Comparison of BER Estimation Methods which Account for Crosstalk

As presented at DesignCon 2009

Co-authored by:

- **Michael Steinberger, Signal Integrity Software, Inc.**
  msteinb@sisoft.com
- **Barry Katz, Signal Integrity Software, Inc.**
  bkatz@sisoft.com
- **Walter Katz, Signal Integrity Software, Inc.**
  wkatz@sisoft.com
- **Todd Westerhoff, Signal Integrity Software, Inc.**
  twesterh@sisoft.com

Signal Integrity Software, Inc (SiSoft)

6 Clock Tower Place, Suite 250 – Maynard, MA 01754

www.sisoft.com

978.461.0449

Abstract

As serial channel data rates continue to increase, crosstalk is becoming a more significant impairment, and accurately accounting for the effects of crosstalk is becoming more important to bit error rate estimation.

There are many ways to incorporated crosstalk into a bit error rate estimate. Using an example channel for which crosstalk is a significant impairment, this paper compares results obtained from four estimation procedures: jitter-based analysis, explicit time domain simulation, semi-analytic BER estimation, and statistical analysis. Results include estimated BER, eye diagrams and bathtub curves. For this example, the more abstract methods yield more consistent results.

Author(s) Biography

Dr. Michael Steinberger is currently responsible for leading the development of SiSoft's serial link analysis products. He has over 30 years experience in the design and analysis of very high speed electronic circuits. Prior to joining SiSoft, Dr. Steinberger worked at Cray Inc., where he designed very high density interconnects and increased the data rate and path lengths to the state of the art. Mike holds a B.S. from the California Institute of Technology and a Ph.D. from the University of Southern California, and has been awarded 13 U.S. patents.

Barry Katz, President and CTO for SiSoft, founded SiSoft in 1995. As CTO, Barry is leading the definition and development of SiSoft's products. He has devoted much of his efforts at SiSoft to delivering a comprehensive design methodology, software tools, and expert consulting to solve the problems faced by designers of leading edge high-speed systems. He was the founding chairman of the IBIS Quality committee. Barry received an MSEE degree from Carnegie Mellon and a BSEE degree from the University of Florida.

Dr. Walter Katz, Chief Scientist for SiSoft, is a pioneer in the development of constraint driven printed circuit board routers. He developed SciCards, the first commercially successful auto-router. Dr. Katz founded Layout Concepts and sold routers through Cadence, Zuken, Daisix, Intergraph and Accel. More than 20,000 copies of his tools have been used worldwide. Dr. Katz developed the first signal integrity tools for a 17 MHz 32-bit minicomputer in the seventies. In 1991, IBM used his software to design a 1 GHz computer. Dr. Katz holds a PhD from the University of Rochester and a BS from Polytechnic Institute of Brooklyn.

Todd Westerhoff, Vice President of Software Products for SiSoft, has over 27 years experience in the modeling and analysis of electronic systems, including 10 years of signal integrity experience. Prior to joining SiSoft, Todd managed a high-speed design group that provided static timing, signal integrity and design rule consultation to various ASIC and system engineering groups within Cisco Systems, Inc. Todd holds a B.E. degree in electrical engineering from the Stevens Institute of Technology in Hoboken, New Jersey.

1.0 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 methods that could be described as being in widespread use are jitter-based analysis and explicit time domain simulation.

This paper compares bit error rate estimates for the same example system using four different estimation methods. These methods are

1. Jitter-based analysis: Crosstalk is treated as a form of determinisitc jitter (DJ) which is separate from the DJ due to intersymbol interference. The peak DJ due to crosstalk is added to the peak DJ due to intersymbol interference to yield a total peak DJ. This is then combined with an assumed level of random jitter (RJ) to produce a bit error rate estimate.
2. Explicit time domain simulation: All channels in the system are simulated simultaneously in the time domain so that the waveform at the receiver decision point is the combined result of the desired signal and the crosstalk aggressors. A bathtub curve is accumulated over the course of the simulation, using the recovered clock to set the timing.
3. Semi-analytic estimation: Each channel in the system is simulated in the absence of crosstalk. An eye diagram is accumulated using the recovered clock to set the timing. The probability density function (PDF) of the crosstalk is calculated separately using convolution engine methods similar to the method described in [1]. The crosstalk PDF is then convolved with the eye diagram, and a bathtub curve and bit error rate estimate derived from that.
4. Statistical analysis: The convolution engine method described in [1] is used to create both a statistical eye diagram and a crosstalk PDF. These are then convolved together to produce a total statistical eye diagram. A bathtub curve is derived from that, and then the PDF of the recovered clock phase is combined with that to produce a bit error rate estimate.

Bit error rate estimates are provided for all four methods, while eye diagrams and bathtub curves are provided for all except the jitter-based analysis. All analyses were performed using SiSoft's Quantum Channel Designer serial channel analysis environment.

This paper also demonstrates that the signal to crosstalk ratio at the receiver is a useful early predictor of the effect of crosstalk on bit error rate for a given channel. Predicted bit error rate is correlated with signal to crosstalk ratio for four different channels, the effect of equalization solution on signal to crosstalk ratio is analyzed, and signal to crosstalk ratio is compared to the parameter insertion loss to crosstalk ratio (ICR) used in recent standards [2].

