

Power Comparison Between High-Speed Electrical and Optical Interconnects for Interchip Communication

Hoyeol Cho, *Student Member, IEEE*, Pawan Kapur, and Krishna C. Saraswat, *Fellow, IEEE*

Abstract—An I/O bandwidth commensurate with a dramatically increasing on-chip computational capability is highly desirable. Achieving this goal using board-level copper interconnects in the future will become increasingly challenging owing to severe increase in high-frequency, skin-effect and dielectric loss, noise due to crosstalk, impedance mismatch, and package reflections. The solutions designed to overcome these deleterious effects require complex signal processing at the interconnect endpoints, which results in a larger power and area requirement. Optical interconnects offer a powerful alternative, potentially at a lower power. Prior work in comparing the two technologies has entailed overly simplified assumptions pertaining to either the optical or the electrical system. In this paper, we draw a more realistic power comparison with respect to the relevant parameters such as bandwidth, interconnect length and bit error rate (BER) by capturing the essential complexity in both types of interconnect systems. At the same time, we preserve the simplicity by using mostly analytical models, verified by SPICE simulations where possible. We also identify critical device and system parameters, which have a large effect on power dissipation in each type of interconnect, while quantifying the severity of their impact. For optical interconnect, these parameters are detector and modulator capacitance, responsivity, coupling efficiency and modulator type; whereas, in the case of electrical system, the critical parameters include receiver sensitivity/offset and impedance mismatch. Toward this end, we first present an optimization scheme to minimize optical interconnect power and quantify its performance as a function of future technology nodes. Next, on the electrical interconnect side, we examine the power dissipation of a state-of-the-art electrical interconnect, which uses simultaneous bidirectional signaling with transmitter equalization and on-chip noise cancellation. Finally, we draw extensive comparisons between optical and electrical interconnects. As an example, for bandwidth of 6 Gb/s at 100 nm technology node, lengths greater than the critical length of about 43 cm yields lower power in optical interconnects. This length becomes lower (making optics more favorable) with higher data rates and lower bit error rate requirement.

Index Terms—Dielectric and skin effect attenuation, electrical interconnect, equalization, interchip communication, modulator, noise modeling, optical interconnect, power comparison, power modeling, simultaneous bidirectional signaling, transimpedance amplifier.

I. INTRODUCTION

DIFFERENT classes of digital systems impose specific requirements on the communication medium. These requirements pertain to the communication length scale and the figure of merit of relevance (bandwidth or latency). The choice of the communication medium is heavily dependent on these factors. For

example, long-haul systems ubiquitously use optical fibers because of low attenuation at high bandwidths. Systems at shorter length scales have traditionally used copper (Cu) interconnects for both latency and bandwidth sensitive applications. However, as the computational bandwidth of the modern integrated circuits (ICs) (measured by the product of the number of transistors and the clock frequency) increases dramatically according to the Moore's law, Cu traces at short distances at least in bandwidth sensitive applications are struggling to keep up, rendering communication bandwidth a bottleneck. This presents a fertile ground for optical medium of communication to penetrate the short distance world, albeit with very different constraints compared to long-haul communication.

In the short distance interconnect world (<100 m), the interconnect hierarchy is typically divided into the following categories in the order of progressively smaller length spans: Cabinet level (1–100 m), backplane level between boards (10 cm–1 m), chip to chip on a board (<10 cm) and on-chip (<2 cm) communication [1]. At the smallest length scales (on-chip), the possibility of insertion of optics, although attractive from latency and possibly power point of view [2] remains inconclusive owing to the existence of cheaper solutions in metal domain at circuits and architectural level. These solutions include repeater insertion (for latency and bandwidth) [3], low swing signaling (for power) [4] and interconnect centric architectures minimizing communication over long distances [5]. The next larger length span constituting communication between two chips on the same board (<10 cm), such as between microprocessor and memory, stresses latency of the communication medium and as such is also not particularly conducive to optics. There are two reasons for this. First, the medium (propagation) latency of optics and off-chip Cu wires which are *LC* in nature is comparable and is based on the speed of light in the respective media. In fact, optics could even be slower due to optoelectronic conversion overhead. Second, the major source of latency in these systems is in the memory itself, where optics cannot help [6].

However, the digital systems which put a premium on communication bandwidth rather than propagation latency can benefit enormously from the choice of optical medium. These systems include multiprocessor, distributed computational systems (e.g., servers) as well as network routers where many boards with multiple chips are communicating with each other or with cross-bar switch boards. The length scales are typically between 10 cm to 1 m and the communication could be between boards connected to a backplane through high-speed connectors. As the demand on aggregate communication bandwidth increases in these systems, because of limited board space, connector density, and/or pin

Manuscript received December 15, 2003; revised June 11, 2004.

The authors are with the Department of Electrical Engineering, Stanford University, Stanford, CA 94305 USA (e-mail: hycho@stanford.edu).

Digital Object Identifier 10.1109/JLT.2004.833531

count, the bandwidth per interconnect is taxed. However, at high bandwidth and large distances, on Cu medium, several factors conspire simultaneously to cause insufficient signal-to-noise ratio (SNR) for reasonable data fidelity. The primary problem is a sharply increasing attenuation over Cu traces which causes signal loss and intersymbol interference (ISI) at the receiver. The attenuation is due to skin effect resistance, which exhibits a square root dependence on frequency, and due to inter- and intralevel dielectric leakage, which increases linearly in frequency. To complicate things further, many sources of noise including connector and trace cross talk (forward and backward), impedance mismatch and package induced reflections become worse, further taxing an already attenuated signal. Simultaneously, the timing reliability dictated by signal and clock uncertainty (skew and jitter) also becomes harder to meet. The commonly deployed mitigating solutions related to attenuation include passive as well as active equalization at transmitter and/or receiver [7], and the use of lower loss tangent dielectrics (example GETEK instead of FR4 board), which tends to be expensive. Even if the data fidelity is met using these solutions, the penalty shows up as excessive power dissipation.

Optical interconnects with low signal attenuation and crosstalk could potentially be very useful in short distance, bandwidth sensitive applications from two standpoints. First, they can provide the high bit rates which Cu may not be able to. Second, even if Cu is able to provide these bit rates, optics may be able to do this at a much lower power and area. These requirements are in stark contrast with long-haul communication requirements, where the primary objective is to achieve lowest possible noise levels in the receiver systems due to scarcity of optical signal, even if it is at the expense of power dissipation. The difference in system requirements ensures that the optical system/component design is not necessarily portable from long-haul to short distance communication and at the same time opens up many different possibilities in terms of choice of optical medium, communication wavelengths and optoelectronic device materials. Already many different backplane optical media have been demonstrated and/or their prospect discussed including polymer waveguides [8], fiber image guides (FIGs) [9], [10], fiber ribbons [11], and free space optical interconnects (FSOI) using lens and mirror system [12]–[15].

Several discussions on the performance of electrical and optical interconnects including their comparisons in the off-chip regime have been published [16]–[19]. Some of these studies ignore the role of end-devices [16], [17]; while most of them do not consider the complexity and the sophistication of the state-of-the-art electrical links. The power modeling of optical interconnects, comparing vertical cavity surface emitting laser (VCSEL) and multi-quantum-well (MQW) modulator, has been detailed at length before [20], [21]. In this paper, we take a more comprehensive view of both Cu and optical systems for short distance, off-chip, bandwidth-sensitive applications. Our primary objective is to compare power dissipation with respect to the relevant parameters such as bandwidth, interconnect length, and bit error rate (BER) by capturing the essential complexity of the two systems. To accomplish this goal we first power optimize the optical system by design [21], [22]. On the electrical side, we choose a sophisticated,

state-of-the-art interconnect for a fair comparison. We model the electrical interconnect attenuation (analytically and using SPICE) including the package effects as well as model various noise sources end-to-end in both systems. In the process, we also identify critical device/system parameters that have the maximum impact on power dissipation in each type of interconnect, while quantifying the severity of their impact. For optical interconnects, these parameters include detector capacitance, coupling efficiency and modulator type, while for electrical interconnect they are receiver sensitivity/offset and impedance mismatch. The rest of the paper is organized as follows. In Section II, we present a power-optimizing method for optical interconnects and quantify the system's performance at future nodes. In Section III we tackle the issues related to power modeling in electrical interconnects. We assume a system, which uses simultaneous bidirectional signaling with transmitter equalization and on-chip noise cancellation. Such interconnects are extremely effective in pin-limited systems. Section IV contains a power comparison between electrical and optical interconnects in terms of the critical length defined as the length beyond which optics becomes more power efficient. This length is characterized as a function of various system requirements (data rate, BER). Finally, we conclude in Section V.

