Analysis of PLL Clock Jitter in High-Speed Serial Links

Pavan Kumar Hanumolu, *Student Member, IEEE,* Bryan Casper, *Member, IEEE,* Randy Mooney, *Member, IEEE,* Gu-Yeon Wei, *Member, IEEE,* and Un-Ku Moon, *Senior Member, IEEE*

***Abstract*—We analyze the effects of transmitter and receiver phased-locked loop (PLL) phase noise, which translates to time-domain clock/data jitter, on the performance of high-speed transceivers. Analytical expressions are derived to incorporate both transmitter and receiver clock jitter into serial link operation. A method to calculate the worst-case noise margin degradation due to clock jitter is discussed in order to obviate impractical time-domain simulations. This analysis relies on the assumption that the channel is linear and time-invariant and, hence, can be characterized by an impulse response. A simple extension to equalized serial links is also presented. The analysis is verified through behavioral simulations using a realistic/measured channel model.**

***Index Terms*—Bit-error rate (BER), eye diagrams, inter-symbol interference, jitter, phase-locked loops (PLL), phase noise, serial links.**

I. Introduction

ADVANCES in integrated circuit (IC) fabrication technology along with innovative circuit design techniques have led to very high-speed digital systems. These systems typically require high speed and efficient communication among multiple ICs. In the case of a large networking system, very high-speed chip-to-chip communication through a system backplane may be required. As the f_T of transistors increase with aggressive technology scaling, the off-chip (I/O) bandwidth can become the major performance bottleneck in the overall system. And as increasing data rates follow technology scaling, limited timing accuracy (i.e., time-domain clock/data jitter) that is bound by the unavoidable use of phase-locked loops (PLLS)and/or delay-locked loops (DLLs) can significantly degrade link performance. While the spectral content of the PLL phase noise is largely unimportant in these systems (unlike in RF/wireless systems), time-domain jitter, namely the probability density function (PDF) of time-domain jitter, plays an important role.[1]

Manuscript received May 1, 2003; revised July 2003. This work was supported by Intel Corporation, Hillsboro, OR. This paper was recommended by Guest Editor M. Perrott.

K. P. Hanumolu and U. Moon are with the Department of Electrical and Computer Engineering, Oregon State University, Corvallis, OR 97331-3211 USA (e-mail: moon@ece.orst.edu).

B. Casper and R. Mooney are with Intel Laboratories, Hillsboro, OR 97124 USA.

G.-Y. Wei is with the Department of Electrical Engineering, Harvard University, Cambridge, MA 02138 USA.

Digital Object Identifier 10.1109/TCSII.2003.819121

[1]The spectral content of the phase noise is still important for the receiver as timing recovery circuit can track low-frequency jitter. Because the scope of this paper is limited to the analysis of PLL jitter in serial links in general, we omit the discussion on this low-frequency jitter tracking of the timing recovery circuit. This kind of jitter tracking is generally considered a common knowledge by the design community.

Several high-speed serial link designs have been proposed recently [1]–[3]. In these systems, limited off-chip bandwidth introduces inter-symbol interference (ISI) [4] which reduces noise margins at the receiver. In addition, timing uncertainty (clock/data jitter) generated by the receiver PLL can lead to suboptimal sampling, thus, reducing the noise margin further. We provide an approach to thoroughly analyze the impact of PLL clock jitter on serial links to identify and understand weaknesses, to verify robustness, and to shed light on new techniques to overcome these problems. In the design phase, transceiver systems typically rely on time-domain simulations involving a long sequence of random data and the performance of serial links is often evaluated using eye diagrams of the received data. There are two problems with this traditional design approach. First, simulation time becomes prohibitively long to evaluate a near worst-case eye diagram. For example, for a serial link with an expected bit-error rate (BER) of 10^{-12}, the input random sequence should be at least 10^{12} long, and preferably, many times longer in order to get an accurate and reasonable statistical measure. Second, it is difficult to properly simulate these serial links with time-domain jitter contributions coming from clock sources at both ends (receiver and transmitter) of the link. In practice, several simplifying assumptions are made regarding the effect of clock jitter on the receive eye diagram. Using these assumptions, the eye diagram generated without clock jitter is modified to obtain an eye diagram with clock jitter. One common way to do this is by closing either side of the eye horizontally by the amount of peak clock jitter. While this method can be helpful in evaluating the effects of jitter at the receiver end, we will show in this paper that this is an overly optimistic approximation of noise margin degradation for transmitter jitter.

Due to the need for integration of clock generators such as PLLs in large digital chips, clock jitter is dominated by power supply and substrate noise, both of which do not scale with technology. Therefore, as data rates increase, bit periods become shorter and the performance of most multigigabit links will be limited by clock jitter. Therefore, it is important to analyze the effects of clock jitter on these high-speed serial links. Casper *et al.* [5] proposed an analysis method to generate a worst-case eye diagram due to ISI and crosstalk. This analysis used simplifying assumptions to accommodate clock jitter. In this paper, we propose an analytical method to incorporate time-domain clock jitter into the design of high speed serial links. This analysis is based on the assumption that jitter is small compared to the clock period. This assumption is valid for well-designed PLLs [6]–[8].

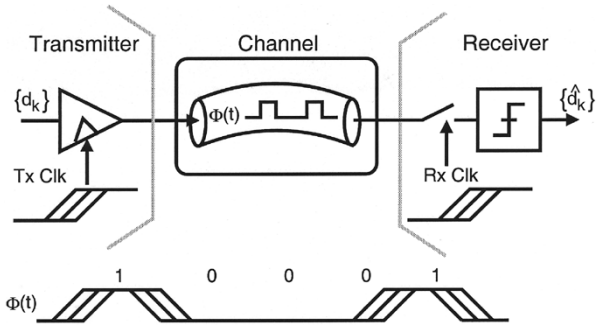

Fig. 1. Simplified block diagram of a serial link.

This paper is organized in the following manner. Section II gives a brief background on the problem of timing jitter in serial links; analysis of receiver sampling jitter is presented in Section III; and the detrimental effect of transmitter jitter on link performance is discussed in Section IV. Both transmitter and receiver jitter are treated together in Section V, and behavioral simulations verifying the analysis are presented in Section VI. Finally, important results are summarized in Section VII.

