Crosstalk-Aware Transmitter Pulse-Shaping for Parallel Chip-to-Chip Links

Mike Bichan
Department of Electrical and Computer Engineering
University of Toronto
Email: mbichan@eecg.utoronto.ca

Anthony Chan Carusone
Department of Electrical and Computer Engineering
University of Toronto
Email: tcc@eecg.utoronto.ca

***Abstract*— Crosstalk and inter-symbol interference are major obstacles towards increasing the data rate of chip-to-chip links. They are especially important in long, single-ended, parallel links. We find that when crosstalk is the dominant source of received noise, a transmitted pulse shape that combines slew-rate limiting with transmitter pre-emphasis is preferred. We present a framework for choosing a transmit pulse-shape that minimizes the combined effects of both crosstalk and inter-symbol interference. When implementing this optimal pulse shape, filter taps that are fractionally spaced are beneficial. As a test, various pulse shapes were injected into a test channel with three parallel microstrip lines, and the received eye diagrams were measured. The crosstalk-aware pulse shape results in an increased eye opening when compared to square and pre-emphasis pulses.**

I. Introduction

Crosstalk between adjacent channels is a severe problem in chip-to-chip communication links. It exists as a result of parasitic capacitance and inductance on printed circuit boards and it is a barrier preventing bit rates for parallel chip-to-chip links from increasing past 5 Gb/s/pin. Even more dramatic are the effects of crosstalk on board-to-board channels and multidrop busses. To extend the useful bandwidth of these channels, techniques must be used to combat crosstalk while also equalizing inter-symbol interference (ISI).

The chip-to-chip communication circuits that achieve the highest per-channel bit rate are serial transceivers. There are no adjacent channels injecting crosstalk onto the desired channel. The fastest of this type of circuit currently achieves a bit rate of 20 Gb/s over short backplane and coaxial cable channels [1], [2]. These circuits use pre-emphasis to compensate for ISI. However, pre-emphasis cannot help reduce crosstalk and can even increase it.

The type of channel considered in this paper is unidirectional, single-ended, and parallel. Circuits designed for this environment must deal with large crosstalk from adjacent signalling paths. For this type of link the maximum reported bit rate is 3 Gb/s [3]. These circuits use slew-rate limiting to minimize each channel's impact on its neighbour. However, slew-rate limiting increases ISI, thus limiting the maximum data rate. Another common way of reducing crosstalk is to explicitly cancel it by sensing the signal in the adjacent channels and subtract the expected crosstalk from the desired channel [4]. This technique is effective, but requires more power as the number of adjacent channels increases.

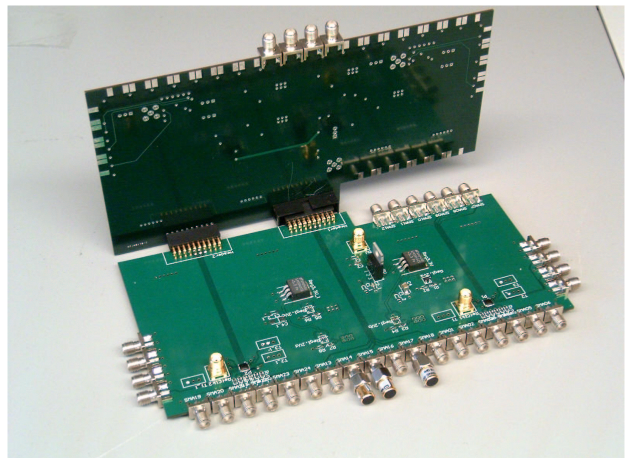

Fig. 1. Board-to-board communication link.

This paper describes the use of a finite impulse response (FIR) filter to shape the transmitted pulses in order to minimize the combined effect of both ISI and crosstalk. Since transmit filters are generally required for pre-emphasis anyway, this approach entails a complexity comparable to conventional high-speed chip-to-chip transceivers. Section II describes the channel model adopted. Section III describes how the transmit pulse shape is optimized for given through and crosstalk channel responses. The optimal pulse shape for a parallel board-to-board channel incorporates both slew-rate limiting and pre-emphasis in order to maximize the received eye-opening-to-crosstalk ratio. In section IV as a proof-of-concept, a 2.7 Gb/s crosstalk-aware pulse shape is generated using a parallel bit error ratio tester (ParBERT) and applied to a single-ended printed circuit board (PCB) bus channel. The resulting channel output is compared to the output produced when square and pre-emphasis pulses are used. The crosstalk aware pulse shape results in an increased eye opening.