2.0 Example System

The demonstration system shown in Figure 1 consists of four high speed serial channels operating mesochronously at 6.25 Gb/s. 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

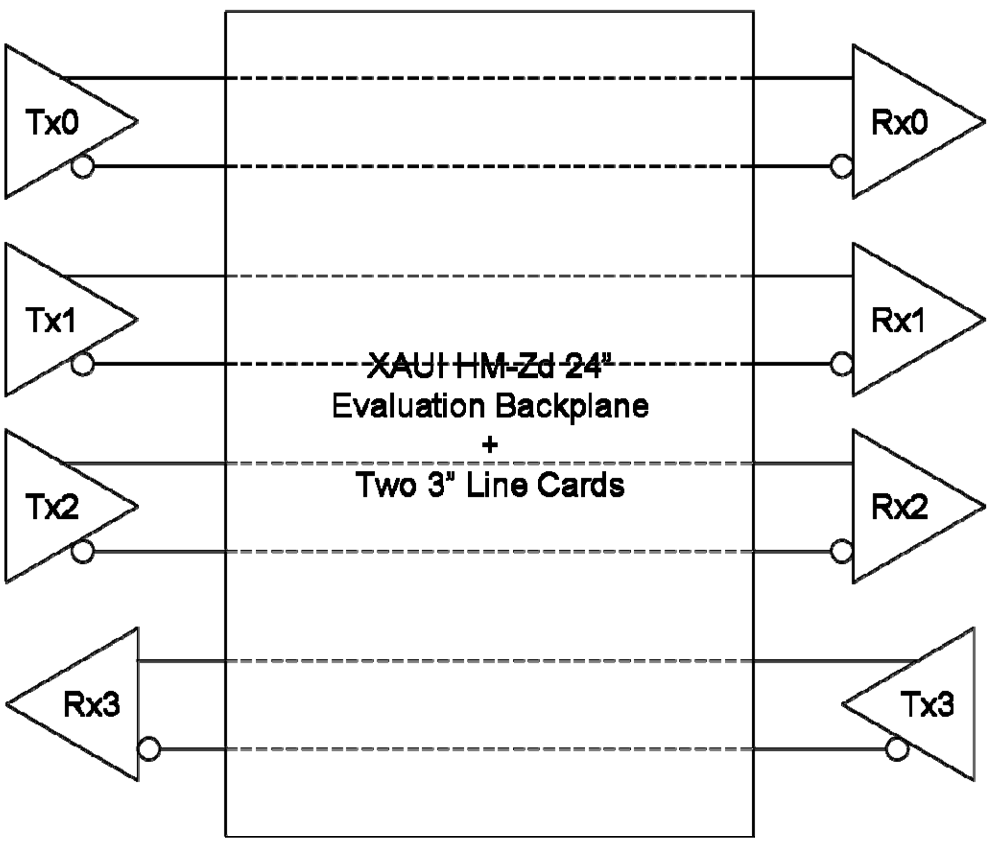

Figure 1: Example system block diagram

2.1 Backplane Model

The backplane model is based on measured S parameters for the XAUI HM-Zd evaluation backplane [3]. 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.

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

To obtain as precise an analysis as possible, we 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.

One of the advantages of the backplane data is that NEXT and FEXT were measured in addition to the desired (THRU) electrical path. One of the challenges, however, was that the data was 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 using this port ordering.

2.2 Equalization

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

The transmit tap weights were -0.025, 0.7, -0.275.

The DFE tap weights converged to -0.027, 0.023, -0.022, -0.006.

This equalization was modeled using transmitter and receiver models conforming to the IBIS AMI (Algorithmic Modeling Interface) standard. The structure and operation of these models is described in [4].

3.0 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. Figure 2 shows the magnitude of the transfer function from each transmitter to the receiver Rx2. These transfer functions include the gain of the transmit and receive buffers, resulting in a gain from Tx2 to Rx2 at lower frequencies.

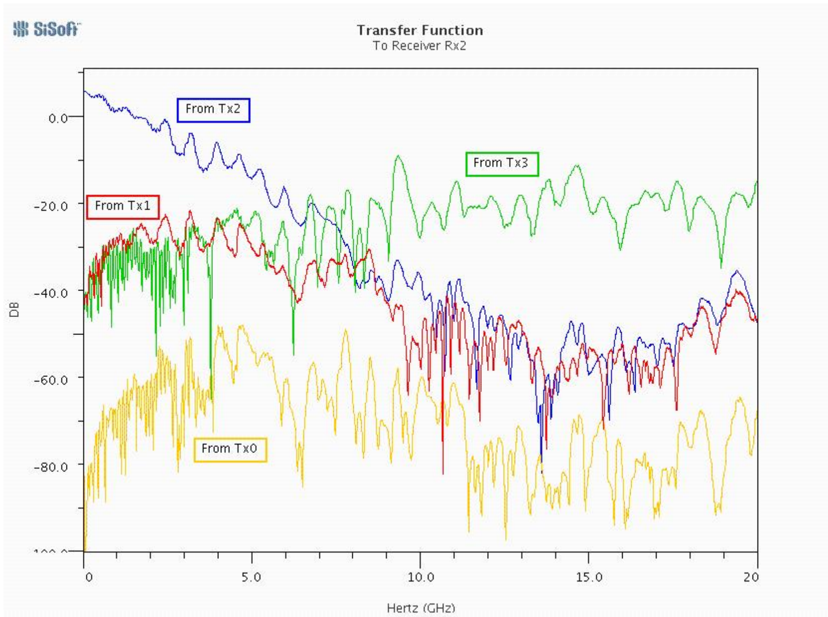

Figure 2: Transfer functions from each transmitter to receiver Rx2

While some of the features of the transfer functions from each transmitter to Rx2 are what one would expect, some are also surprises:

- 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, Figure 3 shows the step response for the primary channel.