II. Serial Link Background

Serial links are predominantly used in high-speed chip-to-chip communication. A generalized model of a serial link is illustrated in Fig. 1. It consists of a transmitter which generates a train of pulses depending on the data symbols to be transmitted, and the pulse width is determined by the timing instant of the transmit clock at both begin/end edges. The receiver is generally a sampler followed by a decision circuit (e.g. high-gain comparator). Even though other receiver architectures based on an integrating amplifier [9] or sense amplifier [10] exist, sampler type receivers are used primarily due to their high-speed advantage [11] in multigigabit links. Communication channels for serial links are typically printed circuit board (PCB) traces or coaxial cables. While loss mechanisms can differ somewhat, linear and time-invariant approximations are valid for most physically realizable channels. Therefore, we characterize a channel by its impulse response. The channel impulse response is typically obtained from measuring the *S*-parameters via a network analyzer, by using field solvers such as advanced design system (ADS) [12], or with some postprocessing of a transient response of the channel measured using time domain reflectometry (TDR).

Nonreturn-to-zero (NRZ) pulses are commonly used as the basis function for discrete data transmission. The response of the channel to the NRZ pulse is defined as the pulse response and is traditionally used to analyze and model the effects of a channel on data transmission and also in the design of equalizers in the case of channels with large attenuation at the frequency of interest. Even though the pulse response is very useful for characterizing the ISI, we will find that it is very difficult to analyze the effects of PLL jitter (especially transmitter jitter) because a pulse is created/defined by two adjacent edges with jitter. Consider the serial link model shown in Fig. 1. Qualitatively, jitter in the transmit PLL can modulate the width of the transmitted NRZ data pulse. This modulation being random, the pulse response of the system displays a level of random variation in accordance with the jitter. This makes the usage of standard deterministic methods difficult. In the case of receiver sampling jitter, several approaches to estimate the signal-to-noise ratio (SNR) loss due to jitter have been proposed [13]. However, it is difficult to translate SNR loss to reduction of noise-margin or degradation of BER in the case of serial links. To circumvent these problems we present a unifying analysis to accommodate both transmitter and receiver sampling jitter to calculate the worst-case noise margin degradations.

Our analysis and discussions are formulated in the context of a two-level (single-bit-per-symbol) NRZ transceiver system, as this is the most common modulation scheme used in serial links today. Some recent implementations employ four-level NRZ signaling (i.e., PAM-4), which doubles the bits-per-symbol rate. While our analysis and conclusions can easily be transferred to this and a variety of other signaling systems, we stay with the common two-level (binary) NRZ signaling scheme to focus our investigations on how PLL jitter impacts transceiver performance.

III. Sampling Clock Jitter in Receiver

The sequence/polarity of bits (symbols) communicated to the receiver by the transmitter can be considered equally likely and independent of each other. We denote these bits by an independent and identically distributed (i.i.d.) sequence $\{d_k\}$. The transmitter produces an output pulse corresponding to data bit d_k and variation in the pulse width is determined by the transmitter clock jitter generated by a PLL.[2] The relationship between clock frequency (of the PLL) and data rate is determined by technology considerations [14]. We begin our analysis by focusing on the effects of jitter on the receiver end and assume that the transmitter clock is "jitter free." (Sections IV and V will consider the effects of jitter only at the transmitter and the combination of jitter on both transmit and receive clocks.) This means that the pulses corresponding to all data bits have equal width. With this assumption, the transmitted pulse train $\phi(t)$, in terms of data bit sequence $\{d_k\}$, can be written as [15]

$$\phi(t) = \sum_{k=-\infty}^{\infty} (d[kT] - d[kT - T]) \cdot u(t - kT) \qquad (1)$$

where T is equal to the bit period and $u(t)$ is the unit-step function such that $u(t) = 0$ for $t \leq 0$ and $u(t) = 1$ for $t > 0$. As mentioned earlier, the channel can be accurately characterized by an impulse response $h(t)$. The output of the channel, y(t), can be evaluated by convolving the input pulse train with the channel impulse response $h(t)$

$$\begin{aligned} y(t) &= \left[\sum_{k=-\infty}^{\infty} (d[kT] - d[kT - T]) \cdot u(t - kT)\right] \otimes h(t) \\ &= \sum_{k=-\infty}^{\infty} [(d[kT] - d[kT - T]) \cdot s(t - kT)] \end{aligned} \qquad (2)$$

where $s(t) = u(t) \otimes h(t)$ is the step response of the channel. A clock-data recovery circuit or a PLL locked to a source-syn-

[2]We generically refer to the clock generator/multiplier as the PLL while fully recognizing that a variety of different architectures (e.g., DLL) are feasible.

chronous clock generates a receiver clock phase that is aligned with the incoming data such that the voltage margin is maximized at the input of the detector. But due to various noise sources (intrinsic device and power supply noise), the receiver clock has jitter associated with each of its edges. This jitter is denoted by the jitter sequence $\{j_{rx}\}$ such that $j_{rx}[n]$ is the jitter associated with the n^{th} sampling edge. Note that we have not yet made any assumptions regarding the properties of the $\{j_{rx}\}$ sequence. With this framework, we can write the sampled channel output as

$$y(nT) = \sum_{k=-\infty}^{\infty} [(d[kT] - d[kT-T]) \cdot s(nT - kT + j_{rx}[nT])]. \tag{3}$$

The sampled step response can be approximated with a first-order Taylor series expansion.[3] It is reasonable to assume that this first-order approximation is valid for the case when $j_{rx}[n]$ is very small compared to bit period T. Therefore, an approximate sampled channel step response can be written as

$$\begin{aligned} s(nT - kT + j_{rx}[nT]) &\approx s(nT-kT) + j_{rx}[nT] \cdot \left.\frac{ds(t)}{dt}\right|_{t=nT-kT} \\ &= s(nT-kT) + j_{rx}[nT] \cdot h(nT-kT). \end{aligned} \tag{4}$$

