Crosstalk Analysis of a System Based on XAUI HM-Zd Evaluation Backplane Data

Barry Katz
Signal Integrity Software, Inc.
Maynard, Massachusetts, USA

Dr. Michael Steinberger
Signal Integrity Software, Inc.
Chippewa Falls, Wisconsin, USA

Todd Westerhoff
Signal Integrity Software, Inc.
Maynard, Massachusetts, USA

***Abstract*—This paper presents the results of a performance analysis of a system based on measured backplane data published for the XAUI HM-Zd evaluation backplane. This analysis traces the impact of near end and far end crosstalk on bit error rate, and compares results from three different methods for estimating bit error rate in the presence of crosstalk.**

***Keywords*—Crosstalk analysis, bit error rate estimation, high speed serial channels**

I. Introduction

While it is clear that the data rate of high speed serial channels will continue to increase for the foreseeable future, one should not expect the number of high speed serial channels in a typical system will decrease. Rather, for many systems, the number of high speed serial channels will increase along with the data rate.

Under these conditions, crosstalk will become an increasingly critical channel impairment. Furthermore, since the data rate, interconnect length, and number of channels is set early in the system design cycle, the signal integrity engineers who support system designers should have a disciplined method for estimating the impact of all channel impairments, including crosstalk, on attainable bit error rate.

There are currently very few crosstalk analyses that are being applied to real systems, and the only method that could be described as being in widespread use is time domain Monte Carlo simulation.

The three methods demonstrated in this paper are time domain Monte-Carlo simulation, semi-analytic bit error rate estimation, and an extension the convolution engine approach [1]. By comparing the results obtained, it will be seen that it is very difficult at best to obtain reliable bit error rate estimates in the presence of significant crosstalk using only time domain simulation methods because such a large sample is required to be statistically significant.

The convolution engine approach is extended by convolving the probability density function (PDF) of the desired signal's eye diagram with the PDF of the combined crosstalk contributors. The computational method itself is a straightforward extension of the existing convolution engine approach, and so the emphasis will be on the results obtained.

The system used as a demonstration vehicle consists of four serial channels operating mesochronously (same frequency, unknown phase) over electrical paths described by the measured S parameter data published for the XAUI HM-Zd evaluation backplane [2]. Channels are operated in both directions (left to right and right to left) so as to exercise both near end and far end crosstalk coupling (NEXT and FEXT).

The equalization scheme chosen is synchronous a four tap transmit de-emphasis (one pre-cursor tap, two post-cursor taps) configuration which is currently in widespread use.

The results presented include bit error rate estimates, "bathtub" curves (probability of error as a function of sampling time), eye diagrams, transfer functions, step responses, and signal to crosstalk ratios.

II. Demonstration System

As shown in Fig. 1, the demonstration system consists of four high speed serial channels operating mesochronously at 6.25 Gb/s.

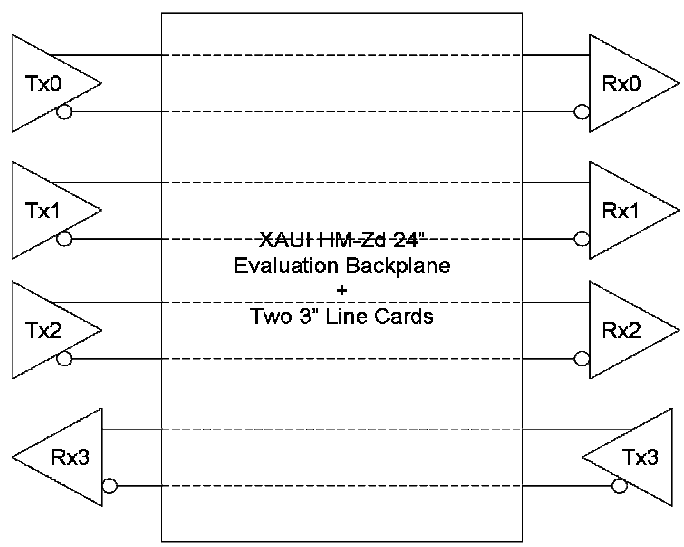

Figure 1. Demonstration system

In Fig. 1, the channels Tx0 to Rx0, Tx1 to Rx1 and Tx2 to Rx2 all propagate from left to right while the channel Tx3 to Rx3 propagates from right to left. Thus, there are a number of different combinations of NE and FE crosstalk:

- Tx0 to Rx0: One instance of FE
- Tx1 to Rx1: Two instances of FE
- Tx2 to Rx2: One instance of FE and one instance of NE
- Tx3 to Rx3: Two instances of NE

The data rate of 6.25 Gb/s is well beyond the data rate this evaluation backplane was originally designed for, and was chosen to create a situation in which the electrical path is to be operated at the limits of its capabilities.

