Equalizing Filter Design for Crosstalk Cancellation

Jihong Ren and Mark Greenstreet
University of British Columbia
{jihong,mrg}@cs.ubc.ca

1 Introduction and Background

Advances in digital integrated circuit (IC) fabrication technology have resulted in an exponential growth for the speed and integration levels of ICs. This has created a corresponding demand for high bandwidth, chip-to-chip interconnect. To meet this demand, designers have resorted to aggressive signaling techniques including multi-level signaling, multi-gigahertz symbol rates, and tight wire spacings [3]. All of these exacerbate crosstalk (electromagnetic coupling) by reducing noise margins, increasing slew rates, and reducing physical isolation. Continued improvements in chip-to-chip interconnect require effective solutions to crosstalk.

This paper demonstrates the effectiveness of equalizing filters for cancelling crosstalk on high-speed buses. Recently high-speed interconnect techniques use single-line equalizing filters to compensate for dispersive losses due to wire resistance, the skin effect, and dielectric losses [1, 2, 4]. To the best of our knowledge, [6] is the only design where an equalizing filter is used for crosstalk cancellation in the context of high speed buses. That paper describes a proprietary design and gives few details of how the filters are derived.

This paper presents a novel method for designing the crosstalk cancelling filters and provides an evaluation comparing it with other design techniques. We show that crosstalk cancellation can double the bandwidth achievable on buses with tight wire spacings. We show that the l_∞ norm is a more appropriate measure for signal integrity in digital designs than the commonly used l_2 norm, achieving roughly 46% improvements in bandwidth. We present a practical method for synthesizing optimal filters based on linear programming and present results from an implementation of our synthesis procedure. Finally, we describe how our filters can be implemented using a practical, look-up table based approach with minimal overheads for area and latency.

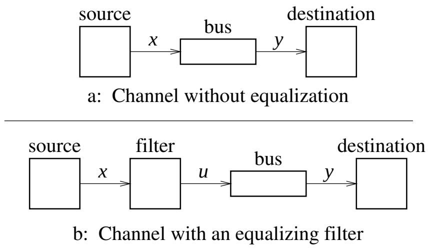

Figure 1. Channels

2 Filter Design

Figure 1 shows the structure we assume for communication channels. Buses have crosstalk, dispersive losses, reflections, and other effects that corrupt digital integrity. Fortunately, all of these phenomena are linear processes. Accordingly, we regard the bus as a time-invariant, linear system and model the bus by its impulse response function. This impulse response gives the coupling from each input of the bus to each output as a function of time. The impulse response can be derived from the electrical or geometrical parameters of the bus, or measured by including high-speed analog to digital converters in the receiver [4, 5].

In theory, the ill effects of crosstalk, etc., could be eliminated by including a filter with a transfer function equal to the inverse of that for the bus. In practice, a perfect inverse is unimplementable due to limitations on high frequency response and voltage swing. Furthermore, for wide buses, filters that consider all interactions require large areas and have high latencies. In practice, it is more practical to consider only the largest contributors to crosstalk, a small neighbourhood around each wire. Thus, the filter design problem is an optimization problem: given limits on the sample-rate, output swing, and filter width, design a filter that achieves the best possible signal integrity. We further simplify the problem by restricting our attention to linear filters.

We use eye height and width to quantify signal integrity. During each sampling interval, a binary signal should be either distinctly high or distinctly low. This allows the receiver to unambiguously determine the value of the bit that was transmitted. The signal can change between sample intervals. We also restrict how high (or low) the signal may go; otherwise, with scaling any eye opening can be made arbitrarily large. Eye height is $1 - \max(\textit{undershoot}, \textit{overshoot})$. Eye width is defined as the time period that high signal transmitted is distinct from low signal transmitted.

Many designers have employed least-squares optimization for the design of equalizing filters [1, 4]. Least squares optimization minimizes the power of the received crosstalk. The main drawback of least-squares optimization is that it uses an *average case* optimality criterion. For digital transmission, we want every bit to be transmitted successfully. Thus our real concern is to minimize the *worst case* error.

Our linear programming formulation is based on the observation that the output of the channel for any wire at any bit time is given by a linear combination of the contributions from each input wire at each bit time. The response of output i to input i after the nominal channel delay is the desired output. All other terms are disturbances. Without loss of generality, we can consider each input to be either $+1$ or -1. We compute the responses to each, individual $+1$ input and then determine the worst-case disturbance by taking the absolute values of each of these responses. This is readily translated into a linear programming problem whose solution gives the coefficients values for the filter.