Using (3) in (4), we can rewrite the sampled channel output as

$$\begin{aligned} y(nT) \approx & \sum_{k=-\infty}^{\infty} [(d[kT] - d[kT-T]) \cdot s(nT-kT)] \\ & + \sum_{k=-\infty}^{\infty} [(d[kT] - d[kT-T]) \\ & \qquad \cdot h(nT-kT)] \cdot j_{rx}[nT] \\ = & (d[nT] - d[nT-T]) \otimes s[nT] \\ & + [(d[nT] - d[nT-T]) \otimes h[nT]] \cdot j_{rx}[nT]. \end{aligned} \tag{5}$$

And rewriting the expression with just the n index

$$\begin{aligned} y[n] &= a[n] \otimes s[n] + (a[n] \otimes h[n]) \cdot j_{rx}[n], \\ &\text{where} \quad a[n] = d[n] - d[n-1]. \end{aligned} \tag{6}$$

The intermediate sequence $a[n]$ is introduced for notational brevity. The first term in (6), is the channel output obtained by sampling the continuous-time channel output with an ideal clock (i.e., no jitter) while the second term represents the equivalent *voltage* noise due to sampling jitter. Qualitatively, the second term in the first-order Taylor series translates the timing jitter into voltage noise depending on the slope of the step response at that instant. Because all practical/realistic channels used in multigigabit serial links are significantly bandwidth limited, the step response of the channel rises/falls quite slowly. This slow rise/fall translates to high accuracy of the first-order Taylor series. Thus, the explicit separation of the jitter noise from the signal in (6) enables us to evaluate worst-case distortion due to ISI and clock jitter independently.

[3]For practical/realistic channels with finite bandwidths, it is reasonable to assume that the first derivative of the step response exists.

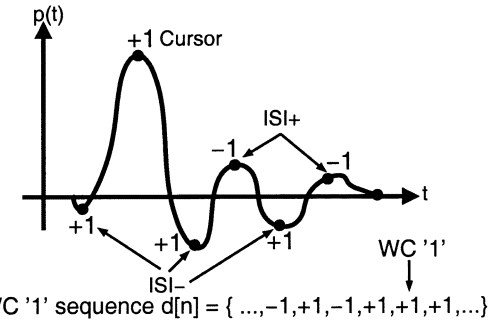

Fig. 2. Example of worst-case ISI calculation using pulse response.

A. Peak Distortion Analysis

The distortion at the detector input is due to ISI and jitter noise. The worst-case distortion (positive pulse example is illustrated) due to ISI alone is given in (7) [5]

$$\begin{aligned} \text{Worstcase ISI noise} &= \sum |ISI_+| + \sum |ISI_-| \\ &= \sum ISI_+ - \sum ISI_- \\ &= \sum_{k=-\infty}^{\infty} y(t-kT)|_{y(t-kT)>0, k\neq 0} \\ &\quad - \sum_{k=-\infty}^{\infty} y(t-kT)|_{y(t-kT)<0, k\neq 0} \end{aligned} \tag{7}$$

where ISI_- and ISI_+ are negative and positive ISI terms. "Worst-case ISI noise" denotes the maximum ISI distortion experienced by the transmitted pulse. A sample pulse response and the corresponding positive and negative ISI terms are shown in Fig. 2. The figure also illustrates the data sequence that causes the worst-case noise. The time-reversal inherent in the convolution can be accounted for by simply reading the sequence from right to left.

In the case of distortion introduced by clock jitter, the worst-case condition can be evaluated by observing the effect of jitter due to the worst-case ISI data pattern illustrated in Fig. 2. The corresponding jitter noise can be evaluated using the second term of (6), $(a[n] \otimes h[n]) \cdot j_{rx}[n]$, by

$$\text{Receiver Jitter Noise} = (\hat{a}[n] \otimes h[n]) \cdot max\,(j_{rx}[n]) \tag{8}$$

where $\hat{a}[n]$ is the worst-case/peak ISI distortion data sequence derived using (7). A specific example of this is illustrated later in Section VI.

B. Sampling Jitter With Linear Receive Equalizer

In the case of severely ISI-limited channels, equalization is used to recover some of the high frequency lost through the channel. An equalizer is typically a filter which inverts the channel response so that the overall response is essentially flat in the band of interest (up to the Nyquist rate of the data), thus, reducing the effects of ISI. There are several issues associated with the design of equalizers and we refer interested readers to [16] for further information. Typically, these equalizers rely on finite impulse response (FIR) filters and are implemented in the analog current-mode domain due to technology limitations [18]. A transceiver model with the linear receive equalizer W,

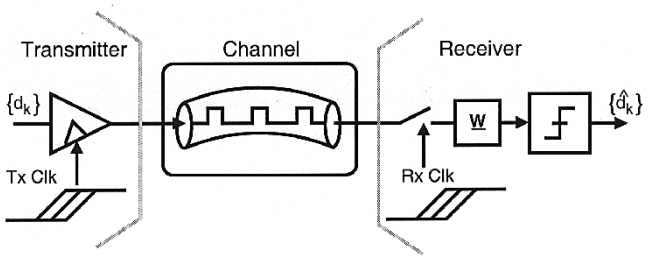

Fig. 3. A simplified block diagram of a serial link with equalizer W.

is shown in Fig. 3. The detector input is simply the sampled channel output convolved with the filter W. This is given by

$$y_{eq}[n] = a[n] \otimes s[n] \otimes W[n] + \{(a[n] \otimes h[n]) \cdot j_{rx}[n]\} \otimes W[n]. \tag{9}$$

The worst-case jitter noise and ISI data patterns can be calculated in a similar way as shown earlier in the case without equalization.

IV. Transmitter Clock Jitter

The transmitter clock determines the pulse width of the transmitted bit or symbol. With transmitter clock jitter, the pulse width of the transmitted data bit can be viewed as being modulated by the jitter.[4] This causes degradation of noise margin at the detector input for the following reasons. First, transmitter clock jitter causes suboptimal sampling at the receiver due to the limited tracking bandwidth of the timing-recovery loop. Second, in the case of equalized serial links, the transmitter jitter degrades equalizer performance. This is because the equalizers are normally optimized for a specific pulse response. Even in the case of adaptive equalizers, the high frequency content of the jitter cannot be tracked due to typically large time constants of adaptation algorithms [17].

