Available online at www.sciencedirect.com
ScienceDirect
ELSEVIER
INTEGRATION, the VLSI journal 40 (2007) 434–446
INTEGRATION the VLSI journal
www.elsevier.com/locate/vlsi

Predictions of CMOS compatible on-chip optical interconnect ☆

Guoqing Chen[a,*], Hui Chen[b], Mikhail Haurylau[a], Nicholas A. Nelson[a], David H. Albonesi[c], Philippe M. Fauchet[a], Eby G. Friedman[a]

[a]*Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY 14627, USA*
[b]*Department of Physics and Astronomy, University of Rochester, Rochester, NY 14627, USA*
[c]*School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14853, USA*

Received 24 January 2006; received in revised form 24 July 2006; accepted 16 October 2006

Abstract

Interconnect has become a primary bottleneck in the integrated circuit design process. As CMOS technology is scaled, the design requirements of delay, power, bandwidth, and noise due to the on-chip interconnects have become more stringent. New design challenges are continuously emerging, such as delay uncertainty induced by process and environmental variations. It has become increasingly difficult for conventional copper interconnect to satisfy these design requirements. On-chip optical interconnect has been considered as a potential substitute for electrical interconnect. In this paper, predictions of the performance of CMOS compatible optical devices are made based on current state-of-the-art optical technologies. Electrical and optical interconnects are compared for various design criteria based on these predictions. The critical dimensions beyond which optical interconnect becomes advantageous over electrical interconnect are shown to be approximately one-tenth of the chip edge length at the 22 nm technology node.
© 2006 Elsevier B.V. All rights reserved.

Keywords: Optical interconnect; On-chip; CMOS compatible; Trends

1. Introduction

In deep submicrometer VLSI technologies, interconnect plays an increasingly important role. Multiple design criteria are considered in the interconnect design process, such as delay, power, bandwidth, and noise. With technology scaling, the device dimensions and clock period continuously decrease. The delay uncertainty caused by process and environmental variations consumes a significant part of the clock period, reducing both performance and yield. It has become increasingly difficult for conventional copper based electrical interconnect to satisfy these requirements. One promising candidate to solve this problem is optical interconnect.

Optical devices are widely used in the telecommunication area, and have been applied as board level interconnects. The concept of on-chip optical interconnect was first introduced by Goodman in 1984 [1]. Since electrical/optical and optical/electrical conversion is required, optical interconnect is particularly attractive for global interconnects, such as data buses and clock distribution networks, where the required signal conversions can be more easily justified. Recently, several comparisons between on-chip electrical and optical interconnects have been described [2,3]. In these papers, the inductive effects of electrical interconnect are ignored, and highly approximate parameters characterizing the optical devices are assumed. The successful realization of on-chip optical interconnect, however, greatly depends upon the development of enhanced CMOS compatible optical devices. Without a reasonable prediction of trends in optical device development, the conclusions presented in [2,3] are less definitive. Furthermore, delay uncertainty is not addressed in these papers. In [4], an optical clock distribution network and an optical

☆This research was supported by the US National Science Foundation under Contract No. CCR-0304574 and Semiconductor Reasearch Corporation through the New York Center for Advanced Interconnect Science and Technology under task 1292.020.
*Corresponding author. Tel.: +1 585 275 1606; fax: +1 585 276 0074.
E-mail address: guchen@ece.rochester.edu (G. Chen).

0167-9260/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.vlsi.2006.10.001

network-on-chip are presented based on heterogeneous integration techniques, which are less suitable for mass IC production.

Based on a practical prediction of optical device development, a more comprehensive comparison between optical and electrical interconnects is performed in this paper for different technology nodes, considering the design criteria of delay uncertainty, latency, power dissipation, and bandwidth density. This comparison is particularly challenging since optical interconnect is a fast developing technology while electrical interconnect is relatively mature. A brief summary of the comparison has been published in [5]. More details describing the device modeling and results are presented in this paper, including a comprehensive analysis of the delay uncertainty in both optical and electrical interconnects. The paper is organized as follows. In Section 2, a delay-optimal design of *RLC* interconnect is presented. In Section 3, an on-chip optical data path is introduced. Predictions of the performance characteristics of next generation optical devices are made based on current technology trends. In Section 4, electrical and optical interconnects are evaluated for different design criteria. Potential challenges to the development of on-chip optical interconnect are discussed in Section 5. Some conclusions are offered in Section 6.

2. Electrical interconnect

Various design methodologies and analysis methods have been proposed to optimize interconnect to satisfy different design criteria. For example, repeaters are widely used in global interconnects to reduce interconnect delay, transition times, and crosstalk noise. The delay model of an *RLC* interconnect with repeaters is reviewed in Section 2.1. Based on this model, an *RLC* interconnect with optimal repeaters to achieve the minimum delay is determined in Section 2.2 for different technology nodes [6]. A model of delay uncertainty is discussed in Section 2.3.

2.1. Delay model

The capacitance and resistance per unit length of the interconnect can be obtained directly from the physical geometries of the global interconnects [6], where the space between adjacent interconnects is assumed to be equal to the minimum interconnect width. The effect of the via impedance is ignored. The interconnect inductance, however, depends upon the distribution of the current return paths, which is difficult to estimate before the physical design of the circuit is completed. The sensitivity of a signal waveform to errors in the on-chip inductance, however, is low. A 30% error in inductance extraction introduces an error of only 9.4% and 5.9% in the delay and rise time, respectively [7]. The magnitude of the on-chip inductance is a slowly varying function of the wire geometry [7]. Based on these two characteristics, a fixed value of 0.5 pH/μm [7] is assumed for all of the technology nodes.

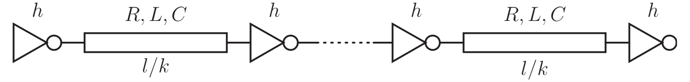

Fig. 1. Repeater insertion in an *RLC* interconnect.

As shown in Fig. 1, a distributed *RLC* interconnect with length l is evenly divided into k segments by uniform repeaters. The repeaters are h times as large as a minimum sized repeater, with the output resistance R_{tr0}/h, output capacitance hC_{d0}, and input capacitance hC_{g0}, where R_{tr0}, C_{d0}, and C_{g0} are, respectively, the output resistance, output capacitance, and input capacitance of a minimum sized repeater.

Repeaters are typically implemented as CMOS inverters [8]. In this analysis, the PMOS transistor is assumed to be twice as large as the NMOS transistor. The repeater output capacitance is assumed to be approximately the same as the input gate capacitance [9]. The sensitivity of the timing model to this assumption is relatively low. The delay model of the interconnect is an extension of the model described in [10] where the repeater output capacitance and input slew effects are considered. By including the repeater output capacitance, the variable ζ [10] used to characterize inductance effects becomes

$$\zeta = \frac{Rl}{2k}\sqrt{\frac{C}{L}} \cdot \frac{R_T C_T (1 + C_{d0}/C_{g0}) + C_T + R_T + 0.5}{\sqrt{1 + C_T}}, \quad (1)$$

where $R_T = kR_{tr0}/(hRl)$ and $C_T = hkC_{g0}/(Cl)$. R_{tr0} can be approximated as $R_{tr0} = KV_{dd}/I_{dn0}$, where K is a fitting parameter, and I_{dn0} is the saturated drain current of a minimum sized NMOS transistor with both V_{gs} and V_{ds} equal to V_{dd}. K can be determined by matching the 50% delay or transition time of the step response of an *RC* equivalent circuit to SPICE simulations. Note that the K for the 50% delay and the K for the transition time are different and are denoted as K_d and K_r, respectively. In this paper, K_d is 0.78 and K_r is 0.55. The corresponding ζ is therefore denoted as ζ_d and ζ_r, respectively. Related transistor parameters can also be found in the ITRS [6]. The delay of a single stage interconnect assuming a step input signal can be obtained by curve fitting,

$$t_d = \frac{e^{-2.3\zeta_d^{1.5}} + 1.48\zeta_d}{w_n}, \quad (2)$$

where $w_n = k/\sqrt{Ll(Cl + C_{g0}hk)}$. In [11], an accurate estimate of the rise time in an *RLC* interconnect is also obtained by curve fitting. The expressions, however, are analytically complicated. In this paper, a simplified piecewise approximation of the rise time is used,

$$t_r = \frac{t_{90\%} - t_{10\%}}{0.8} = \begin{cases} \dfrac{4.4\zeta_r - 1.8}{0.8 w_n}, & \zeta_r > 0.41, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

When $\zeta_r < 0.5$, the interconnect is highly inductance dominant, and (3) can introduce a large error. For those

wire structures considered in this paper, ζ_r is greater than 0.6. The effects of the input transition time on the propagation delay are treated similarly as in [12].

