Modeling and Analysis of High-Speed Links

Vladimir Stojanovic[1,2] and Mark Horowitz[1]
[1]Stanford University, Stanford, CA
[2]Rambus Inc, Los Altos, CA

Abstract

Early work on high-speed links was focused on building CMOS components that could generate, receive, and recover timing of high-speed data. This work rapidly improved data rates, so today's circuits are running into the bandwidth limitations of the electrical wires. The very high symbol rates and very low bit error rate (BER) requirements mean that care is required before applying standard techniques for dealing with band-limited channels. This paper presents the analysis of noise sources, both voltage and PLL timing jitter, in a high-speed link system and their impact on the choice and effectiveness of different communication techniques. We show that the dominant noise sources are colored and bounded, as opposed to standard unbounded Gaussian white noise assumptions. With very low BER requirements, shape of noise probability distributions and their correlations are much more important than just their total power, which contrasts the standard analysis in communication systems.

I. Introduction

During the 90's, the continued scaling of chip performance fostered a large research and development effort in high-speed chip I/O design. This work investigated improving the performance of both high-speed serial links and wider parallel buses. The results of this research are everywhere, and range from high-speed memory interfaces [1,2], and processor interfaces [3], to Gb/s serial links [4,5]. This improvement in I/O performance (which scaled faster than processor frequency) has led to expectations of continued improvements in I/O rates. Unfortunately, the nature of the I/O design problem is changing. Today internal circuits can run at 10's of Gb/s, but the performance of the link is limited by the bandwidth of the channel – the electrical path from one die to the other. The obvious question now is how to continue to scale I/O performance, and what, if anything, will ultimately limit pin bandwidth.

While techniques to deal with band-limited channels are well known, having been used for problems ranging from Ethernet and DSL to disk read channels, tight power budgets and device speed limitations impose significant challenges to implementation of these standard communication algorithms. This paper is an attempt to bridge the gap between high-speed link design and high-speed communication system design. We apply analysis and measurement techniques used in communication system design to the unique problems posed by high-speed channel-limited link design. By analyzing the specific properties of the high-speed link system, and by classifying and analyzing the noise sources, we are able to decide where communication techniques can be applied most cost effectively and how to apply them.

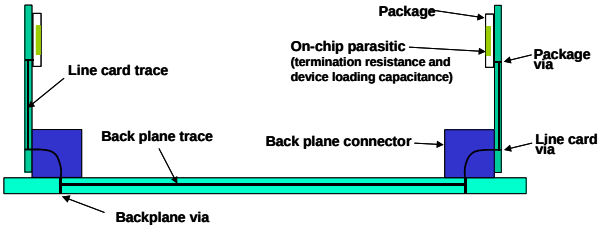

Fig. 1. High-speed system backplane channel [6]. Chips are mounted in packages, which are soldered to boards that plug into the backplane. The channel is the complete path from one die to the other die.

The complexity and performance of the channel depends strongly on application. In this paper we will use a backplane link as our design example, although the analysis method we develop can be applied to any link design. A typical backplane link is shown in Fig. 1, [6], illustrating different channel components. Different trace lengths and backplane material properties, as well as types of connectors, vias and routing layers, cause significant variation in channel transfer function both among different boards and among channels in the same backplane. The typical transfer functions for channels within a single backplane are illustrated in Fig. 2.

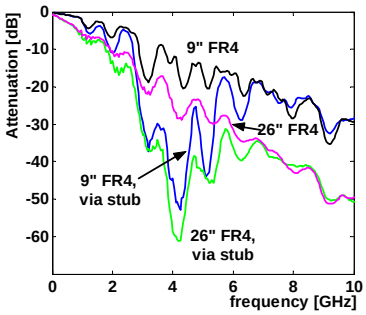

Fig. 2. Frequency response for backplane channels of two lengths, 9" and 26", and for top and bottom traces (i.e. with and without via stub). Note the large notch in the frequency response for the top traces even when the lines are short. Reflections caused by the via stub capacitance cause this notch.

While the channel bandwidth does limit the link's performance, the channel limitation at a few GHz is not that much slower than the speed limitations of the underlying CMOS technology. In this speed and power constrained environment, a typical hierarchical structure with analog

Research supported by the MARCO Interconnect Focus Center and Rambus, Inc.

front-end and digital signal processing back-end, proves to be very inefficient, if not impossible to implement.

To decide what are the most power/cost efficient methods of continuing to scale performance, and to better understand what factors will ultimately limit link performance, we need to create a model that will estimate the noise seen at the receiver from the noise and non-idealities of the system. This model is developed in the next section and contains both deterministic as well as random noise sources. Having elucidated the noise sources, Section III describes how these different noise sources convert to voltage noise at the receiver. This section points out the importance of tracking both correlation between, and bounds on the noise. Section IV then uses this model to compare different communication techniques, and looks at the most important problems that need to be addressed to continue to scale link performance. Our conclusions about the scaling of links are given in Section V.

II. Link Models

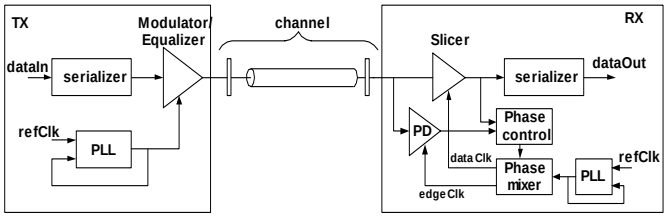

Fig. 3. A high-speed link system showing the transmitter, channel and receiver. The receiver must recover both the data and timing of the in coming signal and often has two detectors connected to the incoming data: the slicer for data, and the phase-detector (PD) for timing.