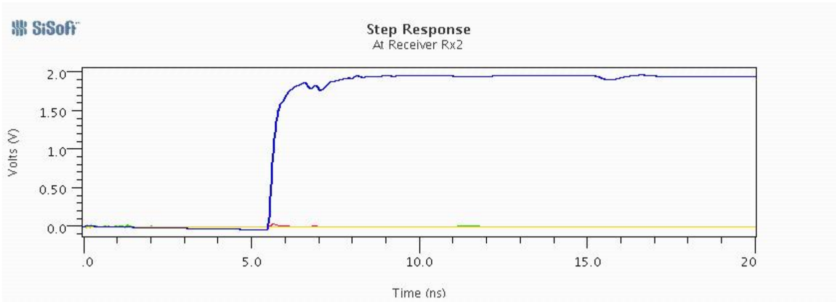

Figure 3: Full scale rendering of step responses at Rx2

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

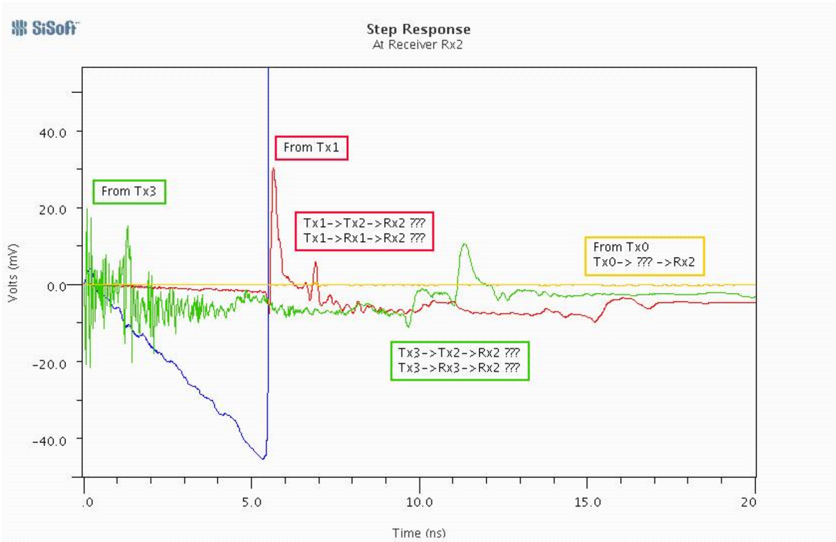

Figure 4: Crosstalk step responses at receiver Rx2

Note: The negative slope at the beginning of the primary step response is a numerical artifact that is a form of Gibbs phenomenon [5]. That is, it is the result of applying the discrete Fourier transform to data that was measured in the continuous frequency domain. This artifact is insignificant for the time domain simulation of most data signals because there are multiple data edges which cancel out the Gibbs phenomenon.

The step responses illustrate 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.

The step responses also illustrate 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. This evidence of multiple reflections and couplings also offers an explanation for the coupling from Tx0 to Rx2, for which there is no direct coupling path. 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.

4.0 BER Estimation Methods

4.1 Global 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.

4.2 Jitter-based Analysis

DJ due to intersymmbol interference was determined from time domain eye diagrams accumulated from one million bits. For Rx2, the peak DJ was 0.27 UI. The slew rate at the inside of the eye diagram was also measured. For Rx2, it was 0.003 mV/pS. Finally, the peak crosstalk was determined from statistical analysis. For Rx2, the peak crosstalk amplitude was 150 mV. From these numbers the peak total DJ was calculated as

$$DJ_{total} = DJ_{eye} + \frac{V_{XT\_max}}{r_{slew}} \qquad \text{(eq 1)}$$

For Rx2, this resulted in a value of 0.58 UI, which is greater than one half UI, indicating that jitter-based analysis predicts a BER of one half for this channel, even if the RJ were zero. Similar results were obtained for the other channels.

4.3 Explicit Time Domain Simulation

Explicit time domain simulations were run by making all transmitters in a time domain simulation 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.

The eye diagram was triggered by the clock recovered by the receiver model, which is the method used in at least one very popular analysis program. While this approach has the advantage of capturing any pattern-dependent effects, it does not include the effect of power supply noise at either the transmitter or receiver. To simulate the effects of power supply noise while maintaining more or less consistent RJ between simulation methods, 3 pS of RJ was injected at the transmitter.

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

4.4 Semi-Analytic BER Estimation

Traditionally, semi-analytic bit error rate estimation has been applied to signals in the presence of additive white Gaussian noise (AWGN). As demonstrated in [6], 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 in [1].

Semi-analytic bit error rate estimates were produced by disabling all transmitters except the transmitter for the desired signal. The eye diagram was then accumulated using exactly the same methods and conditions as described above for explicit time domain simulation.

The convolution engine procedure was 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.

Following the time domain simulation, the crosstalk amplitude PDF from the convolution engine analysis was convolved with the eye diagram to produce an eye diagram with crosstalk. The bit error rate was estimated from the eye diagram using exactly the same method as for the explicit time domain approach described above.

Semi-analytic simulations simulated ten million bits using the same LFSR as the explicit time domain simulations.

4.5 Statistical Analysis

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 [6], this length is more than adequate for any practical applications.

Whereas in the explicit time domain and semi-analytic approaches the DFE had the opportunity to converge before the bit error rate was estimated, for statistical analysis, the DFE tap weights were estimated a priori. As explained in [4], estimate of DFE tap weights does not account for the interaction between the DFE control loop and the clock recovery loop, and so the DFE tap weights for the statistical analysis were slightly different from those for the explicit time domain and semi-analytic methods. This had a negligible effect on the estimated bit error rate, but is visible in the eye diagrams and bathtub curves to follow.

As with the semi-analytic method, 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.

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.

The bit error rate was estimated as the integration of a conditional probability.

$$P_{err} = \int_{t=-T/2}^{T/2} P(err \mid t)\, p(t)\, dt \qquad \text{(eq 2)}$$

Where

- Perr is the total probability of error, or in other words the bit error rate
- P(err|t) is the probability of error given sampling time t. In other words, this is the bathtub curve.
- p(t) is the probability density function of the sampling time, or in other words the clock PDF
- T is the bit period

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.

