

Predictions of CMOS Compatible On-Chip Optical Interconnect

Guoqing Chen, Hui Chen, Mikhail Haurylau, Nicholas Nelson
Philippe M. Fauchet, and Eby G. Friedman
Department of Electrical and Computer Engineering
University of Rochester, Rochester, New York 14627

David Albonesi
Department of Electrical and Computer Engineering
Cornell University, Ithaca, New York 14853

ABSTRACT

Interconnect has become a primary bottleneck in integrated circuit design. As CMOS technology is scaled, it will become increasingly difficult for conventional copper interconnect to satisfy the design requirements of delay, power, bandwidth, and noise. On-chip optical interconnect has been considered as a potential substitute for electrical interconnect in the past two decades. In this paper, predictions of the performance of CMOS compatible optical devices are made based on current state-of-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.

Categories and Subject Descriptors

B.7.1 [Integrated Circuits]: Types and Design Styles
— advanced technologies

General Terms

Design, Performance

Keywords

Optical interconnect, On-chip, CMOS compatible, Trends

This research was supported by the National Science Foundation under Contract No. CCR-0304574.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SLIP'05, April 2–3, 2005, San Francisco, California, USA.

Copyright 2005 ACM 1-59593-033-7/05/0004 ...\$5.00.

1. INTRODUCTION

In deep submicrometer VLSI technologies, interconnect plays an increasingly important role. Multiple design criteria are considered in interconnect design, such as delay, power, bandwidth, and noise. With technology scaling, it has become increasingly difficult for conventional copper based electrical interconnect to satisfy these requirements. One promising candidate to satisfy these performance objectives is optical interconnect.

Optical devices are widely used in the telecommunication area, and are commonly 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. Recently, several comparisons have been made between on-chip electrical and optical interconnects [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.

In this paper, a more comprehensive comparison between optical and electrical interconnects is performed at different technology nodes based on a practical prediction of optical device development. This comparison is particularly challenging since optical interconnect is a young fast-developing technology while electrical interconnect is relatively mature. The paper is organized as follows. In section 2, *RLC*-based delay and power models of electrical interconnect are reviewed. In section 3,

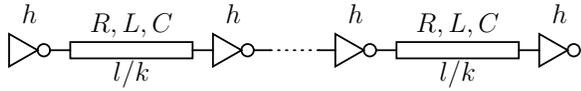


Figure 1: Repeater insertion in an RLC interconnect.

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 interconnect are evaluated for different design criteria. Some conclusions are offered in section 5.

2. ELECTRICAL INTERCONNECT

Repeaters are widely used in submicrometer circuits to reduce interconnect delay, transition times, and crosstalk noise. In this section, an RLC interconnect with repeaters is examined at different technology nodes based on the ITRS [4].

The capacitance and resistance per unit length of the interconnect can be obtained directly from the physical geometries, where the space between adjacent interconnects is assumed equal to the minimum interconnect width. The interconnect inductance, however, depends upon the distribution of the current return paths which are 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, and the magnitude of the on-chip inductance is a slowly varying function of the wire geometry [5]. Based on these two characteristics, a fixed value of $0.5 \text{ pH}/\mu\text{m}$ [5] is assumed for all of the technology nodes.

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 [6]. 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 [7]. 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 [8] where the repeater output capacitance and input slew effects are considered. By including the repeater output capacitance, the variable ζ [8] used to characterize inductance effects becomes

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

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} . Note that the K for the 50% delay and the transition time is different, and the corresponding ζ is therefore denoted as ζ_d and ζ_r , respectively. The related transistor parameters can also be found in the ITRS. 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 [9], 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,

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

When $\zeta_r < 0.5$, the interconnect is highly inductance dominant, and (3) can introduce a large error. The effects of the input transition time on the propagation delay and far end transition time are treated similarly as in an RC interconnect [10].

The interconnect power models used in this analysis are the same as those models described in [10]. Three degrees of freedom are explored in the electrical interconnect design: the wire width, and the number and size of the repeaters. Various combinations are examined to determine the optimal interconnect design with respect to a specific design criterion.

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. 2. The system consists of four primary optical elements: an off-chip laser, an optical modulator, a polymer waveguide, and an optical detector. Because the optical modulator and detector in each data link operate at the same wavelength, there is an inherent conflict in the requirements of the optical material. In

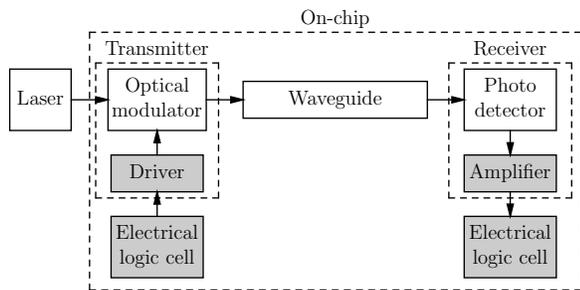


Figure 2: An on-chip optical interconnect data path.

