High-Performance Electrical Signaling

William J. Dally[1], Ming-Ju Edward Lee[1], Fu-Tai An[1], John Poulton[2], and Steve Tell[2]

[1]Computer Systems Laboratory
Stanford University

[2]Department of Computer Science
University of North Carolina - Chapel Hill

Abstract

This paper reviews the technology of high-performance electrical signaling, presents the current state of the art, and projects future directions. We have demonstrated equalized electrical signaling between CMOS integrated circuits at data rates of 4Gb/s. As the factors that determine this signaling rate all scale with improving technology, we expect the data rates of high-performance electrical signaling systems to improve on a Moore's Law curve. The frequency-dependent attenuation of copper wires sets a bandwidth-distance squared (Bd^2) limit on the distance one can signal at a given data rate. Equalizing the channel cancels inter-symbol interference caused by this attenuation and greatly increases signaling distance. In the limit of perfect equalization, distance is ultimately limited by thermal noise in the receiver. At this limit, we calculate that a 4Gb/s system will be capable of operating over 100m of 24-gauge cable without repeaters.

1. Introduction

Data communication or signaling is a dominant power, performance, and cost factor in many digital systems. The connection between processors, caches, and main memory in computer systems; multiprocessor and multicomputer interconnection networks; and high-speed network switches are all critically dependent on signaling technology.

Conventional CMOS signaling has been limited to data rates of 100MHz or less and these data rates have not scaled with improving technology. Costly wide buses and high pin-count chips and modules have resulted from trying to meet the bandwidth demands of modern systems using this slow signaling technology.

We have recently demonstrated equalized CMOS I/O drivers that operate at 4Gb/s over moderate length copper interconnect [DP96, DP97, PDT98]. These designs operate at data rates that are limited by timing jitter, signal rise-time, and receiver aperture time all of which improve as technology scales [DP98b, HYS]. Thus, we expect the data rates of high-performance electrical signaling systems to improve on a Moore's Law curve.

At a given signaling rate, the distance over which one can reliably transmit a signal is limited by the frequency-dependent attenuation of the transmission medium. This limitation can be expressed as a bandwidth-distance squared, Bd^2, product. This Bd^2 product is proportional to the square of the attenuation that can be tolerated by the system.

Only a few dB of frequency-dependent attenuation can be tolerated by an unequalized signaling system before inter-symbol interference overwhelms the signal. Equalizing the channel cancels this inter-symbol interference increasing the tolerable attenuation and quadratically increasing the Bd^2 product. To date we have demonstrated a signaling system that is capable of equalizing up to 10dB of frequency-dependent attenuation. We expect to be able to equalize greater levels of attenuation as this technology evolves.

With perfect equalization, a signaling system is ultimately limited by thermal (Johnson) noise in the receiver. Using conservative assumptions about system performance, we calculate that a Johnson-noise limited system will be capable of communicating over 100m of 24-gauge cable at 4Gb/s without repeaters.

The remainder of this paper discusses the technology of high-performance electrical signaling in more detail. We begin in Section 2 by examining the limitations of conventional CMOS signaling and show how they are overcome by low-swing, incident-wave signaling. In Section 3 we explore the limitations on signaling rate due to electronics and wire. We see that the electronic factors that limit signaling rate all improve as semiconductor technology scales. We also see that wire characteristics place a limit on the Bd^2 product of a signaling technology and that the magnitude of this product is heavily dependent on equalization. We review our experience building an equalized CMOS signaling system in Section 4 and discuss plans for reducing the size and power required by this system.

2. High-Performance Signaling Basics

2.1 Traditional CMOS Signaling

Traditional CMOS signaling systems are limited to data rates of 100Mb/s per wire or less and dissipate large amounts of energy per bit transmitted. More importantly, traditional CMOS signaling rates do not scale with improving semiconductor technology. Because of this, many

modern microprocessors operate their external buses at small fraction of their internal clock rate.