In other analyses not reported here, the second assumption (clock independence) has also proven to be conservative when decision feedback equalization (DFE) is used. This makes sense in that part of the signal being detected has been clocked by the cycles of the recovered clock immediately preceding the sampling time.

5.0 Bit Error Rate Results

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 1.

Table 1: Estimated Bit Error Rates

Channel	Jitter-based Analysis	Explicit Time Domain Simulation	Semi-Analytic BER Estimation	Statistical Analysis
Tx0 to Rx0	5e-4	1.9e-29	3.5e-25	2.4e-25
Tx1 to Rx1	0.5	1.3e-15	2.6e-13	1.9e-14
Tx2 to Rx2	0.5	3.6e-11	1.9e-10	1.4e-10
Tx3 to Rx3	0.5	8.5e-30	2.3e-17	4.8e-16

In Table I, there is good correlation between the bit error rate estimates produced by all methods except jitter-based analysis, especially in the sense that the other three sets of data would result in the same engineering decisions. The correlation between the statistical analysis and semi-analytic results is, however, better than the correlation between either of these results and the explicit time domain simulation 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, or 10^{77} possible messages) while the explicit time domain simulation method uses the smallest sample (ten million bits). In the case of statistical analysis, the large sample is taken for both intersymbol interference and crosstalk, while for semi-analytic estimation only ten million samples of intersymbol interference are produced.

Further insight can be gained from the eye diagrams and bathtub curves from which the above results were derived. Figures 5, 6 and 7 show the eye diagrams for the channel Tx2 to Rx2 for the explicit time domain, semi-analytic, and statistical analysis methods, respectively.

These eye diagrams illustrate that it is very difficult for the explicit time domain method to capture the very rare events that dominate the bit error rate for this channel. By comparison, the eye diagrams for the semi-analytic and statistical analysis methods are nearly identical, largely because they both use a crosstalk PDF which has been determined analytically, and therefore reflects events that would be unlikely to occur in a sample size as small as ten million bits.