II. OPTICAL INTERCONNECT POWER DISSIPATION

The optical interconnect power consists of the transmitter and the receiver powers. In this paper, we assume an off-chip laser source at 1.3 μm wavelength providing light to silicon CMOS driven modulators. This scheme is an alternative to directly modulating hybrid-integrated vertical cavity surface emitting lasers (VCSELs). We chose the modulator-based scheme because VCSELs are prone to serious reliability problems in a harsh, high temperature CMOS environment. The drawback of this scheme is that modulators suffer from relatively low contrast ratio and high insertion loss, which can increase receiver power dissipation. This is especially true in the light of scaling CMOS voltages on a chip. Although, the I/O voltages tend to be higher than normal chip voltages and also tend to scale slower. At high bit rates, VCSEL-based interconnect has about the same total link power with that of MQW modulator-based interconnect [21]. The choice of the wavelength not only gives a larger number of photons for a given optical power compared to 850 nm (hence, larger detector responsivity) but also allows the possibility of monolithic integration of Ge photodetectors directly onto silicon substrate without the danger of noise in silicon circuits. Although, monolithic integration possibility for both detectors and modulators exists with this scheme, our interconnect performance trends in this paper are more representative of the indium phosphite (InP)-based devices that are hybrid-bonded to Si-CMOS. These devices are reversed-biased PIN quantum-well detectors and quantum-well modulators. The hybrid-bonded techniques can already be pushed to yield detector capacitance as low as 50 fF including the bond pad and solder capacitances [21], [23]. The schematic of the analyzed high-speed optical interconnect is shown in Fig. 1. The optical medium could be any of the aforementioned possibilities. Its relevance on power calculation is captured through its attenuation and coupling efficiency from/to end-devices.

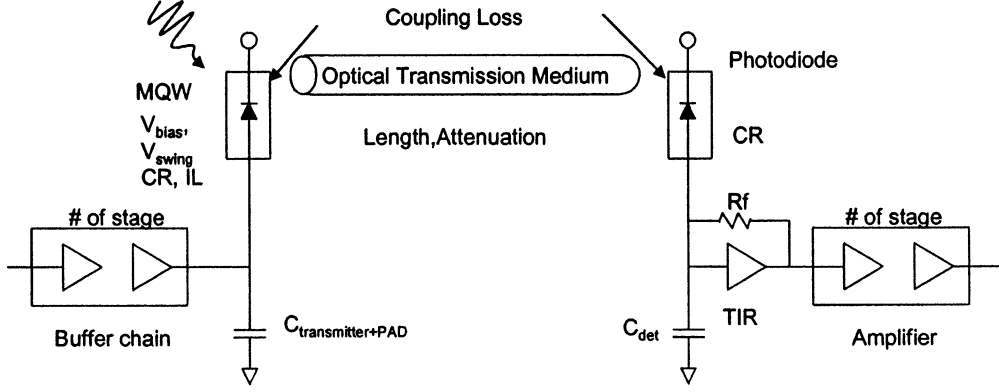


Fig. 1. Schematic showing board-level high-speed optical interconnect.

A. Modulator Power Dissipation

Both the dynamic and the absorption related, static, components of the modulator power dissipation are considered. Modulator is driven by an exponentially sized buffer chain to minimize delay and increase speed. Thus, the dynamic power includes the capacitance of both the modulator and the buffer chain. The static power dissipation depends on the absorbed optical power in the “on” and the “off” state. An ideal modulator should have zero insertion loss (IL) (optical power absorbed during the “on” state) and infinite contrast ratio (CR) (ratio of modulator output optical power in the “on” and the “off” states). However, all modulators exhibit a nonzero IL and only a finite CR. We use the CR and IL to calculate the current in each of the binary state. Next, we multiply it by respective voltage biases and take the average to get power dissipation [21]. Thus, for a one to one transmitter/receiver pair, we end up with

$$P_{\text{static modulator}} = \frac{P_{\text{optrec}}}{\eta} \frac{q}{h\nu} \times \left(\frac{V_{\text{bias}} \left(1 + \text{IL} - \frac{1 - \text{IL}}{\text{CR}} \right) - V_{\text{dd}} \text{IL}}{(1 - \text{IL}) \left(1 - \frac{1}{\text{CR}} \right)} \right). \quad (1)$$

Here, P_{optrec} is the average optical power at the receiver, ν is the frequency of the laser source, V_{bias} is the dc bias applied to the modulator in the highly absorbing state. V_{dd} is the voltage swing (supply voltage of the CMOS generation). In the less absorbing (“on”) state, the modulator is driven to a lower voltage of $V_{\text{bias}} - V_{\text{dd}}$. η , the optical power transfer efficiency, is the optical power at the receiver divided by that at the output of the modulator. The efficiency is usually less than one due to both the coupling losses at the transmitter and the receiver ends as well as the losses incurred by the optical transmission medium. In this paper, the plots in which we have explicitly used interconnect length as an independent variable, correspond to the waveguide medium. Here, the loss was calculated to be about 0.082 dB/cm at wavelength of 1.3 μm using theoretical analysis in [24] and the published values in [21]. As is obvious from (1), we can lower the static modulator power by operating at near zero bias voltage in the highly absorbing state (“0”) and V_{dd} bias in the less absorbing state (“1”). Modulators exhibiting properties close to this ideal modulator can be realized by using resonant cavity around QWM. At low bias, they exhibit absorption such that the front and the effective back mirror reflectivities are the same, leading to maximum passes in the cavity and high total absorption. Increasing the bias, although increases individual QW

absorption, but decreases the effective back mirror reflectivity leading to an asymmetric cavity with less passes, hence, lesser total absorption. We will refer to an ideal modulator along these lines as modulator 1 and the commonly used reflective mode modulator as modulator 2. The bias, CR and IL values for modulator 2 are taken from [21]. An optimization of IL and CR in modulators to further minimize total interconnect power dissipation is possible [22], however not considered in this paper.

B. Receiver Power Dissipation

The optical receiver is assumed to be the photodetector (responsivity ~ 0.5) followed by nonintegrating transimpedance amplifier and gain stages [25]. Its design and power dissipation is detailed in an earlier work [26]. The analytical design included establishing the width, the feedback resistance of the front-end, and the number of subsequent gain stages at a given input optical power (IOP) and detector capacitance. It was constrained by bandwidth, BER through receiver noise, and supply level output swing requirements. The transistor related parameters were taken from the International Technology Roadmap for Semiconductors (ITRS) [27]. The power was found to reduce with higher IOP and lower detector capacitance. Building upon that work, we have now verified the design at 180 nm technology node using BSIM3v3 technology¹ SPICE simulations. As an example, we show the SPICE generated eye-diagram of an analytically designed 4 Gb/s receiver with 20 μW reflectivity difference (RD; receiver optical power difference between the on and the off states) and 100 fF detector capacitance (Table I and Fig. 2). Although, the noise is not simulated here, the clean eye indicates sufficient bandwidth. Table I also shows a reasonable agreement between power dissipation obtained using SPICE and analytically calculated values. A small difference is accounted by two factors. First, the SPICE simulations required an additional gain stage (5 stages) compared to the analytical model (4 stages). This is because in the model we assume the postamplifier gain to be the dc gain (product of transconductance and the output resistance). This turns out to be a slight overestimate at high bit rates, hence an additional stage. Second, we only consider static power dissipation at each stage in the analytical model. This is a good approximation at low swings for most of the gain stages, but toward the last stage as signal becomes large, the static power is reduced. However, this is com-

¹University of California, Berkeley Device Group, CA, USA [Online]. Available: www-device.eecs.berkeley.edu/~ptm/download.html

TABLE I
AN EXAMPLE OF OPTIMIZED DESIGN PARAMETERS AND POWER DISSIPATION

Analytically obtained Optimized design parameters	
Front-end transistor size	94λ (2λ=180nm technology node)
Feedback resistance	375 Ω
Number of post-amplifier stage	4
Power dissipation (bit rate=4Gb/s)	
Analytical modeling	22.86mW
SPICE simulation	18.10mW (4 stages)
(number of post amplifiers)	21.02mW (5 stages)




Fig. 2. Eye diagram@output of transimpedance receiver.

compensated by the increase in the dynamic power dissipation. The favorable comparison with SPICE simulations lends confidence in our analytical models, especially for future technology nodes where SPICE technology files may not exist.