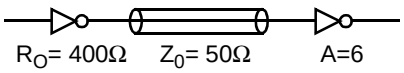

Figure 1: A traditional CMOS signaling system.

Figure 1 shows a traditional CMOS signaling system. CMOS inverters are used as both driver and receiver. The transmission medium, typically a cable or PC-board trace has a characteristic impedance of about 50Ω. The driver typically has an output impedance of 400Ω and the line is unterminated at the receiver. The two power supplies (V_{DD} = 3.3V and GND = 0V) are used to represent logic 1 and 0 respectively.

The traditional CMOS system is slow because the high impedance driver is unable to switch the line on the *incident wave*. Instead, the driver must *ring up* the line as the signal propagates over several round-trips of the line as shown in Figure 2. The figure shows the voltage as a fraction of V_{DD} at the far end of the line as a function of time. The incident wave only switches the line to 10% of V_{DD}. This is doubled by the first reflection at 10ns. Seven traversals of the 4ns line are required to ring the line up to 50% of V_{DD}, and over 15 traversals are needed for the line to converge to within 10% of its final value.

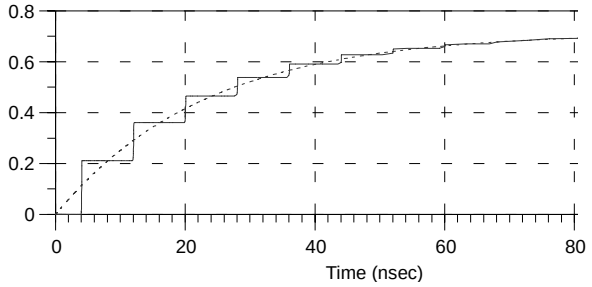

Figure 2: An 8mA CMOS Driver rings up a 4ns line terminated into 1KΩ.

The traditional CMOS system is power hungry, dissipating 1nJ or more to transmit each bit because it requires large signal swings to overcome poor noise isolation. The system uses noisy power supplies as transmit and receive voltage references. It also uses a CMOS inverter which has both a large offset voltage across process corners (typically 300mV) and poor sensitivity (about 500mV) as a receiver. To overcome these noise sources, a large signal swing, typically the full power supply, is required.

2.2 Low-Swing, Incident-Wave Signaling

A signaling system that overcomes the limitations of traditional CMOS signaling is shown in Figure 2. A current-source transmitter drives the line. A typical drive is ±5mA which gives a 250mV signal swing. The line is terminated at both ends into its characteristic impedance. The receiver termination absorbs the incident wave preventing any reflections. The source termination makes the system more tolerant of crosstalk and impedance discontinuities by absorbing any 'stray' waves that arrive back at the source. A high-gain clocked regenerative receive amplifier gives low offset (typically 10mV) and good sensitivity (10mV or better depending on settling time).

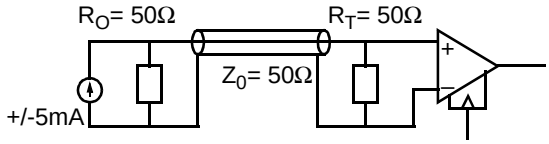

Figure 2: A low-swing, incident-wave signaling system.

By operating using only the incident wave, the system of Figure 2 can operate at a data rate independent of the length of the line. A new bit can be driven onto the source end of the line before the previous bit arrives at the receiver. As discussed below, this results in a data rate that, to first approximation, scales linearly with device speed.

This system operates reliably using a very small voltage swing (250mV) because it isolates the signal from many noise sources. Transmit and receive references are isolated from the noisy power supplies and a clocked receive amplifier greatly improves offset and sensitivity.

3. Limitations on Signaling Rate

The data rate of a signaling system is limited by both the electronics used to generate and receive the signal and the medium over which the signal propagates.