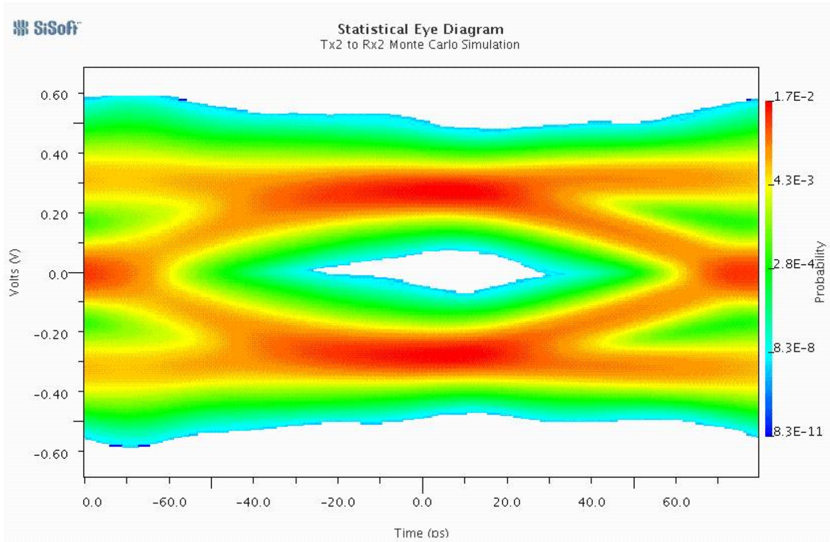

Figure 5: Explicit time domain eye diagram with crosstalk

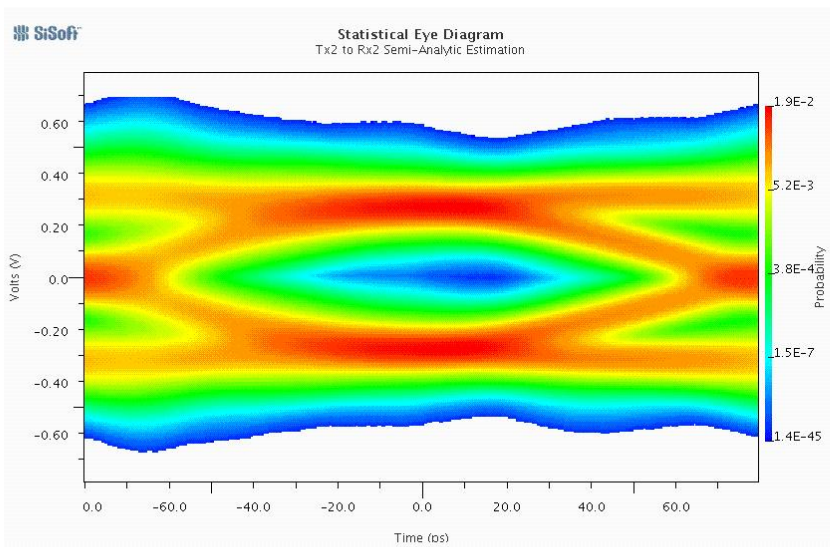

Figure 6: Semi-analytic eye diagram with crosstalk

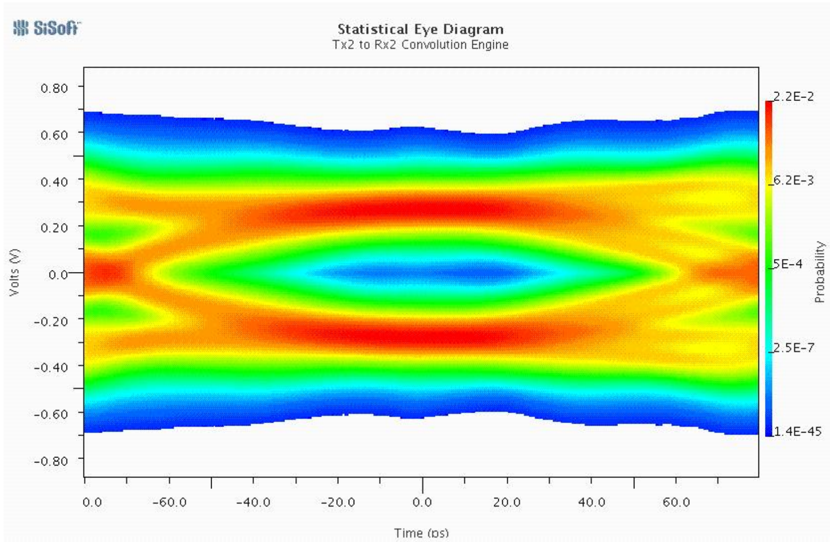

Figure 7: Statistical analysis eye diagram with crosstalk

The same conclusion is suggested by the bathtub curves. Figure 8 shows the bathtub curve, clock PDF, and net BER curves for the statistical analysis case, while Figure 9 is a comparison of the bathtub curves for all three methods.

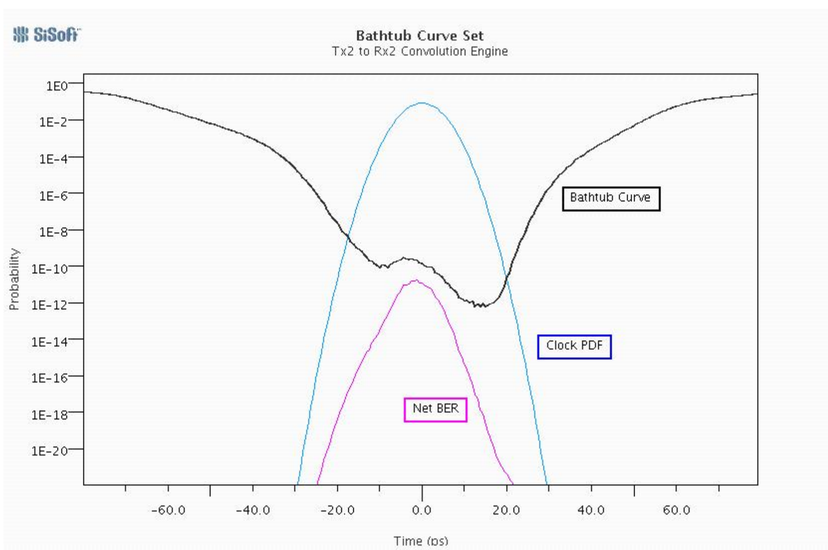

Figure 8: Statistical analysis bathtub curve set

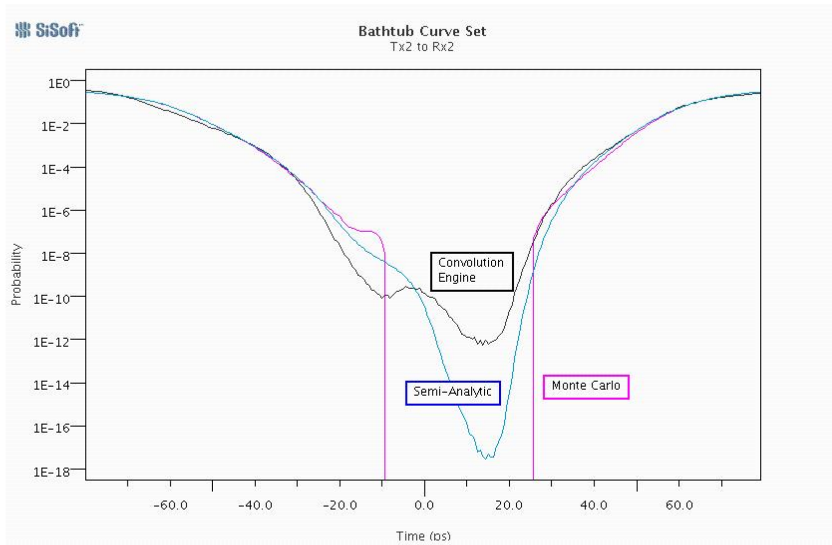

Figure 9: Comparison of bathtub curves

Note in Figure 9 that while the bathtub curves for the statistical analysis and semi-analytic methods have a floor to them, the bathtub curve for the explicit time domain method does not. This is consistent with the eye diagram in Figure 5 and the general discussion of sample size in this section. Especially when using explicit time domain simulation, it is common practice to attempt to compensate for the small sample size by statistical extrapolation. Figure 9 illustrates the hazards of this approach in that it is difficult to see how statistical extrapolation could consistently make the explicit time domain bathtub curve look like the other two curves. Especially since it's been proven [6] that the statistics can not be expected to be Gaussian, it seems rather optimistic to attempt to extrapolate from a probability of 10^{-7} to predict events which have a probability of between 10^{-12} and 10^{-18}.

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

Finally, it is instructive to note that in a mesochronous system (identical frequencies, uncontrolled phase), the total crosstalk can vary considerably with respect to the relative phase between the victim and its aggressors. Figure 10 shows the total crosstalk at Rx2 in the absence of the victim signal.