We will now show that the transmit jitter can be analyzed in a similar framework as shown for receiver sampling clock jitter previously in Section III. Consider (2) repeated below for convenience

$$\begin{aligned} y(t) &= \left[\sum_{k=-\infty}^{\infty} (d[kT] - d[kT-T]) \cdot u(t-kT) \right] \otimes h(t) \\ &= \sum_{k=-\infty}^{\infty} [(d[kT] - d[kT-T]) \cdot s(t-kT)]. \end{aligned}$$

In this equation, the sampling instant kT determines the pulse width of the k^{th} transmitted data pulse/bit. The jitter in the transmitter can be included in the above equation by defining a jitter sequence $\{j_{tx}\}$ such that $j_{tx}[k]$ is the jitter associated with the k^{th} clock edge

$$\begin{aligned} y(t) &= \left[\sum_{k=-\infty}^{\infty} (d[kT] - d[kT-T]) \cdot u(t-kT-j_{tx}[kT]) \right] \otimes h(t) \\ &= \sum_{k=-\infty}^{\infty} [(d[kT] - d[kT-T]) \cdot s(t-kT-j_{tx}[kT])]. \end{aligned} \tag{10}$$

[4]To be precise, the pulse width as well as the position of the pulse are modulated by the jitter. This is because the transmit pulse is affected by the clock jitter on both begin/end edges of the transmit pulse.

Again, first-order Taylor series expansion can be used if $j_{tx}[k] \ll T$, and the approximate channel output can be written as

$$\begin{aligned} y(t) \approx & \sum_{k=-\infty}^{\infty} [(d[kT] - d[kT-T]) \cdot s(t-kT)] \\ & + \sum_{k=-\infty}^{\infty} [(d[kT] - d[kT-T]) \cdot h(t-kT) \cdot j_{tx}[kT])]. \end{aligned} \tag{11}$$

In order to estimate the effects of transmitter clock jitter alone, let us assume for now that the receive sampling clock is jitter free. In that case, the sampled channel output can be written as

$$\begin{aligned} y(nT) &= \sum_{k=-\infty}^{\infty} [(d[kT] - d[kT-T]) \cdot s(nT-kT)] \\ &\quad + \sum_{k=-\infty}^{\infty} [(d[kT] - d[kT-T]) \cdot h(nT-kT) \cdot j_{tx}[kT])] \\ &= \sum_{k=-\infty}^{\infty} [a[kT] \cdot s(nT-kT)] \\ &\quad + \sum_{k=-\infty}^{\infty} [(a[kT] \cdot j_{tx}[kT]) \cdot h(nT-kT)]. \end{aligned} \tag{12}$$

And rewriting this first-order approximated output expression with just the n index

$$y[n] = a[n] \otimes s[n] + (a[n] \cdot j_{tx}[n]) \otimes h[n]. \tag{13}$$

Unlike the receiver sampling jitter of (6), the transmitted data difference sequence $a[n]$ is first modulated by transmitter jitter sequence $j_{tx}[n]$ and then the resulting sequence is convolved with the channel's impulse response $h[n]$.

A. Peak Distortion Analysis

Once again, the peak ISI distortion inherent in the first convolution term in (13) can be calculated using (7). However, the peak distortion due to the transmitter jitter noise is different from that of the receiver sampling jitter. Intuitively, we expect the transmit jitter to be filtered by the channel in some fashion and the second term in (13) reinforces our intuition. Due to the modulation of $a[n]$ by the jitter sequence $j_{tx}[n]$, we can evaluate the peak distortion due to transmitter clock jitter, i.e., the peak distortion of the second term $(a[n] \cdot j_{tx}[n]) \otimes h[n]$, by

$$\text{Transmit Jitter Noise} = (|\hat{a}[n]| \cdot max\,(j_{tx}[n])) \otimes |h[n]| \tag{14}$$

where $\hat{a}[n]$ is the worst-case/peak ISI distortion data sequence derived using (7). It is interesting to note that the peak distortion due to transmitter clock jitter noise can be potentially greater than that of receiver sampling jitter for the similar amount of receiver (j_{rx}) and transmitter (j_{tx}) jitter.

B. Transmitter Jitter With Linear Receive Equalizer

As in the analysis of receiver sampling jitter with receiver equalizer, the transmitter's clock jitter analysis can also be ex-

tended to the serial link with receive equalizer. The input to the detector, with the receive equalizer, is simply the convolution of raw channel output and the FIR equalizer

$$\begin{aligned} y_{\text{eq}}[n] &= \{a[n] \otimes s[n] + (a[n] \cdot j_{tx}[n]) \otimes h[n]\} \otimes W[n] \\ &= a[n] \otimes s[n] \otimes W[n] \\ &\quad + (a[n] \cdot j_{tx}[n]) \otimes h[n] \otimes W[n]. \end{aligned} \tag{15}$$

Equalization is also sometimes used on the transmit side and this is often referred to as transmit pre-emphasis. Transmit pre-emphasis shapes the transmitted pulse so as to make the channel response flat up to the Nyquist rate of the data.[5] Even though this type of equalization is transmit power limited [16], it is commonly applied [19], [20] because of its relatively simple implementation in comparison to building an even more extensive receive equalizer. The jitter analysis for serial links with transmit equalizer/pre-emphasis directly follows from (6) and (13) with a corresponding equalized/pre-emphasized data sequence $\{d_k\}$.

V. Transmitter Jitter and Receiver Jitter

We analyzed transmitter jitter and receiver sampling jitter independently until now. We did this to demonstrate the effect of each of the jitter terms independently. Because both effects of jitter typically appear together in a serial link, we now summarize how the above analysis can be extended to include both the transmitter and receiver jitter. Equation (10) defines the channel output with transmitter jitter and (3) was derived to consider receive sampling jitter. Combining the results of (10) and (3), we can rewrite the sampled channel output which incorporates both of the jitter terms

$$\begin{aligned} y(nT) = \sum_{k=-\infty}^{\infty} &[(d[kT] - d[kT - T]) \\ &\cdot s(nT - kT + j_{rx}[nT] - j_{tx}[kT])]. \end{aligned} \tag{16}$$