A. *Backplane Model*

The backplane model is based on measured S parameters [2]. While data is available for several backplane lengths, the data chosen was for the 24" backplane with two 3" line cards because the goal was to analyze a channel with impairments that are as severe as possible.

As the results will demonstrate, crosstalk coupling is not simply a function of the two paths that are coupled directly. Rather, the crosstalk coupling is affected by the impedance of the surrounding system and, to a lesser extent, by adjacent paths that may not be directly coupled to the victim. The most precise analysis will therefore account for these effects.

To obtain as precise an analysis as possible, we therefore chose to represent the electrical network with a single sixteen port S parameter matrix, and to solve for all of the transfer functions in the system in a single computation rather than combine the results of the computation of multiple subcircuits.

The data for the desired electrical path, FEXT and NEXT were all measured and reported as separate four port measurements. and so it was necessary to assemble the data to produce a complete sixteen port S parameter matrix for the desired interconnect.

The assembly of the sixteen port S parameter matrix required some care and some assumptions. While it was clear which submatrices in general were to receive the THRU, NEXT, and FEXT data, there was no information which definitively stated what the order of the ports/pins would be if multiple channels were placed next to each other in the same backplane.

To choose an ordering for the ports and their associated data, it was assumed that the greatest coupling would occur between adjacent ports. This resulted in a definite order that seemed to make physical sense. The sixteen port S parameter matrix was therefore assembled assuming this port ordering.

B. *Equalization Solution*

The electrical channel has approximately 10 dB of loss at 3.125 GHz, but with about 6dB if gain ripple. The equalization solution chosen therefore used transmit de-emphasis to equalize the main response of the transmission path and four tap DFE to cope with the gain ripple. This type of equalization solution is currently in widespread use. The de-emphasis tap settings were

[-0.025, 0.700, -0.275, 0.000]

The details of these models, along with their operation as IBIS AMI models, are described in [3].

III. Transmission Analysis

Since it is exposed to both NEXT and FEXT, the Tx2 to Rx2 channel is a good example to use for transmission analysis. Fig. 2 shows the transfer function from each transmitter to Rx2.

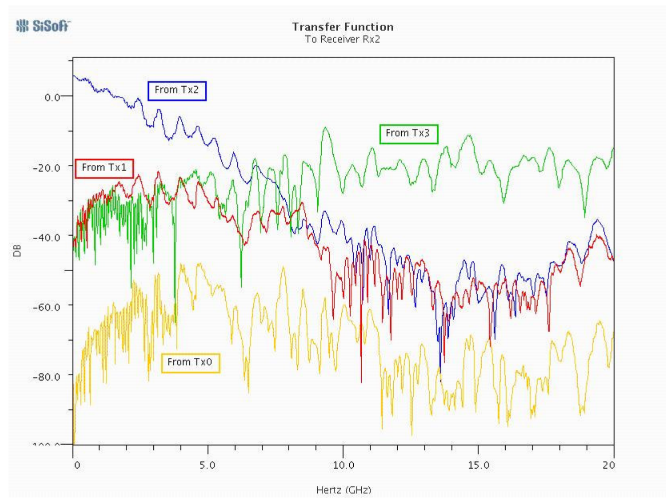

Figure 2. Transfer functions from each transmitter to receiver Rx2

While some of the features of Fig. 2 are what one would expect, some are also surprising. From Fig. 2, one can observe that

- The number of dBs of loss for the primary electrical path increases linearly with frequency, which is typical for a printed circuit board electrical path.
- Since the coupling from Tx3 is NEXT, this coupling does not decrease with frequency because the coupling path is so short.
- Since the coupling from Tx1 is FEXT, one would expect it to decrease with frequency almost as fast as the loss of the primary path; however it is surprising that at higher frequencies, the loss from Tx1 to Rx2 is more or less the same as the coupling from Tx2 to Rx2.
- There is coupling from Tx0 to Rx2 even though the S parameter matrix has no entries for non-nearest neighbor coupling.