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 $1.5 \mu\text{m}$ 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 [11]. The performance and integration ability of optical devices, however, can 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 [12]. One of the few suitable mechanisms for varying the refractive index in pure silicon is the free carrier plasma dispersion effect [13]. There are primarily two electrical structures for changing the carrier concentration in silicon devices. The conventional technique is to inject and deplete carriers in the intrinsic region of a p-i-n diode. One example of this structure is described in [12]. With this approach, a substantial change in the carrier concentration can be obtained over a large volume. The speed of this structure, however, is greatly limited by

Table 2: Prediction of modulator parameters.

Modulator parameter	Δn_{eff}	Delay (ps)	Length (μm)	Capacitance (pF)
[14]	3.75×10^{-5}	114	10^4	72.3
Year 2016	1.5×10^{-4}	15-20	50	0.35

the slow carrier recombination process. An alternative electrical structure is a MOS capacitor. The change in the carrier concentration is achieved by redistribution rather than injection and depletion of carriers, causing a higher modulation speed. The first MOS capacitor based electro-optical modulator was demonstrated by Liu *et al.* [14] and operates at frequencies greater than one gigahertz. 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 GHz to 40 GHz by the year 2016.

Because the device structure used in [14] is a Mach-Zehnder interferometer, the modulator has a large footprint (10 mm long), which also results in an excessive capacitance; hence, increasing the delay and power consumption of the driver circuits. Simulations and early experiments performed by Barrios *et al.* [12] show that optical resonator-based structure can drastically decrease the modulator size down to $10 \mu\text{m}$ to $30 \mu\text{m}$ while maintaining the same operating principle and speed. Further reductions in the modulator size are possible by using photonic bandgap structures. In this paper, a predictive model is generated which combines the advantages of the structures used in [12] and [14], as listed in Table 1.

Based on these considerations, some modulator parameters are predicted for the year 2016, as listed in Table 2. Although the improvement of modulator performance will likely exhibit a step-like behavior, a linear or inverse-linear interpolation (for simplicity) is used to estimate the parameters for the intermediate years. The effective index change Δn_{eff} is determined as

$$\Delta n_{eff} = \sigma \Delta n, \quad (4)$$

where Δn is the index change due to the external modulation, and σ is the fraction of the optical mode power inside the index variation region. For example, in [14], σ is only around 9%. The delay of a resonator-based electro-optical modulator is related to Δn_{eff} and is determined by a predictive analysis of a MOS capacitor modulator [14].

A series of tapered inverters [15] 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 [16]. A minimum sized inverter is used as the first stage. The number of

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

		Electrical structure	
		p-i-n diode	MOS capacitor
Optical Structure	Mach-Zehnder	—	Reference [14] High speed (1 GHz and up) CMOS compatible Large size (10 mm) High power consumption
	Resonator	Reference [12] Low speed (20-30 MHz) CMOS compatible Small size (20 μm) Low power consumption	Predictive model High speed CMOS compatible Small size Low power consumption

stages can be obtained as $N = \log \frac{C_M}{C_{g0}} / \log 3.6$, where C_M is the modulator capacitance. The delay model of each stage is described in [17].

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 could be directly applied to an optical interconnect configuration using an on-chip Raman silicon laser [18], 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 structure will likely diminish this advantage.

Given the operating wavelength of on-chip optical interconnect, there are primarily two candidates for the waveguide material. 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 links, optimized for signal propagation delay and smaller losses [19], low-loss polymer waveguides are better suited [20]. Although polymer waveguides require more area than SOI waveguides, polymer waveguides are fabricated on an additional layer and therefore will not reduce the on-chip silicon resources. In this paper, low-refractive index strip polymer waveguides are assumed with an effective index of 1.4 [20].

3.3 Receivers

The receiver has two components: a photo-detector that converts light into electricity followed by receiver circuits that amplify the analog electrical signal to a digital voltage level. A simplified equivalent circuit model is shown in Fig. 3.

