiver's perspective.

Receive End Adaptive Analog Equalization

An overview of an analog equalization method was previously given in a paper presented by two of the authors last year.[1] Basically, the equalization algorithm works by apply the inverse transfer function of the frequency-dependent attenuation losses over copper medium. This is expressed as,

$$C(f,l) = e^{-k_s l(1+j)\sqrt{f} - k_d l f}$$

where k_s is the skin loss, and k_d is the dielectric loss. Key to using this function is to note that it is exponentially linear with length. A simple exponential approximation can then be used to obtain an equation that is linear with length.

$$C(f,l) \approx 1 - l[k_s(1+j)\sqrt{f} + k_d f]$$

Hence a single response shape is used, but the gain is varied to match the length and attenuation combination of the copper trace. Generally, the relationship between skin and dielectric losses is fairly consistent for traces with different physical characteristics. Although there may be minor differences in the attenuation shape, in general it is close enough that it works for a wide variety of media, trace thickness, and trace width. Good performance can be achieved for any length up to a maximum attenuation. This attenuation is usually specified at half the baud rate since the majority of the energy lies below this frequency.

Channel Degradations

Unfortunately, in the backplane, the channel rarely follows the attenuation characteristics exactly. Capacitances and inductances due to channel discontinuities and multiple reflections between discontinuities add distortions to the channel. Moreover, these distortions will change in a random fashion from channel to channel as described in the channel section. It is therefore impossible to tune the channel response for a particular channel since it can vary significantly even for two adjacent channels of the same length.

Gennum can, however, predict the degradation in the signal based upon the signal characteristics. The figure of merit is the channel ripple, which is defined as the maximum variation from the ideal attenuated channel at half the data rate. Of course, some interpretation must be made since it is rare that a minimum will occur precisely at half the baud rate. Figure 8 shows a comparison of simulation and measured results.

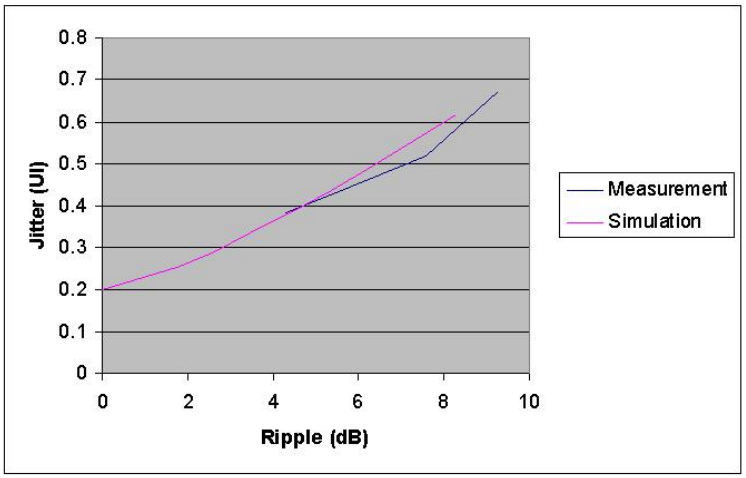

Figure 8 - Impact of Various Source Jitter Components on Equalized Link Performance

Crosstalk

Crosstalk generally occurs in the connector, although poor daughtercard or backplane design can cause crosstalk between traces. Both near end crosstalk (NEXT) and far end crosstalk (FEXT) contribute to the signal, although generally FEXT is considered to have less of an impact since it suffers the same channel attenuation as the desired signal. However it must still be taken into account.

Figure 9 to Figure 11 shows a near-end crosstalk response in both time and frequency domain.