Additional insight can be gained by examining the step responses generated by these transfer functions. Since these step responses are entirely equivalent to time domain transmission (TDT) measurements one would make in the lab, they can also be used for the same purpose: to determine the delay associated with various segments of the response.

For the sake of reference, Fig. 3 shows the step response for the primary channel. Note: The negative slope at the beginning of the primary step response is an artifact of the numerical processing that investigation to date has suggested is due to Gibb's phenomenon. This slope does not affect the pulse response or time domain simulations because the slope from a rising edge is canceled by the slope from the falling edge which follows it.

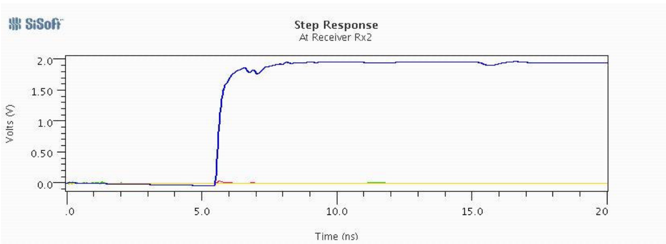

Figure 3. Transmitter Tx2 to receiver Rx2 step response

Fig 4. is an expanded view of the step responses to show the crosstalk responses clearly.

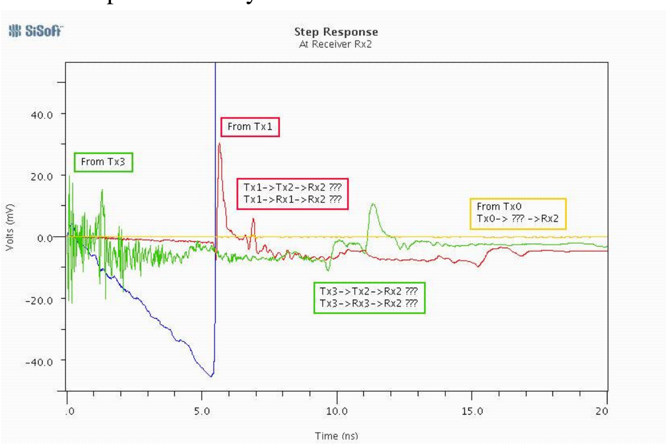

Figure 4. Crosstalk step responses at receiver Rx2

Fig. 4 illustrates clearly that, because it arrives before the primary step, the crosstalk from Tx3 is NEXT whereas the FEXT from Tx1 can arrive no sooner than the primary step.

Fig. 4 also illustrates some secondary crosstalk responses. Because these secondary paths result from a combination of several reflections and couplings, there is more than one way that any one of these responses could occur. For example, the response at about 6.8 nS could either occur because:
- The signal from Tx1 was near-end coupled to Tx2, reflected by Tx2 and then transmitted to Rx2 or
- The signal from Tx1 was transmitted to Rx1, reflected off Rx1, and was near end coupled to Rx2.
Either of these explanations would account for the total delay of 6.8 nS. A similar set of possible explanations applies to the response from Tx3 at about 11.5 nS.

This evidence of multiple reflections and couplings also offers an explanation for the coupling from Tx0 to Rx2. It could occur either because the signal from Tx0 was coupled to Tx1, reflected off Tx1 and was coupled to Rx2, or because the signal from Tx0 was coupled to Rx1, reflected off Rx1, and was coupled to Rx2.

These results also serve to illustrate the importance of considering the entire circuit at one time when calculating crosstalk. Crosstalk is not the result of a crosstalk coupling mechanism acting in isolation, but rather the result of a crosstalk coupling mechanism interacting with the impedances that surround it. When performing a system-level analysis, it is therefore desirable to include as much of the circuit environment as possible in a single measurement or network analysis rather study each of a number of separate crosstalk coupling paths in isolation.

IV. Performance Analysis

A. *Methods*

1) Convolution Engine

The convolution engine procedure [1] was applied to the desired signal, thus obtaining the PDF of the amplitude as a function of position in the eye. This PDF, then, is one way to represent the effects of inter-symbol interference. For the results reported here, the convolution engine took into account all messages of length 256 bits. As demonstrated in [4], this length is more than adequate for any practical applications.