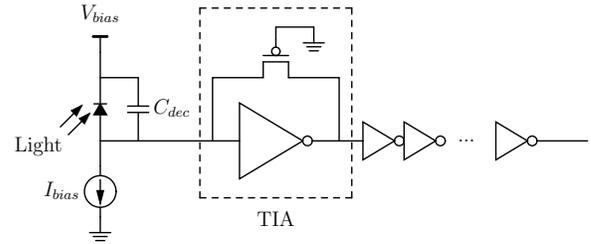


Figure 3: Circuit model of an optical receiver.

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 [21]. 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 $\tau_{dec} = 0.315T_r$. In 2002, an interdigitated Ge p-i-n detector fabricated on a Si substrate with a 3 dB bandwidth of 3.8 GHz was demonstrated [22]. There have been several other papers published on SiGe detectors and these detectors exhibit similar performance levels, such as [23, 24]. The bandwidth or delay of most of these detectors are limited by the carrier transit time, which can be improved through device optimization.

Based on a model proposed by Averine *et al.* [21], the trend in the performance of future detectors is projected. In Fig. 4, the MSM detector response time as a function of electrode width is plotted for different detector sizes and is compared with some experimental results. The spacing between the electrodes is assumed to be equal to the electrode width. As shown in Fig. 4, an optimal electrode width exists that produces a minimum response time. When the electrode is too narrow,

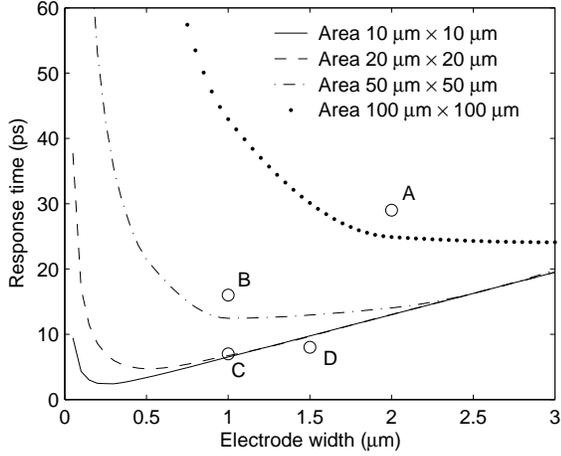


Figure 4: Detector response time versus electrode width. A ($100\ \mu\text{m} \times 100\ \mu\text{m}$) [21], B ($50\ \mu\text{m} \times 50\ \mu\text{m}$) [21], C ($20\ \mu\text{m} \times 20\ \mu\text{m}$) [26], and D ($10\ \mu\text{m} \times 10\ \mu\text{m}$) [23].

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 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 [25].

The photo-current I_{ph} from the photo-detector is pre-amplified by a transimpedance amplifier (TIA). The TIA 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 [7]. 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} [7]. For the receiver circuits, the static power dominates and is

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

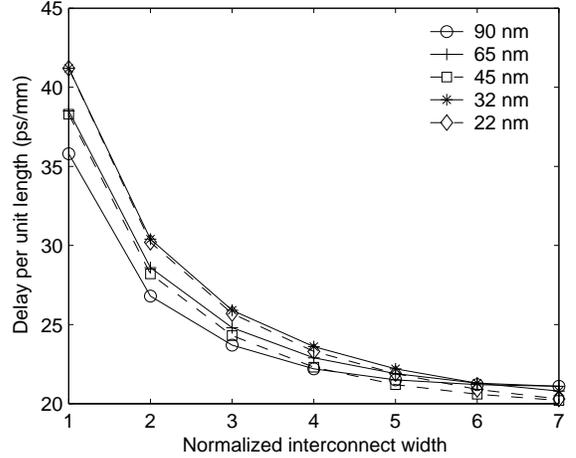


Figure 5: Minimum delay per unit length as a function of interconnect width.

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 which is determined by the output swing requirement. The delay of the receiver circuit τ_{cct} is obtained by approximating the circuit as a one pole system, $\tau_{cct} = 0.7/(2\pi\Delta f_{req})$, where Δf_{req} is the bandwidth requirement.

4. COMPARISON BETWEEN ELECTRICAL AND OPTICAL INTERCONNECTS

In this section, electrical interconnects are compared with optical interconnects for different design criteria.

4.1 Delay

The optimal number and size of repeaters along an RLC interconnect can be determined to achieve the minimum delay [8]. This minimum delay can be further decreased by increasing the wire width [27]. The achievable minimum delay per unit length for different wire widths is illustrated in Fig. 5. The interconnect widths are normalized to the minimum wire width W_{min} as predicted by the ITRS. As shown in Fig. 5, scaling has only a small effect on the delay of interconnects with repeaters, consistent with the conclusions from [28]. 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 a constant — 20 ps/mm for all technology nodes of interest.