2.2. Delay optimal design

The optimal number and size of repeaters along an *RLC* interconnect can be determined to achieve the minimum delay [10]. This minimum delay can be further decreased by increasing the wire width [13]. Three degrees of freedom are, therefore, explored in electrical interconnect design: the wire width, and the number and size of the repeaters. Various combinations are examined to determine the optimal interconnect design to achieve the minimum delay. The minimum delay per unit length for different wire widths is illustrated in Fig. 2. The interconnect widths are normalized to the minimum wire width W_{min} as predicted in the ITRS [6]. The achievable minimum delay of an *RC* interconnect at the 90 nm technology node is also plotted in Fig. 2 to illustrate the effect of inductance on delay. As shown in Fig. 2, at smaller interconnect widths (W_{min} to $4W_{min}$), the interconnect is dominated by the interconnect resistance, making the effect of inductance negligible. For wider interconnects (typically the global interconnects), the effect of inductance on the delay can be significant. It can also be seen from the figure that scaling has only a small effect on the delay of interconnects with repeaters, consistent with the conclusions from [14]. The decrease in delay with increasing wire width slows when the wire width exceeds $3W_{min}$. The minimum achievable delay per unit length is approximately in the range of 20–22 ps/mm for all technology nodes of interest. Further increasing the wire width greater than $7W_{min}$ only produces small delay differences; the optimal wire width, therefore, is chosen as $7W_{min}$ for each technology node.

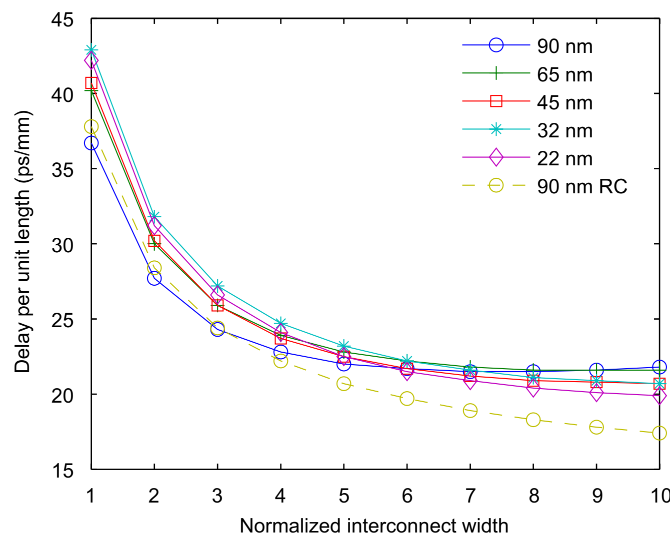

Fig. 2. Minimum delay per unit length as a function of interconnect width.

2.3. Delay uncertainty model

Once the delay optimal design is determined, the delay uncertainty caused by process and environmental variations is analyzed through the Monte Carlo method. Since signal delay is directly related to the power and ground voltage levels, a specific model [15] is used to analyze the effect of power/ground noise on the delay. In this paper, the 50% delay is based on the effective power/ground voltage levels (rather than the ideal power/ground), which corresponds to the second delay definition among the four types of buffer delays described in [15]. The delay uncertainty caused by power/ground noise consists of two components, one component t_{dif} is due to the differential mode noise (i.e., $\Delta V_{dd} - \Delta V_{ss}$), which affects the delay by changing the effective driver resistance; the other component t_{com} is due to the common mode noise (i.e., $\Delta V_{dd} + \Delta V_{ss}$), which affects the delay by changing the effective switching threshold [15]. Changing the variable V_{dd} in the delay expressions only addresses the differential mode noise. The effect of the common mode noise, however, needs to be considered explicitly [15],

$$t_{com} = -\frac{\alpha t_r}{2(1+\alpha)V_{dd_ideal}}(\Delta V_{dd} + \Delta V_{ss}), \qquad (4)$$

where α is the velocity saturation index of an MOS transistor.

3. On-chip optical data path

Introducing optical interconnects into VLSI architectures requires compatibility with CMOS technology. This requirement significantly limits the choice of materials and processes available for fabricating optical components. One of the most significant issues in optical interconnect is the absence of an efficient silicon-based laser that can be monolithically integrated. Only configurations that utilize an external laser as a light source are considered. This type of architecture, however, increases optical losses due to light propagation and coupling.

A diagram of an optical interconnect system is shown in Fig. 3. The system consists of four primary optical elements: an off-chip laser, an optical modulator, a waveguide, and an optical detector. Coupling among the

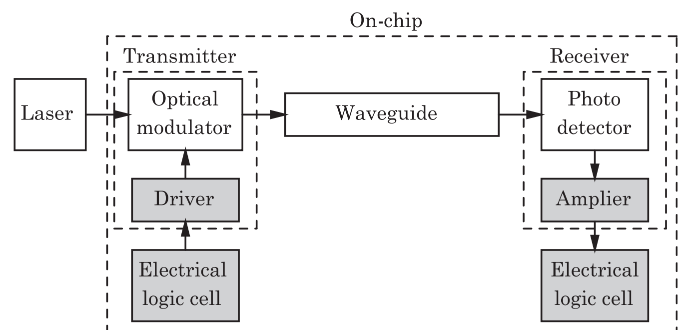

Fig. 3. An on-chip optical interconnect data path.

modulator, waveguide, and photodetector can be implemented through waveguide gratings or 45° mirrors. Because the optical modulator and detector in each data path operate at the same wavelength, there is an inherent conflict in the requirements of the optical material. In contrast to a modulator, which requires negligible optical loss, the principle of detector operation relies on the absorption of light. Considering compatibility with a CMOS technology, a practical solution is a 1.3–1.5 μm wavelength light source with a silicon modulator and a SiGe or Ge photo-detector. Unlike electrical devices, optical devices are not readily scalable due to the light wavelength constraint [16]. The performance and integration ability of optical devices, however, will likely be further improved by technology invention and structural optimization. Although various device models exist for these optical elements, a specific design is selected to satisfy the on-chip requirements. Based on this specific design, trends in the development of optical transmitters, waveguides, and optical receivers are described, respectively, in the following subsections.

3.1. Transmitters

A transmitter is composed of an electro-optical modulator and a driver circuit. The design of a fast and cost efficient CMOS compatible electro-optical modulator is one of the most challenging tasks on the path towards realizing on-chip optical interconnects. In a modulator, conversion between electrical and optical signals is performed in two steps. First, certain optical properties of the medium, e.g., the refractive index or absorption coefficient, are changed by the electrical signals. Second, the optical signals are modulated, either in amplitude or in phase, by varying the optical properties.

Unstrained bulk crystalline silicon, unfortunately, does not exhibit a linear Pockels effect, and refractive index changes due to the Kerr effect are very weak [17]. One of the few suitable mechanisms for varying the refractive index in pure silicon is the free carrier plasma dispersion effect [17]. There are primarily two electrical structures for changing the carrier concentration in silicon devices. One technique is to inject and extract carriers in the intrinsic region of a p-i-n diode. This structure is described in [18]. With this approach, a substantial change in the carrier concentration can be obtained over a large volume. The speed of this structure, however, is limited by the carrier extraction process. To enhance the speed, a high voltage is needed to extract carriers. An alternative electrical structure is a MOS capacitor. A change in the carrier concentration is achieved by redistribution rather than injection and extraction of carriers, causing a higher modulation speed. The first MOS capacitor based electro-optical modulator was demonstrated by Liu et al. [19] and operates at frequencies up to 10 GHz [20]. By design optimization and technology improvements, such as thinning the gate oxide and using an epitaxial over-growth technique, the bandwidth of the modulator is expected to increase to 30–40 GHz by the year 2016.

Because the device structure used in [20] is a Mach–Zehnder interferometer, the modulator has a large footprint (~10 mm long), resulting in an excessive capacitance; hence, increasing the delay and power consumption of the driver circuits. Simulations and early experiments performed by Barrios et al. [21] show that an optical resonator-based structure can drastically decrease the modulator size down to 10–30 μm while maintaining the same operating principle and speed. Further reductions in the modulator size are possible by using photonic bandgap structures. It is important that either an active or passive temperature compensation method similar to that published in [22] is used to maintain stable device operation, particularly in high quality factor resonators.

In this paper, a predictive modulator model is used that combines the advantages of the structures used in [18] and [20], as listed in Table 1. The performance of a modulator significantly depends on the dimensions and related physical parameters. For example, an increase in the extinction coefficient can degrade the modulator bandwidth. To optimize the performance of a modulator, a comprehensive closed-form model [23] is used to determine a proper trade-off among all of the physical parameters characterizing a MOS modulator. An optimized modulator structure with a gate oxide thickness of 1.01 nm, width of 0.89 μm, length of 10 μm, and capacitance of 304 fF is determined from this model and used in the analysis for all of the technology nodes. The driving voltage of the modulator is set to 1.2 V.