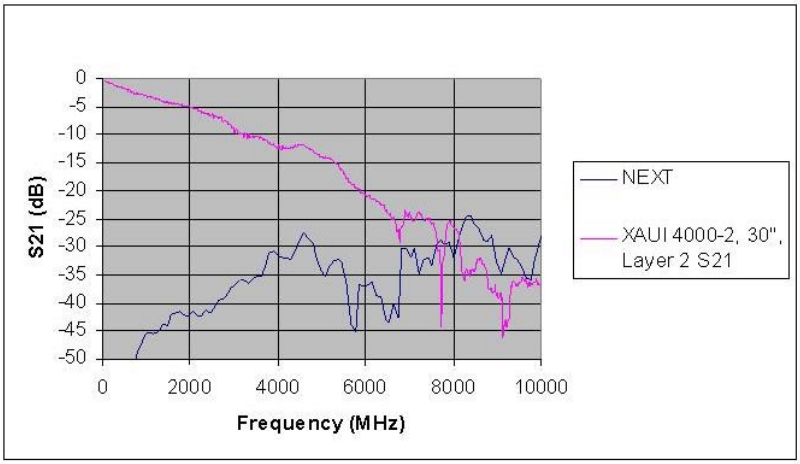

Figure 9 - Crosstalk in S-domain

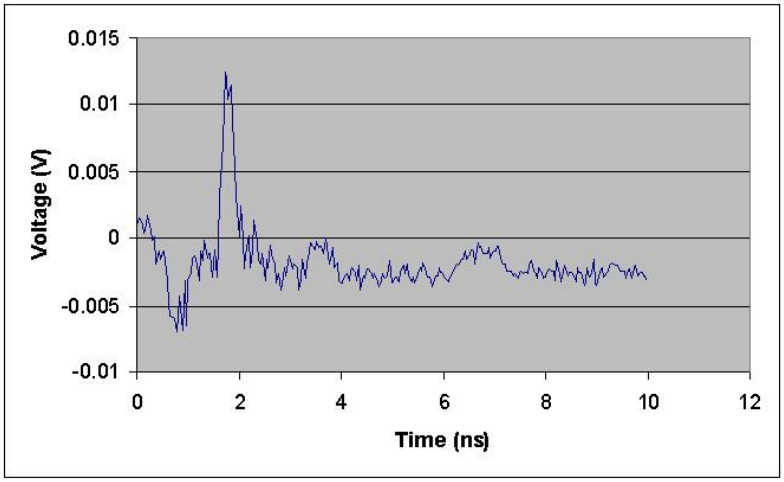

Figure 10 – Step Response of Crosstalk in Time Domain

Based only on the step response of the crosstalk , one would expect that the peak-to-peak contribution of the crosstalk would be on the order of 20mV, the peak-to-peak value shown in Figure 10. However, the durations of the peaks are longer than a symbol period at 5Gbps, and hence the maximum crosstalk effect will be higher. This is illustrated in Figure 11, which shows multiple overlapping symbols. The peak-to-peak ripple is around 120mV, or six times higher than the step response.

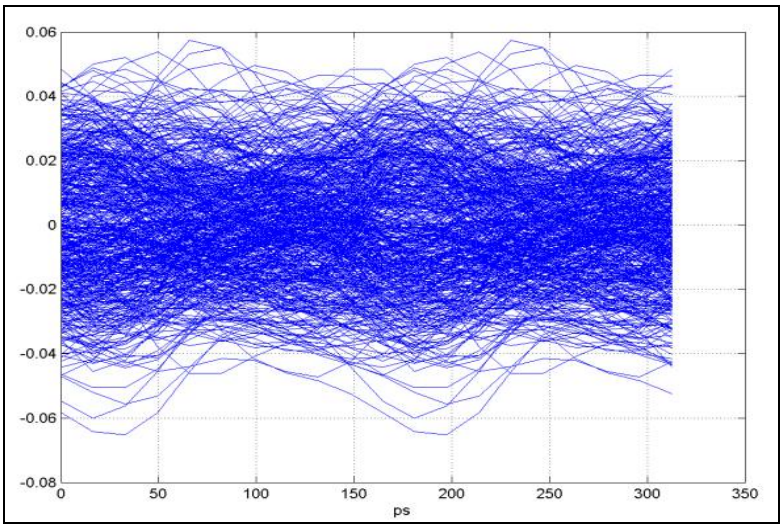

Figure 11 - Impact of Multiple Symbols on Crosstalk Noise Levels at 5Gbps

It is clear from a comparison of Figure 9 and Figure 10 that it is not enough to simply measure a step response for the crosstalk and assume that the peak-to-peak levels apply. Additionally, since all channels are linear, multiple crosstalk sources add linearly. Similarly, the crosstalk in an adjacent channel is proportional to the peak-to-peak transitions in that channel (with some additional effect from the transition times).

Impact of Crosstalk on Jitter

The worst-case effect of crosstalk on jitter can be simply calculated. Since the system is linear, the crosstalk noise appears as an envelope around the received signal without crosstalk. This is illustrated in Figure 12.