II. Chip-to-Chip Channel Impairments

The channel considered here is shown in Figure 1. This board-to-board channel consists of two 10 cm sections of 50 Ω microstrip connected by a simple through-hole parallel board connector. This channel is representative of the type of channel found in the backplane environment in which multiple

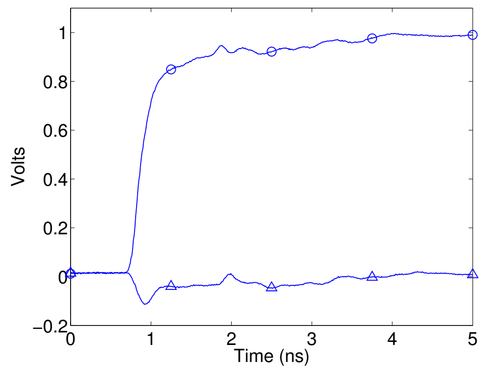

Fig. 2. Measured through (◯) and crosstalk (△) responses of the board-to-board channel.

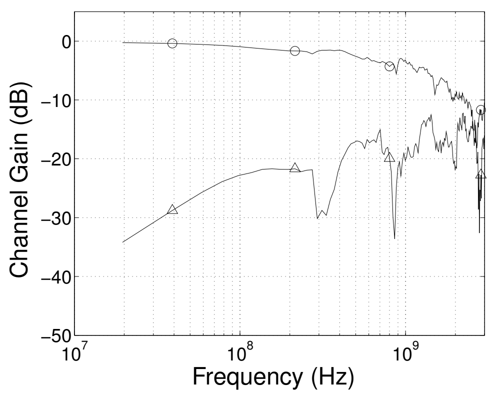

Fig. 3. Frequency response of the through channel (◯) and crosstalk channel (△) computed from the step response.

daughtercards are connected to a motherboard by a parallel board connecter. Both the proximity of the microstrip lines to one another, and the board connectors themselves, are responsible for significant crosstalk.

The measured step response of this channel is shown in Figure 2. This figure indicates that the crosstalk response is not insignificant compared with the through response and therefore cannot be ignored.

The frequency response of the channel can be computed from the step response and is shown in Figure 3. The differentiating nature of the crosstalk channel can be seen from this figure. An ideal differentiator has a frequency response with a slope of 20 dB/decade which is what this channel displays at frequencies from 10 MHz to 100 MHz. Above 100 MHz the through and crosstalk responses are attenuated by the length of the channel.

For frequencies above 2 GHz the through channel experiences severe attenuation which causes ISI. Traditionally, ISI is reduced by amplifying the high frequency content of the signal. For the frequency response shown, however, the through and crosstalk responses have similar amplitude for frequencies above 2 GHz. In this case, amiplifying the high frequency content may not improve the signal-to-noise ratio (SNR) because the crosstalk is also being amplified. Clearly, a tradeoff must be made between ISI and crosstalk when

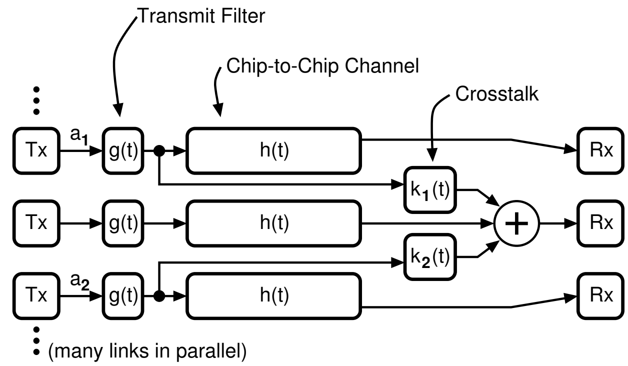

Fig. 4. System model of a parallel link

equalization is being performed.

III. Optimal Pulse Shape

This section describes a procedure for finding the optimal transmit filter for the general parallel link shown in Figure 4. Let $a_i(t)$ be the aggressor data patterns with bit period T_{bit} and $k_i(t)$ be the impulse responses of the corresponding far-end crosstalk (FEXT) channels. Assuming all transmitters employ the same pulse-shaping filter with impulse response $g(t)$, the received crosstalk is:

$$c(t) = \sum_i a_i(t) * g(t) * k_i(t) \tag{1}$$

Since $a_i(t) = \sum_k \delta(t - kT_{bit}) \cdot d_i(k)$ where $d_i(k)\varepsilon\{\pm 1\}$ is the binary transmitted data, the worst-case crosstalk is given by:

$$c_{max} = \sum_i \max_{0 \le \tau_i < T_b} \sum_k |(g * k_i)(kT_{bit} + \tau_i)| \tag{2}$$

where τ_i is the skew between channels that leads to the worst-case crosstalk. If h is the through response of the channel, the worst-case eye opening is given by a peak distortion analysis [5]:

$$h_{min} = (g * h)(\tau_{pk}) - \sum_{k \neq 0} (g * h)(\tau_{pk} + kT_{bit}) \tag{3}$$