As an example, figure 2 compares eye diagrams for a bus with no equalization, and with FIR filters optimized by the l_2 and l_∞ criteria. Both filters have four taps, four samples per bit, and consider the value input on the wire itself and its seven nearest neighbours to the left and seven nearest to the right. The bus model corresponds to 1 oz. copper PC board traces, that are 5 cm long, 3 mils wide with 9 mil spacing and a 42Ω characteristic impedance.

This bus has a tight wire spacing, as we expect for future interconnects. With no filter, the minimum bit time for which 50% eye height can be achieves is 687ps. Independent pre-emphasis for each wire gives little improvement (680ps could be achieved) showing that crosstalk is the primary signal integrity issue for this bus. With l_2 optimization, the minimum bit time can be reduced to 512ps (a 4 tap, 5 wide filter), adding more taps or filters to the l_2 results in a slight *increase* in the minimum bit-time, showing that the optimizer compromises worst-case performance to reduce the average-case error. With l_∞ optimization, a 349ps minimum is achieved with the 4-tap, 8-wide filter described above. Thus, our l_∞ approach provides 46% better performance than the commonly used l_2 method and 96% better performance than a channel without a filter. We are aware of no other published research describing l_∞ design of crosstalk cancelling filters. Our results demonstrate the effectiveness of this approach.

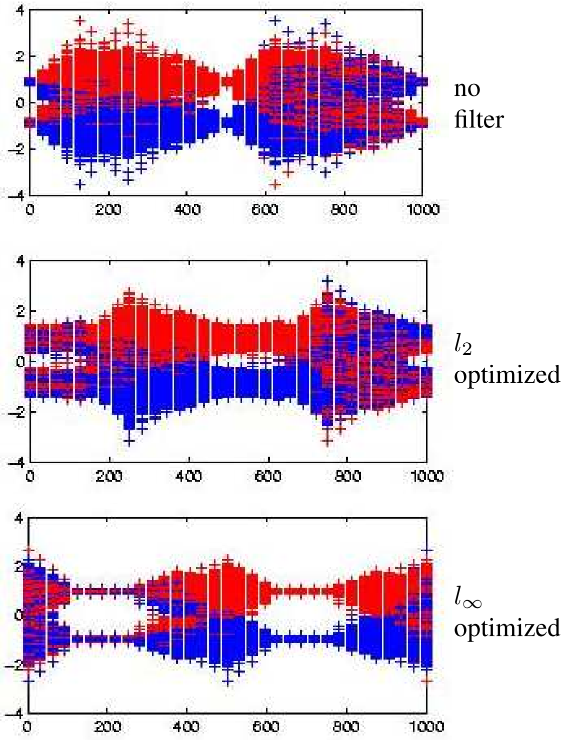

Figure 2. Eyes with and without filters

3 Hardware Implementation

Figure 3 shows an hardware implementation suitable for the filters described in this paper. Our equalizing filters is a multi-input, multi-output filters. We construct a separate filter for each output wire. Figure 3(a) shows our design for the filter for one output wire. It uses an interleaved DAC as described in [5]. The clock generator produces phases for enabling each DAC. A current summing circuit combines the DAC outputs to produce the filter output, v. For simplicity, we show a design where the interleaving factor for the DAC is the same as the oversampling rate of the filter. By using a separate filter for each DAC, the DACs are incorporated into the channel, and filter coefficients can be adjusted to compensate for variations between the DACs.

The filters for each DAC channel can be typical FIR designs. To further simplify the design, we precompute the convolutions of individual bits with the filter coefficients and store the results of these convolutions rather than the coefficients themselves. The filter for a single channel multiplies each of these coefficients by the corresponding bits for the wire and its neighbours. As the inputs to the filters are bits (0 or 1), these multiplications are straightforward.

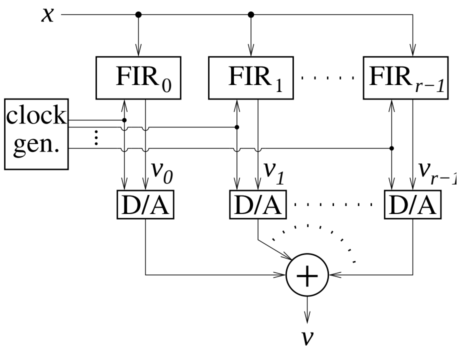

Figure 3. An Interleaved Equalizing Filter

These products are combined using an adder tree to produce the input to the DAC.