A series of tapered inverters [24] is used to drive the modulator. If the inverter output capacitance is equal to the input gate capacitance, the optimal size ratio between two neighboring inverters is 3.6 [25]. A minimum sized inverter is used as the first stage. The number of stages can be obtained as $N = \log(C_M/C_{g0})/\log 3.6$, where C_M is the

Table 1
Predictive model of future silicon based electro-optical modulators

	Electrical structure	
Optical structure	p-i-n diode	MOS capacitor
Mach–Zehnder	–	Ref. [20] High speed (up to 10 GHz) CMOS compatible Large size (10 mm) High power consumption
Resonator	Ref. [18] CMOS compatible Small size (38 μm) Low power consumption	Predictive closed-form model [23] High speed CMOS compatible Small size Low power consumption

modulator capacitance. The delay model of each stage is described in [26].

Although the transmitter analysis presented here is for a specific structure, the analysis can be directly applied to other silicon-based optical interconnect systems. For example, the power and delay model can be applied to an optical interconnect configuration using an on-chip Raman silicon laser [27], which behaves as a resonator.

3.2. Waveguides

The performance of optical waveguides is primarily limited by the wavelength of the utilized light and the choice of optical material. Although novel waveguide platforms, such as photonic crystal waveguides, can potentially reduce the waveguide pitch, optical losses in such structures will likely diminish this advantage.

Given the operating wavelength of on-chip optical interconnect, there are primarily two candidates for the waveguide material [28]. For applications requiring dense and short waveguide arrays, a silicon-on-insulator (SOI) structure is more beneficial due to the smaller waveguide pitch. For longer paths, optimized for signal propagation delay and smaller losses [29], low-loss polymer waveguides are better suited [30]. The core index and cladding index in the polymer waveguide are assumed to be 1.6 and 1.1, respectively.

3.3. Receivers

The receiver has two components: a photo-detector that converts light into current followed by an amplifier that converts the analog current signal into a digital voltage signal. A simplified equivalent circuit model is shown in Fig. 4.

In this paper, interdigitated SiGe p-i-n or metal semiconductor metal (MSM) detectors are considered due to the fast response and reasonable quantum efficiency of these structures. The signal rise time (response time) of the detector can be expressed as $T_r = \sqrt{T_{tr}^2 + T_{RC}^2}$, where T_{tr} is the time required for the photo-generated carriers to drift to the electrical contact, and T_{RC} is the RC response time of the detector [31]. The 3 dB bandwidth of a detector is $\Delta f_{dec} = 0.35/T_r$. Based on a one pole approximation, the delay of the photo-detector is related to the rise time as $t_{dec} = 0.315T_r$. In 2002, an interdigitated Ge p-i-n detector fabricated on a Si substrate with a 3 dB bandwidth of 3.8 GHz at a 1.3 μm wavelength has been demonstrated [32]. Several other papers have been published on SiGe detectors [33,34]. These detectors exhibit similar performance levels. The bandwidth or delay of most of these detectors is limited by the carrier transit time, which can be improved through device optimization.

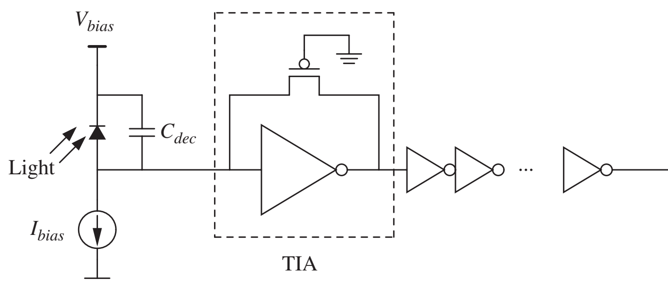

Fig. 4. Circuit model of an optical receiver.

Based on a model proposed by Averine et al. [31], the trend in the performance of future detectors is projected. In Fig. 5, the MSM detector response time as a function of electrode width is plotted for different detector sizes and is compared with experimental results. The spacing between the electrodes is assumed to be equal to the electrode width. As shown in Fig. 5, an optimal electrode width exists that produces a minimum response time. When the electrode is too narrow, the response time is dominated by T_{RC}. When the electrode is too wide, the response time is dominated by T_{tr}. In this paper, the electrode width is assumed to be optimized for minimum response time. The minimum response time of a detector decreases as the detector area becomes smaller. The detector response time is expected in the near future to drop significantly, from tens of picoseconds to a few picoseconds. The reason for this decrease is that present detectors are generally bulky, and a longer time is required for the carriers to transit. Effort has been placed on making smaller detectors. Once efficient coupling between the waveguides and the detectors is realized, the size of the detector is expected to significantly decrease, greatly reducing the response time. This trend, however, is expected to slow and eventually saturate due to fundamental limitations in material properties [35].

The photo-current I_{ph} from the photo-detector is pre-amplified by a transimpedance amplifier (TIA). The TIA

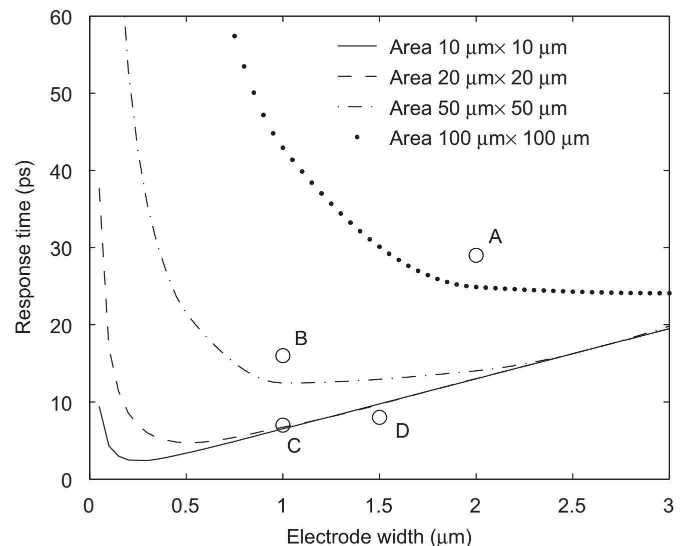

Fig. 5. Detector response time versus electrode width: A (100 μm × 100 μm) [31], B (50 μm × 50 μm) [31], C (20 μm × 20 μm) [36], and D (10 μm × 10 μm) [33].

consists of an inverter and a feedback resistor, implemented as a PMOS transistor. Additional minimum sized inverters are used to amplify the signal to a digital level. A current source I_{bias} is used to bias the input DC current to zero. All of the inverters are assumed to be biased at $V_{dd}/2$. The size of the inverter and feedback transistor in the TIA is determined by bandwidth and noise constraints [9], where the receiver circuit is treated as a one-pole system and the noise model described in [37] is used. The bandwidth requirement of the receiver is assumed to be 0.7 times the bit rate, and the bit error rate (BER) is assumed to be 10^{-15} [9]. For the receiver circuits, the static power dominates and is

$$P_{\text{rec}} = W_{\text{TIA}} I_{\text{dsat0}} V_{dd} + (I_{\text{bias}} V_{dd} + I_{\text{ph}} V_{\text{bias}})/2 + N_{\text{inv}} I_{\text{dsat0}} V_{dd}, \tag{5}$$

where I_{dsat0} is the saturation drain current of a minimum sized inverter biased at $V_{dd}/2$. W_{TIA} is the size of the TIA normalized to a minimum sized inverter. N_{inv} is the number of additional inverter stages determined by the output swing requirements [9]. The delay of the receiver amplifier t_{amp} is obtained by approximating the circuit as a one pole system, $t_{\text{amp}} = 0.7/(2\pi\Delta f_{\text{req}})$, where Δf_{req} is the bandwidth requirement. The input optical power is assumed to be $200\,\mu\text{W}$ [38]. The detector is assumed to have a responsivity of 0.32 A/W and a dark current of $0.5\,\mu\text{A}$.

4. Comparison between electrical and optical interconnects

In this section, different criteria used in the design of the two interconnect systems described in Sections 2 and 3 are compared. The interconnect length is 10 mm. In Section 4.1, the delay uncertainty of the two kinds of interconnects is analyzed. Delay uncertainty can significantly affect the actual circuit delay due to the pipeline nature of synchronous systems, which is discussed in Section 4.2. The power and bandwidth density are compared in Sections 4.3 and 4.4, respectively. The critical length above which optical interconnect is beneficial is determined in Section 4.5.

4.1. Delay uncertainty

Delay uncertainty is caused by process and environmental variations. Variations in the environment include power/ground noise, temperature fluctuations, and crosstalk coupling. In this paper, all of the variations are assumed to be random with a normal distribution and independent unless explicitly indicated.