The figure of merit optimized in this work is “eye-opening-to-crosstalk” ratio:

$$\text{E2C} = \frac{\text{crosstalk-free eye opening}}{\text{maximum possible crosstalk}} \tag{4}$$

Given h and k_i, E2C can be computed for any prospective transmit filter $g(t)$. We maximize E2C by performing an exhaustive search over all possible transmit filters $g(t)$. Since the FEXT channels bewtween neighboring links are symmetric and c_{max} is generally dominated by the nearest aggressors, on a wide bus the optimal pulse shapes will be nearly identical for all links. Therefore, we can assume that using the same transmit filter for each channel will not be far from optimal.

The transmitter is assumed to be an FIR filter with four bits per tap. The FIR filter is not necessarily baud-rate; different tap spacings are considered. The number of taps is varied from two to six. Finally, the peak output swing of the transmit filter

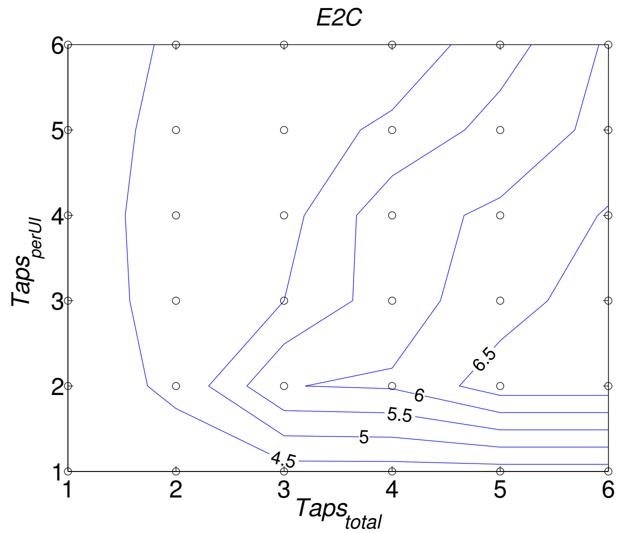

Fig. 5. Contour plot of simulated E2C against number of taps and taps per UI.

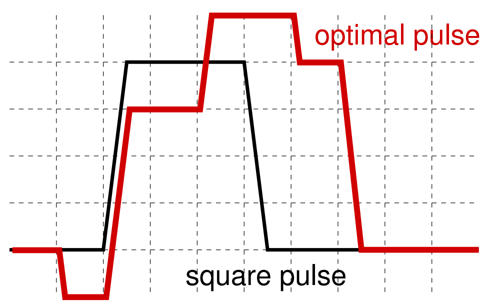

Fig. 6. Comparison of optimal and regular NRZ pulses.

is also constrained as it would be in practice by the supply voltage.

As we vary the total number of taps ($Taps_{total}$) and taps per unit interval ($Taps_{perUI}$), there is an optimal pulse shape for each ($Taps_{total}$,$Taps_{perUI}$) pair. A two-dimensional plot in Figure 5 shows the highest figure of merit possible for each filter configuration. The highest figure of merit occurs for a filter with 6 taps and 3 taps per unit interval (UI). The optimal pulse shape generated by this procedure is shown in Figure 6, alongside a square pulse for comparison, for the case of a filter with at most 6 taps.

The pulse shown in Figure 6 incorporates both slew-rate limiting and pre-emphasis in order to compensate for imperfections in the board-to-board channel. Figure 5 shows that for 3- to 5-tap filters, a tap spacing of $T_{bit}/2$ is preferable to a baud-spaced filter. For a 6-tap filter, 3 taps per UI is preferred. This is because a total filter span of roughly 2 UI is sufficient to cancel most ISI for this particular channel. So, as additional taps are added, the link's performance benefits more from the increased time resolution offered by smaller tap spacings than from increasing the filter span beyond 2 UI. The smaller tap spacings can be used to further mitigate the crosstalk.

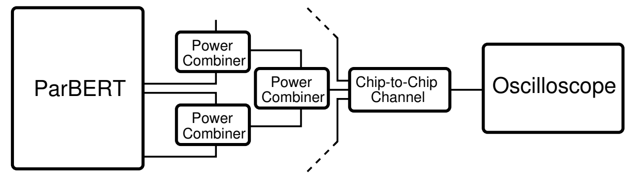

Fig. 7. Test setup for the hardware demonstration of the transmit filter.