The convolution engine procedure was then applied separately to each crosstalk aggressor, producing a PDF of that aggressor's crosstalk amplitude as a function of position in the eye. The PDFs for all the crosstalk aggressors were then convolved together to get a PDF for the total crosstalk, This PDF was then saved for use in semi-analytic bit error rate estimation, as described below.

Finally, the crosstalk PDF was convolved with the PDF for the desired signal to obtain a PDF for the complete signal. This PDF was then used to estimate the bit error rate.

2) Monte-Carlo

Monte-Carlo time domain simulations were run by making all transmitters active at the same time and accumulating a persistent eye diagram at the receiver in essentially the same way that a digital sampling oscilloscope would in a real system. That is, for each sample, the amplitude bin at the sampled amplitude and the sample's time position in the eye was incremented.

Since all transmitters were active, each sample explicitly included the effect of crosstalk, and the probability of error for a given sample was equal to either one or zero, depending on whether or not the sample was on the same side of the decision threshold as the sample for that bit in the center of the eye.

To remain consistent with common practice, the Monte-Carlo simulations simulated ten million bits using a linear feedback shift register (LFSR) data pattern with a register length of 31 bits.

3) Semi-Analytic

Traditionally, semi-analytic bit error rate estimation has been applied to signals in the presence of additive white Gaussian noise (AWGN). As demonstrated in [4], however, crosstalk cannot usually be accurately represented as a Gaussian process. Crosstalk was therefore assumed to be an additive process whose PDF was that calculated using the convolution engine procedure described above.

Semi-analytic bit error rate estimates were produced by disabling all transmitters except the transmitter for the desired signal, and then using the crosstalk amplitude PDF from the convolution engine analysis to estimate the bit error rate.

The semi-analytic simulations simulated ten million bits using the same LFSR as the Monte-Carlo simulations.

4) Other assumptions

For all bit error rate estimates, it was assumed that the decision circuit has a minimum latch overdrive of 25 mV. In other words, it was assumed that if the differential voltage at the decision point was less than 25 mV when the clock edge occurred, the decision circuit output would not switch soon enough to output the correct bit.

It was also assumed that the recovered clock driving the decision circuit has a Gaussian distributed phase noise process with a standard deviation of 3 pS. This standard deviation was chosen to represent the phase noise due to not only reference clock phase noise and pattern-dependent jitter, but the effects of power supply noise at the transmitter and receiver as well. The primary effect of power supply noise is to modulate the delay of the clock and data paths in both the transmitter and receiver, thus modulating the clock to data delay at the decision point. While the standard deviation of 3 pS is consistent with a SerDes design in which the power to the critical clock and data paths has been carefully filtered, power supply noise is still the primary contributor to the overall sampling clock phase noise.

In all cases, the bit error rate was estimated as the integration of a conditional probability. This calculation is based on the assumptions that the data bits are independent of each other, and that the sampling time is independent of the data signal being sampled, While the second assumption is not absolutely true in that pattern dependent jitter is one of the sources of sampling clock phase noise, it is a good approximation for two reasons:

1. Pattern dependent jitter is comparatively small contributor to the overall sampling clock phase noise.

2. The bandwidth of the clock recovery loop is much smaller than the data rate, and therefore the pattern dependent jitter at a given bit time is the aggregate result of many previous, independent bits.

Given that the bathtub curve is the probability of error as a function of sampling time $P(err \mid t)$, that the PDF of the sampling time is $p(t)$, and that the sampling time is independent of the data signal, the probability of error $P(err)$ is

$$P(err) = \int P(err \mid t)\, p(t)\, dt \qquad (1)$$

B. Results

1) Bit Error Rate

The bit error rate for each of the four channels was estimated using each of the three estimation techniques. The results are given in Table I.

TABLE I. ESTIMATED BIT ERROR RATES

Channel	Convolution Engine	Monte-Carlo Time Domain Simulation	Semi-Analytic Bit Error Rate Estimation
Tx0 to Rx0	2.4e-25	1.9e-29	3.5e-25
Tx1 to Rx1	1.9e-14	1.3e-15	2.6e-13
Tx2 to Rx2	1.4e-10	3.6e-11	1.9e-10
Tx3 to Rx3	4.8e-16	8.5e-30	2.3e-17

In Table I, there is good correlation between the bit error rate estimates produced by all three methods, especially in the sense that all three sets of data would result in the same engineering decisions. The correlation between the convolution engine and semi-analytic results is, however, better than the correlation between either of these results and the Monte Carlo results. This observation is consistent with the fact that the convolution engine takes the largest sample of statistics into account (all messages of length 256) while the Monte-Carlo method uses the smallest sample (ten million bits).