The total hardware required for our filters is quite small. As an example, we consider the the 4-tap, eight-wide filter with four filter taps per data bit as described in the previous section. We use 12-bit data paths to provide generous guard bits for 8-bit DACs. Our design requires 156 one-bit full-adders for each of the four interleaved DACs. Thus, a filter for each output pad can be constructed with fewer than 5000 transistors. Furthermore, the latency is very small. For the filter considered here, the adder tree has 15 inputs and a depth of four adders. Thus, it is reasonable to estimate that the filter adds less than $1ns$ to the latency of the channel for an implementation in a 0.13μ process. The significant bandwidth advantages, the small per-pad transistor count, and the low added latency demonstrate that crosstalk cancelling filters are a practical way to use on-chip computation resources to improve chip-to-chip signal integrity and bandwidth.

4 Conclusion

This paper explores the effectiveness of equalizing filters in crosstalk cancellation for high-speed, off-chip buses. It demonstrates that linear programming provides effective methods for designing crosstalk canceling equalizing filters that greatly increase the bandwidth of high-speed digital buses. For 5 cm long 32-bit wide PCB buses with (75 μm wire width and 225 μm separation), the channel with a crosstalk cancelling filter can operate at 2.9GHz and achieve 50% eye height. Without crosstalk cancellation, such buses can only operate at 1.45GHz.

In this paper, we described and compared filter design methods based on l_2 and l_∞ optimization methods. These correspond to least-squares and linear programming based optimization respectively. Although l_2 optimization has received the most attention for single-wire pre-emphasis, it cannot guarantee worst-case performance. The l_∞ metric corresponds to the traditional eye-height measure of signal integrity and provides guarantees of worst-case performance. In fact, the worst-case input is derived as a by-product of the optimization process. Our sample designs indicate that the l_∞ approach significantly outperforms the more common l_2 method for buses where crosstalk is a primary signal integrity concern.

Scaling trends of the VLSI technology favor the equalizing filter approach. Long buses cost more and support lower data rates. The cost of the bus justifies added circuitry on the chip. The lower data rate provide more time for the filtering operations. Furthermore, improvements in chip fabrication are producing smaller and faster circuits for implementing the filter while buses remain big and slow. This also contributes to the favorability of adding more sophisticated equalizing filters. Given these scaling trends, we expect that crosstalk cancelling equalizing filters will be essential for chip-to-chip interconnect in the near future. Our design methods provide a practical and effective way for designing these filters.

Acknowledgements

Uri Ascher has provided insightful suggestions with respect to the numerical optimization procedures used in this work. Robert Drost provided valuable comments on filter design and channel modeling. Ken Steiglitz provided feedback about equalizing filter design. The System-On-Chip group at UBC has provided valuable feedback throughout this research. This work has been supported in part by NSERC.

References

[1] W.J. Dally and J.W. Poulton. Transmitter equalization for 4-GBPs signaling. *IEEE Micro*, 1:48–56, 1997.

[2] A. Fiedler, R. Mactaggart, et al. A 1.0625Gbps transceiver with 2x-oversampling and transmit signal pre-emphasis. In *Proc. of ISSCC97*, pages 238–239, 1997.

[3] M. Horowitz, C. Ken, and S. Sidiropoulos. High speed electrical signalling: Overview and limitations. *IEEE Micro*, 18:12–24, 1998.

[4] V. Stojanovic, G. Ginis, and M.A. Horowitz. Transmit pre-emphasis for high-speed time-division-multiplexed serial-link transceiver. *IEEE Transactions on Communications*, 38:551–558, 2001.

[5] C.-H. Yang, V. Stojanovic, et al. A serial link based on 8 GSa/s A/D and D/A converters in 0.25μm cmos. *IEEE J. of Solid-State Circuits*, 36(12):1684–1692, December 2001.

[6] J.L. Zerbe, R.S. Chau, et al. A 2Gb/s/pin 4-PAM parallel bus interface with transmit crosstalk cancellation, equalization and integrating receivers. In *IEEE Int'l. Solid State Circuits Conf.*, pages 430–432, 2001.