Once again, we can approximate the step response using first-order Taylor series approximation for two variables (i.e., when $j_{tx}[k] \ll T$ and $j_{rx}[k] \ll T$)

$$\begin{aligned} &s\left(nT - kT + j_{rx}[nT] - j_{tx}[kT]\right) \\ &\quad \approx s(nT - kT) + j_{rx}[nT] \cdot \left.\frac{ds(t)}{dt}\right|_{t=nT-kT} \\ &\qquad - j_{tx}[kT] \cdot \left.\frac{ds(t)}{dt}\right|_{t=nT-kT} \\ &\quad = s(nT - kT) + j_{rx}[nT] \cdot h(nT - kT) \\ &\qquad + j_{tx}[kT] \cdot h(nT - kT). \end{aligned} \tag{17}$$

Putting (16) and (17) together, we can write the channel output as

$$y[n] \approx a[n] \otimes s[n] + (a[n] \otimes h[n]) \cdot j_{rx}[n] + (a[n] \cdot j_{tx}[n]) \otimes h[n]. \tag{18}$$

Since we did not use any specific properties of the jitter sequence, (18) is valid for any jitter sequences $j_{tx}[n]$ and $j_{rx}[n]$. The correlation between $j_{tx}[n]$ and $j_{rx}[n]$, if any, is determined by the clocking scheme and the system. By defining the jitter sequences accordingly, the trade-offs between various clocking schemes (e.g. mesochronous, source synchronous, and embedded clocking [4]) can be analyzed using (18). The properties of individual jitter sequence depend on the type of clock source used and the system architecture of the serial link. In most situations, it would be reasonable to assume for the worst case that the transmitter and the receiver jitter properties are uncorrelated. However, any amount of observed correlation between the transmit and receive jitter would result in an overall improvement of the system.

[5]In practice, the low-frequency content of the pulse is attenuated due to the limited transmitter power.

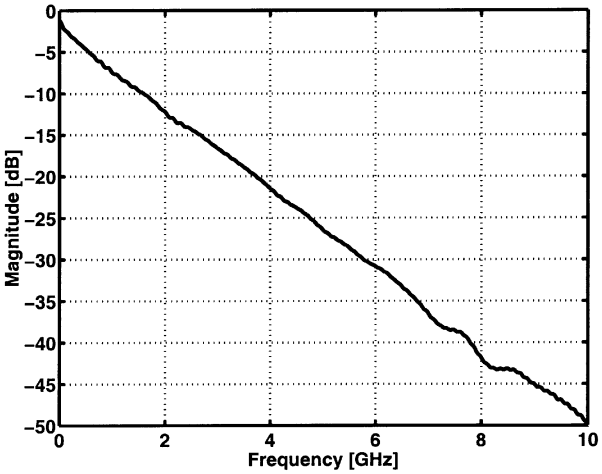

Fig. 4. Magnitude response.

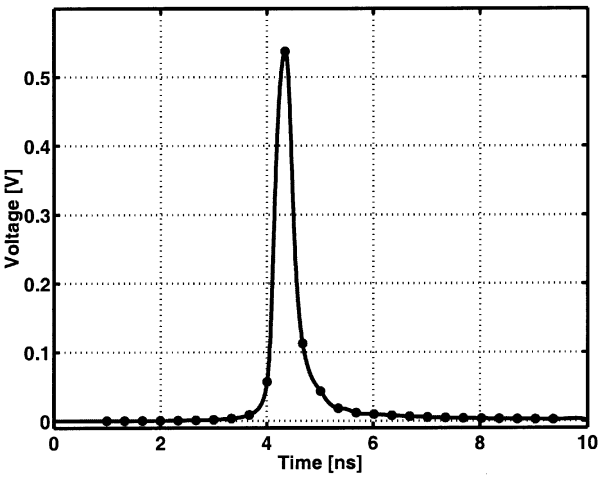

Fig. 5. Pulse response to a 333 ps (3 Gb/s) pulse.

VI. Behavioral Simulations

The analysis presented thus far is verified using behavioral simulations in MATLAB. The channel model used is a 20″ FR4 trace with two connectors. The impulse response is derived from measured differential scattering parameters. The transmitter and receiver bandwidths are modeled using single pole transfer functions. The effective frequency response and the pulse response for a 3-Gb/s pulse are shown in Figs. 4 and 5, respectively. As is clear from the pulse response, this particular channel does not require any equalization for 3-Gb/s data transmission. Because the key idea here is to demonstrate

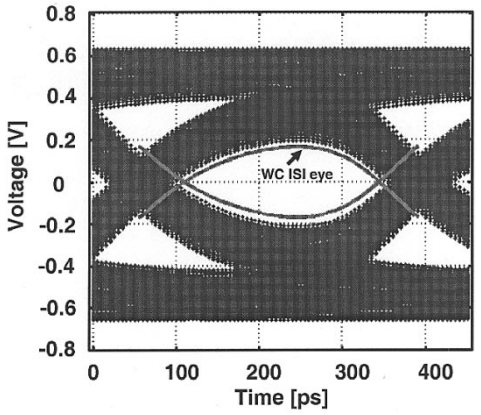

Fig. 6. Eye diagram without jitter simulated with 2000 data bits.

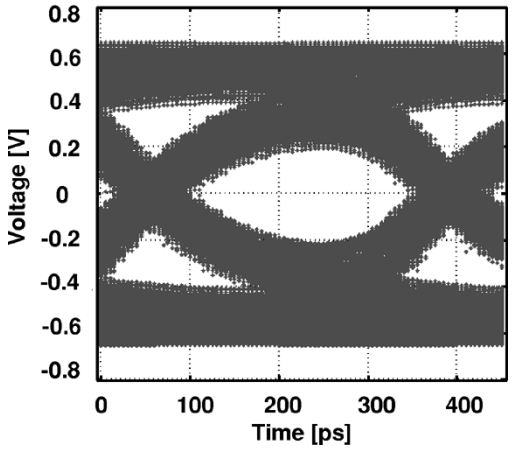

Fig. 7. Calculated eye diagram using (6).

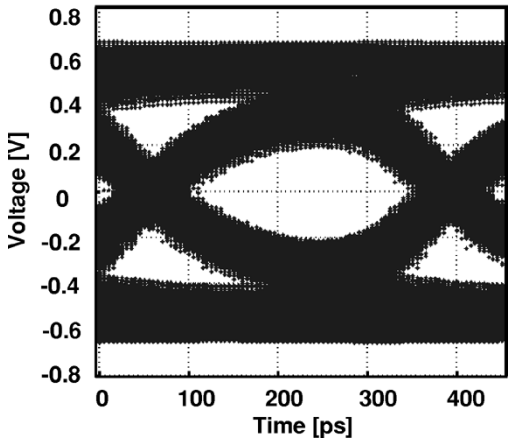

Fig. 8. Simulated eye diagram with 2000 random bits.