It should also be noted that whereas the convolution engine analysis was completed in a few seconds, the time domain simulations for the semi-analytic and Monte Carlo methods took a little over an hour each to run.

2) Signal to Crosstalk Ratios

In the analysis of digital transmission in the presence of AWGN, it is common practice to state the signal to noise ratio as energy per bit divided by the noise spectral density. In contrast to AWGN, however, crosstalk has bounded energy but highly non-uniform spectral density. It is therefore more appropriate to state the ratio of signal energy to crosstalk energy directly.

Signal energy for each channel was calculated by integrating the square of the pulse response. Similarly, the crosstalk energy from each aggressor was calculated by integrating the square of its pulse response into the victim channel.

Furthermore, the appropriate point at which to calculate the signal to crosstalk ratio is the decision point of the receiver, and to include the effects of equalization as well as channel loss in the calculation.

Table II contains the signal to crosstalk ratio from each aggressor to each victim. For the sake of later comparison, each entry in Table II shows the signal to crosstalk ratio including the effects of equalization above the signal to crosstalk ratio without equalization. The diagonal elements of Table II show that total signal to crosstalk ratio for the channel.

TABLE II. SIGNAL TO CROSSTALK RATIOS (DB)

From\To	Rx0	Rx1	Rx2	Rx3
Tx0	**26.0**	26.0 29.3	59.4 63.1	62.4 67.0
Tx1	26.1 29.3	**23.0**	26.0 29.3	44.8 49.2
Tx2	57.6 62.2	26.1 29.3	**21.7**	23.7 27.4
Tx3	77.3 81.2	58.8 61.8	23.5 27.2	**23.5**

Not surprisingly, there is a close correlation between the signal to crosstalk ratios in Table II and the bit error rate estimates in Table I.

From Tables I and II, it is apparent that to meet typical bit error rate requirements, the total equalized signal to crosstalk ratio should be at least 23 dB. Given that an ideal NRZ receiver would require a signal to noise ratio of 17dB to achieve a 1e-12 bit error rate (i.e., BPSK receiver with Eb/N0 of 17 dB), this suggests that the implementation margin for this system is about 6 dB.

It is recommended that, especially in the early phases of system design, an estimated unequalized signal to crosstalk ratio of 26a dB or greater can be considered safe, while any signal to crosstalk ratio less than that should be cause for concern, and link performance should be validated through laboratory measurements.

For the remainder of this section, results will be be given only for the channel Tx2 to Rx2. This choice is based on two factors:

- Tx2 to Rx2 is the channel with the highest error rate.
- Tx2 to Rx2 is exposed to both NEXT and FEXT, and so allows both mechanisms to be illustrated in a single example.

3) Bathtub Curves

Fig. 5 shows the bathtub curve, clock PDF, and net BER curve (i.e., the integrand in equation 1) for the convolution engine bit error rate estimate.

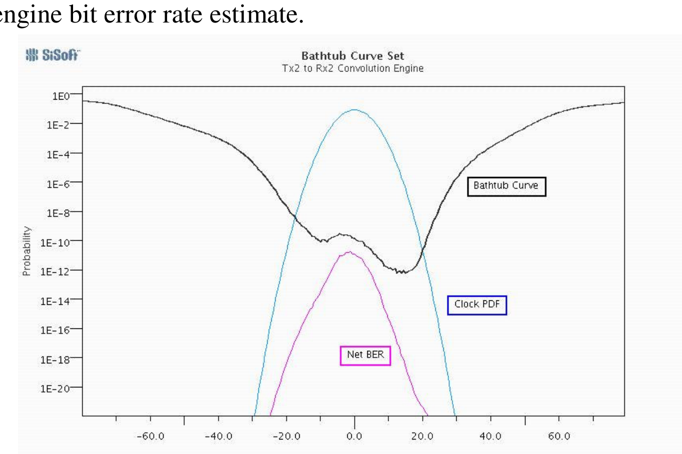

Figure 5. Bathtub curve set for Convolution Engine BER estimate

Note that this bathtub curve has a floor to it, and that there is therefore no time at which it would be possible to sample the data without incurring errors. This is typical of channels with significant amounts of crosstalk.