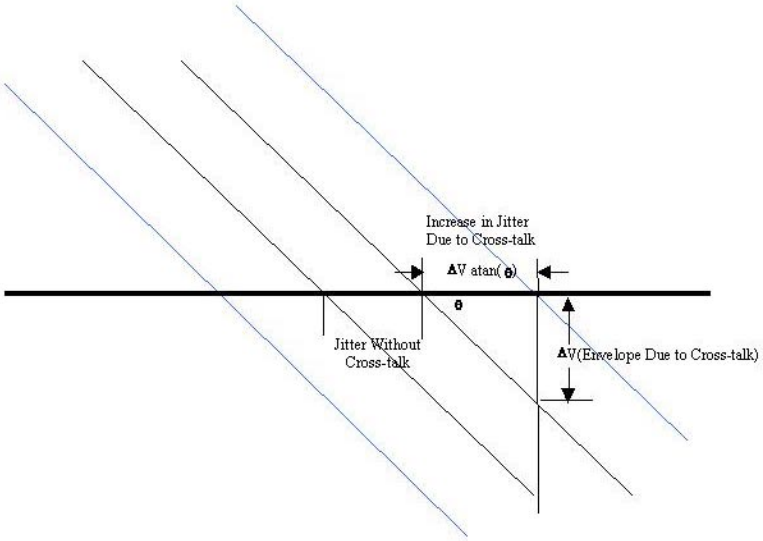

Figure 12 - Jitter Contribution Due to Crosstalk

The additional jitter contributed by crosstalk is dependent on the vertical contribution due to the crosstalk, and the slope of the signal at the threshold (in mV/ps). Note that when considering the impact of crosstalk on a signal going through a receive end equalizer; the crosstalk *after* the equalizer must be used.

Interaction Between the Channel, Equalizer, and Jitter

The focus of this paper is to develop an understanding of the interaction between the passive and active components by presenting various channels to the Gennum equalizer and analyzing the jitter content at its output. The basic set-up is shown in Figure 13. Three different test platforms were used that consisted of various materials, trace conditions, layer connections, and board thicknesses. See the following section on Test Platforms for a further description of the boards.

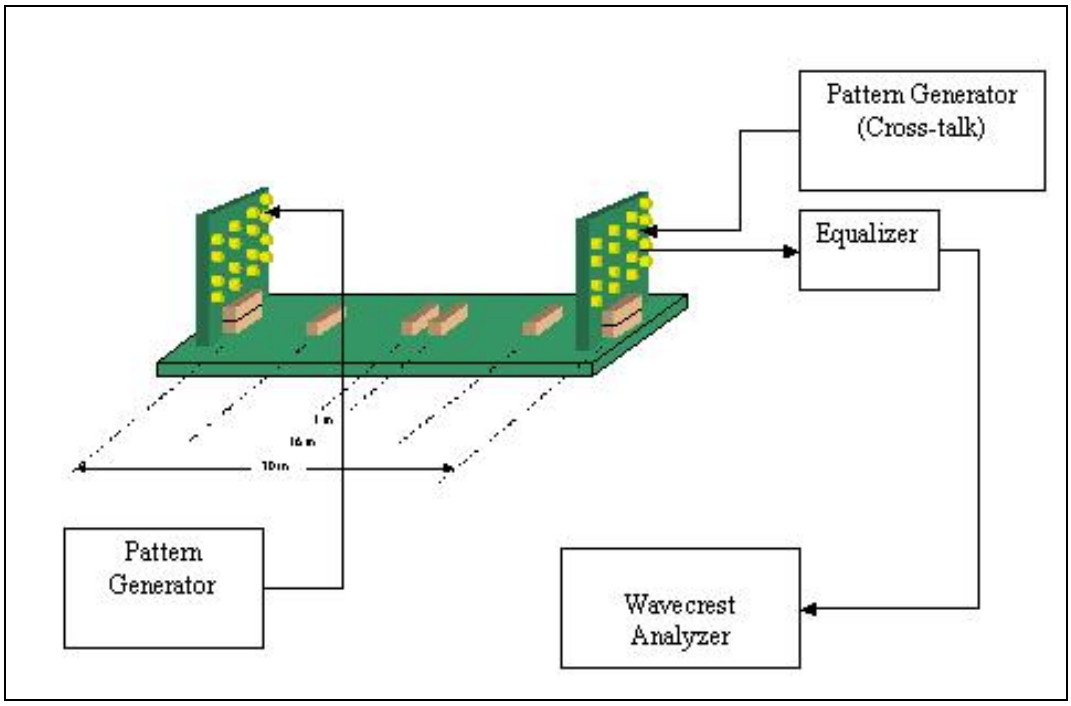

Figure 13 – Measurement Setup

Test Platforms

For the purpose of this paper two basic platforms were chosen. Material variations of each platform were also used in order to include the effect of material changes.