3.1 Electronic Limits on Signaling Rate

As illustrated in Figure 3, electronics limits signaling rate due to rise-time, aperture time, and timing uncertainty [DP98b]. The time for a bit cell, t_{bit}, must be made long enough for the signal to slew from one logic level to another, t_r, for the receiver to sample this signal while stable, t_a, and to tolerate jitter between the signal and the sampling clock, $2t_u$.

All three of these factors are related to the basic time constant of the semiconductor technology, τ_n, the time for a minimum-sized nFET to discharge the gate of an equal-sized nFET [DP98b]. This time constant is given approximately by

$$\tau_n = \frac{V_{DD} C_g}{I_{DSS}}. \qquad (1)$$

For a typical 0.35μm (drawn gate length) CMOS process, τ_n = 10ps, and this time constant scales linearly with gate length.

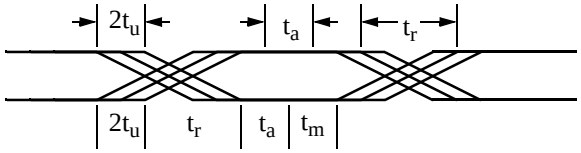

Figure 3: Abstract eye diagram showing limitations on signaling rate.

The transition time of the signal, t_r, is determined by the time required to slew the current source from one extreme to the other. Using a current source switched by logic gates it is easy to generate a slew time of $10\tau_n$, 100ps in an 0.35μm process. The RC time constant of the terminated line (25Ω) and the capacitance of the driver (about 1pF) is smaller than the driver slew time.

The aperture time of the receive amplifier, t_a, can be made as small as $2\tau_n$ by using a gate-isolated sense amplifier [DP98]. Using a pass-gate to sample the arriving signal gives an aperture time of about $10\tau_n$.

The peak-to-peak jitter, $2t_u$, is caused by jitter in the sampling clock, jitter in the receive clock, and delay variations in the signal path. All three sources of jitter are primarily due to power-supply modulation of delay and crosstalk-induced delay variation. Both of these factors scale with delay and hence with process gate length. Using circuits with good power supply rejection and layout techniques to avoid crosstalk, peak-to-peak jitter of less than $5\tau_n$ can be achieved.

As gate lengths shrink, variation in device parameters, such as V_T is expected to increase. This variation will lead to a fixed variation in, for example, the delay elements used to build timing circuits. Left uncompensated, this effect would give fixed-pattern jitter term that does not scale linearly with process technology. Fortunately, it is possible to measure these variations and compensate for them.

With careful design, an incident-wave signaling system can achieve a bit cell width of 15-30τ_n (150-300ps in 0.35μm technology). Most importantly, because there is no need to ring up the signal line, this signaling rate scales linearly with the basic time constant of the technology, τ_n, independent of the wire length. As gate lengths continue to shrink according to the SIA roadmap [SIA97] the speed of signaling systems can improve at the same rate.

3.2 Wire Limits on Signaling Rate

While the speed of the transmitter and receiver improve as the gate length of semiconductor technology shrinks, the bandwith of the wire used to connect the transmitter to the receiver remains constant. For interconnects that traverse any appreciable distance, the bandwidth of this wire is rapidly becoming the limiting factor. We can express this limitation as a bandwidth-distance squared product with units of bits×m^2/s [MO97].

Frequency-dependent attenuation due to the skin-effect resistance of the wire is the major factor limiting wire bandwidth. Impedance discontinuities are not a limiting factor. In a well designed system, the transmission path from transmitter to receiver has a constant impedance with no appreciable discontinuities to attenuate the signal and cause inter-symbol interference. With some insulating materials, dielectric absorption also causes frequency-dependent attenuation; however this effect can be mitigated by using a low-loss dielectric.