A simplified block diagram of the high-speed link system is shown in Fig. 3. The transmitter chip serializes the data and modulates it into the channel using a phase-locked loop (PLL) to provide the needed timing reference. The channel filters the data, and adds interference from neighboring signals. The receiver must extract the timing information and the data value from the incoming bits. Given the high symbol rates and low complexity requirements of these links, historically the transmitter used a simple binary non-return-to-zero modulation (NRZ), and the receiver samples the data and immediately decides on the value of the received data symbol. The operation of these systems is usually visualized by using an "eye diagram" which is shown in Fig. 4. For timing recovery most systems sample each bit twice, once in the middle of the eye to get the data value, and then once at the edge of the eye to get timing information.

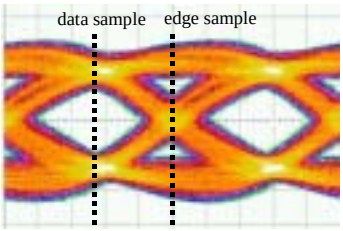

Fig. 4. Non-return to zero, eye diagram.

This eye diagram clearly shows the two types of noise that must be considered in these systems – the vertical fuzz is the voltage noise seen at the receiver, and the horizontal fuzz is the timing noise. Some of the "noise" is really deterministic, it is either from other signals, or time delayed versions of the signal being transmitted, while other noise components are from truly random sources.

A. Deterministic Noise

Often the largest "noise" seen at the receiver is not really noise at all – it is either delayed versions of the signal itself, caused by uncorrected dispersion of the channel (which is called inter-symbol interference, or ISI) or it is delayed versions of other signals caused by cross-talk. Channel dispersion comes from fundamental loss mechanisms in the wire, including skin-effect and dielectric loss, as well as filters formed by unintended transmission line impedance discontinuities caused by via stubs and impedance mismatches. These latter filters cause reflections, which cause resonances in the frequency response. Crosstalk or co-channel interference mainly occurs in dense connectors, as far and near-end cross-talk (FEXT and NEXT), [7]. Since the channel is the complete path from one chip to another, it is important to include the effects of the chip parasitics on the link as well.

The channel that we will use for most of the examples in this paper is shown in Fig. 5, which gives both the response in the frequency and time domain. Since deterministic noise is not really noise at all, but residual error, once the pulse response and input sequences are known, it is easy to calculate the probability density functions for the resulting error voltages. In fact, it is not difficult to even find the patterns that lead to the worst-case noise.

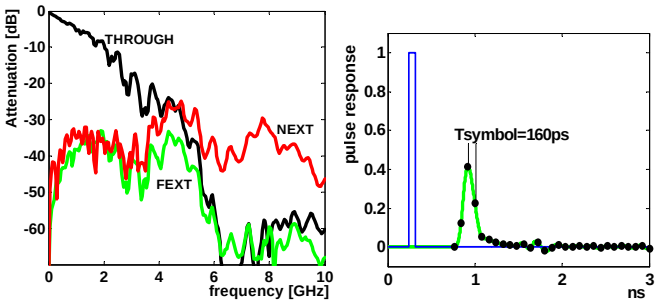

Fig. 5. Frequency response and crosstalk of a typical link, with pulse response at 6.25Gb/s. NEXT is the noise reflected back to the transmitter, and FEXT is the noise coupled to the receiver.

Since the ISI of this channel is so large, the residual error would overwhelm the signal if no correction was done. As a result, most of the examples will consider the effect of a 5-tap linear equalizer that greatly reduces the ISI.

B. Random Noise Sources and Modeling

Noise sources in high-speed environment arise from different device components, and occur both in timing and voltage domains. We will try to examine each of the noise sources in some detail, particularly focusing on its distribution and

spectral content since both can be critical parameters for determining the significance of the noise.

True random voltage noise has traditionally been neglected in high-speed links, since it is assumed that thermal and transistor device noise is very small with respect to signal magnitude. However, with increased bandwidth, these sources are slowly rising in importance, but are still small compared to other noise sources. For example, input referred noise for a receiver with 5GHz noise bandwidth will have a sigma of roughly 0.3mV, which is roughly 40dB down from the equalized signal level at the receiver.

The larger sources of noise arise from coupling from other signals, or though the power supply or substrate connections through finite power-supply rejection ratio (PSRR). Power supply and substrate noise are difficult to characterize, and despite a number of attempts to measure on-chip noise, [8], for most systems the distribution and spectral properties remain unquantified. For the purposes of our analysis, we will assume that due to the large amount of on-chip capacitance this noise is band-limited to frequencies smaller than on-chip clock frequency, and since it is really the superposition of a finite number of different events, we model it with a bounded Gaussian distribution.

In addition to these dynamic noise sources, we need to address static noise sources, which can be dominant in the noise budget of high-speed links. Receiver input offset can be as large as tens of mV, and even when corrected will be a few mV [9]. Other examples of static noise in more complex receivers are quantization errors in transmitter and receiver (for example quantization of equalizer coefficients or D/A and A/D steps) and estimation error of equalizer coefficients. These errors are uniformly distributed with σ determined by the size of the quantization steps. For our example we will use 10mV steps, and 10% estimation errors, which give σ of a few mV.

Supply and device noise also cause uncertainty in the timing signals used in the link. The effect of this noise on the link PLLs has been extensively studied [10-12]. While Hajimiri in [10] and Demir in [11] mostly focused on the performance of the voltage controlled oscillators and cyclo-stationary aspects of phase noise, caused by device noise, Mansuri in [12] has illustrated the impact of the most significant noise sources in a PLL. Since this noise is critical in most loops we repeat some of this analysis next.