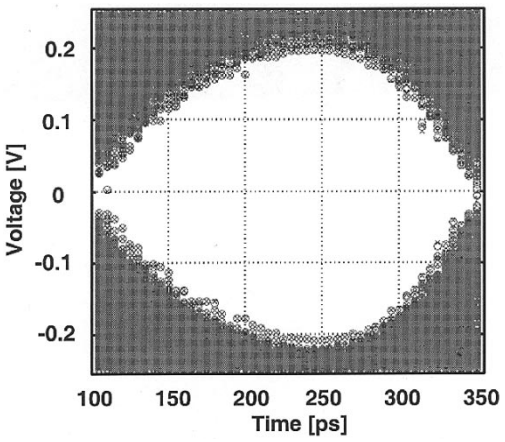

Fig. 9. Calculated (marker 'o') and simulated (marker 'x') zoomed-in eye diagrams with 2000 data bits.

the accuracy of the presented analysis, for simplicity, the jitter for both the transmitter and receiver is assumed to be white Gaussian noise. We refer the interested readers to [21] and [22] for a more detailed model of PLL generated jitter.

Shown in Fig. 6 is a simulated receiver eye diagram without sampling jitter. The worst-case eye diagram calculated using (7) is superimposed for comparison. It can be seen that the simulated eye diagram approaches the worst-case eye with a data stream length of 2000 bits. Figs. 7 and 8 illustrate the eye diagrams in the presence of 5 ps rms receiver sampling jitter. Fig. 7 is the "calculated" eye diagram, by which we mean the eye generated using (6) for a data stream length of 2000 data bits. On the other hand, the simulated eye diagram shown in Fig. 8 is generated using the same data and jitter sequence in a *time-domain* simulation.[6] The calculated and simulated eye diagrams are overlayed in order to compare them in Fig. 9. Note that the simulated and calculated eye diagrams are virtually identical, meaning the error in the first-order Taylor series approximation is acceptable.[7]

We now present simulation results that demonstrate noise margin degradation due to receiver clock jitter. The simulated opening of the eye diagram with 30 k data bits[8] and 5 p/s rms (σ) jitter is shown in Fig. 10. Note that the eye is still wide-open due to only a limited amount of data bits used in the simulation. We now evaluate the receiver eye based on the analysis presented in the preceding sections. First, the worst-case ISI data pattern is calculated using (7) and jitter noise generated due to this data pattern is calculated using (8). The worst-case eye is then obtained by subtracting the jitter noise from the worst-case ISI eye. The calculated worst-case eye diagrams using 3σ and 7σ amounts of peak jitter are also shown in Fig. 10. The noise margin degradation is minimal at the center of the eye and maximum near the zero-crossing. This makes intuitive sense because the center of the eye is reasonably flat (slope is zero) and, hence, any jitter at the optimal sampling point only results in a small voltage-margin degradation. However, due to the larger slope at the edges, jitter translates to larger voltage margin degradation at the edge of the eye. Also, notice that even with 30 k data bits, the simulated eye is not close to the calculated worst-case eye even with 3σ jitter. This reinforces the fact that it is generally very difficult to find the absolute worst-case margin from time-domain simulations. For this reason, time-domain simulation is seldom used to estimate BER in practice. Commonly used

[6]In a time-domain simulation the data sequence is convolved with the impulse response and the output is sampled appropriately to maximize the voltage margin. Jitter is incorporated by randomly varying the optimal sampling point.

[7]First-order Taylor series approximation can be directly validated by comparing the approximate step response with actual step response. However, it is not clear how this error translates to the error in the eye diagram.

[8]30 k bits as opposed to say 2 million bits are chosen for simulation as a direct tradeoff between long data stream versus long simulation time.

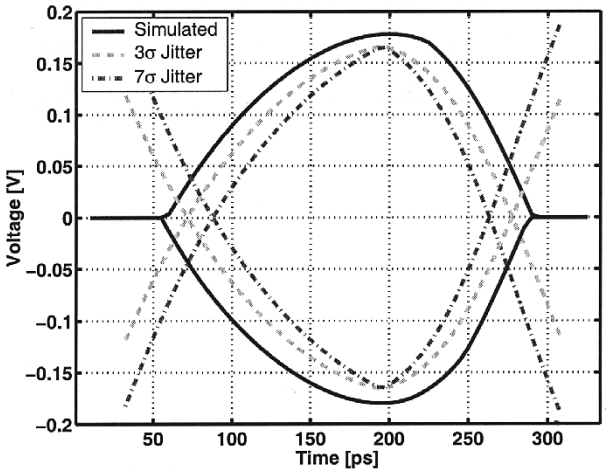

Fig. 10. Simulated eye diagram with 30 k bits and worst-case eye diagram for 3σ and 7σ values of peak receiver jitter.

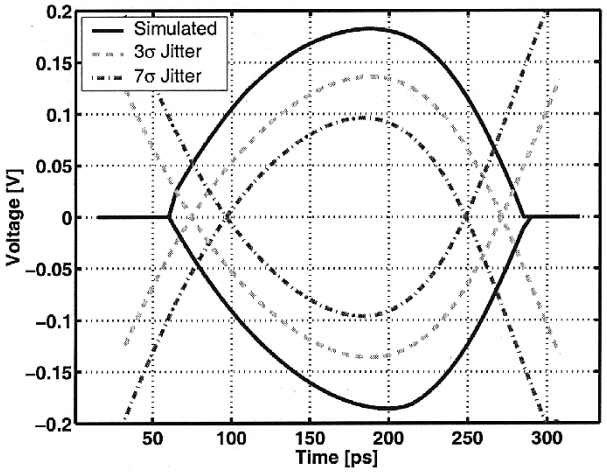

Fig. 11. Simulated eye diagram with 30 k bits and worst-case eye diagram for 3σ and 7σ values of peak transmitter jitter.

methods incorporate the effects of jitter into the worst-case ISI eye by shifting the ISI eye edges horizontally toward the center of the eye by the peak jitter amount. Even though this method results in a worst-case eye, it provides little insight and is not applicable to the transmitter jitter.