The skin-effect attenuation in dB at frequency f of a wire with length d is given by

$$A(d,f) = A_1 d\left(\frac{f}{f_1}\right)^{\frac{1}{2}} \quad (2)$$

where A_1 is the attenuation of a 1m wire at frequency f_1. For example, the attenuation of a 100Ω, 24AWG twisted pair is 0.45dB/m at 1GHz (A_1 = 0.45dB/m, f_1 = 1GHz). At 10GHz a 1m wire has an attenuation of 1.4dB/m and a 10m wire has an attenuation of 14dB. Attenuation is also proportional to the inverse of the wire radius, so one can get lower attenuation by using fatter wires. However, for the remainder of this discussion we will use the 24AWG cable for our illustrative example.

Because this attenuation is frequency dependent it causes inter-symbol interference in addition to reducing signal energy. At an attenuation of 6dB, the unattenuated low frequency components of the signal completely swamp the high frequency components. Because of this effect, transmission without equalization is usually limited to an attenuation of 2dB or less. If our tolerable attenuation in dB is A_T and our bandwidth in bits/s is $B = 2f$, then we can rewrite Equation (2) as

$$Bd^2 = 2f_1\left(\frac{A_T}{A_1}\right)^2 \quad (3)$$

For an unequalized 24AWG pair, the Bd^2 product is limited to 4×10^{10} bits×m^2/s. Without equalization this cable can carry 40Gb/s over 1m or 400Mb/s over 10m.

Equalizing the signal, by pre-emphasizing the high-frequency components of the signal before transmission, eliminates the inter-symbol interference caused by the frequency-dependent nature of the attenuation. To date we have demonstrated signaling systems that can equalize up to 10dB of frequency-dependent attenuation. From Equation (3) the Bd^2 product for a 24-gauge cable with this limited equalization is 10^{12} bits×m^2/s, 1Tb/s over 1m or 10Gb/s over 10m. Better equalization can improve this capacity further.

With perfect equalization, the amount of attenuation that can be tolerated, and hence the Bd^2 product is limited by the noise floor. A receiver limited by Johnson noise in a 50Ω environment at 300°K has a noise floor in dBV of

$$N = 10\log_{10}(f) - 180 \quad (4)$$

Assuming a transmit signal level of 1V and a minimum SNR of 20dB for reliable signaling, we can from Equations (3) and (4) derive

$$Bd^2 < 2f_1\left(\frac{-N-20}{A_1}\right)^2$$
$$Bd^2 < 10^{10}[163 - 10\log_{10}(B)]^2 \quad (5)$$

In this case Bd^2 is not a constant since the bandwidth determines both A and A_T. Distance rolls off just slightly faster than the square root of bandwidth:

$$d < 10^5 B^{-1/2}[163 - 10\log_{10}(B)] \quad (6)$$

Equation (6) gives a maximum distance of 230m at 1Gb/s, 63m at 10Gb/s, and 17m at 100Gb/s using 24AWG pair. This approximately quadratic rolloff shows that the $B^{1/2}$ term dominates Equation (5). These distances are not hard physical limits but rather estimates of what can be accomplished with good engineering. One may be able to build receivers with a better noise temperature than 300°K, with ideal coding, one can recover more than one bit per symbol with a 20dB SNR, and one could operate with higher transmit power levels.

The distance limits calculated by Equation (5) assumes that we can build an equalized channel to correct for a frequency-dependent attenuation equal to the expression in square brackets. This ranges from 73dB at 1Gb/s to 53dB at 100Gb/s. To date we have implemented Gb/s-rate equalizers that correct for up to 10dB of frequency dependent attenuation. Closing the gap between 10dB and 70dB requires solving many engineering problems to build more precise equalizers. Also, to achieve this limit we must cancel many sources of receiver offset and interference to build a high-speed receiver that is limited by thermal noise. The fact that most radio receivers operate at the thermal limit, however, suggests that this goal is feasible.