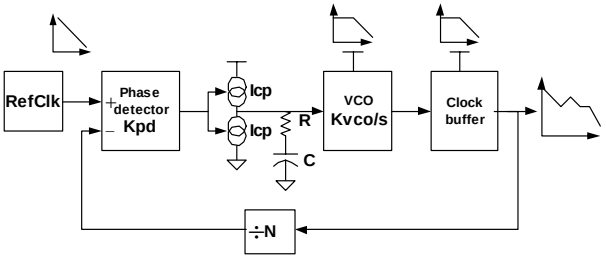

Fig. 6. The most significant noise sources for a PLL: Input clock jitter, VCO supply noise and clock buffer supply noise. Each is transformed through the loop to have a different spectral response.

A typical 2nd order PLL is shown in Fig 6, which also shows the dominant noise sources. Approximating the PLL as a second order system, we can obtain the noise transfer function from the power supply of the VCO and clock buffer to the output of the PLL as well as the transfer from reference clock noise to the output of the PLL. Fig. 7a gives the transfer function of the noise.

We see that VCO supply noise is band-pass filtered to the output, while clock buffer supply noise is high-pass filtered, and reference clock is low-pass filtered. To find the magnitude of the jitter, we need both the power spectrum of the supply noise, and the sensitivity of the VCO and clock buffers to the supply noise. Even if we assume the supply noise is white, the PLL jitter cannot be assumed white, and hence autocorrelation of jitter samples has to be taken into account in noise analysis. In addition, since the supply noise is bounded, the approximation of the jitter distribution should be bounded as well.

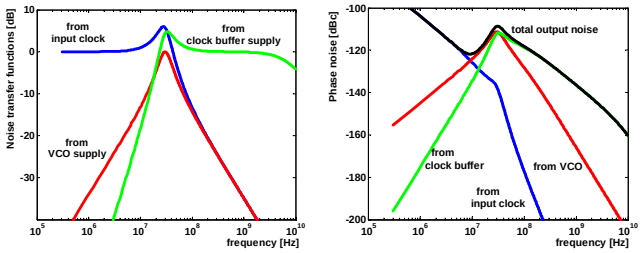

Fig. 7. a) Noise transfer functions from different supplies to the output of the PLL, b) Power spectral densities of PLL phase noise components assuming the supply noise is uniformly distributed with 20mV p-p and filtered by on-chip supply network with bandwidth of 100MHz. For our example design the VCO sensitivity is 0.15ps/mV, and the buffer sensitivity is 0.75ps/mV.

C. Clock and Data Recovery (CDR)

Since the receiver is actively tracking the input data's timing, we need to model how the CDR loop responds to data jitter to understand jitter's effect on the whole system. This is tricky since one of the most popular CDR techniques involves the use of identical samplers for data slicing and phase detection, in order to cancel the sampler's setup time. This binary type of phase detector results in bang-bang control loop, that is non-linear and in general very hard to analyze. In literature, there have been two separate methodologies to characterize this type of CDR system. Communications camp used Markov chain analysis of the loop [13-15], while IC designers usually linearize the loop and treat it as linear control system [16,17]. Both approaches have their limitations.

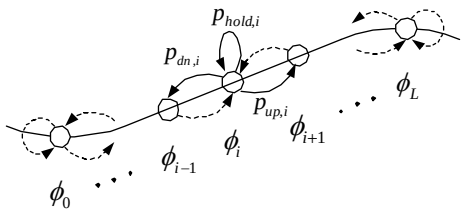

Fig. 8. First-order Markov chain phase-state model. Each state represents a different phase position, and the arcs are the probability of transition, given that position.

As illustrated in Fig. 8, possible phase positions of the

recovered clock are shown as states in a Markov chain, with probabilities $p_{hold,i}$, $p_{up,i}$ and $p_{dn,i}$ to hold, advance or retard the phase ϕ_i. The probabilities are generated by filtering (using moving-window, random-walk filter [18], or accumulate-reset filter [14]) the phase update information. Once the state-transition probabilities are found from the statistics of input data, which are shown in Fig. 9a, steady-state phase probabilities, which are shown in Fig. 9b, can be obtained.

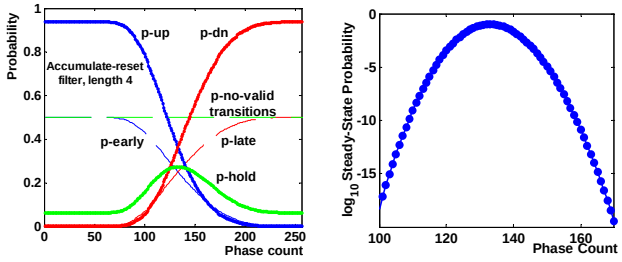

Fig. 9. a) The raw input probabilities (p-early, p-late, p-no-valid transition) are converted by a filter to state transition probabilities (p-up, p-dn, p-hold) for each possible phase location. Note that in a random NRZ data stream 50% of all data does not contain a transition to lock to. b) A plot of the steady-state phase probabilities for the input statistics shown in a).

Unfortunately real systems can depend on more than just the previous state of the system, which violates the first-order Markov chain model. In channels heavily affected by ISI and colored noise, first-order Markov chain can be applied only on accumulate-reset filter in cases where the length of the strong correlation of ISI or noise is similar to the length of the filter. For random-walk filters, or longer correlation lengths, higher-order Markov chains have to be used, [15].[1]

To avoid the problem with Markov models, one can linearize the loop, for example using first-order $\Sigma\Delta$ approximation, [16,17], in which the phase detector of a second-order CDR loop can be replaced with a white noise source with variance of dither jitter. This enables a frequency domain analysis of the loop, and creates the jitter tolerance mask for the receiver CDR loop. Loop delay can be correctly taken into account [19], which is essential for low-frequency jitter tracking, because of the peaking in the CDR transfer function that occurs from excessive delay.