Platform #1 / #2 - The XAUI HM-Zd Interoperability Backplane

The platform is conceptually shown in Figure 14. It consists of 2 line cards that provide SMA access and the Z-PACK HM-Zd based backplane. Each line card is 0.093" (nominal) thick, consists of 14 layers, and Platform #1 is fabricated using Nelco 4000-2 material. There are four signal layers distributed throughout the entire stackup where the 100Ω differential geometries are based on 0.006" (nominal) wide traces. The trace length from the SMA to the Z-PACK HM-Zd connector is two inches. The backplane is 0.200" (nominal) thick, consists of 14 layers, and is also fabricated using 4000-2 material. There are four signal layers distributed throughout the entire stackup where the 100Ω differential geometries are based on 0.010" (nominal) wide traces. On the backplane there are three sets of trace lengths – 1", 16", and 30". Thus, for the platform there are overall system lengths of 5", 20", and 34". To examine the impact of better materials for the backplane on overall performance, Tyco Electronics fabricated the backplane only in Nelco 4000-6 (Platform #2).

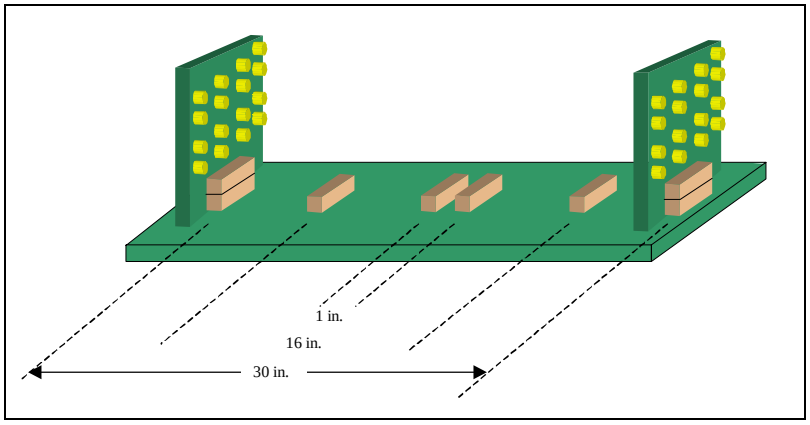

Figure 14 – XAUI HM-Zd Interoperability Platform

Platform #3 - The HM-Zd QuadRoute Backplane

Tyco Electronics developed the QuadRoute design technique based on its Z-PACK HM-Zd connector. By utilizing the unique footprint of this connector, which has 2.5 mm spacing between signal columns, two differential pairs, based on 4.5 mil wide traces, can be routed between the signal columns with enough isolation to minimize crosstalk between the two pair. The use of the QuadRoute technique reduces signal layer count in half, which can economically justify the use of higher performance materials to reduce the predictable losses associated with dielectric losses. Furthermore, this routing density can help minimize overall board thickness, which helps minimize connector footprint noise and the associated unpredictable losses caused by the connector / PWB interface. This results in predictable behavior for all signal layers throughout the entire backplane.

The HM-Zd QuadRoute backplane totals 20 layers, including 8 signal layers, and is 0.125" thick. The routing density that is enabled by the QuadRoute technique gives this backplane the same capacity as a backplane that has 16 signal layers, which would be 0.250" thick. As discussed in detail, the associates stub capacitance of such a backplane thickness will drive performance variability across the stackup of the backplane. The minimal stub capacitance of the 0.125" thick backplane results in minimal performance variability.

Platform Channel Performance

Figure 15 illustrates the performance of the layer extremes of all three platforms for the 30 inch link. For Platform #3, the HM-Zd QuadRoute Backplane, there is minimal difference between the two layer extremes, as compared to the other two platforms. Reduced channel ripple is also evident for the QuadRoute backplane. Furthermore, the QuadRoute backplane doesn't go below –30 dB for the top layer connection all the way up to 6 GHz, which neither of the two XAUI backplanes achieved.