C. Power Dissipation Minimization

Fig. 3 illustrates our optical interconnect power minimization methodology. The increase in the optical power increases the modulator power but decreases the receiver power as discussed earlier. This leads to an optimal laser power at which total interconnect power (receiver and modulator) is minimized. The figure demonstrates the minimization for two different losses and for modulators 1 and 2. The receiver power does not change with laser power on continuous bases beyond a certain point as it goes into gain-limited regime [24]. As expected, modulator 2 yields larger power dissipation than modulator 1. Also, the receiver power is dominant over the modulator power. A higher loss (6 dB in Fig. 2) in optical power through either less efficient coupling and/or greater attenuation in the optical interconnect (longer lengths), results in a larger receiver power dissipation. This is due to a lower reflectivity difference between the on and the off states at the receiver with larger optical power loss, which is tantamount to digital signal-to-noise ratio (DSNR) deterioration. An increase in receiver power, in turn, results in a larger optimal laser power and total power dissipation. The numerator of the DSNR is the product of the reflectivity difference and responsivity of the photodetector (R_s) and is given by

$$\begin{aligned} I_{\text{on}} - I_{\text{off}} &= R_s(\text{RD}) = R_s(P_{\text{on}} - P_{\text{off}}) \\ &= R_s \eta P_l (1 - \text{IL}) \left(1 - \frac{1}{\text{CR}}\right) \end{aligned} \quad (2)$$

where $I_{\text{on}}(P_{\text{on}})$ and $I_{\text{off}}(P_{\text{off}})$ is the photodetector current (receiver optical power) in the on and the off states, respectively, and P_l is the laser power. Thus, in the high loss (low RD) case, the noise in the receiver has to be reduced to maintain same DSNR. This costs power dissipation. Fig. 4 exhibits both the optimum laser power and the resulting minimum power dissipation as a function of loss for two different bit rates (2 and 6 Gb/s). Increase in the power dissipation with bit rate is almost entirely due to a larger power dissipation in the receiver at higher bit rates. Fig. 4 also shows the reduction in power dissipation with technology scaling (100 and 50 nm) due to improvement in receiver transistors. The detector capacitance of 250 fF in this plot is somewhat pessimistic. Capacitance approaching 50 fF has been demonstrated [22]. We will later quantify the impact of lowering the detector capacitance on power dissipation.

III. ELECTRICAL INTERCONNECT POWER DISSIPATION

For the electrical interconnect, we choose a link capable of simultaneous bidirectional signaling. This scheme, by enabling full-duplex channels, provides higher aggregate bandwidth over smaller number of pins, a scenario particularly useful in pin-limited chips. However, it uses a more complicated detection scheme, where a transmitter replica is fed as a reference to the differential amplifier receiver to isolate the received and the transmitted signal (Fig. 5). We also choose low swing current mode, bipolar, differential signaling scheme. The rationale for these choices was driven by the achievement of maximum noise immunity. Differential signaling dramatically reduces signal return crosstalk and facilitates a noise-free receiver [28]. High impedance of current mode signaling, further, enables noise immunity to power supply [28].

The board trace dimensions are chosen to yield characteristic impedance Z_0 of 45 Ω (Fig. 6) on a high performance GETEK board. This board, although expensive compared to usual FR4 board, provides lower dielectric loss, hence, lower signal attenuation (loss tangent is 0.01). Striplines, as opposed to microstrips, are used to eliminate forward crosstalk. Flip-chip package model with lower parasitic inductance and capacitance is chosen (Fig. 5, [28]). Complete ISI cancellation is assumed using a transmitter side preemphasis equalization with multi-tap FIR filter [29]. The rise time is taken to be a third of the bit period. This is assumed to provide a reasonable compromise between smaller rise time requirement for adequate timing margin and larger rise time for lower noise. An on-chip cancellation circuit to reduce reverse channel crosstalk due to package

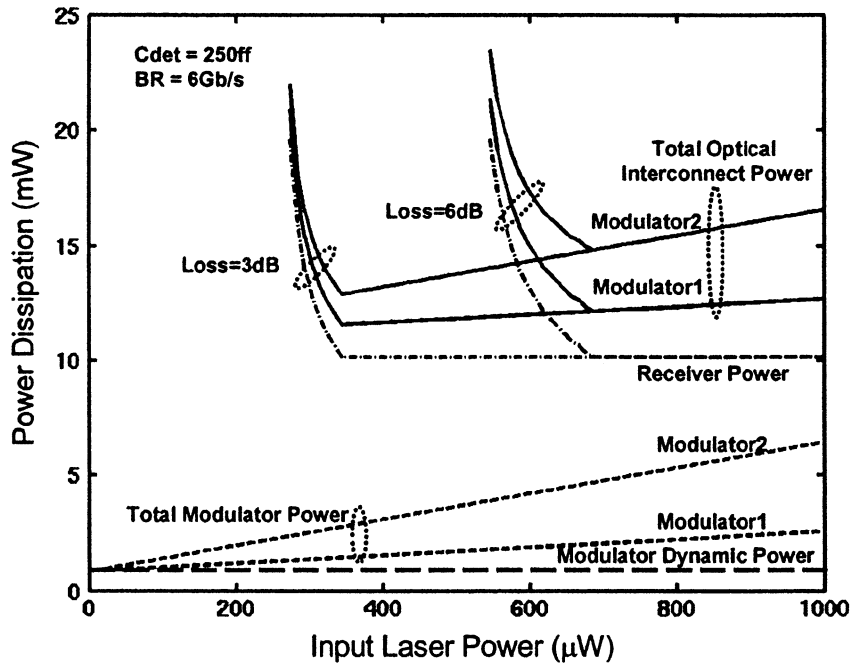


Fig. 3. Power optimization method for optical interconnect.

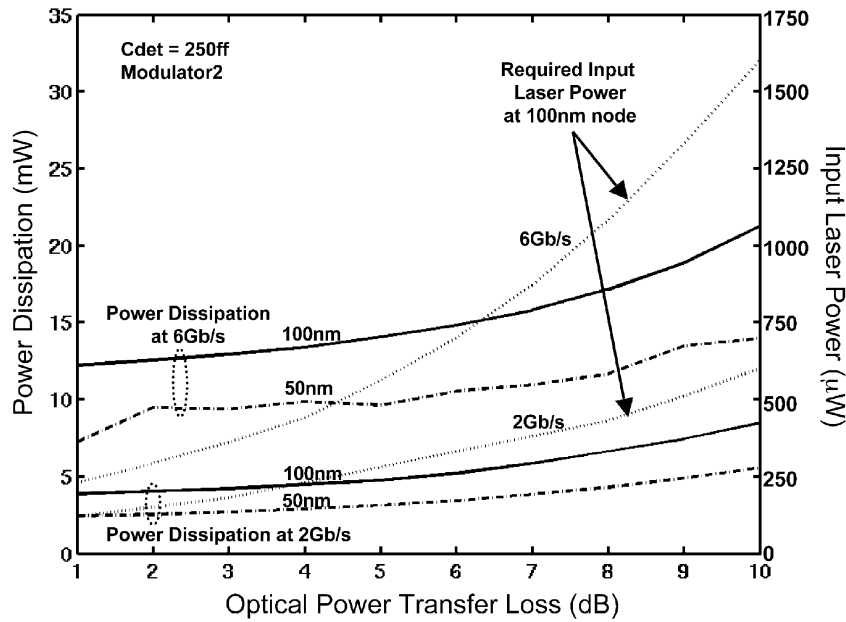


Fig. 4. Optical interconnect power dissipation and optimized input laser power for 100 nm/50 nm technology nodes with bandwidths.

reflections is also assumed. The principle behind this concept is the following. Since we know the exact value and the time at which reflection from package parasitic (and connectors, if applicable) arrive back at the receiver, we can synchronously generate the same voltage using an additional source and feed it into the reference of the receiving differential amplifier to partially cancel it. In short, full consideration was given to maximize electrical interconnect performance with sophisticated schemes for fair comparison with its optical counterpart.