As the next section will show, much of the jitter is due to factors related to the bit stream. To analyze this jitter our analysis will use a Markov model and we will assume that strong correlation of noise and ISI exists within a window covered by the CDR accumulate-reset loop filter, and that the weak correlation of the residual ISI due to long latency reflections and associated colored noise can be assumed uncorrelated (but with accurate bounded distributions). This approach works well with different edge selection algorithms present in multi-level and other more advanced signaling techniques.

[1] Another issue with Markov model is it ignores the latency of the CDR feedback loop, which leads to dither jitter. For systems where the phase steps are small, the input jitter is larger than the dither jitter and the latter can be ignored.

III. Link Performance Analysis

Driven by requests for essentially error-free operation, link performance analysis has traditionally used peak-to-peak noise values to estimate voltage and timing budget. Thus designers needed to show that with worst-case voltage and timing noise the link eye was still large enough to overcome the true statistical noise. As designs entered multi-Gb/s speeds and ISI started to become a major concern, these methods proved too pessimistic, and made the resulting circuit designs either impossible or with power/area costs that were undesirable. Clearly some statistical analysis was needed. Unfortunately assuming all noise sources and ISI to be Gaussian and independent, and using the resulting distribution to estimate bit-error rates (BER) leads to predictions that are also far off since most of the noise sources in the link are actually bounded [20,21].

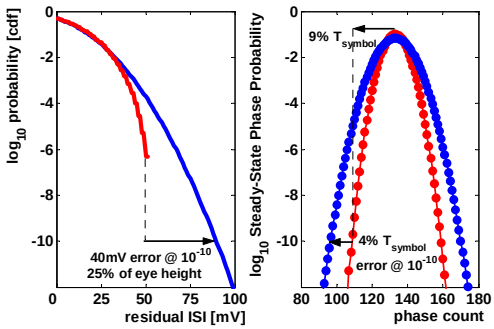

Fig. 10. Error of Gaussian approximation of residual ISI distribution, a) True voltage noise distribution for this channel vs. Gaussian model b) Steady-state phase distribution from Markov CDR model with input noise as true residual ISI distribution vs. its Gaussian approximation.

We can demonstrate the problem using our channel model, compensated by a simple 5-tap linear equalizer. Different bit sequences will have different residual errors, and one can compute the probability of an error for a given voltage margin. This data is given in Fig. 10a along with an approximation using Gaussian distribution with the same variance. In cases when systems operate at relatively high BER (~10^{-1} - 10^{-5}), like many standard communication systems, the approximation holds well. However, for target BERs of 10^{-15}, the error is large. Fig. 10b gives the output of our Markov model for receive timing giving the timing sample data for the same system, comparing the steady-state phase distributions for true residual ISI distribution and its Gaussian approximation.

To get a more complete picture of the link performance, timing noise sources have to be included in the statistical voltage analysis. Casper in [20] treats transmitter and receiver jitter as one random noise source, and conditions the BER for each phase offset with the Gaussian distribution of the combined jitter sources, in order to obtain the final BER number. Correctly observing that transmitter and receiver jitter cannot be just combined, Ahmad in [21] uses exhaustive combinatorial simulation to find the distribution of Gaussian transmitter jitter at the end of the channel. This approach requires both excessive run-time if length of the pulse response

is long, and approximations in discretizing the Gaussian distribution.

We will use a simple decomposition, shown in Fig. 11, to allow us to analyze the effect of transmitter jitter on link performance. With jitter a transmitter creates two signals: the output without jitter and an error signal caused by the jitter. The error signal is the difference between the actual pulse and the ideal pulse, and is simply a narrow pulse at each transition with a width equal to the jitter.

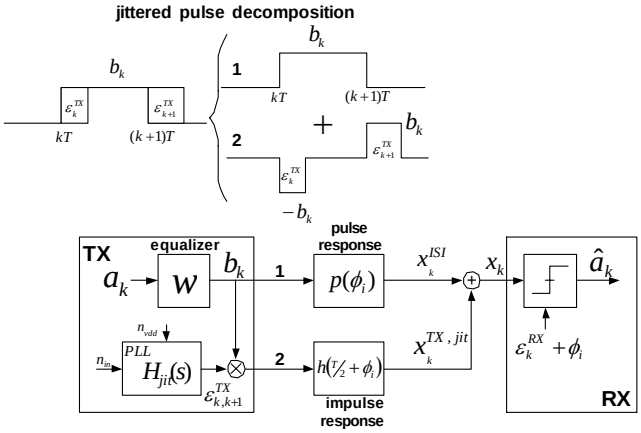

Fig. 11. Jittered pulse decomposition. A pulse transmitted with jitter is converted to a pulse with no jitter, plus an error term where the widths of the error pulses are equal to the jitter ε_k^{TX}. Since the error pulses are narrow they are represented by impulses located at the transitions (½ a bit from the data sample point) and after passing through the channel act as additive noise to the input data.