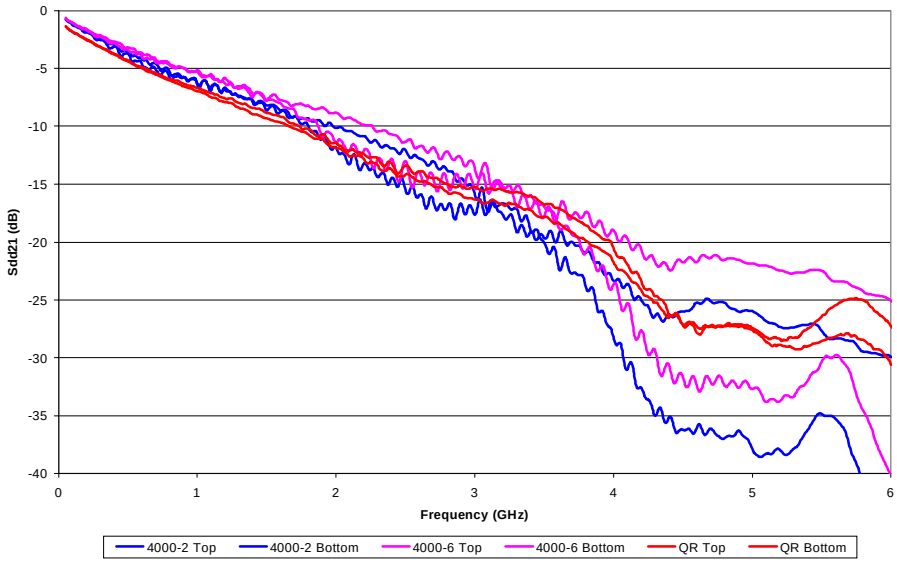

Figure 15 – S_{DD21} Data for Layer Extremes For All Platforms

Test Procedure

Three different patterns were chosen, PRBS 2^7-1, 2^{10}-1, and CJTPAT to cover a wide range of traffic characteristics. The basic test setup is shown in Figure 13. Measurements were taken at 5Gbps over a 30" link (+2x2" on line cards) on all three boards, going through the stack-up from top layer to bottom layer. The Wavecrest equipment was used to obtain the total jitter out to a Bit Error Ratio (BER) of 10^{-12} to give more quantitive results than eye diagrams.

Crosstalk impacts were briefly investigated on one link at 5Gbps. Measurements were taken over the 30" link on Platform #1, the XAUI 4000-2 board. The victim signal was input to IN_C1, and the aggressor signal was input to OUT_C1 on the receiving daughter card to achieve the more damaging near-end crosstalk (NEXT). The selected pins neighbor each other in adjacent signal columns and are representative of typical pinout conditions in a system environment. The aggressor signal had the same data rate as the desired signal, and was set at 4 times the desired signal to emulate the effect of multiple aggressors.

Results

The total jitter out of the pattern generator for the three data patterns used is given in Table 1.

Table 1: Total Jitter (10^{-12} BER) for Data Patterns

Data Pattern	Total Jitter (ps)	Total Jitter (UI)
2^7-1	38.96	0.19
2^{10}-1	40.28	0.20
CJTPAT	43.08	0.21

The additive jitter is defined as the total jitter at a BER of 10^{-12} at the output of the equalizer, minus the BER at 10^{-12} out of the pattern generator. No data recovery is performed on the data, as this will hide the margin available to the data recovery unit. Figure 16, Figure 17, and Figure 18 show the additive jitter at for a 30" backplane link for each of the three platforms.