4. Experimental Signaling Systems

To demonstrate the concepts of equalized, high-speed, incident-wave signaling we started a research program in 1995 to build an equalized 4Gb/s transceiver in 0.5µm CMOS technology. To date we have demonstrated one test chip and are in the process of designing a second. In this section we report on our experience with these experiments.

4.1 Equalized Signaling Test Chip

Figure 4 shows a die plot of our first prototype equalized 4Gb/s transceiver chip which was designed in the Spring of 1996. The transmitter along with a test pattern generator is at the bottom of the chip. The receiver and a test pattern checker occupy the top portion of the chip. The design of this chip along with simulation results are presented in [DP96, DP97]. Actual test results are reported in [PDT98].

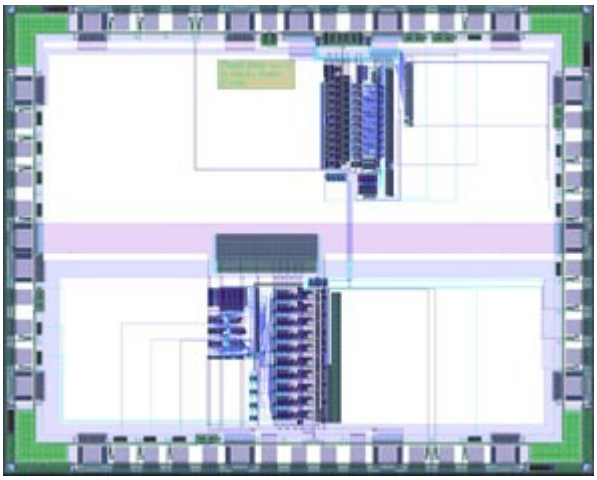

Figure 4: Die plot of 4Gb/s transceiver test chip

A block diagram of the transmitter portion of this test chip is shown in Figure 5. The transmitter operates on 10 bits in parallel at 400MHz and creates a 4Gb/s output signal by multiplexing ten independent signals together under control of a precision 10-phase clock.

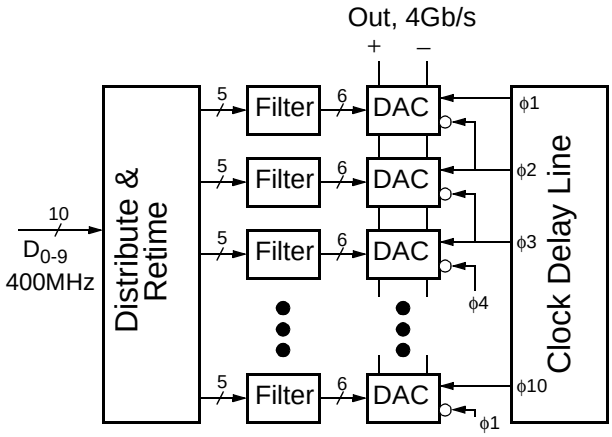

Figure 5: Block diagram of 4Gb/s transmitter

Equalization is performed using a bank of ten five-tap digital filters. Each filter examines a data bit along with the four immediately preceding bits and computes a 4-bit output (represented as a pair of three bit numbers). The filters are implemented as SRAM look-up tables for speed and efficiency. Each filter's output is used to drive a digital-to-analog converter (DAC). The DACs are sequenced

on to the line by a clock generator to generate the output signal.

Output waveforms from the transmitter are shown in Figure 6. The figure shows an isolated 250ps *one* pulse in a field of *zeros* being transmitted across 1m of 30AWG twisted-pair cable. The waveform on the left is without equalization. With 10dB of attenuation there is no eye opening and the one is undetectable. With equalization, as shown on the right, a clean 250ps eye opening is visible.