Similar to the receiver sampling jitter case, the simulated and calculated worst-case eye diagrams with transmitter jitter are shown in Fig. 11. Again, the simulated eye is not close to the calculated worst-case eye even with 3σ jitter. It is interesting to note that the noise margin degradation due to transmitter jitter is severe all across the eye unlike the receiver jitter case, where degradation is minimal at the center of the eye and maximum near the zero-crossing. This is consistent with (15), which showed that the jitter is shaped along with the data symbols by the band-limited channel. The calculated eye-diagram incorporating both the transmitter and receiver jitter is shown in Fig. 12. Eye diagrams calculated using zero jitter (i.e., only worst-case ISI), transmitter jitter alone, and receiver jitter alone, are also shown. It is clear that the transmitter and receiver jitter degrade both the voltage margin and timing margin. However, the transmitter jitter has more adverse effect on both the voltage and timing margins.

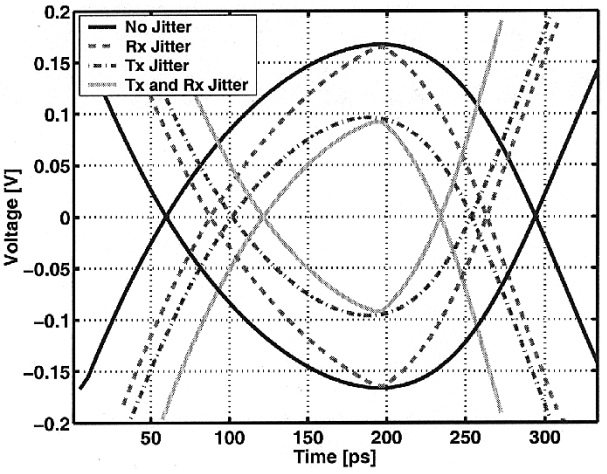

Fig. 12. Comparison of the effect of transmitter (Tx) and receiver (Rx) jitter on the worst-case ISI eye. 7σ peak transmitter jitter and 7σ peak receiver sampling jitter is assumed.

VII. Conclusion

The analysis and net effects of receiver and transmitter clock jitter on high-speed serial links are presented. In particular, the effect of transmitter clock jitter and receiver sampling jitter on the worst-case ISI condition is analyzed. Based on the linear time-invariant assumptions of the channel and using the first-order Taylor series approximation, analytical expressions representing the detector input for various conditions are derived. The noise due to jitter was decoupled from the expression of the channel output without jitter. This enables efficient calculation of the noise margin degradation due to jitter. Mathematical expressions useful for calculating the receive and transmit jitter degradations are summarized. Behavioral simulations indicate a good match between the calculation and simulation. This analysis enables efficient calculation of the worst-case margin without indulging in prohibitively large simulation.

References

[1] B. Casper, A. Martin, J. Jaussi, J. Kennedy, and R. Mooney, "8 Gb/s differential simultaneous bidirectional link with 4 mV 9 ps waveform capture diagnostic capability," in *ISSCC Dig. Tech. Papers*, Feb. 2003, pp. 78–79.

[2] B. Lee, M. Hwang, S. Lee, and D. Jeong, "A 2.5–10 Gb/s CMOS transceiver with alternating edge sampling phase detection for loop characteristic stabilization," in *ISSCC Dig. Tech. Papers*, Feb. 2003, pp. 76–77.

[3] R. Farjad-Rad, C. Yang, and M. Horowitz, "A 0.4-μm CMOS 10-Gb/s 4-PAM pre-emphasis serial link transmitter," *IEEE J. Solid-State Circuits*, vol. 34, pp. 580–585, May 1999.

[4] W. Dally and J. Poulton, *Digital Systems Engineering*. Cambridge, U.K.: Cambridge Univ. Press, 1998.

[5] B. Casper, M. Haycock, and R. Mooney, "An accurate and efficient analysis method for multi-Gb/s chip-to-chip signaling schemes," in *IEEE VLSI Circuits Symp. Tech. Papers*, Jun. 2002, pp. 54–57.

[6] J. Maneatis, "Low-jitter process-independent DLL and PLL based on self-biased techniques," *IEEE J. Solid-State Circuits*, vol. 31, pp. 1723–1732, Nov. 1996.

[7] I. Novof, J. Austin, R. Kelkar, D. Strayer, and S. Wyatt, "Fully integrated CMOS phase-locked loop with 15 to 240 MHz locking range and ± 50 ps jitter," *IEEE J. Solid-State Circuits*, vol. 30, pp. 1259–1266, Nov. 1995.

[8] I. Young, J. Greason, and K. Wong, "PLL clock generator with 5 to 10 MHz of lock range for microprocessors," *IEEE J. Solid-State Circuits*, vol. 27, pp. 1599–1607, Nov. 1992.