We consider the power dissipated in the termination resistors related to current swing requirement, as this power is becoming increasing fraction of the total power [30]. Also, this power critically depends on the attenuation and noise characteristics of interconnects, attributes where there is a stark difference between

electrical and optical media. The idea is to model the attenuation and noise sources in the electrical interconnect as a function of the bit rate and length. From this, we backtrack the minimum current swing required at the transmitter for an adequate signal-to-noise ratio (SNR) at the receiver. Finally, from current swing we calculate the power dissipation.

The BER in an electrical interconnect system is approximated by [28]

$$BER = e^{-\frac{V_{SNR}^2}{2}} \quad \text{where } V_{SNR} = \frac{V_{nm}}{V_{Gaussian}}. \quad (3)$$

Here, V_{SNR} is the voltage SNR. V_{nm} , the net noise margin, is given by the difference of half the signal swing and the sum off all worst-case noise sources at the receiver. $V_{Gaussian}$ is the standard deviation of all the statistical noise sources, which are

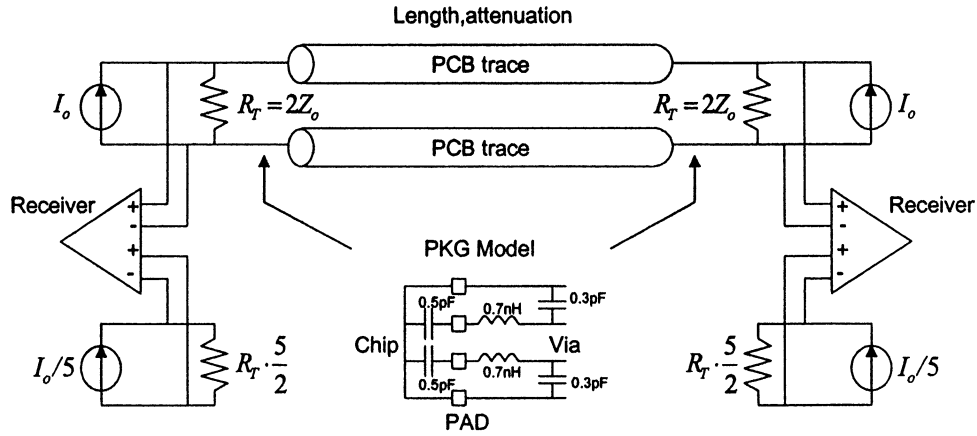



Fig. 5. Schematic showing board-level electrical interconnects.

L	302 nH/m
C	148 pF/m
Z_o	45 Ω
R_o	4.71 Ω /m
R_s	$1.313 \times 10^{-3} \Omega / m \cdot \sqrt{Hz}$
G_o	0
G_d	$9.299 \times 10^{-12} \Omega / m \cdot Hz$



$$R = R_o + R_s \sqrt{f} \text{ (Skin Effect)}$$

$$G = G_o + G_d f \text{ (Dielectric Loss)}$$

Fig. 6. Summaries of the board trace parameters and the corresponding SPICE parameters used for dielectric and skin effect loss.

assumed to be uncorrelated. From (3), we have the condition on the net required noise margin for adequate BER

$$V_{\text{nmreq}} = V_{\text{Gaussian}} \sqrt{2 \ln \left(\frac{1}{\text{BER}} \right)}. \quad (4)$$

The net available noise margin at the receiver depends on two factors: the attenuated signal swing and the sum of all worst-case noise sources. The attenuation in the signal swing is modeled extensively and will be described subsequently. The worst-case noise sources are of two kinds: Proportional to signal swing or independent of it (fixed sources). The first type of proportional noise source is attenuated by the trace just as the signal because it is acquired at the transmitter end. Its proportionality constant is denoted by K_A . This includes trace crosstalk, impedance mismatch, and package reflections. The second type of proportional noise source is acquired at the receiver-end, hence, is not attenuated by the board trace (denoted by K_U). K_U is present only in the case of simultaneous bidirectional signaling due to the opposite direction transmitter at the receiver end. It includes reverse channel crosstalk, package reflections, and transmitter replica mismatch. Finally, the fixed noise sources (V_{NF}) arise due to the receiver offset and its sensitivity.

The effect of transmitter-end preequalization with multtap filters is to increase the voltage fraction from $1 - 2A$ to A ,

where A is the attenuated fraction of the signal at the receiver at a particular bit-rate [29]. The postequalization swing at the receiver is, then, A times the swing at the transmitter (V_{swtrans}), resulting in a gross noise margin of half this value. To meet the BER, the available net noise margin should be greater than the required net noise margin, hence

$$V_{\text{swtrans}} \left(\frac{A}{2} - AK_A - K_U \right) - V_{\text{NF}} > V_{\text{nmreq}}. \quad (5)$$

Thus, from (4) and (5), the minimum swing required is

$$\begin{aligned} V_{\text{swtrans}} &= 2 \frac{V_{\text{nmreq}} + V_{\text{NF}}}{A(1 - 2K_A) - 2K_U} \\ &= 2 \frac{V_{\text{NF}} + V_{\text{Gaussian}} \sqrt{2 \ln \left(\frac{1}{\text{BER}} \right)}}{A(1 - 2K_A) - 2K_U}. \end{aligned} \quad (6)$$

For differential, bipolar current mode signaling with parallel termination ($2Z_0$), and a current swing from $-I_0$ to $+I_0$. (Fig. 5), the transmitter side voltage swing is given by $2(I_0/2)(2Z_0) = 2I_0Z_0$. Thus, the required I_0 (one way swing) is given by

$$I_0 = \frac{V_{\text{NF}} + V_{\text{Gaussian}} \sqrt{2 \ln \left(\frac{1}{\text{BER}} \right)}}{Z_0 [A(1 - 2K_A) - 2K_U]}. \quad (7)$$

TABLE II
SUMMARY OF NOISE SOURCES IN ELECTRICAL INTERCONNECTS

Proportional Noise			
Attenuation (1-A@PCB Length=50cm, 6Gb/s)		0.239	
Noise sources due to unidirectional signaling (Attenuated, K_A)			
Near-end cross-talk	Impedance mismatch	Transmitter impedance mismatch	Package (LC)
0.0005	0.025	0.025	0.075
Noise sources due to bidirectional signaling (Un-attenuated, K_U)			
Reverse-channel crosstalk	Reverse crosstalk (same interconnect)		Transmitter offset
	Impedance mismatch	OCC*	
0.022	0.01	0.05	0.05
Fixed noise (V_{NF}) (Un-attenuated)			
Gaussian noise ($V_{gaussian}$)		5mV	
Low-end	Receiver offset	17.4mV	Receiver
High-end		8.0mV	sensitivity
			20mV
			0.8mV

* OCC: On chip cancellation

The component of total power which is dissipated in the termination resistance for one-way signaling is the sum of the power dissipated in the two termination resistances and the power dissipated in the replica transmitter circuit to cancel the opposite side transmitter signal. For further power minimization, we used a scaled current (factor of 5) and increased impedance by the same factor in the replica circuit. Thus,

$$\begin{aligned}
 P_{\text{term}} &= \left(\frac{I_0}{2}\right)^2 (2Z_0) + \left(\frac{I_0}{2}\right)^2 (2Z_0) + \left(\frac{I_0}{5}\right)^2 (5Z_0) \\
 &= 1.2I_0^2 Z_0
 \end{aligned} \tag{8}$$

where I_0 is given by (7). The other sources of power dissipation in this link are the transmitter and receiver logic circuit power, equalization power, and the power due to additional transmitter for canceling the near-end LC tank package reflections. The transmitter logic power includes the dynamic power due to exponentially sized buffer chain. For low-end receiver logic circuit, we estimate the tail current of the differential amplifier to be about 100 μA . A two-stage amplifier was needed to amplify the input signal to the power rails. The power dissipation in the high-end receiver with very low sensitivity and offset is assumed to be the same, which is optimistic. Amongst the sources considered, the termination resistance power is found to be dominant. The transmitter logic as well as the cancellation circuit power expectedly tracks the power in the termination resistance. Equalization power is neglected, as subsequent taps are scaled version of the main transmitter current. The power due to clock and timing circuits for clock recovery is not considered in this paper. These components are also omitted in the case of optical interconnects for fair comparisons. We believe that omitting the power dissipation in the interconnect system analysis will not greatly impact the conclusion of our analysis because this clock recovery circuit is the common component for both optical and electrical interconnects.