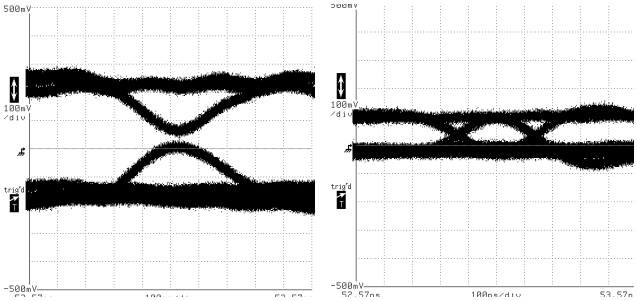

Figure 6: Measured waveforms from 4Gb/s test chip with equalization off (left) and on (right).

A demultiplexing receiver, also operating at 400MHz is used to receive the signal as shown in Figure 7. The line is received by a bank of 20 clocked sense amplifiers. Each amplifier is allowed 1.25ns, about 10 time constants, to settle giving very good sensitivity. The amplifiers sample the 10 data bits as well as the 10 potential edges between the bits. The edge samples are used to control the 20-phase clock. By centering the odd samples on the edges, the even samples remain centered on the eye openings.

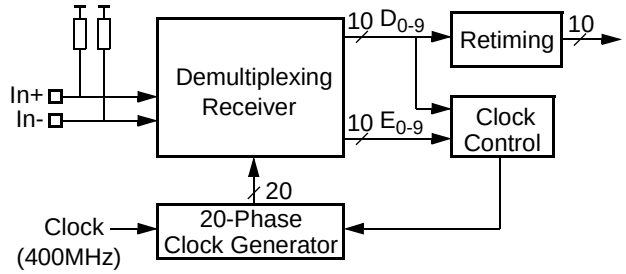

Figure 7: Block diagram of 4Gb/s receiver

4.2 Input Multiplexed Architecture

While our initial test chip was successful in demonstrating equalized signaling at 4Gb/s rates it had two significant shortcomings: (a) the transmitter was very large, measuring 1mm x 0.5mm, and (b) the source-coupled delay lines used for transmitter and receiver timing consumed considerable power. While our first design is adequate for a chip with up to a few tens of transceivers, a smaller, lower-power transceiver is required to enable hundreds of 4Gb/s transceivers to be placed on a chip.

After experimenting with several alternatives, we arrived at the input-multiplexed transmitter architecture using analog equalization to achieve our size and power goals. A block diagram of this transmitter is shown in Figure 8. Equalization is performed by a two-tap analog filter that is realized by having the main transmitter and an equalizing duplicate transmitter sum their output currents on the line. The equalizing duplicate operates one bit delayed and with programmable weight. Longer analog FIR filters can be constructed by building additional duplicates each delayed one bit from the previous duplicate.

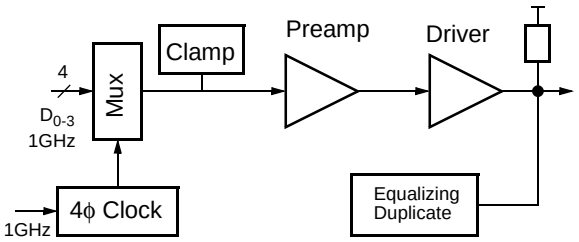

Figure 8: Input-multiplexed transmitter architecture with analog equalization

The input-multiplexed transmitter of Figure 8 is considerably smaller than the output-multiplexed transmitter of Figure 7. While layout of the new transmitter is not yet complete, we estimate that it will fit in an area of 0.1mm × 0.1mm, about 2% of the area of the old transmitter. This reduction in area is due to three factors. First, there is only a single copy of the transmitter circuitry as opposed to ten copies in the previous design. Second, using an analog equalizing filter eliminates the costly look-up table RAMs. Finally, operating the transmitter circuitry at 1GHz rather than 400MHz allows us to operate with 4:1 multiplexing instead of 10:1 multiplexing reducing the required size of the multiplexer and clock delay line.