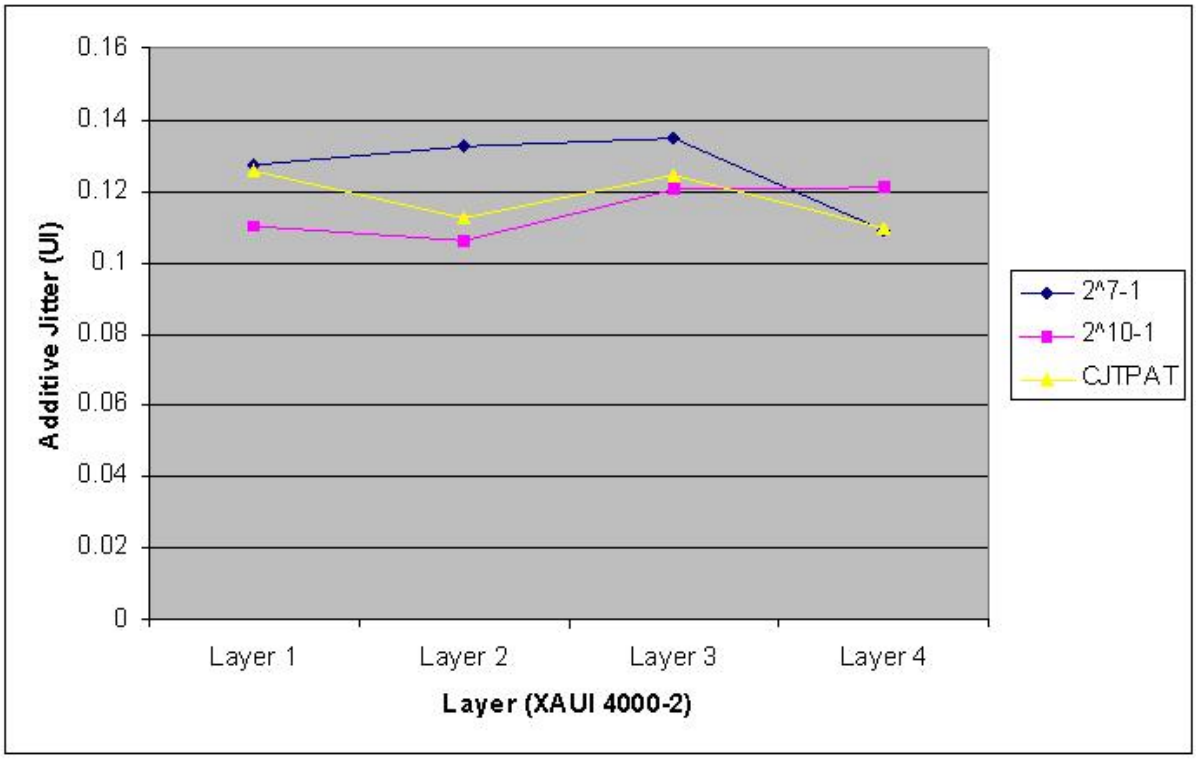

Figure 16: Additive Jitter After Equalizer for 30" Link over Platform #1

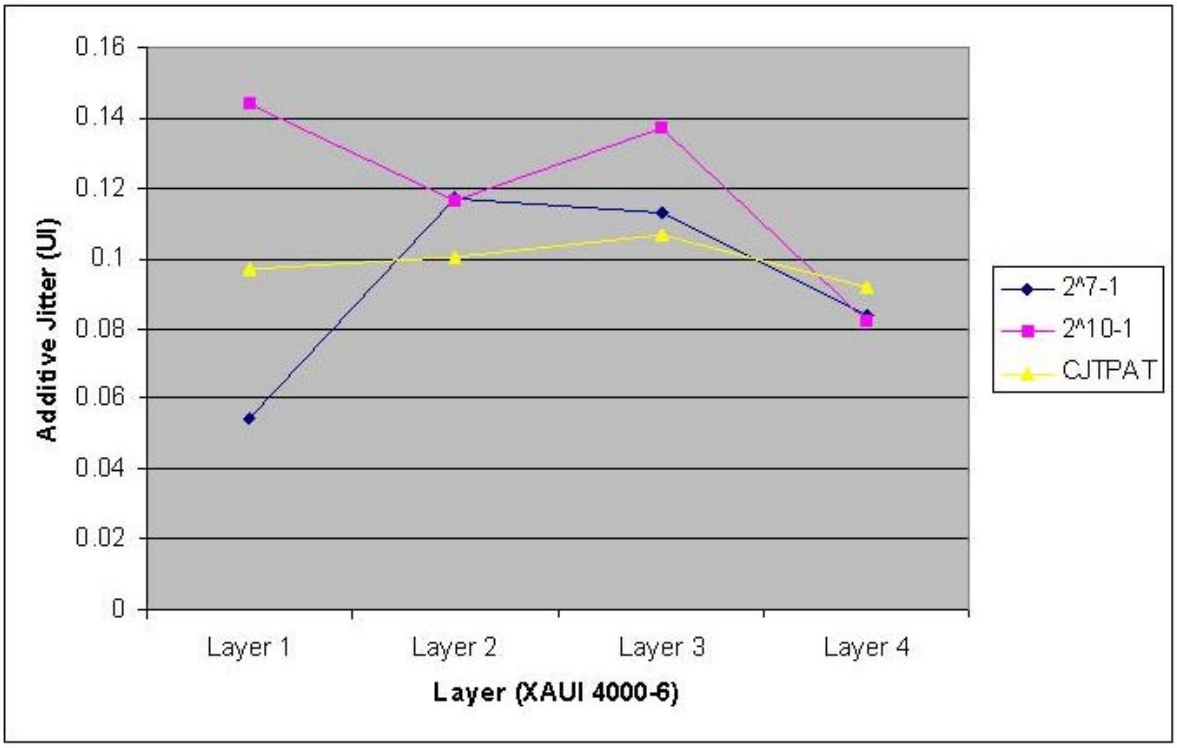

Figure 17: Additive Jitter After Equalizer for 30" Link over Platform #2

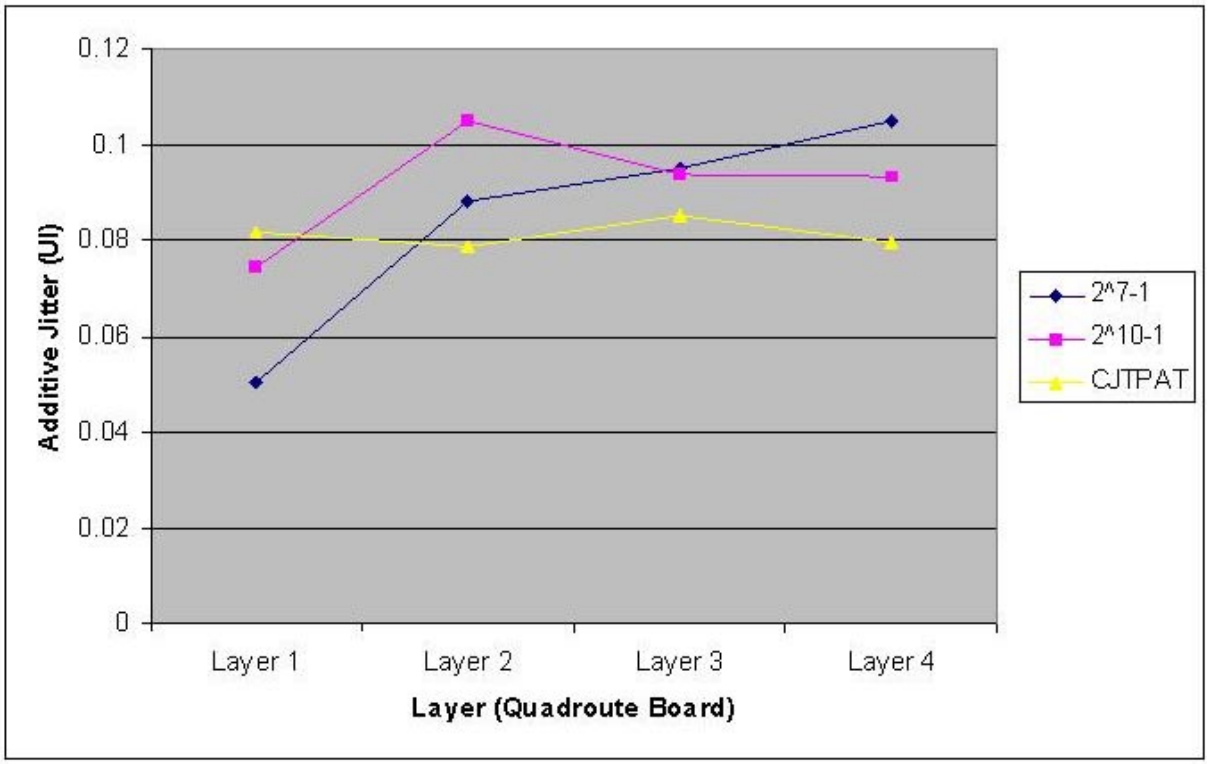

Figure 18: Additive Jitter After Equalizer for 30" Link over Platform #3

As mentioned previously, crosstalk results were taken. The difference in TJ was less than 0.01UI, and is well within the measurement repeatability.

Discussion of Results

The maximum additive jitter of all the measurement was just over 0.14UI. Clearly, with the XAUI transmit jitter allowance of 0.35UI, transmission over these channels at 5Gbps using receive end analog equalization methods is easily doable, even considering crosstalk. The layer effect appears to be minimal at this data rate with actually some unexpectedly low additive jitters on the top layer of the XAUI 4000-6 and Quad Route boards for the PRBS patterns. The CJTPAT pattern is relatively insensitive to layering. These anomalies may be a chance combination of data pattern affect on the adaptation algorithm, plus channel and equalizer frequency responses that match extremely well.

Another trend to note is that the additive jitters for the Quad Route board average about 0.03UI less than the other two boards, despite having a higher attenuation.

The increase in jitter due to crosstalk at this data rate is negligible. Figure 9 shows a crosstalk margin (S21 of desired signal minus S21 of crosstalk signal) of around 30dB at the 2.5GHz half data rate frequency. With aggressor levels of 4 times the desired signal, or 12dB, the margin for this case is closer to 20 dB. This measurement illustrates that crosstalk margins much less than 20dB (including all crosstalk degradations) are supportable.