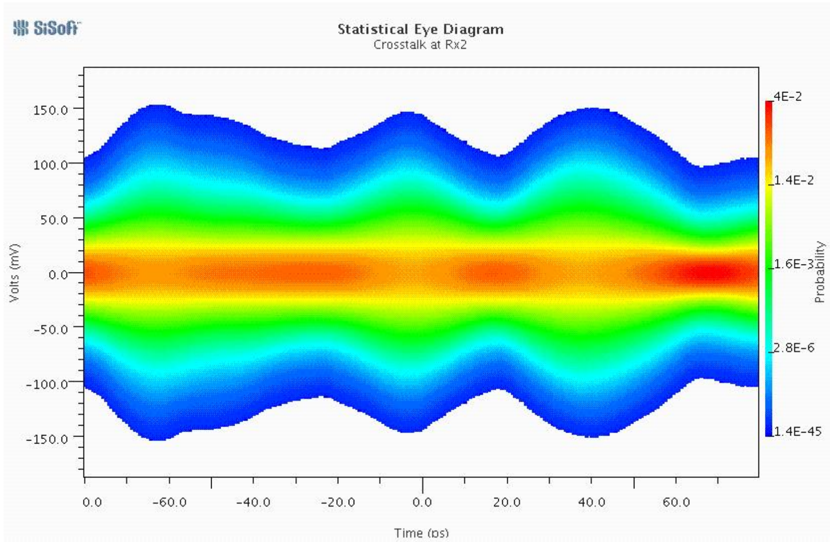

Figure 10: Crosstalk eye at Rx2

6.0 Signal-to-Crosstalk Ratio

Especially early in the definition of a system, it is very useful to have a parameter which reliably predicts performance (i.e., bit error rate) without requiring a lot of effort or a great deal of detailed information. The insertion loss to crosstalk ratio (ICR) defined in [2] is one such parameter. This section suggests an even simpler parameter: signal to crosstalk ratio.

In the analysis of digital transmission in the presence of AWGN, it is common practice to state the signal to noise ratio at the input to a band limited receiver as energy per bit divided by the noise spectral density (Eb/N0), and to present the bit error rate as a function of Eb/N0. Eb/N0 is calculated early in the system definition phase and used to determine whether a given system is likely to meet requirements. The following two considerations argue against applying this approach directly to high speed serial channels, however:

1. The receiver may be band limited, but the bandwidth and passband characteristics are not usually very well controlled.
2. In contrast to AWGN, crosstalk has bounded energy but highly non-uniform spectral density.

It is therefore more appropriate to state the ratio of signal power to crosstalk power at the receiver decision point directly.

Signal power at the receiver decision point can be calculated as

$$P_s = R_s \int p_s^2(t)dt \qquad \text{(eq 3)}$$

Where

- Ps is the signal power

- Rs is the signal data rate
- ps(t) is the signal pulse response at the receiver decision point

A similar equation can be applied to each crosstalk aggressor. The resulting signal to crosstalk ratio is

$$SXT = \frac{P_s}{P_{xt}} = \frac{R_s \int p_s^2(t)dt}{\sum_i R_i \int p_{xt,i}^2(t)dt} \quad \text{(eq 4)}$$

6.1 Effect of Equalization

One of the questions that often arises is the interaction between crosstalk and the equalization solution. For example, how is crosstalk affected if more equalization is put at the transmitter or receiver?

If one is willing to make four assumptions, the answer is simple: More or less equalization can be put at either the transmitter or receiver without affecting the impact of crosstalk on bit error rate. These assumptions also result in a relatively simple method for estimating the signal to crosstalk ration, and therefore the effect on bit error rate. These assumptions are:

1. Each victim runs parallel to its aggressors, so that victims and aggressors have essentially the same path loss.

2. Victims and aggressors propagate in the same direction.

3. Victims and aggressors have the same data rate, signal strength, and equalization solution.

4. The equalization solution does a good job of equalizing the channel.

This set of conditions is illustrated in Fig. 11 for one victim and one aggressor

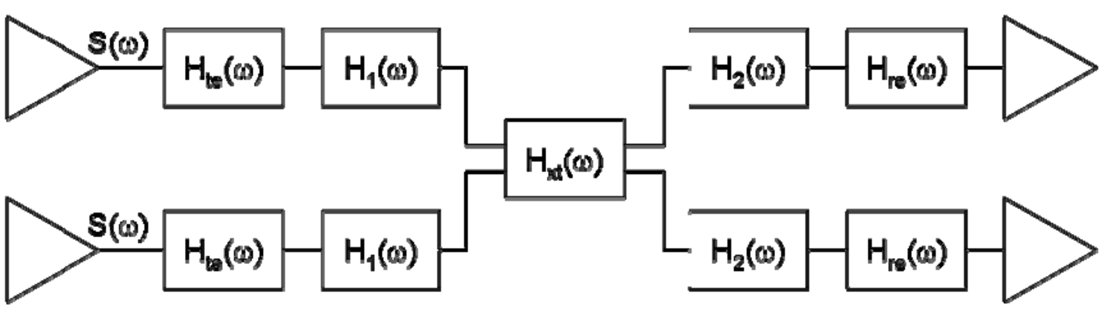

Figure 11: FEXT block diagram

Where

- $S(\omega)$ is the Fourier transform of the data signal.
- $H_{te}(\omega)$ is the transfer function of the transmit equalization.

- $H_1(\omega)$ is the transfer function of the electrical path leading up to the crosstalk coupling section.
- $H_{xt}(\omega)$ is the transfer function of the crosstalk coupling section from aggressor to victim. It is assumed that the crosstalk section is transparent to the primary transmission path.
- $H_2(\omega)$ is the transfer function of the electrical path from the crosstalk coupling section to the receiver
- $H_{re}(\omega)$ is the transfer function of the receive equalizer.