Since jitter pulses are much shorter than symbol time, and the channel is usually band-limited significantly below the Nyquist frequency, we represent these pulses as impulses at the transitions.[2] The ideal transmitter output is unaffected by jitter, and can be modeled by the channel response, ***p***, to the pulse of symbol time duration. At the output of the channel, the two signals are uncorrelated, provided that there is no correlation between the transmitted data and jitter. The resulting expression is:

$$x^{ISI}(kT+\phi_i+\varepsilon_k^{RX}) = \sum_{j=-sbS}^{sbE} b_{k-j}\, p(jT+\phi_i) \qquad (1a)$$

$$x^{jitter}(kT+\phi_i+\varepsilon_k^{RX}) = \sum_{j=-sbS}^{sbE} b_{k-j}\left[h(jT+\frac{T}{2}+\phi_i)(\varepsilon_k^{RX}-\varepsilon_{k-j}^{TX}) - h(jT-\frac{T}{2}+\phi_i)(\varepsilon_k^{RX}-\varepsilon_{k+1-j}^{TX})\right] \qquad (1b)$$

where ϕ_i represents the offset of the steady-state CDR phase, and ε_k^{RX}, jitter of the selected phase ϕ_i of the receiver PLL, and b_k the value of the transmitted symbol. Equation (1a) gives the output that results from the channel ISI, and receiver jitter, while (1b) provides the effect of transmitter jitter. If a transmitter filter is used, then the values of the b_k will not be independent, and instead will be correlated through the FIR filter with coefficients, w_i.

$$b_k = \sum_{i=-pre}^{post} w_i a_{k-i} \qquad (2)$$

[2] This is just a zero-order Taylor expansion of the error. We can extend this analysis method to higher order expansions.

Equation (1b) shows that the effect of transmitter and receiver jitter can be quite different. In fact they are the same only in the extreme case where all of the transmitter jitter samples within the length of the impulse response are roughly the same: the whole waveform is effectively shifted. This situation resembles the impact of the receiver jitter. In case of shorter jitter correlation, the impact of transmitter jitter is far worse than that of receiver jitter, as shown in Fig. 12. The jitter correlation at the transmitter strongly depends on the filtering of the power supply noise. For white supply noise, the resulting voltage noise looks similar to Fig. 12a, not 12b. The ratio of jitter mapped to voltage from transmitter is much bigger than from receiver, when jitter is assumed white as shown in Fig 12a. This points out the importance of understanding the jitter spectrum since for our assumptions about filtered power supply noise, the correlation of jitter is strong enough to make transmitter and receiver jitter look similar.

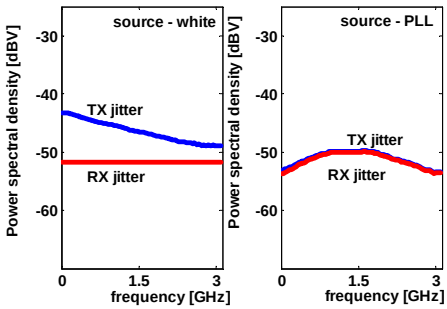

Fig. 12. Power spectral density (PSD) of voltage noise at sample time, due to transmitter and receiver jitter. The PSDs are multiplied by the Nyquist frequency, to obtain the noise power at each frequency, hence the unit [dBV] instead of customary [dB V^2/Hz].

When we compare this noise distribution to other noise sources, like equalizer quantization noise, tap value estimation error, and residual ISI on the equalized channel we see that the jitter noise is comparable to some of these noise sources. The size of the residual ISI depends on the length of the equalizer used. For this plot a 5-tap linear equalizer (large for current high-speed links) is used. Power spectral densities of all three sources of error are shown both for data and edge samples, in Fig. 13.

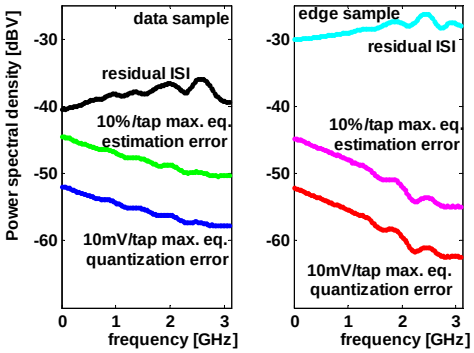

Fig. 13. Power spectral densities of quantization, estimation and residual ISI errors, at data and edge samples for a 6.25GHz NRZ signal using a 5-tap linear equalizer.

Besides power spectral densities, it is also illustrative to see the impact of different noise sources on CDR performance. Shown in Fig. 14 are steady-state phase probabilities, if we consider ISI, transmitter and receiver jitter separately. Although the receiver jitter causes the biggest steady-state phase standard deviation, 2.4% of symbol time, followed by ISI, 1.5% and transmitter jitter, 1.4%, the distribution due to ISI is wider for most of the error range of interest, so it dominates the low BER analysis.

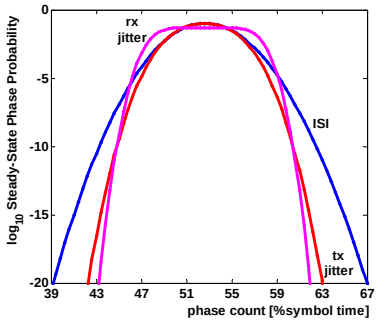

Fig. 14. Steady-state phase probabilities, dependence on ISI, transmitter and receiver jitter.

IV. Communications Techniques

There are many techniques that have been developed to deal with bandwidth limited channels. The simplest approach is equalization to flatten out the frequency response of the channel and to reduce the dispersion. Other techniques that are being discussed are decision feedback equalization, and multi-level signaling. This section uses our noise model to evaluate these different techniques.

A. Equalization

Many multi-Gb/s link designs use some kind of equalization, based on linear filtering, in order to achieve desired data rates over severely band-limited channels. Since linear equalization effectively "amplifies" noise, it is important to evaluate how equalizer affects the noise terms. The largest random noise in our links comes from jitter, so we look at that first.