Table II summarizes noise sources in electrical interconnect assuming 5% mismatch between termination resistances and the characteristic impedance of the printed circuit boards (PCB) trace [31]. We have ignored the noise due to connector crosstalk, which will make the electrical interconnects power results slightly optimistic. The reverse crosstalk was estimated with SPICE simulations. The Gaussian noise was assumed to be 5 mV [32]. We considered two types of electrical receivers for each bit rate. The high-end receiver has an offset of 8 mV and sensitivity of 0.8 mV [33]. When using PCB for multigigabit data rate, attenuation due to both the skin effect loss and dielectric loss become extremely important and it is imperative to model it accurately. Attenuated/remaining fraction of the signal at the receiver due to these two effects is given by

$$A = e^{-R/2Z_0} e^{-GZ_0/2} \tag{9}$$

$$R = \frac{R_{\text{dc}}}{2} \left(\frac{f}{f_s}\right)^{1/2}, G = 2\pi f C \tan \delta_D \tag{10}$$

where R_{dc} is constant dc resistance, f_s is the frequency, where the skin depth is equal to the height of the conductor, and $\tan \delta_D$ is the loss tangent. We have calculated attenuation for our dimensions using the comprehensive analytical model (accounts for frequency dependent loss tangent) developed in [34] (analytical model²) and have compared the results with SPICE simulations (including package effect). The SPICE models use a constant loss tangent. As shown in Fig. 7, the frequency dependency in the loss tangent causes more attenuation at higher bit rates in the analytical model compared to SPICE without package effect. However, including package effect with SPICE increases attenuation above the analytical model. The figure also decouples the skin-effect and the dielectric loss and clearly shows that dielectric loss becomes more limiting at high frequency.

An interesting point that follows from (6) is that there exists a minimum allowed attenuated signal (maximum attenuation)

²Coded with MATLAB version 6.5.0.

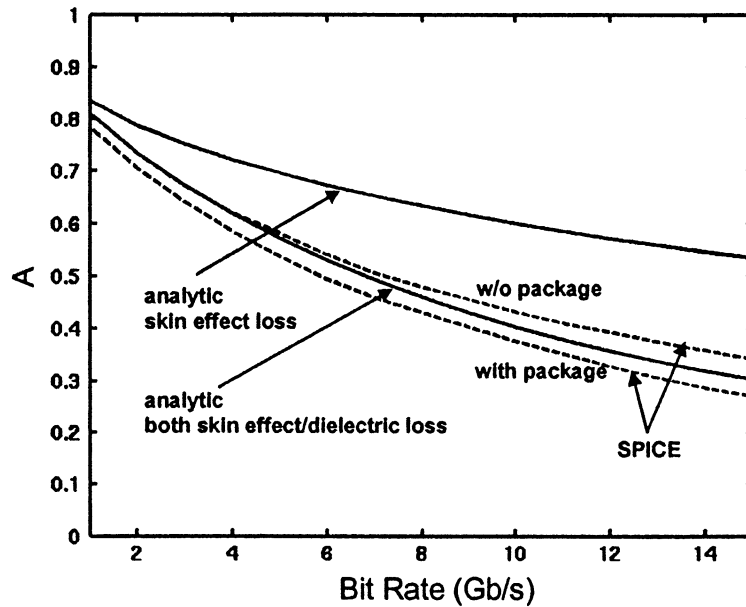


Fig. 7. Attenuation comparison between SPICE and modeling with skin effect, dielectric loss, and package effect.

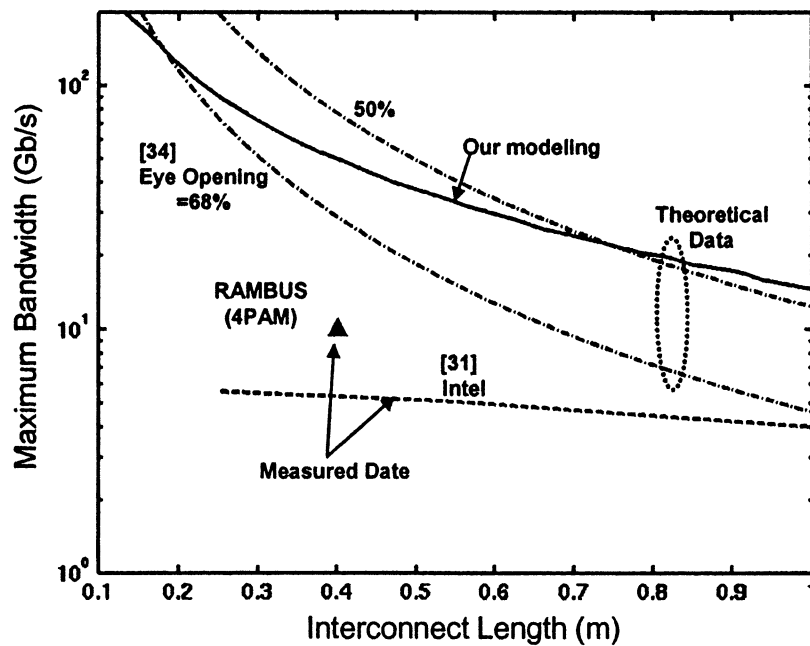


Fig. 8. Maximum bandwidth of electrical interconnect with simultaneous bidirectional signaling.

for simultaneous bidirectional signaling, beyond which, the denominator in (6) becomes negative and noise margin will never be met. This limit is given by

$$A_{\min} = \frac{2K_U}{1 - 2K_A}. \quad (11)$$

Since attenuated signal limit in (11) depends on both bit rate and length [(9) and (10)], it follows that for a given interconnect length, there is a maximum allowed bit rate and vice versa. This bit rate-length contour is plotted for our case in Fig. 8. However, the power dissipation will become prohibitively high much before this limit is reached. Fig. 8 also compares the maximum

bit rate limit with other calculations in the literature as well as some experimental results. Reference [35] calculates the maximum bit rate for a single-ended system by using (9) to ultimately obtain the step response of interconnects. Using this step response and an arbitrarily chosen signal swing requirement at the receiver, one obtains the minimum permissible bit time for signal to rise to the chosen value, hence the maximum bit rate. In practice, the swing requirement would depend on the noise in the system. This yields the Area/length² limit for electrical interconnects [35]. The maximum bit rate for two different swing requirements (50% and 68% of the transmitter side voltage) is shown. In contrast, our calculation in the figure relies on explicit noise estimation and is for bidirectional signaling.

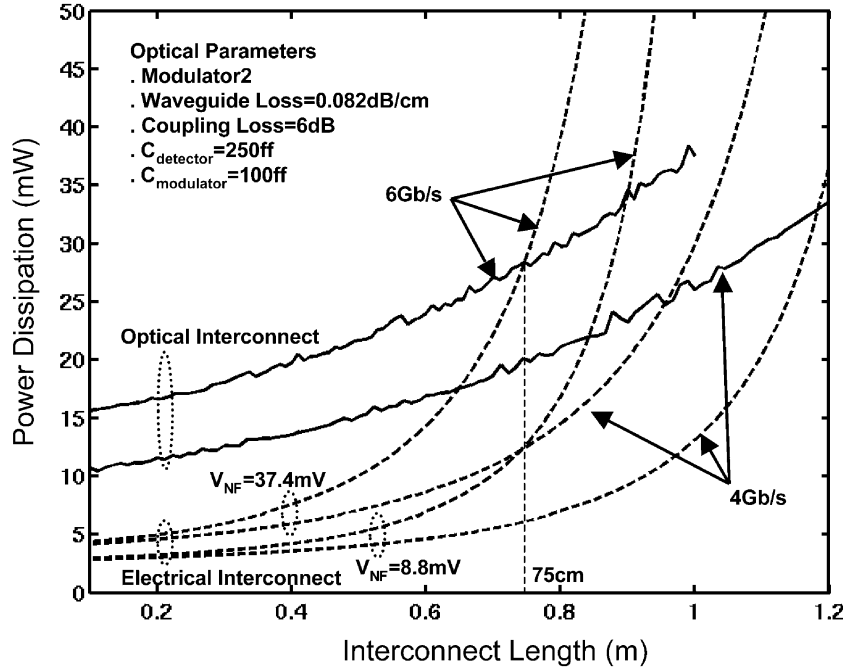


Fig. 9. Power comparison between electrical and optical interconnects for modulator 2.