The Fourier transform of the signal at the receiver decision point is therefore

$$S_r(\omega) = S(\omega)H_{te}(\omega)H_1(\omega)H_2(\omega)H_{re}(\omega) \qquad \text{(eq 5)}$$

and the crosstalk at the receiver decision point is

$$S_{rxt}(\omega) = S(\omega)H_{te}(\omega)H_1(\omega)H_{xt}(\omega)H_2(\omega)H_{re}(\omega) \qquad \text{(eq 6)}$$

Define

$$\Psi(\omega) \equiv S(\omega)H_{te}(\omega)H_1(\omega)H_2(\omega)H_{re}(\omega) \qquad \text{(eq 7)}$$

Then

$$S_r(\omega) = \Psi(\omega) \qquad \text{(eq 8)}$$

$$S_{rxt}(\omega) = \Psi(\omega)H_{xt}(\omega) \qquad \text{(eq 9)}$$

The signal to crosstalk ratio R_{xt} is therefore

$$SXT = \frac{\int |\Psi(\omega)|^2 \, d\omega}{\int |\Psi(\omega)|^2 |H_{xt}(\omega)|^2 \, d\omega} \qquad \text{(eq 10)}$$

If a reasonable approximation is chosen for the function PSI, then equation 10 will provide an accurate estimate of the signal to crosstalk ratio.

The first crude approximation is to assume that the channel is unequalized. A second choice would be to use the channel equalization conditions described in [6], and a third choice would be to use the actual equalized response of the channel. The next section will present results using the first and third of these choices.

If the victim and aggressor propagate in opposite directions, as illustrated in Fig. 12, the analysis is more complex, and in general the situation should be avoided for the sake of system performance.

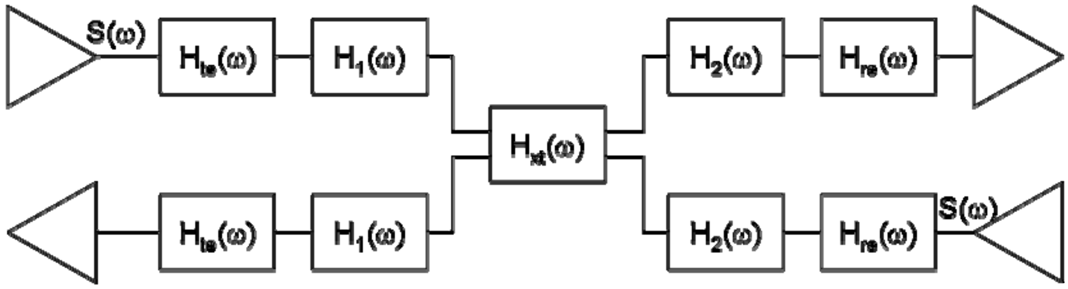

Figure 12: NEXT block diagram

In this case,

$$S_{rxt}(\omega) = S(\omega)H_{te}(\omega)H_2(\omega)H_{xt}(\omega)H_2(\omega)H_{re}(\omega) \qquad \text{(eq 11)}$$

$$S_{rxt}(\omega) = \Psi(\omega)H_{xt}(\omega)\frac{H_2(\omega)}{H_1(\omega)} \qquad \text{(eq 12)}$$

$$SXT = \frac{\int \left|\Psi(\omega)\right|^2 d\omega}{\int \left|\Psi(\omega)\right|^2 \left|H_{xt}(\omega)\right|^2 \left|\frac{H_2(\omega)}{H_1(\omega)}\right|^2 d\omega} \qquad \text{(eq 13)}$$

The risk is that if the crosstalk coupling mechanism is considerably closer to the receiver than the transmitter, then there will be a strong aggressor attacking a weak victim. In equation 12, this phenomenon would cause the factor

$$\frac{H_2(\omega)}{H_1(\omega)} \qquad \text{(eq 14)}$$

to have a magnitude much greater than one, with the result that the signal to crosstalk ratio could decrease quite a bit.

6.2 Results

Signal energy for each channel was calculated by integrating the square of the pulse response at the receiver. Similarly, the crosstalk energy from each aggressor was calculated by integrating the square of its pulse response into the victim channel. Since the data rates for all channels were the same, the signal to crosstalk ratio was the ratio of these two energies.

Table 2 contains the signal to crosstalk ratio from each aggressor to each victim. For the sake of later comparison, each entry in Table 2 shows the signal to crosstalk ratio including the effects of equalization and below that the signal to crosstalk ratio without equalization, as was promised in the previous section. The diagonal elements of Table 2 show that total signal to crosstalk ratio for the channel together with the statistical analysis bit error rate estimate for that channel.

Table 2: Signal to Crosstalk Ratios (dB)

From/To	Rx0	Rx1	Rx2	Rx3
Tx0	**26.0** **2.4e-25**	26.0 29.3	59.4 63.1	62.4 67.0
Tx1	26.1 29.3	**23.0** **1.9e-14**	26.0 29.3	44.8 49.2
Tx2	57.6 62.2	26.1 29.3	**21.7** **1.4e-10**	23.6 27.4
Tx3	77.3 81.2	58.8 61.8	23.5 27.2	**23.5** **4.8e-16**

Not surprisingly, there is a close correlation between the signal to crosstalk ratios in Table 2 and the bit error rate estimates in Table 1.

The greater apparent signal to crosstalk ratio for the unequalized case is due to the fact that equalization tends to reduce the response at low frequencies compared to frequencies near one half the data rate. Since crosstalk is typically lower at these low frequencies, equalization reduces the signal power more than it does the crosstalk power, resulting in a lower signal to crosstalk ratio. Note that the equalized signal to crosstalk ratio will be the more reliable predictor of system performance.

From Tables 1 and 2, 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 26 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.

6.3 Comparison to ICR

The ICR equations from IEEE 802.3ap [2] were applied to the transfer functions for the example system. Figure 13 shows the ICR (blue), fitted ICR (red) and ICR mask (black) for Rx2. This figure demonstrates that the ICR fitting equations are very effective with this data.

Figure 14 shows the fitted ICR for all four channels, along with the ICR mask. In Figure 14, none of the four channels in the example system are in the “high confidence” region of the ICR graph, and yet at least Tx0 to Rx0 and Tx3 to Rx3 are predicted to have a very low bit error rate. It therefore appears that the ICR mask is at least somewhat conservative.