Process variations include both die-to-die and within-die variations. Temperature variations also exhibit a similar behavior. The average on-chip temperature is different from die-to-die due to the ambient environment and local circuit activity. Since different on-chip blocks typically dissipate different amounts of power, the temperature is also non-uniform within a die [39]. The within-die variations at different locations are generally correlated according to the separation distances. Since the focus of this paper is on comparing electrical and optical global interconnects which cross different regions of an IC, within-die spatial correlation effects need to be considered. The spatial correlation coefficient is modeled as [40]

$$\rho_{\text{cor}}(x) = \begin{cases} 1 - \dfrac{x}{X_L}(1-\rho_B), & x \leqslant X_L, \\ \rho_B, & x > X_L, \end{cases} \tag{6}$$

where x is the separation distance between the variations, ρ_B is the correlation coefficient of the die-to-die variations, and X_L is related to the gradient of the systematic within-die variations. In this paper, ρ_B is assumed to be 0.5 for process and temperature variations, which means the variations are caused equally by die-to-die and within-die variations. X_L is assumed to be 10 mm for MOS device variations, temperature variations, and power/ground noise, 2 mm for the interconnect height and ILD thickness, and 5 mm for the interconnect width and spacing [41].

The process variations considered in the MOS transistors are: L_{eff}, T_{ox}, N_{ch}, and R_{sd}. By using the analytic models of V_{th} [42,43], μ [44], and I_{dsat} [45], the variation of the transistor performance can be analyzed.

For electrical interconnects, process variations occur in the following primary geometric parameters: wire width W, wire height H, space between wires in the same level S, and thickness of the interlevel dielectric oxide T_{ILD}. The variation in resistance R is primarily determined by geometric parameters and the temperature of the interconnect,

$$R = \frac{\rho_0 l}{WH}[1 + \beta(T - T_0)], \tag{7}$$

where β is the temperature coefficient of resistivity, and is $0.0036/^\circ\text{C}$ at $20\,^\circ\text{C}$ [46]. The interconnect capacitance can be expressed as $C_{\text{gnd}} + 2\eta C_{\text{c}}$, where C_{gnd} is the ground capacitance and C_{c} is the coupling capacitance to the neighboring wires. C_{gnd} and C_{c} can be determined from the closed-form expressions provided in [47]. The switching factor η is used to model the Miller effect due to switching activities on neighboring wires [48]. Although the switching patterns are discrete, the signal skews and transition times are continuous and also significantly affect the switching factor. Furthermore, by staggering the repeaters [49], the coupling effect inside a bus structure can be significantly reduced. Based on these considerations, η is assumed to exhibit a continuous normal distribution with $3\sigma = 1$ rather than a discrete distribution, where σ is the standard deviation. Since the sensitivity of the inductance L on the wire geometry is low [50], the effect of process variations on the inductance is ignored. The variation of L is primarily due to neighboring coupling effects, which change the current return paths. It is shown in [51] that, for multiple interconnect systems, the effective wire inductance can be

described as

$$L = L_{\text{self}} + \sum_i \xi_i M_i, \quad (8)$$

where L_{self} is the partial self inductance and M_i is the partial mutual inductance between wire i and the wire of interest. ξ_i is a coefficient which depends upon the signal switching patterns and wire capacitances [51]. To simplify the problem, the inductance is modeled in this paper as $L = L_0(1+\xi)$, where L_0 is a typical value of inductance of $0.5\,\text{pH}/\mu\text{m}$ [7] and ξ is used to model the coupling effect on the effective inductance.

The interconnect is decomposed by repeaters into a number of segments. In each segment, process, temperature, and power/ground variations are assumed to be uniformly distributed, i.e., the ρ_{cor} inside a segment has a value of one. The ρ_{cor} between different segments is determined by (6). η and ξ, however, are assumed to be uniform along the total length of the interconnect, since in a bus structure wires often experience the same neighboring coupling environment over the total length of the line.

Those parameters considered to vary and the corresponding 3σ values are listed in Table 2, which are extracted from [6,52,53]. The absolute variations in the optical devices are assumed to be the same as in a similar electrical structure, since the same microelectronic technology will be applied to fabricate these optical devices. For example, the variations in a MOS modulator are assumed to be the same as in an MOS transistor. For optical interconnect, only the polymer waveguide is considered for process variations. Among these parameters, three kinds of correlations are considered. First, power and ground are inversely correlated with a correlation coefficient of −0.5. The 3σ value of V_{dd} and V_{ss} is 5.8% of the nominal V_{dd} such that $V_{dd} - V_{ss}$ has a 3σ value of 10%. Second, assuming a fixed wire pitch, the interconnect width and space are inversely correlated with a coefficient of −1. This behavior is also true for the electrode finger width W_{finger} and space S_{finger} in an interdigitated photo-detector. The interconnect height H and T_{ILD} are assumed to be correlated with a coefficient of −0.5 [52]. Third, ξ is assumed to be correlated to η with a correlation coefficient of −0.3. This assumption is based on the observation that oppositely switching neighbors result in a greater effective capacitance (i.e., a greater η), and a smaller effective inductance (i.e., a smaller ξ), since neighboring wires provide nearby current return paths. Unlike the effective capacitance, which is only related to the immediate neighbors, the effective inductance depends on the neighboring wires over a long distance, making the correlation between η and ξ fairly weak.

For an optical interconnect system, although the waveguide crosses a long distance, geometric variations

Table 2
Parameters and 3σ variations

Year		2004	2007	2010	2013	2016
Technology node		90 nm	65 nm	45 nm	32 nm	22 nm
MOS transistor	L_{eff} (nm)	37 ± 10%	25 ± 10%	18 ± 10%	13 ± 10%	9 ± 10%
	T_{ox} (nm)	2.0 ± 4%	1.3 ± 4%	1.1 ± 4%	1.0 ± 4%	0.9 ± 4%
	N_{ch} ($\times 10^{18}\,\text{cm}^{-3}$)	1.55 ± 5%	2.74 ± 5%	4.00 ± 5%	5.85 ± 5%	10.0 ± 5%
	R_{sd} ($\Omega - \mu$m)	180 ± 10%	162 ± 10%	135 ± 10%	107 ± 10%	79 ± 10%
Electrical interconnect	H (μm)	0.431 ± 15%	0.319 ± 15%	0.236 ± 15%	0.168 ± 15%	0.125 ± 15%
	T_{ILD} (μm)	0.431 ± 15%	0.319 ± 15%	0.236 ± 15%	0.168 ± 15%	0.125 ± 15%
	W (μm)	1.44 ± 3%	1.02 ± 3%	0.72 ± 3%	0.49 ± 3%	0.35 ± 3%
	S (μm)	0.205 ± 20%	0.145 ± 20%	0.103 ± 20%	0.07 ± 20%	0.05 ± 20%
	η	1 ± 1	1 ± 1	1 ± 1	1 ± 1	1 ± 1
	ξ	0 ± 0.5	0 ± 0.5	0 ± 0.5	0 ± 0.5	0 ± 0.5
Optical modulator	W_{mod} (μm)	0.89 ± 3%	0.89 ± 2.1%	0.89 ± 1.5%	0.89 ± 1.0%	0.89 ± 0.7%
	H_{mod} (μm)	0.1 ± 36%	0.1 ± 26%	0.1 ± 20%	0.1 ± 14%	0.1 ± 10%
	N_d ($\times 10^{18}\,\text{cm}^{-3}$)	7.4 ± 5%	7.4 ± 5%	7.4 ± 5%	7.4 ± 5%	7.4 ± 5%
	N_a ($\times 10^{18}\,\text{cm}^{-3}$)	5.4 ± 5%	5.4 ± 5%	5.4 ± 5%	5.4 ± 5%	5.4 ± 5%
	$T_{\text{ox_mod}}$ (nm)	1.01 ± 8%	1.01 ± 5%	1.01 ± 4.3%	1.01 ± 4.0%	1.01 ± 3.6%
	V_{bias} (V)	1.2 ± 10%	1.2 ± 9.2%	1.2 ± 8.3%	1.2 ± 7.5%	1.2 ± 6.7%
Polymer waveguide	H_{wav} (μm)	1.5 ± 4%	1.5 ± 3%	1.5 ± 2%	1.5 ± 2%	1.5 ± 1%
	W_{wav} (μm)	1.5 ± 3%	1.5 ± 2%	1.5 ± 1%	1.5 ± 1%	1.5 ± 1%
Detector	W_{finger} (μm)	0.505 ± 8%	0.105 ± 28%	0.105 ± 20%	0.095 ± 15%	0.085 ± 12%
	S_{finger} (μm)	0.505 ± 8%	0.105 ± 28%	0.105 ± 20%	0.095 ± 15%	0.085 ± 12%
Environment variation	V_{dd} (V)	1.2 ± 0.07	1.1 ± 0.064	1.0 ± 0.058	0.9 ± 0.052	0.8 ± 0.046
	V_{ss} (V)	0 ± 0.07	0 ± 0.064	0 ± 0.058	0 ± 0.052	0 ± 0.046
	T (°C)	100 ± 50	100 ± 50	100 ± 50	100 ± 50	100 ± 50