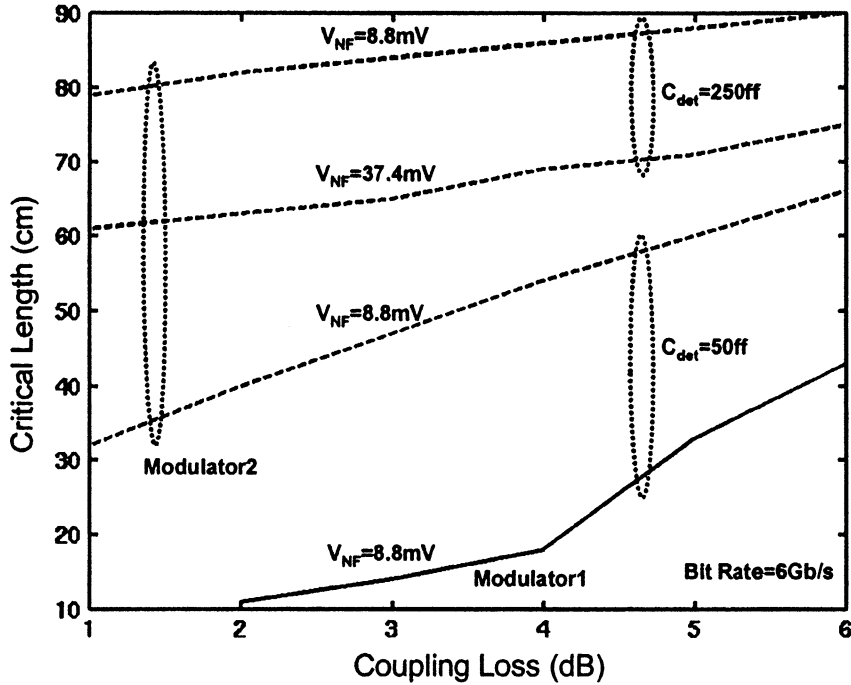


Fig. 10. Critical length in terms of design parameters.

IV. COMPARISON BETWEEN ELECTRICAL AND OPTICAL INTERCONNECTS

Fig. 9 compares the electrical and optical interconnect power versus length at 4 and 6 Gb/s. Electrical interconnect consists of two sets of curves corresponding to different receiver offset/sensitivity. Electrical interconnect power rises with length and bit rate due to a larger attenuation, and a greater impact of unattenuated as well as fixed noise sources. At higher bit rates a smaller rise time would further increase noise sources such as parasitic reflections from package LC tanks. The optical interconnect power dissipation also rises with length owing to a greater loss

in optical power resulting in a smaller η (Section II), albeit this rise is slower than that for electrical interconnect. Beyond a critical length, optical interconnect yields lower power. This critical length reduces at higher bit rates. The figure also explicitly quantifies the impact of improving fixed noise, V_{NF} (receiver sensitivity/offset) on electrical interconnect power dissipation.

In Fig. 10, we explicitly quantify the impact of critical device/system parameters in respective interconnects, on the critical length. Specifically, for optical interconnect, we consider the role of detector/modulator capacitance, coupling loss and ideal modulator 1, whereas, for electrical interconnect, we examine

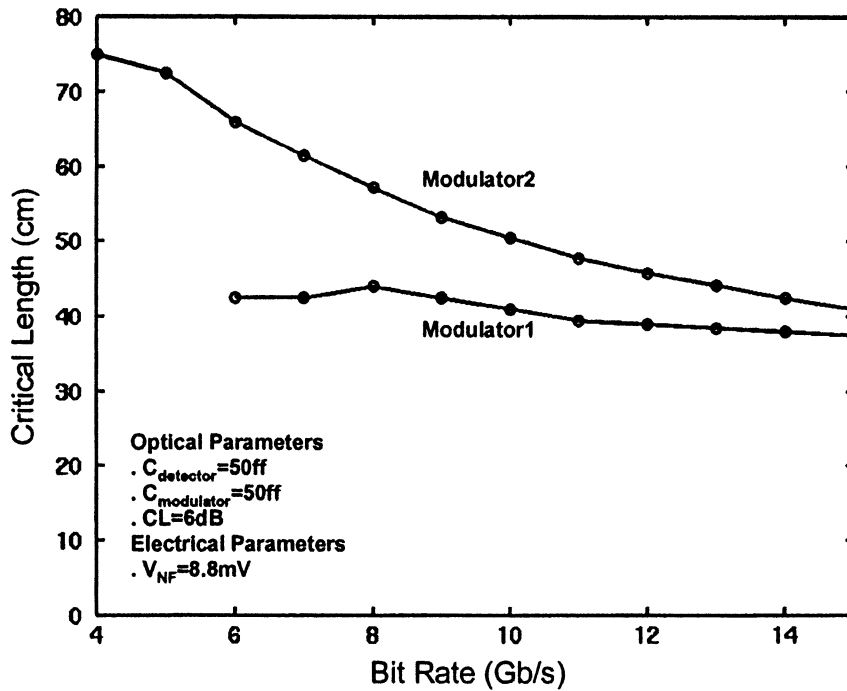


Fig. 11. Critical length in terms of bit rate.

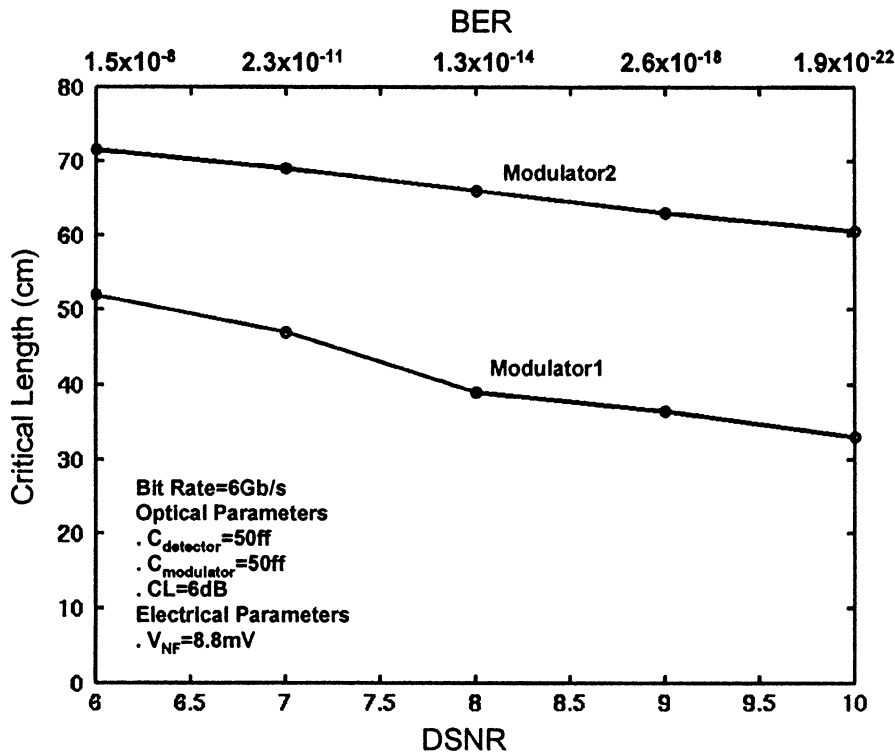


Fig. 12. Critical length in terms of BER.

the role of receiver sensitivity/offset on the critical length. The critical length with modulator 1 is found to be 44 and 20 cm when compared with high and low-end electrical receivers, respectively, with 6 dB coupling loss and 50 fF detector capacitance (low-end with mod 1 not shown in the graph). Both coupling loss and detector capacitance play a pivotal role in dictating critical length. For example, bringing down the detector capacitance from 250 to 50 fF with 3 dB coupling loss reduces critical length from about 80 to 45 cm with modulator 2.