[9] S. Sidiropoulos and M. Horowitz, "A 700 Mbps/pin CMOS signaling interface using current integrating receivers," in *Proc. VLSI Circuits Symp.*, June 1996, pp. 142–143.
[10] M.-J. E. Lee, W. Dally, and P. Chiang, "A 90 mW 4 Gb/s equalized I/O circuit with input offset cancelation," in *ISSCC Dig. Tech. Papers*, Feb. 2000, pp. 252–253.
[11] H. Johansson and C. Svensson, "Time resolution of NMOS sampling switches used on low-swing signals," *IEEE J. Solid-State Circuits*, vol. 33, pp. 237–245, Feb. 1998.
[12] "Advanced design systems," in *ADS Design Guides*: Agilent Technologies, 2000.
[13] M. Shinagawa, Y. Akazawa, and T. Wakimoto, "Jitter analysis of high-speed sampling systems," *IEEE J. Solid-State Circuits*, vol. 25, pp. 220–224, Feb. 1990.
[14] C. Yang and M. Horowitz, "A $0.8\ \mu$m CMOS 2.5 Gb/s oversampling receiver and transmitter for serial links," *IEEE J. Solid-State Circuits*, vol. 31, pp. 2015–2023, Dec. 1996.
[15] H. Tao, L. Toth, and J. Khoury, "Analysis of timing jitter in bandpass sigma-delta modulators," *IEEE Trans. Circuits Syst. II*, vol. 46, pp. 991–1001, Aug. 1999.
[16] E. Lee and D. Messerschmitt, *Digital Communication*. Norwell, MA: Kluwer, 1999.
[17] J. Proakis, *Digital Communications*. New York: McGraw-Hill, 2000.
[18] R. Farjad-Rad, C. Yang, and M. Horowitz, "A 0.3-μ CMOS 8-Gb/s 4-PAM serial link transceiver," *IEEE J. Solid-State Circuits*, vol. 35, pp. 757–764, May 2000.
[19] A. Fiedler, R. Mactaggart, J. Welch, and S. Krishnan, "A 1.0625 Gbps transceiver with 2x-oversampling and transmit signal pre-emphasis," in *ISSCC Dig. Tech. Papers*, Feb. 1997, pp. 238–239.
[20] R. Gu, J. Tran, H. Lin, A. Yee, and M. Izzard, "A 0.5–3.5 Gb/s low-power low-jitter serial data CMOS transceiver," in *ISSCC Dig. Tech. Papers*, Feb. 1999, pp. 352–353.
[21] N. Da Dalt, M. Harteneck, C. Sandner, and A. Wiesbauer, "Numerical modeling of PLL jitter and the impact of its nonwhite spectrum on the SNR of sampled signals," in *Proc. Southwest Symp. Mixed-Signal Design*, 2001, pp. 38–44.
[22] ——, "On the jitter requirements of the sampling clock for analog-to-digital converters," *IEEE Trans. Circuits Syst. I*, vol. 49, pp. 1354–1360, Sept. 2002.

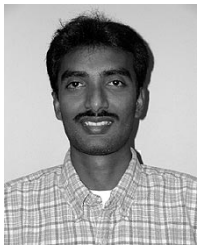

Pavan Kumar Hanumolu (S'99) received the B.E. degree (Hons.) in electrical and electronics engineering and the M.Sc. degree (Hons.) in mathematics from the Birla Institute of Technology and Science, Pilani, India, in 1998, and received the M.S. degree in electrical and computer engineering from the Worcester Polytechnic Institute, Worcester, MA, in 2001. He is currently working toward the Ph.D. degree in electrical engineering at the Oregon State University, Corvallis.

From 1998 to 1999, he was a Design Engineer with Cypress Semiconductor, India, working on phase locked loops for low-voltage differential signaling (LVDS) interfaces. His current research interests include equalization, clock and data recovery for high-speed I/O interfaces, data converters and low-voltage mixed-signal circuit design.

Bryan Casper (S'97–M'98) received the B.S. and M.S. degrees in electrical engineering from Brigham Young University, Provo, UT.

In 1998, he joined Intel Labs, Hillsboro, OR, where he is currently a Circuits Researcher involved in research, design, validation, and characterization of high-speed mixed signal circuits and I/O systems. He is married with two children.

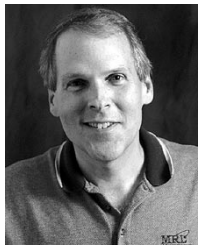

Randy Mooney (M'88) received the M.S.E.E. degree from Brigham Young University, Provo, UT.

He is currently a Senior Principal Engineer with Intel Labs, Hillsboro, OR, leading the Microprocessor Research Labs (MRLs) high-speed signaling team. From 1980 to 1992, he was with the Standard Products Division, Signetics Corporation, Orem, UT, where he developed products in Bipolar, CMOS, and BiCMOS technologies, concentrating on components for use in bus driving applications. In 1992, he joined the Supercomputer Systems Division, Intel Corporation, Hillsboro, OR, where he was involved in the development of interconnect components for parallel processor communications; responsible for the development of signaling technology for these components; and developed a method of simultaneous bidirectional signaling that was used for the Intel Teraflops Supercomputer. His current work is focused on multigigabit, differential, serial stream based interfaces, and pt-to-pt memory interfaces.

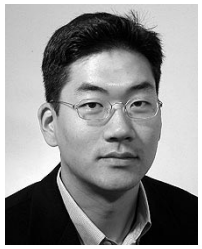

Gu-Yeon Wei (S'97–M'01) received the BS, MS, and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1994, 1997, and 2001, respectively.

In August 2000, he joined Accelerant Networks, Inc., Portland, OR, where he was a Member of the design team for a 5-Gb/s backplane transceiver. Since January 2002, he has been an Assistant Professor in the Department of Electrical Engineering, Harvard University. His research interests include high-speed low-power link designs; mixed-signal circuits for communications; low-noise circuits for bio-sensor applications; and energy-efficient design strategies for future advanced CMOS technolgies.

Un-Ku Moon (S'92–M'94–SM'99) received the B.S. degree from University of Washington, Seattle, the M.Eng. degree from Cornell University, Ithaca, New York, and the Ph.D. degree from the University of Illinois, Urbana-Champaign, all in electrical engineering, in 1987, 1989, and 1994, respectively.

From February 1988 to August 1989, he was a Member of Technical Staff at AT&T Bell Laboratories, Reading, PA. From August 1992 to December 1993, he taught a course in microelectronics at the University of Illinois. From February 1994 to January 1998, he was a Member of Technical Staff at Lucent Technologies Bell Laboratories, Allenton, PA. Since January 1998, he has been with Oregon State University, Corvallis. His research interests include analog and mixed analog-digital integrated circuits. His past work includes highly linear and tunable continuous-time filters, telecommunication circuits including timing recovery and analog-to-digital converters, and switched-capacitor circuits.

Prof. Moon received the National Science Foundation CAREER Award in 2002, and the Engelbrecht Young Faculty Award from Oregon State University, College of Engineering, in 2002. He has been an Associate Editor of the IEEE TRANSACTIONS ON CIRCUIT AND Systems—II: ANALOG AND DIGITAL SIGNAL PROCESSING since January 2003, and serves as a Member of the IEEE Custom Integrated Circuits Conference Technical Program Committee and the Analog Signal Processing Program Committee of the IEEE International Symposium on Circuits and Systems.