Conclusions

Moving current backplane designs beyond XAUI speeds of 3.125 Gbps per lane while maintaining the desired BER level is a formidable task. It requires understanding the passive and active elements, their interaction, as well as utilizing the proper techniques to analyze their synergistic performance. This paper has brought together passive and active component expertise as well as testing expertise. The industry as a whole is recognizing that this is necessary and is bringing together the various disciplines required to help solve the common problem of multi gigabit serial transmission.

This paper has focused on applying analog adaptive equalization techniques to backplane testbeds that are representative of real world backplanes. Using Wavecrest jitter analysis equipment, which is based on industry-compliant jitter analysis techniques, it has been demonstrated that the analog adaptive equalization technique utilized by Gennum can extend current backplanes designed for XAUI speeds, based on the Tyco Electronics Z-PACK HM-Zd connector, to 5Gbps serial data rates. At this speed the adaptive equalization technique is using frequency content up to 2.5 GHz. This is important because this part of the spectrum has low ripple in comparison to the response of the channel at higher frequencies. Nonetheless, there is still enough ripple to impact the total amount of jitter measured at the output of the equalizer.

In the development of tomorrow's standards, the total jitter and random jitter is typically specified. Multiple platforms measured throughout the entire backplane stackup have permitted multiple deterministic jitter test scenarios by altering the amount of data dependent jitter for each test channel. There are several interesting points to note:

- From the three patterns used, CJTPAT pattern proved the most stable in relation to layer connection effect.
- Based on CJTPAT (and with most points for the other two patterns), Platform 3 had the best performance (approximately 0.08UI), followed by Platform 2 (approximately 0.10UI), followed by Platform 1 (approximately 0.12UI).
- 2^7-1 and 2^10-1 PRBS patterns yielded the most variation with Platforms 2 and 3, but minimal variation with Platform 1. This is seen to be a random occurrence where the adaptation algorithm and the data patterns combine to give an equalization response that very closely matches the ideal attenuation response of the channel.

One might infer that the associated material losses can account for the performance ranking of the platforms. This would not agree with the measured S_{DD21} responses of the systems, however, which showed Platform 3 having the highest attenuation for the 0 to 2.5 GHz frequency range. Thus, it is not necessarily the absolute magnitude of the measured S_{DD21} response, but also the quality of the curve. From Figure 15 it is observed that Platform 3 has the least amount of channel ripple while also having the best low to high frequency relation, which helps to minimize total ISI. Thus, the designer may consider using higher channel attenuation to minimize channel ripple to optimize the overall performance of interconnects based on this adaptive equalization technique.

Furthermore, although limited crosstalk testing was performed, efforts were made to exaggerate the effect of the crosstalk by increasing the strength of the intended aggressor. No noticeable difference in jitter performance was observed. This

implies that the Z-PACK HM-Zd connector has sufficient crosstalk performance to support 5 gigabit per second operation. Further testing with multiple aggressors is necessary.

The allowable XAUI transmit jitter is 0.35UI. Adding the channel/equalizer maximum additive jitter of 0.14UI to this gives 0.49UI. The XAUI maximum jitter amplitude tolerance is 0.65UI, giving a margin of 0.16UI. This is a fairly significant margin and indicates that the link can be made robust in an operational environment. It is also an indication that operation to higher data rates such as 6.25Gbps is feasible.

Significant attention has been given to the overall jitter performance of the measured Active Interconnects. The reason is simple - it is necessary in order to support the desired level of BER that is desired. Links are being designed with a higher maximum eye opening in order to provide margin to the user. This margin is needed to account for other sources of noise, such as the power supply, device process variation, manufacturing tolerances of the associated PWB's, including material properties variation and impedance variation, and operating condition variation, including temperature and power supply variation. By using frequency domain techniques to develop the best signal to noise ratio and jitter analysis in time domain to verify the performance to low BERs, the system designer can engineer an Active Interconnect to achieve the desired level of BER under all operating conditions.

---

[1] John D'Ambrosia, Michael Fogg, Ken Lazaris-Brunner, Bharat Tailor, "Reliable Serial Backplane Data Transmission at 10 Gbps," DesignCon, 2002.

[2] Jitter Fundamentals, www.Wavecrest.com

[3] J. Wilstrup, "A method of serial data jitter analysis using one-shot time interval measurements", ITC Proceeding, 1998.