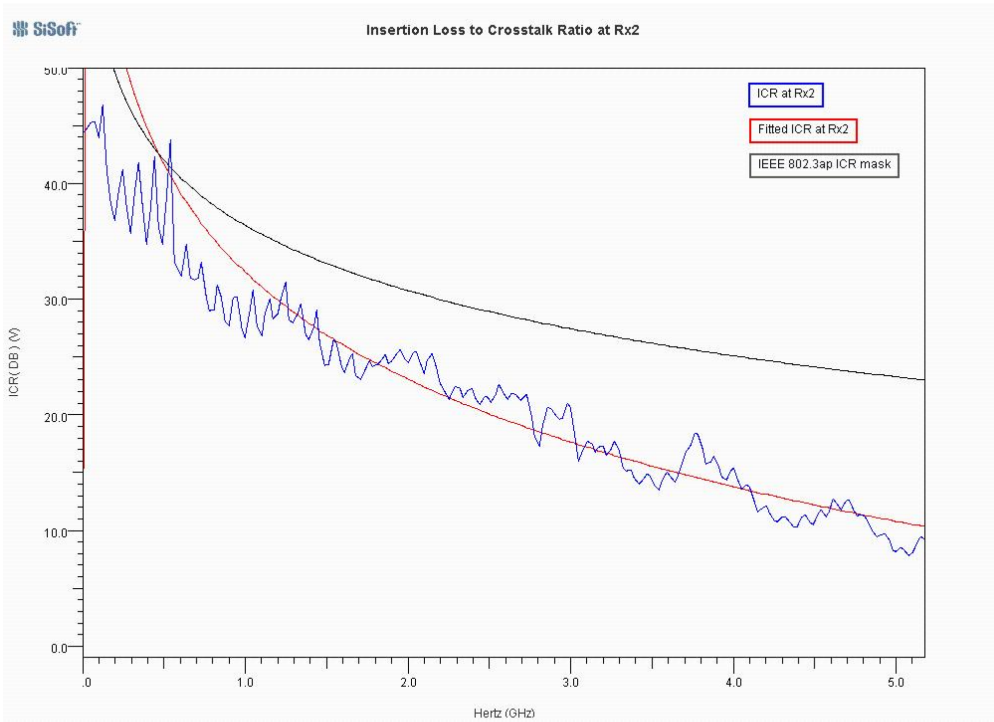

Figure 13: Raw and fitted ICR at Rx2

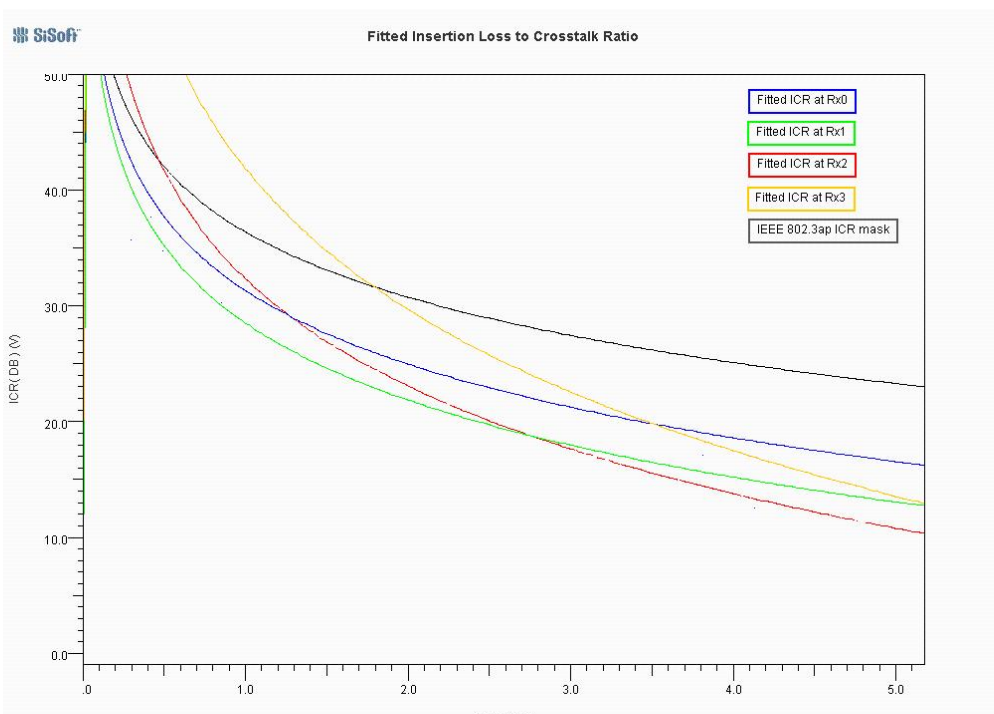

Figure 14: Fitted ICR and ICR mask for all four channels in the example system

7.0 Conclusions

This paper has presented estimates of the bit error rate of an example system using four different estimation methods: jitter-based analysis, explicit time domain simulation, semi-analytic estimation, and statistical analysis. The conclusions of this study are

1. While jitter-based analysis is excessively conservative, at least for the system analyzed, the other three methods all 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 statistical analysis method should be used whenever that method can be applied.
3. Insertion loss to crosstalk energy ratio is a reliable indicator of the effect of crosstalk on BER. An unequalized insertion loss to crosstalk ratio of at least 26dB for satisfactory BER is a useful rule of thumb.
4. 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.
5. 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] IEEE Draft P802.apTM/Draft 3.3, Annex 69B.4.6, copyright IEEE, January 26, 2007

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

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

[5] Gold and Rader, Digital Processing of Signals, McGraw Hill, copyright 1969, pg. 220, 226.

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

For more information on this subject, please contact us at:

Signal Integrity Software, Inc (SiSoft)

6 Clock Tower Place, Suite 250 – Maynard, MA 01754

www.sisoft.com

978.461.0449