are assumed to be uniform across the total length. This assumption overestimates the delay uncertainty, since those independent components of variations in different parts of the waveguide can average out, producing a smaller delay uncertainty. This overestimation, however, does not affect the conclusions of this paper, since delay uncertainty caused by the waveguide is small as compared with other parts of the system. As described in Section 3.3, the receiver amplifier is designed to satisfy the target bandwidth and noise constraints. The design margins of the bandwidth and noise are assigned such that the target requirements can be satisfied with process and dynamic environmental variations. The design of the amplifier will be more challenging in future technology nodes due to parameter variations. The input optical power needs to be increased to produce an effective circuit. Although parameter variations in different parts of an optical interconnect may be correlated, the effects of these variations on delay uncertainty are different due to the different operational mechanisms. In this paper, the delay uncertainty generated at different parts of the optical data path is assumed to be independent, resulting in the following expression for the standard deviation of the total delay:

$$\sigma_{\text{optical}} = \sqrt{\sigma_{\text{drv}}^2 + \sigma_{\text{mod}}^2 + \sigma_{\text{wav}}^2 + \sigma_{\text{dec}}^2 + \sigma_{\text{amp}}^2}. \quad (9)$$

The terms in (9) represent standard deviations of different optical components, ‘‘drv’’ is the driver circuit of the modulator, ‘‘mod’’ is the modulator, ‘‘wav’’ is the waveguide, ‘‘dec’’ is the detector, and ‘‘amp’’ is the amplifier in the receiver.

Based on these assumptions, the delay uncertainty of both the electrical and optical interconnect is analyzed. The delay distribution of a 10 mm electrical interconnect at the 45 nm technology node is shown in Fig. 6. The delay uncertainty (defined as from -3σ to 3σ) is about one half of the nominal delay. The delay and 3σ values for different parts of a 10 mm optical data path are listed in Table 3. The delay of the transmitter and receiver is determined as explained in Sections 3.1 and 3.3, respectively. The signal delay in the waveguide is treated as a light propagation delay. The delay uncertainty of the optical interconnect is dominated by the modulator, as listed in Table 3. The dimensions of the modulator are not scaled; the manufacture process, however, can be controlled more accurately with technology improvements, reducing the delay uncertainty of the modulator. The total delay uncertainty of the optical interconnect, therefore, is expected to be lower in future technology nodes. The delay uncertainty of the electrical interconnect, in contrast, is expected to slowly increase in future technology nodes due to the larger number of inserted repeaters. A comparison of the standard deviation of the delay of electrical and optical interconnect is shown in Fig. 7.

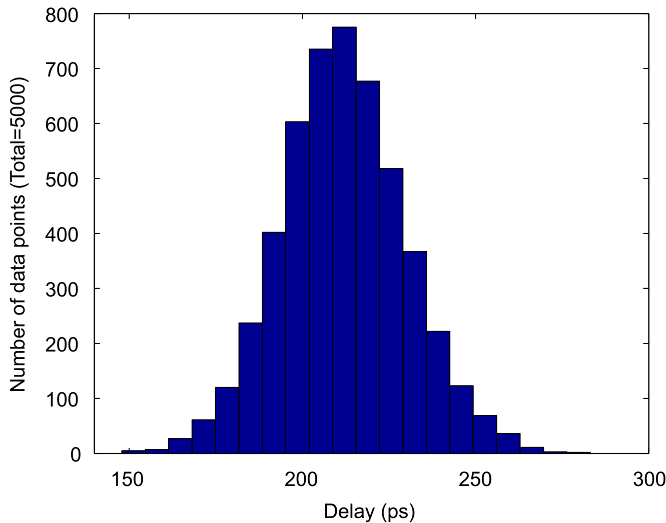

Fig. 6. Delay distribution of a 10 mm electrical interconnect at the 45 nm technology node.

4.2. Delay

As shown in Fig. 2, the minimum delay of electrical interconnect is about 20–22 ps/mm. This minimum delay, however, may not be achievable due to the effect of delay uncertainty. A timing diagram of the data and clock waveforms is shown in Fig. 8. In the figure, T_{un} is the delay uncertainty, T_{setup} and T_{hold} are the minimum setup and hold requirements at the receiving register, respectively, and T_{bit} is the bit period which is the same as the clock

Table 3
Delay (ps) and 3σ value of a 10 mm optical data path

Year	2004	2007	2010	2013	2016
Technology node	90 nm	65 nm	45 nm	32 nm	22 nm
Modulator driver	37.3 ± 20.9%	26.5 ± 20.4%	16.6 ± 23.5%	10.3 ± 29.1%	5.2 ± 40.4%
Modulator	40.0 ± 67.0%	40.0 ± 51.0%	40.0 ± 41.0%	40.0 ± 32.0%	40.0 ± 27.0%
Polymer waveguide	49.3 ± 1.1%	49.3 ± 0.8%	49.3 ± 0.5%	49.3 ± 0.2%	49.3 ± 0.1%
Detector	2.5 ± 5.6%	1.1 ± 21.9%	0.6 ± 14.1%	0.5 ± 9.3%	0.4 ± 7.1%
Amplifier	34.0 ± 10.6%	13.5 ± 23.8%	8.7 ± 17.6%	5.7 ± 15.8%	3.4 ± 15.0%
Total optical	163.1 ± 17.3%	130.4 ± 16.4%	115.2 ± 14.7%	105.8 ± 12.5%	98.3 ± 11.2%

period $T_{\rm clk}$. The clock signal is assumed to be properly skewed in order for the data to be correctly latched. In this paper, the timing budget assigned to $T_{\rm setup}$ and $T_{\rm hold}$ is assumed to be 20% of the clock period, i.e., the delay uncertainty cannot exceed 80% of the clock period. Note that this 20% clock period timing budget also includes the delay uncertainty of the register, clock jitter, and the effect of the rise time (i.e., the rise time of the data cannot be excessively large such that the data waveform will maintain the correct value over a specific period). If this requirement is not satisfied, additional pipeline registers are inserted such that the timing requirements of each stage are satisfied. The delay of the interconnect considering delay uncertainty is

$$T_{\rm total} = m(T_{\rm max} + T_{\rm setup} + T_{C-Q}), \tag{10}$$

where m is the number of register stages, T_{C-Q} is the time required for the data to leave the register after the clock signal arrives, and $T_{\rm max}$ is the maximum interconnect delay in a stage and is the summation of the nominal stage delay

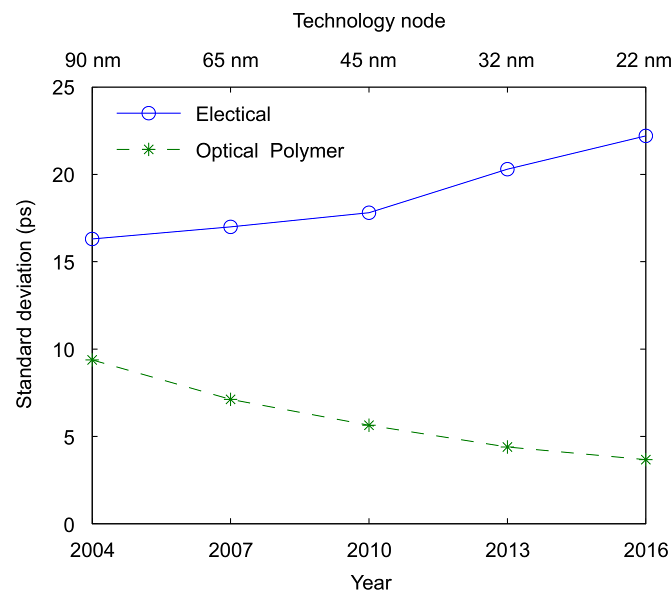

Fig. 7. Comparison of standard deviation of delays of electrical and optical interconnects.

and 3σ of the delay uncertainty of each stage. $T_{\rm setup} + T_{C-Q}$ is also assumed to be 20% of the clock period.