Fig. 6 compares the bathtub curves from the three different methods. This comparison follows the same pattern as the comparison of the bit error rate estimates.

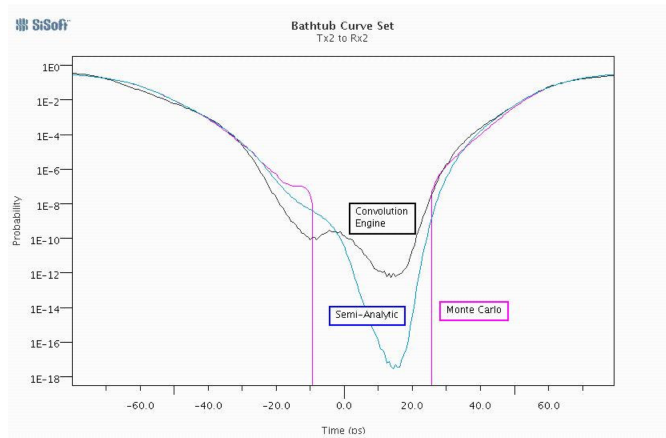

Figure 6. Comparison of bathtub curves

Fig 6 also serves to illustrate the hazards of statistical extrapolation in that it's not clear how one would extrapolate the Monte Carlo bathtub curve to produce a curve which more closely resembled the other two bathtub curves.

4) Eye Diagrams

The statistical eye diagrams generated by the three methods also serve to illustrate the difference in the sample sizes used in bit error rate estimation process.

In decreasing order of sample size, Fig. 7 is the eye diagram produced by the convolution engine, Fig. 8 is the eye diagram produced by the semi-analytic method, and Fig. 9 is the eye diagram produced by the Monte-Carlo simulation.