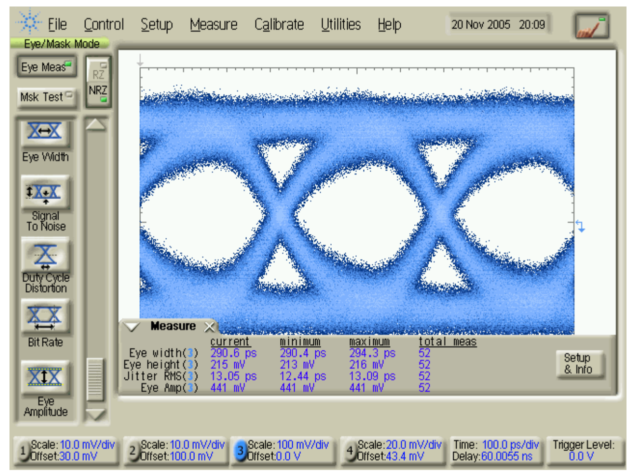

Fig. 8. Measured eye diagram at 2.7 Gb/s with a PRBS sequence of length $2^{31}-1$. This figure shows the channel output corresponding to a square pulse input with no aggressors.

IV. Transmit Filter Proof-of-Concept

This section presents a proof-of-concept hardware demonstration of a crosstalk-aware transmit filter. A ParBERT is used to imitate the function of a transmit filter, and the output is applied to a board-to-board channel. The ParBERT has nine output modules with a maximum bit rate of 2.7 Gb/s. For each channel, three of the 2.7 Gb/s module outputs have been combined together with power combiners as shown in Figure 7, to produce the transmitted pulse shape. The swing and relative phase of each module can be controlled in software, and so any pulse shape can be reproduced with this setup.

When pseudo-random square pulse data is fed into this channel with no adjacent agressor signals, the result is the eye diagram shown in Figure 8. When there are two adjacent signals causing crosstalk onto the desired signal, the received eye diagram looks like the one shown in Figure 9. It is clear from these figures that FEXT is a major problem in this channel.

Using the optimal pulse shape described in the previous section, the eye diagram shown in Figure 10 is received. According to the oscilloscope, the eye opening has increased from 39 mV to 101 mV by using the crosstalk-aware pulse shape.

A bathtub plot comparing the square and optimal pulses is shown in Figure 11. It is clear that the crosstalk-aware pulse improves the bit-error rate (BER) performance of the link substantially. Figure 12 shows a similar comparison, but

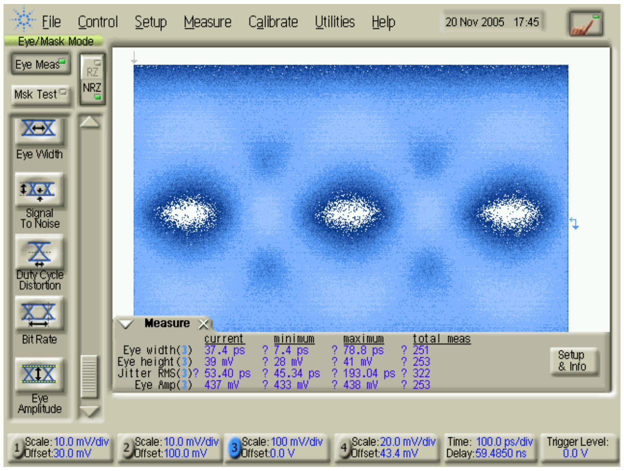

Fig. 9. Measured eye diagram at 2.7 Gb/s with a PRBS sequence of length $2^{31}-1$. This figure shows the output of the chip-to-chip channel for square pulse input with two agressors.

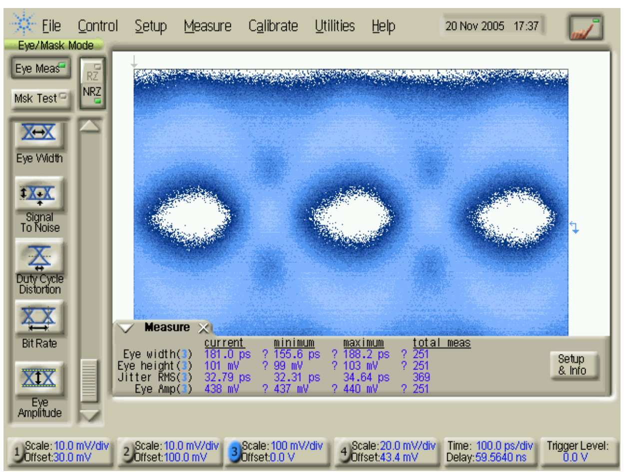

Fig. 10. Measured eye diagram at 2.7 Gb/s with a PRBS sequence of length $2^{31}-1$. This figure shows the output of the chip-to-chip channel for optimal pulse input with two aggressors.