Expression (10) can also be used to determine the delay of the optical interconnect. Since no register-like storage device can be inserted into an optical data path, the delay uncertainty determines an upper bound on the channel bandwidth,

$$B_{\rm optical} = \frac{1}{T_{\rm bit}} \leqslant \frac{0.8}{T_{\rm un}}. \tag{11}$$

From Table 3, note that the clock frequency, as predicted in the ITRS [6], can be achieved by optical interconnect for each technology node except the 22 nm node. For the 22 nm node, the highest bandwidth as determined by (11) is 36.3 Gb/s. The actual delay of the electrical and optical interconnect is compared in Table 4. The delay at the 22 nm node is determined with a clock frequency of 36.3 GHz. As listed in Table 4, the actual delay of the electrical interconnect remains approximately fixed for all of those technology nodes. The delay of the optical interconnect, however, decreases with future technology nodes due to the increasing performance of the electrical circuits in the modulator driver and the receiver amplifier.

4.3. Power

The power dissipated by the electrical interconnect includes dynamic power, short-circuit power, and leakage power. The electrical interconnect power models used in this analysis are the same as those models described in [12]. The clock frequency used in the power model is the local clock frequency as predicted in the ITRS (the clock frequency at the 22 nm technology node is assumed to be 36.3 GHz). The switching factor is assumed to be 0.15 [12]. The power of the registers can be estimated by scaling a typical master–slave D flip-flop where the result is negligible as compared to other interconnect power components.

The power consumed by the optical interconnect is almost independent of the interconnect length, since the length is sufficiently short such that the optical power loss in the waveguide is negligible. In this paper, only electrical power is evaluated for the optical data path, as listed in

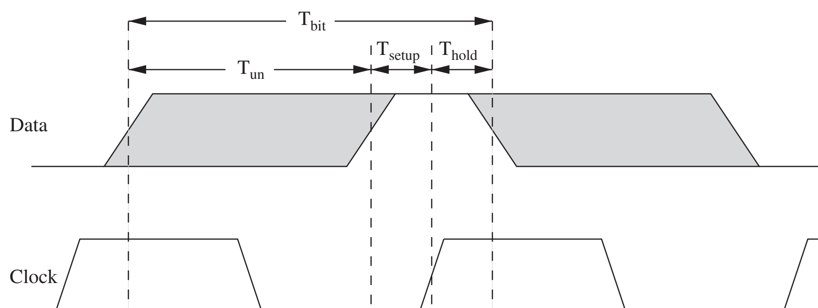

Fig. 8. A timing diagram of data and clock waveforms.

Table 4
Delay comparison between electrical and optical interconnects

Year		2004	2007	2010	2013	2016
Technology node		90 nm	65 nm	45 nm	32 nm	22 nm
Electrical	Delay (ps)	311.9	313.2	291.3	312.0	317.8
	# of register stages	1	2	2	4	7
	# of clock cycles	2	3	5	7	12
Optical-polymer	Delay (ps)	238.9	173.3	145.4	127.7	114.8
	# of register stages	1	1	1	1	1
	# of clock cycles	1	2	2	3	4
Optical-SOI	Delay (ps)	291.6	226.0	198.1	180.4	167.5
	# of register stages	1	1	1	1	1
	# of clock cycles	2	2	3	4	6

Table 5
Power consumption (mW) in optical and electrical interconnects

Year	2004	2007	2010	2013	2016
Technology node	90 nm	65 nm	45 nm	32 nm	22 nm
Transmitter	0.9	1.9	3.4	5.9	11.2
Receiver	0.5	0.5	0.3	0.3	0.3
Total optical	1.4	2.4	3.7	6.2	11.5
Electrical	9.8	16.9	21.7	33.4	45.3

Table 5. The power consumed by the transmitter dominates the power of the receiver, which is in contrast to the assumption made in [2]. The reason for this difference is that the modulator assumed in this analysis is CMOS compatible. The size as well as the capacitance of the modulator is large, requiring a large driver circuit. The power consumed by a 10 mm electrical interconnect is also listed for comparison in Table 5. The power consumption of both the electrical and optical interconnects increases due to the higher signal switching frequencies and greater leakage current. Optical interconnect consumes less power than electrical interconnect for all of the technology nodes.

4.4. Bandwidth density

Bandwidth density is an effective criterion for evaluating the ability to transmit data through a unit width. The maximum bit rate for a single interconnect is assumed to be the clock rate (one bit is transmitted per clock period). Similar to the power estimation, the clock frequency at the 22 nm technology node is assumed to be 36.3 GHz. Another parameter related to the bandwidth density is the interconnect pitch. As illustrated in Fig. 2, the optimal electrical interconnect width is $7W_{\min}$, corresponding to a pitch of $8W_{\min}$. For optical interconnects, the waveguide size should be larger than the optical mode size. Based on this limitation, the polymer waveguide pitch is assumed to be 4 μm, and the SOI waveguide pitch is assumed to be 0.85 μm. By introducing wavelength division multiplexing (WDM) [16], the bandwidth of the optical interconnect can be significantly improved. The bandwidth density of different interconnects is compared in Fig. 9. For optical interconnect with WDM, the channel number in a waveguide is assumed to be one at the 90 nm technology node, and to increase by four for each new technology node.

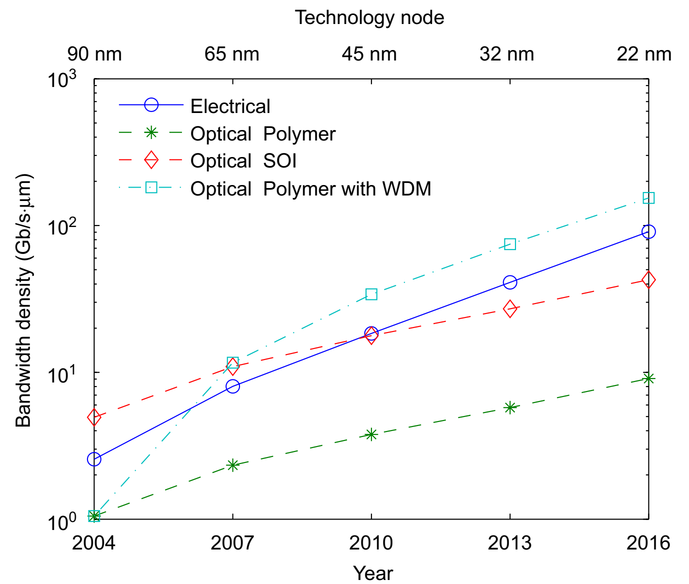

Fig. 9. Comparison of bandwidth density of electrical and optical interconnects. It is assumed that the area of the WDM components can be minimized in the realization of on-chip optical interconnect. The area penalty of WDM is therefore not included in the bandwidth density calculation.

4.5. Discussion

The critical length beyond which optical interconnect overcomes electrical interconnect is plotted in Fig. 10 for different design criteria. The lengths are normalized to the chip edge dimension. As shown in Fig. 10, the critical length is approximately one-tenth of the chip edge length at the 22 nm technology node.

A direct area comparison of on-chip optical and electrical interconnects might not be legitimate due to the different chip layers used by the two systems. With the use of a polymer waveguide in optical interconnect, an entirely new layer is required. Electrical interconnects, however, are implemented on traditional metal layers. The large optical transmitter and receiver are located at the two ends of the waveguide; in contrast to electrical repeaters, which are distributed along the interconnect. Via congestion issues, therefore, are avoided in optical interconnects.

As compared with [3], the results obtained in this analysis are optimistic for optical interconnect. The reason for this optimism is the choice of device models used in this analysis. Rather than a nitride waveguide [3], a polymer waveguide is assumed, increasing the light speed in the waveguide. Furthermore, a more aggressive WDM scheme is applied here, four additional channels per technology node rather than one additional channel per two technology nodes [3]. Another difference in this analysis is that a CMOS-compatible modulator is assumed, which is shown to be one of the most challenging elements in the optical data path. An additional advantage of optical interconnect is the smaller crosstalk noise as compared with electrical interconnect.

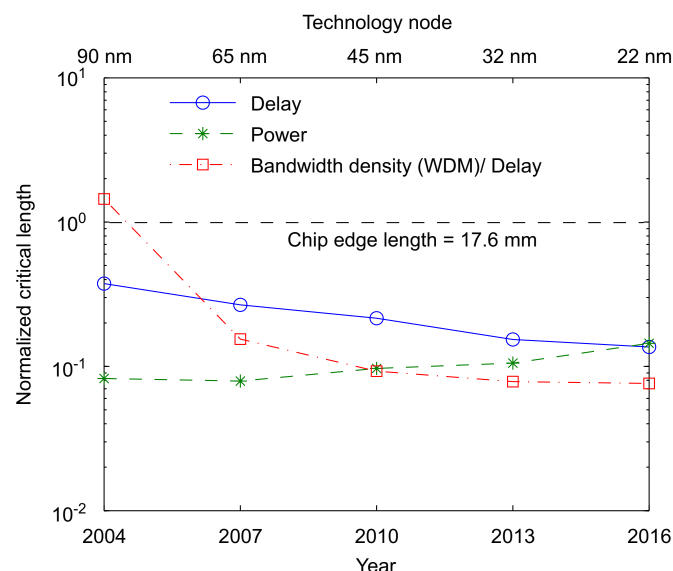

Fig. 10. Normalized critical length beyond which optical interconnect (a polymer waveguide) is advantageous over electrical interconnect.