Most high-speed links use transmitter equalization, since it is easier to implement the needed FIR filters on the digital data before transmission, than on the analog received data. We use a 5-tap linear equalizer in the transmitter to de-emphasize the low frequency components, and flatten the frequency response. The effect of the equalizer on the jitter noise is shown in Fig. 15 for both white and colored jitter at the output of the PLL, for the channel in Fig. 16. The equalizer has a small effect on the effective voltage noise when then transmitter jitter is white because the jitter for each bit is uncorrelated with the previous bit. The FIR filter is not really equalizing the noise since each noise value is independent. Unlike the transmitter, the receiver sees the effect of the equalized channel, so the jitter noise is reduced. The colored noise has very small energy at high-frequencies, making the correlation length in the transmitter longer than the filter length. Thus, for the colored noise case, both transmitter and receive jitter cause similar voltage noise, and have a similar reduce due to equalization.

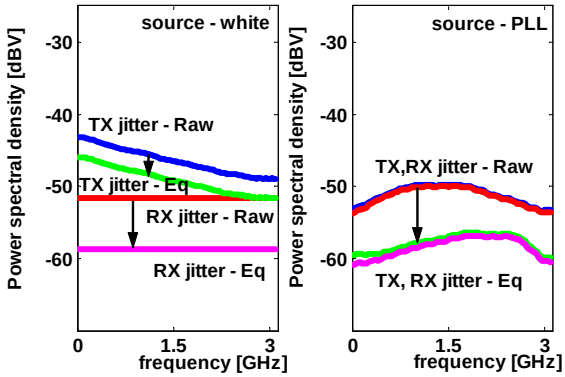

Fig. 15. PSDs of transmitter and receiver jitter mapped to the data sample point through equalized and unequalized (Raw) channel, for cases of white jitter PSD and colored jitter PSD as output of the PLL.

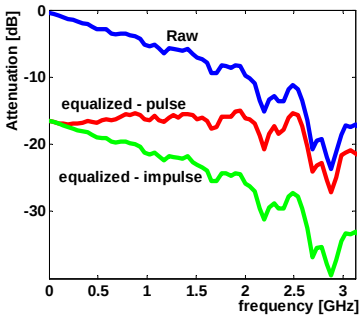

Fig. 16. Channel frequency response (Raw), and frequency response of the equalized 160ps pulse, and impulse.

B. Decision Feedback Equalization (DFE)

Rather than building a linear filter to equalize the pulse response, DFE uses the history of the received symbols to cancel the trailing ISI that is present in the channel. If the channel pulse response is known, we can subtract the residual signals from all the bits we have already seen, leaving only the signal from the bit of interest. In communications, DFE has been used heavily instead of linear filtering, to circumvent the problem of noise amplification, [22]. These systems implement DFE in two basic ways: either in the digital domain, or as a mixed-signal circuit. For the digital domain the input is initially digitized at a high enough resolution and rate, and the DFE is implemented in a DSP. Modems and Gigabit Ethernet [23] are implemented this way. For mixed-signal implementation, the input quantizer only needs to resolve the input. The output of this quantizer, plus older bits feed an FIR filter that then drives a DAC whose output is subtracted from the input signal. This approach has been used in analog read channels [24].

Both approaches are, however, impractical in multi-Gb/s high-speed links. The digital approach requires very fast and accurate A/D converters. While high-speed A/D converters have been created [25], they would be too expensive (area/power) to use in a practical link. The second approach, suffers from latency problems. For a 6.25Gb/s binary link, you have 160ps to resolve the input, drive the DAC, and have the DAC outputs settle to the required precision.

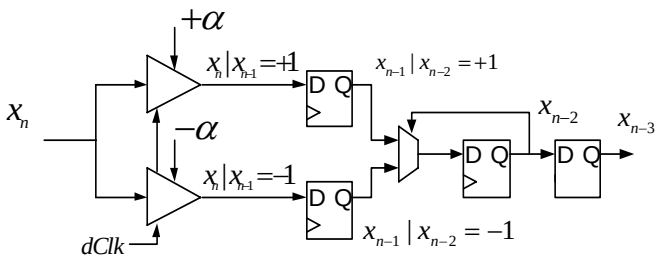

Fig. 17, Parallel DFE or Partial-response DFE for 1+αD channel.

An interesting solution to this latency problem has been independently proposed by Parhi and Kasturia [26,27], and applied to long-haul fiber optical links. Since we cannot run the feedback loop fast enough, we unroll it once and make two decisions each cycle. One comparator decides the input if the previous output was a 1, and the other comparator decides the input if the previous bit was a -1. Once we know the previous bit, we select the correct comparator output. This algorithm is shown in Fig. 17. Instead of just one data sampler, for 2PAM signaling the receiver has two samplers that are offset by $\pm\alpha$, anticipating the impact of trailing ISI α, from a previously sent symbol of value of ± 1. This method can be applied to two or more taps of feedback, however, the number of required receivers is M^L, where M is the number of signal levels and L is the number of feedback taps, and each receiver nominally has $M-1$ comparators. Usually only a small amount of unrolling is needed. Corrections for bits that are far enough away from the current sample don't have a latency problem. This makes it easy for DFE to correct for long latency interference caused by reflections from connectors, vias, transmit/receive parasitic capacitances and other termination mismatches. To prevent the complexity of the resulting DFE and parasitic output capacitance from becoming a problem, only limited taps are used, [28].