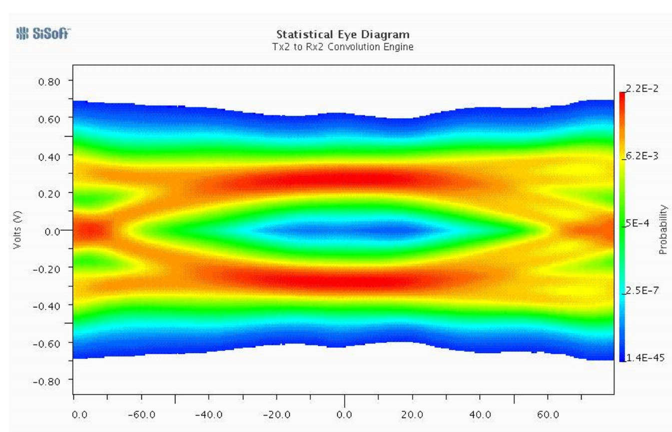

Figure 7. Convolution engine eye diagram

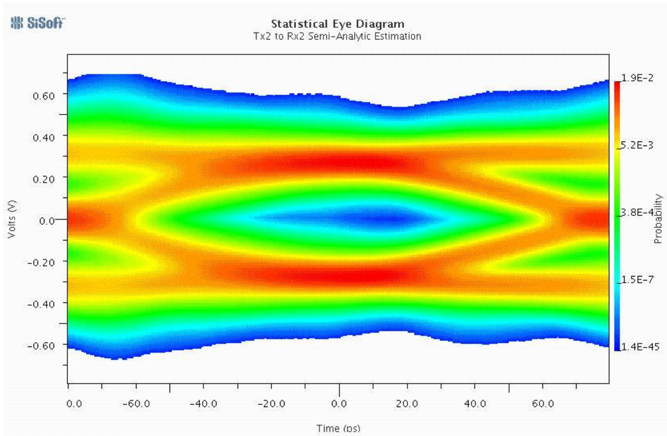

Figure 8. Semi-analytic eye diagram

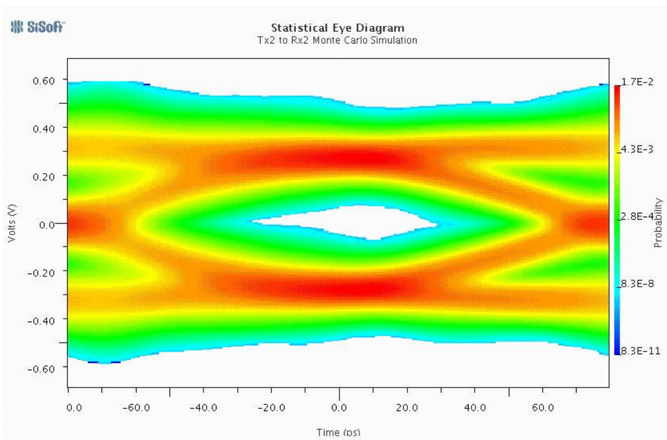

Figure 9. Monte-Carlo eye diagram

5) Crosstalk Characteristics

It can sometimes also be instructive to examine the eye diagram of the crosstalk in the absence of the desired signal. For the Tx2 to Rx2 channel, this is shown in Fig. 10.

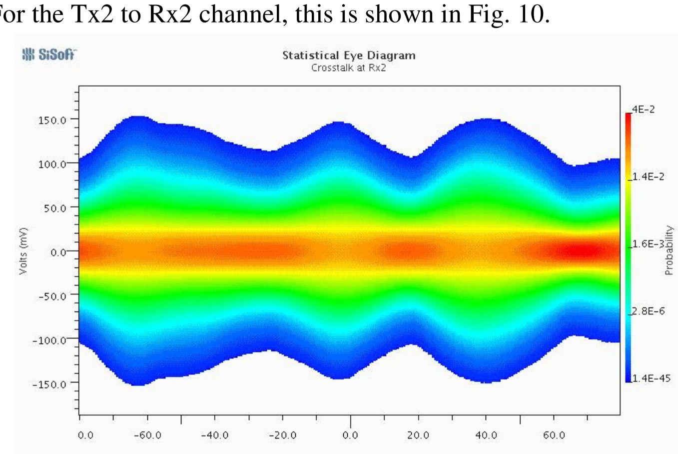

Figure 10. Crosstalk eye diagram at Rx2

Fig. 10 shows significant variation of crosstalk PDF with sample time. This variation of crosstalk PDF with sampling time can affect the performance of mesochronous systems in a very visible way. For example, if the relative phase of the crosstalk at Rx2, were to be delayed by 65 pS, the crosstalk would affect the bit error rate far more than if the crosstalk were advanced by 65 pS. Such shifts can occur either through small changes in routing or changes in printed circuit board material, even if the two materials have ostensibly the same dielectric constant.

V. CONCLUSIONS

This paper has presented estimates of the ultimate performance of the XAUI evaluation backplane using three different estimation methods: convolution engine, semi-analytic BER estimation, and Monte Carlo simulation. The conclusions of this study are

1. All three methods can produce useful BER estimates which are consistent with each other.
2. Because it takes the most statistics into account and is the most efficient computationally, the convolution method should be used whenever the method can be applied.
3. There are hazards associated with statistical extrapolation, and so this technique should only be used when there are no other options.
4. Signal to crosstalk energy ratio is a reliable indicator of the effect of crosstalk on BER. An unequalized signal to crosstalk ratio of at least 26dB for satisfactory BER is a useful rule of thumb.
5. In the system analyzed, near end crosstalk was the most serious impairment. This is consistent with the fact that near end crosstalk can couple a strong aggressor into a weak victim. Configurations with significant near end crosstalk should therefore be avoided in system designs.
6. In mesochronous systems, the crosstalk energy can vary significantly due to small changes in the relative delay between aggressor and victim. Analyses should therefore be performed for the worst case relative delay.

REFERENCES

[1] Anthony Sanders, Mike Resso, John D'Ambrosia, "Channel Compliance Testing Utilizing Novel Statistical Eye Methodology", DesignCon 2004 Proceedings.

[2] John D'Ambrosia and Adam Healey, "Channels for Consideration by the Signaling Ad Hoc", IEEE P802.3ap Signaling Ad Hoc, September 17, 2004.

[3] Steinberger, Westerhoff, and White, "Demonstration of SerDes Modeling using the Algorithmic Model Interface (AMI) Standard", DesignCon 2008 Proceedings, February 2008.

[4] Michael Steinberger, "Exploration of Deterministic Jitter Distributions", DesignCon 2008 Proceedings, February 2008.