5. Potential challenges in optical interconnects

Although significant progress had been made towards the development of on-chip optical interconnects over the past decade, there still remains a number of issues that need to be solved [28]. The first problem is the large footprint and power consumption of the optical components and driving circuits, particularly the modulator. This characteristic permits only a limited number of electrical connections that can be replaced with optics. A solution to this problem can be found by using alternative optical platforms, such as photonic bandgap structures [54] or ring resonators [18], which result in more compact optical components. These enhancements come at the price of stricter fabrication tolerances. A second problem is the generation of sufficient optical power to maintain optical operation. Although a state-of-the-art detector requires only 200 μW of optical power at the input [38], and passive optical loss during the light propagation can be as low as 25%, the number of required detectors can exceed 100–1000, even for simple optical interconnect systems. For example, a 64-bit optical interconnect system with 20 point-to-point optical connections requires 0.5 W of optical power at the IC input. Generation of this optical power requires multiple off-chip lasers and optical couplers. Using optical interconnects in multiple fan-out applications will further increase the input optical power requirement. Thus, both efficient light sources and detectors are crucial for the development of future on-chip optical interconnects. Finally, a set of integrated silicon-compatible WDM components needs to be developed in order to fully exploit the inherent advantages of optical interconnects.

6. Conclusions

A prediction of the performance characteristics of future CMOS compatible optical devices is described in this paper. Based on this prediction, electrical and optical on-chip interconnects are compared for various design criteria at different technology nodes. Critical lengths beyond which optical interconnect becomes advantageous in terms of delay, power, and bandwidth density/delay are presented. With technology scaling, these lengths are well below expected die size dimensions. The delay uncertainty of both electrical and optical interconnects is shown to significantly affect the actual signal delay. Delay uncertainty also provides an upper bound on the optical channel bandwidth.

References

[1] J.W. Goodman, et al., Optical interconnects for VLSI systems, Proc. IEEE 72 (7) (1984) 850–866.

[2] P. Kapur, K.C. Saraswat, Comparisons between electrical and optical interconnects for on-chip signaling, in: Proceedings of the IEEE International Interconnect Technology Conference, June 2002, pp. 89–91.

[3] M.J. Kobrinsky, et al., On-chip optical interconnects, Intel Technol. J. 8 (2) (2004) 129–141.
[4] I. O'Connor, F. Gaffiot, On-chip optical interconnect for low-power, in: E. Macii (Ed.), Ultra-Low Power Electronics and Design, Kluwer, Dordrecht, 2004.
[5] G. Chen, et al., On-chip copper-based vs. optical interconnects: delay uncertainty, latency, power, and bandwidth density comparative predictions, in: Proceedings of the IEEE International Interconnect Technology Conference, June 2006, pp. 39–41.
[6] International Technology Roadmap for Semiconductors, Semiconductor Industry Association, 2003.
[7] Y.I. Ismail, E.G. Friedman, Sensitivity of interconnect delay to on-chip inductance, in: Proceedings of the IEEE International Symposium on Circuits and Systems, May 2000, pp. 403–406.
[8] V. Adler, E.G. Friedman, Repeater design to reduce delay and power in resistive interconnect, IEEE Trans. Circuits Syst. Part II 45 (5) (1998) 607–616.
[9] P. Kapur, Scaling induced performance challenges/limitations of on-chip metal interconnects and comparisons with optical interconnects, Ph.D. Dissertation, Stanford University, Stanford, California, 2002.
[10] Y.I. Ismail, E.G. Friedman, Effects of inductance on the propagation delay and repeater insertion in VLSI circuits, IEEE Trans. Very Large Scale Integration (VLSI) Syst. 8 (2) (2000) 195–206.
[11] N.H. Mahmoud, Y.I. Ismail, Accurate rise time and overshoots estimation in *RLC* interconnects, in: Proceedings of the IEEE International Symposium on Circuits and Systems, May 2003, pp. 477–480.
[12] G. Chen, E.G. Friedman, Low power repeaters driving *RC* and *RLC* interconnects with delay and bandwidth constraints, IEEE Trans. Very Large Scale Integration (VLSI) Syst. 14 (2) (2006) 161–172.
[13] M.A. El-Moursy, E.G. Friedman, Optimum wire sizing of RLC interconnect with repeaters, Integration VLSI J. 38 (2) (2004) 205–225.
[14] R. Ho, K.W. Mai, M.A. Horowitz, The future of wires, Proc. IEEE 89 (4) (2001) 490–504.
[15] L.H. Chen, M. Marek-Sadowska, F. Brewer, Buffer delay change in the presence of power and ground noise, IEEE Trans. Very Large Scale Integration (VLSI) Syst. 11 (3) (2003) 461–473.
[16] G.P. Agrawal, Fiber-Optic Communication Systems, Wiley, New York, 1997.
[17] R.A. Soref, B.R. Bennett, Electrooptical effects in silicon, IEEE J. Quantum Electron. 23 (1) (1987) 123–129.
[18] Q. Xu, et al., Micrometer-scale silicon electro-optic modulator, Nature 435 (2005) 325–327.
[19] A. Liu, et al., A high-speed silicon optical modulator based on a metal–oxide–semiconductor capacitor, Nature 427 (2004) 615–618.
[20] L. Liao, et al., High speed silicon Mach-Zehnder modulator, Opt. Express 13 (8) (2005) 3129–3135.
[21] C.A. Barrios, V.R. Almeida, M. Lipson, Low-power-consumption short-length and high-modulation-depth silicon electrooptic modulator, J. Lightwave Technol. 21 (4) (2003) 1089–1098.
[22] S.M. Weiss, M. Molinari, P.M. Fauchet, Temperature stability for silicon-based photonic band-gap structures, Appl. Phys. Lett. 83 (10) (2003) 1980–1982.
[23] M. Haurylau, et al., Closed-form model of a capacitor-based electro-optical modulator, IEEE Trans. Electron Devices (2006), in press.
[24] B.S. Cherkauer, E.G. Friedman, A unified design methodology for CMOS tapered buffers, IEEE Trans. Very Large Scale Integration (VLSI) Syst. 3 (1) (1995) 99–111.
[25] N. Hedenstierna, K.O. Jeppson, Cmos circuit speed and buffer optimization, IEEE Trans. Comput. Aided Des. Integrated Circuits Syst. 6 (2) (1987) 270–281.
[26] T. Sakurai, A.R. Newton, Alpha-power law mosfet model and its applications to CMOS inverter delay and other formulas, IEEE J. Solid-State Circuits 25 (2) (1990) 584–594.
[27] O. Boyraz, B. Jalali, Demonstration of a silicon raman laser, Opt. Express 12 (21) (2004) 5269–5273.
[28] M. Haurylau, et al., Optical interconnect roadmap: challenges and critical directions, IEEE J. Sel. Top. Quantum Electron. (2006), in press.
[29] Y.A. Vlasov, S.J. McNab, Losses in single-mode silicon-on-insulator strip waveguides and bends, Opt. Express 12 (8) (2004) 1622–1631.
[30] L. Eldada, L.W. Shacklette, Advances in polymer integrated optics, IEEE J. Sel. Top. Quantum Electron. 6 (1) (2000) 54–68.
[31] S.V. Averine, Y.C. Chan, Y. Lam, Geometry optimization of interdigitated schottky-barrier metal-semiconductor-metal photodiode structures, Solid-State Electron. 45 (3) (2001) 441–446.
[32] J. Oh, et al., Interdigitated Ge p-i-n photodetectors fabricated on a Si substrate using graded SiGe buffer layers, IEEE J. Quantum Electron. 38 (9) (2002) 1238–1241.
[33] D. Buca, et al., Metal–germanium–metal ultrafast infrared detectors, J. Appl. Phys. 92 (12) (2002) 7599–7605.
[34] J. Oh, S.K. Banerjee, J. Campbell, Metal–germanium–metal photodetectors on heteroepitaxial Ge-on-Si with amorphous Ge Schottky barrier enhancement layers, IEEE Photonics Technol. Lett. 15 (5) (2003) 745–747.
[35] L. Pavesi, D.J. Lockwood, Silicon Photonics, Springer, New York, 2004.
[36] S.Y. Chou, Y. Liu, 32 GHz metal–semiconductor–metal photodetector on silicon, Appl. Phys. Lett. 61 (15) (1992) 1760–1762.
[37] J.J. Morikuni, et al., Improvements to the standard theory for photoreceiver noise, J. Lightwave Technol. 12 (4) (1994) 1174–1184.
[38] B. Yang, et al., 10-Gb/s all-silicon optical receiver, IEEE Photonics Technol. Lett. 15 (5) (2003) 745–747.
[39] A.H. Ajami, K. Banerjee, M. Pedram, Modeling and analysis of nonuniform substrate temperature effects on global ULSI interconnects, IEEE Trans. Comput. Aided Des. Integrated Circuits Syst. 24 (6) (2005) 849–861.
[40] P. Friedberg, et al., Modeling within-die spatial correlation effects for process-design co-optimization," in: Proceedings of the IEEE International Symposium on Quality Electronic Design, March 2005, pp. 516–521.
[41] Y. Cao, et al., Effective on-chip inductance modeling for multiple signal lines and application to repeater insertion, IEEE Trans. Very Large Scale Integration (VLSI) Syst. 10 (6) (2002) 799–805.
[42] Z. Liu, et al., Threshold voltage model for deep-submicrometer MOSFET's, IEEE Trans. Electron Devices 40 (1) (1993) 86–95.
[43] G.F. Niu, G. Ruan, R.M.M. Chen, Further comments on 'threshold voltage model for deep-submicrometer MOSFET's' and its extension to subthreshold operation, IEEE Trans. Electron Devices 43 (12) (1996) 2311–2312.
[44] N. Arora, MOSFET Models for VLSI Circuit Simulation—Theory and Practice, Springer/Wien, NY, 1993.
[45] N.G. Einspruch, VLSI Electronics Microstructure Science—Advanced MOS Device Physics, Academic Press, San Diego, California, 1989.
[46] C.K. Hu, J.M.E. Harper, Copper interconnect: fabrication and reliability, in: Proceedings of the VLSI Technology, Systems, and Applications, June 1997, pp. 18–22.
[47] S. Wong, G. Lee, D. Ma, Modeling of interconnect capacitance, delay, and crosstalk in VLSI, IEEE Trans. Semicond. Manuf. 13 (1) (2000) 108–111.
[48] A.B. Kahng, S. Muddu, E. Sarto, On switch factor based analysis of coupled *RC* interconnects, in: Proceedings of the IEEE Great Lakes Symposium on VLSI, June 2000, pp. 79–84.
[49] A.B. Kahng, et al., Interconnect tuning strategies for high-performance ICs, in: Proceedings of the IEEE Design, Automation and Test in Europe Conference, February 1998, pp. 471–478.
[50] Y.I. Ismail, E.G. Friedman, On the extraction of on-chip inductance, J. Circuits Syst. Comput. 12 (1) (2003) 31–40.
[51] Y. Cao, et al., Switch-factor based loop *RLC* modeling for efficient timing analysis, IEEE Trans. Very Large Scale Integration (VLSI) Syst. 13 (9) (2005) 1072–1078.
[52] Y. Cao, et al., Design sensitivities to variability: extrapolation and assessments in nanometer VLSI, in: Proceedings of the IEEE ASIC/SOC Conference, September 2002, pp. 411–415.
[53] V. Venkatraman, W. Burleson, Robust multi-level current-mode on-chip interconnect signaling in the presence of process variations, in:

Proceedings of the IEEE International Symposium on Quality Electronic Design, March 2005, pp. 522–527.
[54] M. Soljacic, et al., Photonic-crystal slow-light enhancement of nonlinear phase sensitivity, J. Opt. Soc. Am. B 19 (9) (2002) 2052–2059.

Guoqing Chen received his B.S. (with honor) and M.S. degrees in Electronic Engineering from the Tsinghua University, Beijing, China, in 1998 and 2001, respectively. He is now working toward his Ph.D. degree in the Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY.

In the summers of 2004 and 2005, he was with Manhattan Routing, Inc., New York City, NY, where his work was focused on the development of CAD tools for the timing closure stage in the IC design process. His research interests include high-speed interconnect design and modeling, signal integrity, clock distribution networks, and low-power circuit design.

Hui Chen received her Bachelor of Engineering degree and Master of Engineering degree from the Tianjin University, China in 1996 and 1999, respectively. She is now working toward her Ph.D. degree in the Department of Physics and Astronomy, University of Rochester, Rochester, NY.

Her current research interests are in the areas of silicon nanocrystal light emission and optical gain, and silicon-based on-chip optical interconnects.

Mikhail Haurylau is currently employed with KLA-Tencor Corp. He received his B.S. degree in Microelectronic Engineering from the Belarusian State University of Informatics and Radioelectronics, Minsk, Belarus, in 2001, and his Ph.D. degree in the Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY, in 2006.

His research interests are in the areas of tunable photonic bandgap structures, silicon-based electro-optic modulators and on-chip optical interconnects.

Nicholas A. Nelson received his M.S. degree in Electrical and Computer Engineering from the University of Rochester in 2002 and is currently working towards his Ph.D. degree in the Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY. His current research interests are in the area of unique multiprocessor-on-chip organizations and power amplifier control logic.

David H. Albonesi is an Associate Professor of Electrical and Computer Engineering at the Cornell University and a Member of the Computer Systems Laboratory. He received his Ph.D. from the University of Massachusetts, Amherst in 1996.

His current research interests include adaptive, power-efficient, and reliability-aware computer architectures, multithreaded and multicore processors, and microarchitectures exploiting new technologies. Dr. Albonesi has received a National Science Foundation CAREER Award and IBM Faculty Partnership Awards in 2001, 2002, and 2003. He holds eight US patents and is a Senior Member of the IEEE.

Philippe M. Fauchet is a Distinguished Professor of Electrical and Computer Engineering at the University of Rochester where he is also a Professor of Optics, Biomedical Engineering, and Materials Science, and a Senior Scientist at the Laboratory for Laser Energetics. He received his Engineer's Degree from the Faculté Polytechnique de Mons in Belgium in 1978, an M.S. in Engineering from the Brown University in 1980, and, in 1984, a Ph.D. in Applied Physics from the Stanford University, where he stayed as an IBM Post-doctoral Fellow and Acting Assistant Professor for 1 year before he took a faculty position in Electrical Engineering at the Princeton University. Dr. Fauchet's current research interests include nanoscience and nanotechnology with silicon quantum dots, silicon-based LEDs and lasers, semiconductor optoelectronics and devices, optical interconnects, silicon-based biosensors, the silicon/biology interface, and optical diagnostics. In 1998, he founded the Center for Future Health, a multidisciplinary laboratory where engineers, physicians, and social scientists develop affordable technology that can be used at home to keep individuals independent and healthy. Dr. Fauchet received an IBM Faculty Development Award in 1985, an NSF Presidential Young Investigator Award in 1987, an Alfred P. Sloan Research Fellowship in 1988, and the 1990–1993 Prix Guibal & Devillez for his work on porous silicon. He has chaired many conferences and workshops, is the author of over 350 publications, has edited nine books, and is a Fellow of the Optical Society of America, the American Physical Society, and the Institute of Electrical and Electronic Engineers.

Eby G. Friedman received his B.S. degree from the Lafayette College in 1979, and M.S. and Ph.D. degrees from the University of California, Irvine, in 1981 and 1989, respectively, all in Electrical Engineering.

From 1979 to 1991, he was with the Hughes Aircraft Company, rising to the position of manager of the Signal Processing Design and Test Department, responsible for the design and test of high-performance digital and analog IC's. He has been with the Department of Electrical and Computer Engineering at the University of Rochester since 1991, where he is a Distinguished Professor, the Director of the High-Performance VLSI/IC Design and Analysis Laboratory, and the Director of the Center for Electronic Imaging Systems. He is also a Visiting Professor at the Technion—Israel Institute of Technology. His current research and teaching interests are in high-performance synchronous digital and mixed-signal microelectronic design and analysis with application to high-speed portable processors and low-power wireless communications.

He is the author of more than 300 papers and book chapters, six patents, and the author or editor of eight books in the fields of high-speed and low-power CMOS design techniques, high-speed interconnect, and the theory and application of synchronous clock and power distribution networks. Dr. Friedman is the Regional Editor of the *Journal of Circuits, Systems and Computers*, a Member of the editorial boards of the *Analog Integrated Circuits and Signal Processing, Microelectronics Journal, Journal of Low Power Electronics*, and *Journal of VLSI Signal Processing*, Chair of the *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* steering committee, and a Member of the technical program committee of a number of conferences. He previously was the Editor-in-Chief of the *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, a Member of the editorial board of the *Proceedings of the IEEE and IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, Program and Technical chair of several IEEE conferences, and a recipient of the University of Rochester Graduate Teaching Award and a College of Engineering Teaching Excellence Award. Dr. Friedman is a Senior Fulbright Fellow and an IEEE Fellow.