Decision-feedback equalization is not entirely free of noise amplification or equivalently, signal attenuation, since leading ISI, Fig. 5b, can only be eliminated using linear filter preceding the feedback stage.

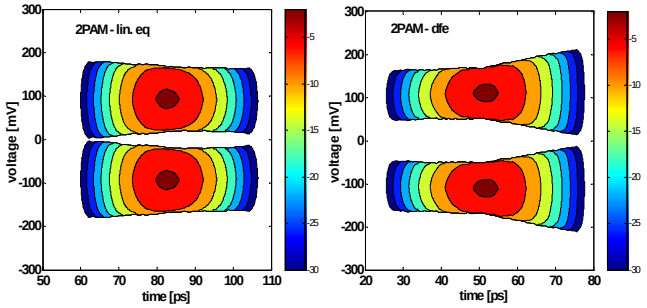

Fig. 18. Probability distributions of sampled signal due to distributions of ISI and CDR phase. The plots are contours of log(P) for 6.25Gb/s link, with ±500mV p-p transmitter output using a) 2PAM with linear equalizer, 1 leading and 4 trailing taps, b) one tap parallel DFE with linear feed-forward filter as in a). The spacing between the top and bottom contours of the same value is the voltage margin in the system.

We can compare the performance of DFE and linear equalizer solutions using our noise model, and the results are shown in Fig 18. The margin for an error rate of 10^{-x} is the smallest vertical distance between the –x contours in the plots. One can clearly see the larger margins available for DFE.

C. Multi-level Signaling

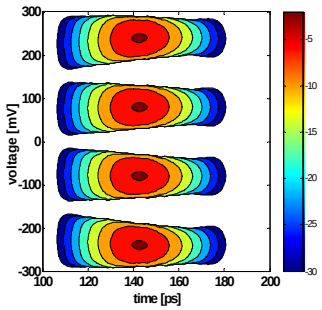

Fig. 19. Probability distribution map of sampled signal for 4PAM signaling with linear equalizer.

Signal modulation has been effectively used in communications, to combat band-limited channels, improving the bit rate by transmitting more bits/Hz rather than increasing the signaling rate and hence the required bandwidth. For example, rather than signaling at 5GHz to achieve 10Gb/s with 2PAM signaling, we can use 4PAM and decrease the Nyquist frequency to 2.5GHz while achieving the same bit rate. The choice between the two schemes is between the signal loss of 3x due to increased number of levels, and equalized channel attenuation at two Nyquist frequencies, [28,29]. Actually, the trade off is a bit more complicated because of different distribution of ISI and jitter in 2PAM and 4PAM cases, as well as the impact of crosstalk. Given the decrease in the received signal relative to the maximum power transmitted into the channel, multi-level modulation scheme is generally more sensitive to residual errors from ISI and cross-talk than normal binary signaling.

We can again use our noise models to look at the margins for a 4PAM system, and the result is shown in Fig 19. This plot is generated for the same data rate (so ½ the symbol rate) as Fig. 18, using the same 5-tap linear equalizer. For this channel the margin is better for 4PAM than the 2 PAM system, but they are slightly worse than the margins for the DFE system. We can compare the margin of these three approaches for different channels formed by changing the length of the backplane trace. This data is shown in Table 1.

TABLE I

Voltage margins [mV] at target BER=10^{-12} at 6.25Gb/s, for 2PAM, 2PAM with DFE and 4PAM signaling, over 3, 10 and 20" backplanes. Transmitter peak output swing is ± 500mV, and receiver sensitivity ± 10mV.

Eq/Mod type vs. BP length	3"	10"	20"
2PAM	32	17	19
2PAM w. DFE	79	49	44
4PAM	10	37	31

Short channels have less loss and more noise factors, so not surprisingly 2PAM systems are much better than 4 PAM systems in these situations. As the backplane get longer, the attenuation increases, and now 4PAM is better than 2PAM with linear filters, but it is still not better than 2PAM with a simple DFE.

V. Conclusion

During the 90's link performance was scaled dramatically by treating I/O wires as transmission lines, and using on-chip parallelism (multiplex transmitters and de-multiplex receivers) and PLLs to rapidly scale the bit rate so it was comparable or even faster than the internal clock rate. We have scaled links so well, that we are now running into the intrinsic frequency limitations of the wires. In this new regime of operation, we need to borrow many tools from communication theory, but apply them to our environment where we are running Gsymbols/sec, and have different error requirements.

In these bandwidth limited systems, the most important parameter to track is the effective noise that the receiver sees. Since most of the noise in this system is not unbounded Gaussian noise, it is critical to model both the bounded nature of the noise, and its correlation. We have presented one such model for noise that includes the effects of ISI and the clock CDR loop, and show how using white Gaussian noise leads to both overestimates of the noise, and changes the relative importance of different factors. Preliminary evaluations using this model indicate that both DFE and 4 PAM are interesting alternatives to simple NRZ linear equalization for high-speed links, and can give much larger margins. The main limitations on link performance are the deterministic errors – so more complex equalization and DFE should help. After these issues are solved, it looks like both CDR jitter and equalization accuracy will be the next limits that need to be addressed.

Acknowledgments

The authors would like to acknowledge the help and support of J. Zerbe, R. Kollipara, A. Ho, F. Chen and J. Wei of Rambus, and E. Alon of Stanford University. V. Stojanovic also thanks I. Stojanovic for full-hearted help on the paper and technical discussions.

References

[1] M.M Griffin *et al*, "A process-independent, 800-MB/s, DRAM byte-wide interface featuring command interleaving and concurrent memory operation," *IEEE J. Solid-State Circuits*, vol: 33, no. 11, Nov. 1998, pp. 1741 –1751.