The new design is also considerably more power-efficient than the previous design. While we have not completed our power studies at this time, we anticipate considerable gains in three areas. First, by moving the multiplexer to the input the clock load, and hence the clock power, is reduced by an order of magnitude. Second, this design uses a clock circuit based on regulated CMOS delay stages, as opposed to the source-coupled delay stages in the previous design, that dissipates about a third the power for a given clock load. Finally, all of the power previously dissipated in the digital filters has been eliminated.

5. Conclusion

Incident-wave electrical signaling systems currently operate at data rates up to 4Gb/s. We expect these signaling rates to improve on a Moore's law curve over time because all of the electronic factors limiting performance scale with improving semiconductor technology.

At a given signaling rate, signaling distance, without repeaters, is limited by the frequency-dependent attenuation of the transmission medium. This limitation can be

described by a bandwidth-distance squared constant for a given medium and allowable attenuation. Without equalization, frequency-dependent attenuation must be kept to less than 2dB giving a Bd^2 constant of 4×10^{10} bits×m^2/s (3m at 4Gb/s) for a typical 24 gauge cable.

Equalization can flatten the attenuation profile of a signaling system, greatly increasing the amount of frequency-dependent attenuation that can be tolerated. We have demonstrated a 4Gb/s equalized signaling system that can tolerate up to 10dB of frequency dependent attenuation improving the Bd^2 constant to 10^{12} bits×m^2/s (16m at 4Gb/s). We expect high-data rate equalization to improve to the point that the tolerable attenuation is limited by Johnson noise in the receiver to levels between 50 and 70dB depending on signaling rate. At these attenuation levels, the Bd^2 limit is between 2.5 and 5×10^{13} (79 to 111m at 4Gb/s).

High-speed equalized transceivers can be built in a fraction of a square mm using a standard CMOS process and operate at power levels of about 100mW each. One can place hundreds of these transceivers on a standard CMOS chip much in the way one uses conventional CMOS I/O drivers today. The cost per transceivers is about the same as the cost of a conventional CMOS I/O driver.

We expect the low cost, high and improving performance, and high level of integration of equalized CMOS signaling to make it the technology of choice for applications within its Bd^2 envelope. Even outside this envelope, CMOS signaling with repeaters or parallel CMOS signaling at a lower per-line bit rate are likely to be cost-effective alternatives to a single high-speed optical link over moderate distances.

6. Acknowledgments

We thank Mark Horowitz and his students for many helpful suggestions. This work is sponsored in part by the Defense Advanced Research Projects Agency (DARPA) under DARPA Order E253 and monitored by the Army under contract number DABT63-96-C-0039 and in part by Intel Corporation.

7. References

[DP96] William J. Dally and John Poulton, “Equalized 4Gb/s CMOS Signaling”, *Proceedings of Hot Interconnects IV,* pp. 29-39, 1996.

[DP97] William J. Dally and John Poulton, “Equalized 4Gbps CMOS Signaling,” *IEEE Micro,* pp. 48-56, Jan-Feb 1997.

[PDT98] John Poulton, William J. Dally, and Steve Tell, “A Tracking Clock-Recovery Receiver for 4Gbps Signaling,” *IEEE Micro,* pp. 25-26, Jan-Feb 1998.

[DP98]William J. Dally and John W. Poulton, *Digital Systems Engineering*, Cambridge University Press, 1998.

[HYS] Mark Horowitz, Chih-Kong Ken Yang, and Stefanos Sidiropoulos, “High-Speed Electrical Signaling: Overview and Limitations,” *IEEE Micro,* pp. 12-24, Jan-Feb 1998.

[MO97] D. A. B. Miller and H. M. Özaktas, “Limit to the Bit-Rate Capacity of Electrical Interconnects from the Aspect Ratio of the System Architecture,” Special Issue on Parallel Computing with Optical Interconnects, *J. Parallel and Distributed Computing* 41, pp. 42-52, 1997.

[SIA97] Semiconductor Industry Association, *The National Technology Roadmap for Semiconductors: Technology Needs*, 1997.