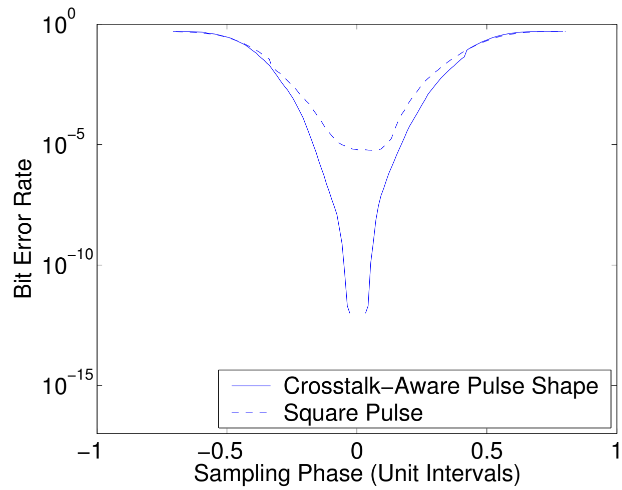

Fig. 11. Bathtub plot comparing crosstalk-aware and square pulses.

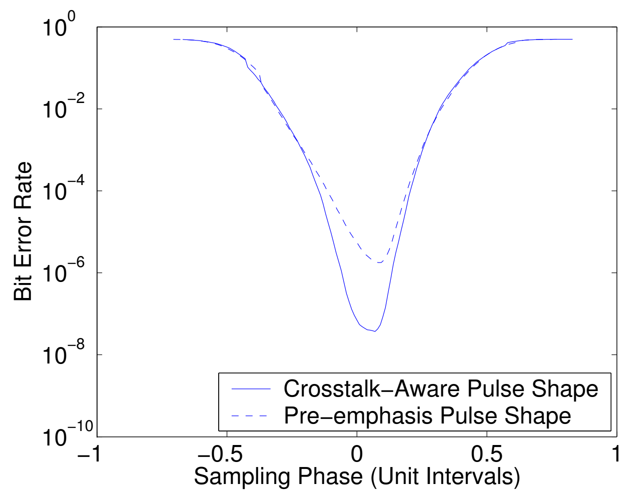

Fig. 12. Bathtub plot comparing crosstalk-aware and pre-emphasis pulses. BER is higher than in Figure 11 because a smaller signal swing was used in this measurement for both pulse shapes.

instead of a square pulse a pre-emphasis pulse is used. The pre-emphasis pulse used here has 2 taps, with the first tap having an amplitude 50% greater than the second tap.

V. Conclusion

The impact of crosstalk can be reduced by choosing a transmitted pulse shape that incorporates both slew-rate limiting and pre-emphasis. This technique uses less power than the alternative, namely explicit crosstalk cancellation, especially when a large number of parallel links are present. In addition, measured results show that this optimal pulse results in a lower received BER than a simple, two-tap pre-emphasis pulse of the type that is commonly employed in chip-to-chip transceivers.

Acknowledgment

The authors would like to thank Intel Corporation for their generous support.

References

[1] Y. Hur, M. Maeng, C. Chun, F. Bien, H. Kim, S. Chandramouli, E. Gebara, and J. Laskar, "Equalization and near-end crosstalk (NEXT) noise cancellation for 20-Gb/s backplane serial I/O interconnections," *IEEE Trans. Microwave Theory Tech.*, vol. 53, no. 1, pp. 246–255, Jan. 2005.

[2] P. Chiang, W. J. Dally, M.-J. E. Lee, R. Senthinathan, Y. Oh, and M. A. Horowitz, "A 20-Gb/s 0.13-μm CMOS serial link transmitter using an LC-PLL to directly drive the output multiplexer," *IEEE J. Solid-State Circuits*, vol. 40, no. 4, pp. 1004–1011, Apr. 2005.

[3] S.-J. Bae, H.-J. Chi, H.-R. Kim, and H.-J. Park, "A 3Gb/s 8b single-ended transceiver for 4-drop DRAM interface with digital calibration of equaliztion skew and offset coefficients," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2005, pp. 520–521.

[4] J. F. Buckwalter and A. Hajimiri, "Cancellation of crosstalk-induced jitter," *IEEE J. Solid-State Circuits*, vol. 41, no. 3, pp. 621–632, Mar. 2006.

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