In Fig. 11, the critical length is shown in terms of bit rate of the system. For modulator 2, this length gradually reduces to about 40 cm at 15 Gb/s, while the slope is more shallow for modulator 1. These results are for 6 dB coupling loss. If the coupling loss is reduced further to 3 dB, the critical lengths could come down by significant amount as shown in Fig. 10 for 6 Gb/s. We see an apparent saturation of critical length at high bit rates. This is because we tend to underestimate electrical interconnect power at high bit rates. With bit rate increase, we account for the impact

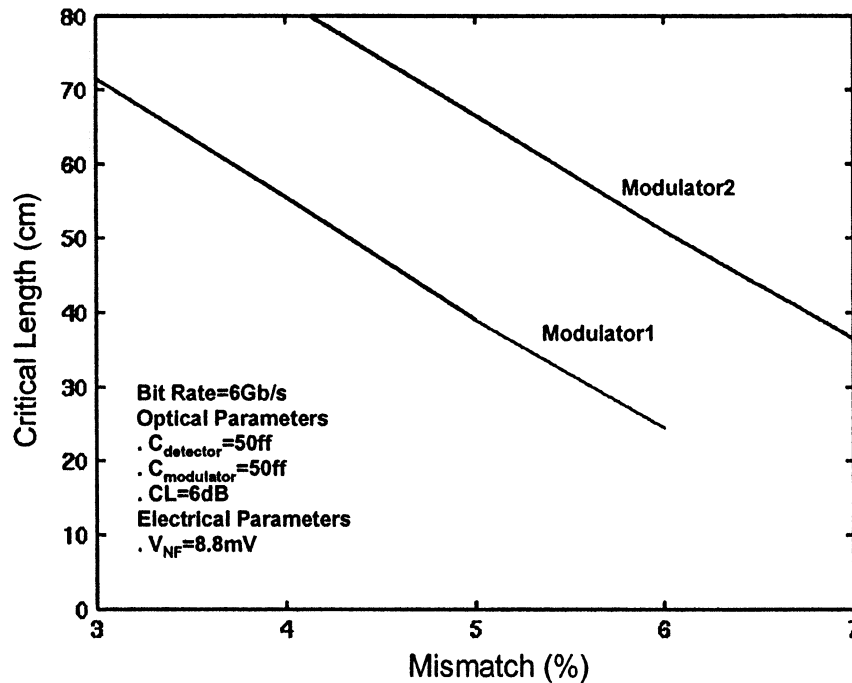


Fig. 13. Critical length in terms of mismatch in terminator.

of worsening trace attenuation on power dissipated in the termination resistances, but neglect the power increase in the electrical transmitter and receiver due to complexity in its modeling. If this power increase was incorporated, the zero length power dissipation in Fig. 9 would increase with bit rate rather than remaining constant. This would yield a lower crossover point (critical length) between electrical and optical curves than what we have calculated, with the difference accentuated at higher bit rates.

Next, we plot the critical length as a function of BER requirement (implicitly in terms of digital signal-to-noise ratio, DSNR) in Fig. 12. Different BER is demanded in different system applications and high BER can be tolerated if explicit error correction schemes are utilized. For example, in communication application BER between 10^{-15} and 10^{-12} is deemed sufficient, whereas, server systems, particularly if they are not deploying error correction, require BERs less than 10^{-15} . From Fig. 12, it is clear that for small BER values, the critical lengths are smaller and optical interconnects have advantage over electrical interconnects. Thus, optical interconnects are more power-favorable for systems where data reliability criteria is demanding.

Finally, we examine the sensitivity of critical length on the mismatch between termination impedances and the characterization impedance of the PCB trace as this constitutes a significant noise source in electrical interconnects (Fig. 13). The critical length is found to substantially increase with small reduction in the impedance mismatch.

V. CONCLUSION

We have done extensive power dissipation comparison between electrical and optical interconnects for bandwidth sensitive applications in 10 cm to 1 m range of interconnects. This comparison, among other things, included introducing a sophisticated power optimization scheme for optical interconnects and

detailed noise and attenuation modeling in electrical interconnects. Based on this modeling, power dissipation was calculated as a function of length and bandwidth. We find that beyond a critical length, within the application range, power optimized optical interconnects dissipate lower power compared to the state-of-the-art high-speed electrical signaling scheme. We have further quantified the impact of various device and system components in electrical and optical interconnects on the critical length. This falls under three categories. 1. On the optical side, we have explicitly quantified the impact of detector/modulator capacitance, coupling loss and modulator type on the critical length. Whereas, it is also implicitly possible to conclude from (2) that the impact of detector responsivity would be similar to that of coupling loss. 2. On the electrical side, we have characterized critical length as a function of receiver sensitivity/offset and impedance mismatch. 3. On the system demand side, we have studied the critical length as a function of bandwidth and BER. This gives both optical device designers as well as electrical circuit designers a framework to assess the system's level (power) impact of various figures of merit of the devices. When compared with the high-end electrical receiver (8.8 mV fixed noise), the critical length is found to be about 43 cm with low optical coupling losses and close to ideal modulator at the bit rate of 6 Gb/s. At higher bit rates and lower BER, the critical length reduces and optics becomes more power favorable. These trends can be fundamentally thought of in terms of a twofold tradeoff between electrical and optical interconnects. Optical interconnects are superior because they have lower attenuation and lower noise (no crosstalk, etc). Whereas, their downside is that they need extra power for conversion from electronics to optics and vice versa. Since the power penalty is fixed, whereas, the power advantage is length and bit rate-dependent, the optical interconnects become beneficial at longer lengths. Finally, including several factors, which were ignored in this electrical interconnect analysis, such as equalization power, especially, as

more taps are required at larger lengths and bit rates, incomplete residual ISI cancellation even with more taps, and rise time reduction induced greater package ringing at higher bit rates, will further reduce the critical lengths.