[2] H. Yoon *et al*, "A 2.5-V, 333-Mb/s/pin, 1-Gbit, double-data-rate synchronous DRAM," *IEEE J. Solid-State Circuits*, vol. 34, no. 11, Nov. 1999, pp. 1589 –1599.

[3] T.P. Thomas and I.A. Young, "Four-way processor 800 MT/s front side bus with ground referenced voltage source I/O," *IEEE Symposium on VLSI Circuits*, June 2002, pp. 70 –71

[4] K. Lee *et al*, "1.04 GBd low EMI digital video interface system using small swing serial link technique,"*IEEE J. Solid-State Circuits*, vol. 33, no. 5 , May 1998, pp. 816 –823.

[5] C.-K. Yang and M. Horowitz, "A 0.8um CMOS 2.5Gb/s oversampling receiver and transmitter for serial links," *IEEE J. Solid-State Circuits*, December 1996, pages 2015-2023.

[6] R. Kollipara *et al*, "Design, Modeling and Characterization of High-Speed Backplane Interconnects," *DesignCon 2003*.

[7] S. Sercu and J. De Geest, "BER Link Simulations," *DesignCon 2003.*

[8] M. Badaroglu *et al*, "Methodology and experimental verification for substrate noise reduction in CMOS mixed-signal ICs with synchronous digital circuits," *IEEE J. Solid-State Circuits*, vol. 37, no. 11, Nov 2002, pp. 1383 – 1395.

[9] W. Ellersick *et al*, "GAD: A 12-GS/s CMOS 4-bit A/D converter for an equalized multi-level link," *IEEE Symposium on VLSI Circuits*, June 1999.

[10] A. Hajimiri, "Noise in phase-locked loops," *Southwest Symposium on Mixed-Signal Design*, 2001. pp. 1 –6

[11] A. Demir, *et al*, "Phase noise in oscillators: a unifying theory and numerical methods for characterization," *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, vol. 47, no. 5, May 2000, pp. 655 –674.

[12] M. Mansuri and C-K.K. Yang, "Jitter optimization based on phase-locked loop design parameters," *IEEE J. Solid-State Circuits*, vol. 37, no. 11, Nov 2002, pp. 1375 –1382

[13] J.K. Holmes, "Performance of a First-Order Transition Sampling Digital Phase-Locked Loop Using Random-Walk Models," *IEEE Transactions on Communications*, vol.com-20, no. 2, April 1972.

[14] A.E. Payzin, "Analysis of a Digital Bit Synchronizer," *IEEE Transactions on Communications*, vol.com-31, no. 4, April 1983.

[15] D.W. Paranchych and N.C. Bealieu, "Use of second order Markov chains to model digital symbol synchronizer performance," *IEE Proceedings – Communications*, vol. 143, no. 5, October 1996, pp. 250-258.

[16] I. Galton, "Analog-input digital phase-locked loops for precise frequency and phase demodulation," *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 42, no. 10, October 1995, pp. 621-630.

[17] R. Walker *et al*, "A 2.488Gb/s Si-Bipolar Clock and Data Recovery IC with Robust Loss of Signal Detection," *IEEE International Solid-State Circuits Conference*, Feb.1997, San Francisco.

[18] J.R. Cessna and D.M. Levy, "Phase Noise and Transient Times for a Binary Quantized Digital Phase-Locked Loop in White Gaussian Noise," *IEEE Transactions on Communications, vol.com-20, no. 2, April 1972.*

[19] J. Kim, *Design of CMOS Adaptive-Supply Serial Links, Ph.D. Thesis*, Stanford University, December 2002.

[20] B.K. Casper *et al*, "An accurate and efficient analysis method for multi-Gb/s chip-to-chip signaling schemes," *IEEE Symposium on VLSI Circuits*, June 2002, pp. 54-57.

[21] B. Ahmad, "Performance Specification of Interconnects," *DesignCon 2003*.

[22] J. Proakis and M. Salehi, *Communication Systems Engineering*, Prentice Hall 1994.

[23]] E.F. Haratsch and K. Azadet, "A pipelined 14-tap parallel decision-feedback decoder for 1000BASE-T Gigabit Ethernet," *IEEE International Symposium on VLSI Technology, Systems, and Applications*, 2001, pp. 117 - 120

[24] M.Q. Le *et al*, "An Analog DFE for Disk Drives Using a Mixed-Signal Integrator," *IEEE J. Solid-State Circuits*, May 1999, pp. 592-598.

[25] K. Poulton *et al*, "A 20GS/s 8b ADC with a 1MB Memory in 0.18μm CMOS," *IEEE International Solid-State Circuits Conference*, Feb. 2003, San Francisco.

[26] K.K. Parhi, "High-Speed architectures for algorithms with quantizer loops," *IEEE International Symposium on Circuits and Systems*, vol. 3, May 1990, pp. 2357-2360

[27] S. Kasturia and J.H. Winters, "Techniques for high-speed implementation of nonlinear cancellation," *IEEE J. Selected Areas in Communications*, vol. 9, no. 5, Jun 1991, pp. 711-717.

[28] J. Zerbe *et al*, "Design, Equalization and Clock Recovery for a 2.5-10Gb/s 2-PAM/4-PAM Backplane Transceiver Cell," *IEEE International Solid-State Circuits Conference*, Feb. 2003, San Francisco.

[29] J.T. Stonick *et al*, "An adaptive pam-4 5-Gb/s backplane transceiver in 0.25-μm CMOS," *IEEE J. Solid-State Circuits*, vol. 38, no. 3, March 2003, pp. 436-443.