REFERENCES

- [1] Y. Li, E. Towe, and M. W. Haney, "Special issue on optical interconnections for digital systems," *Proc. IEEE*, vol. 88, pp. 723–727, June 2000.
- [2] P. Kapur and K. C. Saraswat, "Comparisons between electrical and optical interconnects for on-chip signaling," in *Int. Interconnect Technol. Conf.*, 2002, pp. 89–91.
- [3] H. B. Bakoglu, *Circuits, Interconnections, and Packaging for VLSI*. Reading, MA: Addison-Wesley, 1990.
- [4] R. Ho, K. Mai, and M. Horowitz, "Efficient on-chip global interconnects," in *Symp. VLSI Circuits*, 2003, pp. 271–274.
- [5] K. Mai, T. Paaske, N. Jayasena, R. Ho, W. J. Dally, and M. Horowitz, "Smart memories: A modular reconfigurable architecture," in *Proc. 27th Int. Symp. Computer Architecture*, 2000, pp. 161–171.
- [6] D. Hwang, T. Sze, A. Landin, R. Lytel, and H. L. Davidson, "Optical interconnects: Out of the box forever?," *J. Select. Topics Quantum Electron.*, vol. 9, no. 2, pp. 614–623, Mar. 2003.
- [7] R. Mooney, "Scaling methods for electrical interconnects to meet the performance requirements of microprocessor platforms," in *14th Annual Workshop on Interconnects Within High Speed Digital Syst.*, 2003.
- [8] R. T. Chen, L. Lin, C. Choi, Y. J. Liu, B. Bihari, L. Wu, S. Tang, R. Wickman, B. Picor, M. K. Hibbs-brenner, J. Bristow, and Y. S. Liu, "Fully embedded board-level guided-wave optoelectronic interconnects," *Proc. IEEE*, vol. 88, pp. 780–793, June 2000.
- [9] Y. Li, J. Ai, and J. Popelek, "Board-level 2-D data-capable optical interconnect circuits using polymer fiber-image guides," *Proc. IEEE*, vol. 88, pp. 794–805, 2000.
- [10] T. May, A. G. Kirk, D. V. Plant, J. F. Ahadian, C. G. Fonstad, K. L. Lear, K. Tatah, M. S. Robinson, and J. A. Trezza, "Interconnection of a two-dimensional array of vertical-cavity surface-emitting lasers to a receiver array by means of a fiber image guide," *Appl. Opt.*, vol. 39, no. 5, pp. 683–689, Feb. 2000.
- [11] K. P. Jackson, "High-density, array, optical interconnects for multi-chip modules," in *Multi-Chip Module Conf.*, Mar. 1992, pp. 142–145.
- [12] G. A. Keeler, B. E. Nelson, D. Agarwal, and D. A. B. Miller, "Optical interconnects using short optical pulses," in *Lasers and Electro-Opt. Soc. 12th Ann. Meet.*, 1999.
- [13] Y. Liu, E. M. Strzelecka, J. Nohava, M. K. Hibbs-brenner, and E. Towe, "Smart-pixel array technology for free-space optical interconnects," *Proc. IEEE*, vol. 88, pp. 764–768, 2000.
- [14] D. V. Plant and A. G. Kirk, "Optical interconnects at the chip and board level: Challenges and solutions," *Proc. IEEE*, vol. 88, pp. 806–818, 2000.
- [15] S. C. Esener, "Implementation and prospects for chip-to-chip free-space optical interconnects," in *Int. Electron Devices Meet.*, 2001.
- [16] A. V. Mule, A. Naemi, E. N. Glytsis, T. K. Gaylord, and J. D. Meindl, "Toward a comparison between chip-level optical interconnection and board-level interconnection," in *Interconnect Technol. Conf.*, June 2002, pp. 92–94.
- [17] A. Naemi, A. V. Mule, and J. D. Meindl, "Partition length between board-level electrical and optical interconnects," in *Interconnect Technol. Conf.*, June 2003, pp. 230–232.
- [18] D. A. B. Miller, "Rationale and challenges for optical interconnects to electronic chips," *Proc. IEEE*, vol. 88, pp. 728–749, 2000.
- [19] M. R. Feldman, S. C. Esener, C. C. Guest, and S. H. Lee, "Comparison between optical and electrical interconnects based on power and speed considerations," *Appl. Opt.*, vol. 27, no. 9, pp. 1742–1751, May 1988.
- [20] M. Yoneyama, K. Takahata, T. Otsuji, and Y. Akazawa, "Analysis and application of a novel model for estimating power dissipation of optical interconnections as a function of transmission bit error rate," *J. Lightwave Technol.*, vol. 14, pp. 13–22, 1996.
- [21] O. Kibar, D. A. A. Blerkon, C. Fan, and S. C. Esener, "Power minimization and technology comparison for digital free-space optoelectronic interconnects," *J. Lightwave Technol.*, vol. 17, pp. 546–555, Apr. 1999.
- [22] P. Kapur and K. C. Saraswat, "Minimizing power dissipation in chip to chip optical interconnects using optimal modulators and laser power," in *Interconnect Technol. Conf.*, June 2003, pp. 224–226.
- [23] G. A. Keeler, N. C. Helman, P. Atanackovic, and D. A. B. Miller, "Cavity resonance tunneling of asymmetric Fabry-Perot MQW modulators following flip-chip bonding to silicon CMOS," *Optics in Computing*, 2002.
- [24] P. K. Tien, "Light waves in thin films and integrated optics," *Appl. optics*, vol. 10, no. 11, pp. 2395–2413, 1971.
- [25] A. V. Krishnamoorthy and D. A. B. Miller, "Scaling optoelectronic-VLSI circuits into the 21st century: A technology roadmap," *IEEE J. Select. Topics Quantum Electron.*, vol. 2, pp. 55–76, Apr. 1996.
- [26] P. Kapur and K. C. Saraswat, "Power dissipation in optical clock distribution network for high performance ICs," in *Int. Interconnect Technol. Conf.*, 2002, pp. 151–153.
- [27] K. C. Saraswat, *International Technology Roadmap for Semiconductor*, 1999 ed. San Jose, CA: Semiconductor Industry Assoc., 1999.
- [28] W. J. Dally and J. H. Poulton, *Digital Systems Engineering*. Cambridge, MA: Cambridge Univ. Press, 1998.
- [29] W. J. Dally and J. Poulton, "Transmitter equalization for 4 Gbps Signaling," *IEEE Micro*, pp. 48–56, 1997.
- [30] K. Y. K. Chang, J. Wei, C. Huang, S. Li, K. Donnelly, M. Horowitz, Y. Li, and S. Sidiropoulos, "A 0.4–4 Gb/s CMOS quad transceiver cell using on-chip regulated dual-loop PLLs," *J. Solid-State Circuits*, vol. 38, no. 5, pp. 747–754, May 2003.
- [31] G. Ahn, D. K. Jeong, and G. Kim, "A 2-Gbaud 0.7-V swing voltage-mode driver and on-chip terminator for high-speed NRZ data transmission," *IEEE J. Solid State Circuits*, vol. 35, pp. 915–918, June 2000.
- [32] B. K. Casper, H. Haycock, and R. Mooney, "An accurate and efficient method for multi-Gb/s chip-to-chip signaling schemes," in *Symp. VLSI Circuits*, 2002, pp. 54–57.
- [33] M. J. E. Lee, W. Dally, and P. Chiang, "A 90 mW 4 Gb/s equalized I/O circuits with input offset cancellation," in *Int. Solid-State Circuits Conf.*, 2000, pp. 252–253.
- [34] C. Svensson and G. D. Dermer, "Time domain modeling of lossy interconnects," *Trans. Adv. Packag.*, vol. 23, no. 2, pp. 191–196, May 2001.
- [35] D. A. B. Miller and H. M. Ozaktas, "Limit to the bit-rate capacity of electrical interconnect from the aspect ratio of the system architecture," *J. Parallel Distrib. Comput.*, vol. 41, 1997.



Hoyeol Cho (S'02) received the B.S. and M.S. degrees in electrical engineering from Korea University, Seoul, in 1991 and 1993, respectively. He is currently working toward Ph.D. degree in electrical engineering at Stanford University, Stanford, CA.

In 1993, he joined Samsung Electronics, Kiheung, Korea, where he was engaged in the design of Window RAM, Synchronous DRAM, Synchronous Graphic RAM, and Double Data Rate SDRAM. In 1997, he transferred to the System LSI Division, where he was involved in the development of merged memory logic (MML) including DRAM, SRAM, and flash memory. Recently, his research interests are focused on the performance comparison between electrical and optical interconnects.



Pawan Kapur was born in Kanpur, India. He received the B.S. degree (*summa cum laude*) in physics and mathematics from Moravian College, Bethlehem, PA, in 1995. He received the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1998 and 2002, respectively.

He is currently a Research Associate with Stanford University and is involved in research related to interconnect requirements for future high-performance integrated circuits, which includes exploring limitations of both electrical interconnects and assessing alternative technologies, such as optical interconnects. He also looks at viability of optical interconnects for off-chip applications. His current research projects include Si/Ge compatible optical devices such as modulators, 3-D integration and systems for heterogeneous technologies, CMOS transistor scaling issues related to power and delay, and novel cell structures to facilitate scaling of FLASH memories.



Krishna C. Saraswat (S'71–M'75–SM'85–F'89) received the B.E. degree in electronics in 1968 from the Birla Institute of Technology and Science, Pilani, India. He received the M.S. and Ph.D. degrees in electrical engineering and high voltage MOS devices and circuits, respectively, in 1969 and 1974, respectively, from Stanford University, Stanford, CA.

He remained at Stanford, working as a researcher and was appointed Professor of Electrical Engineering in 1983. He serves as the Chair of Stanford's Materials Council, and as the Associate Director of the NSF/SRC Engineering Research Center for Environmentally Benign Semiconductor Manufacturing. He also serves on the leadership council of the MARCO/DARPA-funded Focus Center for Materials, Structures, and nano-Devices. Since January 2004, he has an honorary appointment of an Adjunct Professor with the Birla Institute of Technology and Science, Pilani, India. His research interests are in new and innovative materials, structures, and process technology of silicon and germanium devices, and interconnect for nanoelectronics. During 1969–1970, he was with Texas Instruments, working on microwave transistors. He joined Stanford University as a Research Associate in 1975 and later became a Professor of Electrical Engineering in 1983. For the next 15 years, he worked on modeling of CVD of silicon, conduction in polysilicon, diffusion in silicides, contact resistance, interconnect delay and 2-D oxidation effects in silicon. He pioneered the technologies for aluminum/titanium layered interconnects, CVD of tungsten silicide MOS gates, CVD tungsten MOS gates, and tunable work function SiGe MOS gates. During the late 1980s, he became interested in the economics and technology of single wafer manufacturing. He developed equipment and simulators for single wafer thermal processing, deposition and etching and technology for the *in situ* measurements and real-time control. Jointly with Texas Instruments, a *microfactory* for single-wafer manufacturing was demonstrated in 1993. Since the mid-1990s, he has been working on new materials, devices, and interconnects for scaling MOS technology to sub-10 nm regime. He has pioneered several new concepts of 3-D ICs with multiple layers of heterogeneous devices. His group has recently demonstrated the first high-performance germanium MOSFETs with high-k dielectrics. Since 2000, he has also been doing research on environmentally benign semiconductor manufacturing. He has graduated more than 48 doctoral students and has authored or coauthored more than 400 technical papers, of which six have received *Best Paper Award*.

Dr. Saraswat is a Member of both The Electrochemical Society and The Materials Research Society. He received the Thomas Callinan Award from The Electrochemical Society in 2000 for his contributions to the dielectric science and technology. He is the recipient of the 2004 IEEE Andrew Grove Award for seminal contributions to silicon process technology. He received two gold medals for academic excellence during his education in India.