The delay distribution of a 1 cm optical data path is listed in Table 3. The delay of the transmitter and the 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. As listed in Table 3, the delay of the transmitter is much greater than that of the receiver. With improvements in technology,

Table 3: Delay (ps) distribution in a 1 cm optical data path as compared with the electrical interconnect delay.

Year	2004	2007	2010	2013	2016
Technology node	90 nm	65 nm	45 nm	32 nm	22 nm
Modulator driver	83.7	45.8	25.8	16.3	9.5
Modulator	114.0	52.1	30.4	20.0	14.3
Waveguide	46.7	46.7	46.7	46.7	46.7
Photo-detector	1.4	0.5	0.3	0.3	0.2
Receiver amplifier	37.5	16.9	10.4	6.9	4.0
Total optical	283.3	162.0	113.6	90.2	74.7
Electrical	200.0	200.0	200.0	200.0	200.0

both the transmitter and the receiver delay are expected to decrease. By the year 2007 (the 65 nm technology node), optical interconnect is expected to operate faster than electrical interconnect.

4.2 Power

For electrical interconnects, the power should be evaluated under specific design requirements, such as delay and bandwidth. A minimum sized wire without repeaters consumes minimum power, however, this configuration is not practical for global interconnect due to the significant delay and low bandwidth of the line. The power-delay product (PDP), therefore, is often used as an effective design criterion [27]. For each wire size, a local minimum PDP can be obtained by adjusting the repeater size and number. From simulations, the global minimum PDP can be achieved with a wire size ranging from $4W_{min}$ to $5W_{min}$ for different technology nodes.

The power consumed by the optical interconnect is almost independent of the interconnect length, since the length is sufficiently short 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 Table 4. 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. Note that there is a significant power decrease from the 90 nm technology node to the 65 nm technology node, which reflects the expected improvements in modulator structures from Mach-Zehnder to a resonator. The power consumed by a 10 mm delay-optimized electrical interconnect is also listed for comparison in Table 4. Although the optimized delay of electrical interconnect is almost constant with technology scaling, the power increases significantly due to the higher signal switching frequency and greater leakage current. Optical interconnects are also beneficial from a power perspective.

Table 4: Power consumption (mW) in an optical data path.

Year	2004	2007	2010	2013	2016
Technology node	90 nm	65 nm	45 nm	32 nm	22 nm
Transmitter	177.5	18.4	8.6	6.0	5.0
Receiver	0.4	0.3	0.2	0.3	0.3
Total optical	177.9	18.7	8.8	6.3	5.3
Electrical	7.5	12.7	15.8	22.8	31.2

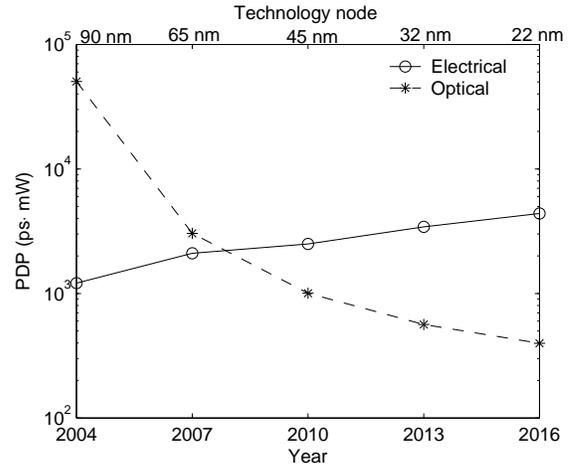


Figure 6: Comparison of the PDP of electrical and optical interconnects (for a length of 1 cm).

The PDP of an electrical interconnect is compared with an optical interconnect in Fig. 6. Note the crossover point between the 65 nm technology node and the 45 nm technology node.

4.3 Bandwidth density

The 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). With proper design, this bandwidth can be achieved in both electrical and optical interconnect. The bandwidth density, therefore, is only determined by the interconnect pitch. As illustrated in Fig. 5, the minimum delay of an electrical interconnect can be achieved with an interconnect width of $7W_{min}$, 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 waveguide pitch is assumed to be $4 \mu\text{m}$, much larger than the electrical interconnect pitch. Single wavelength optical interconnects, therefore, are not beneficial if a high bandwidth density is desired. The bandwidth of optical interconnect, however, can be significantly improved by introducing wavelength division multiplexing

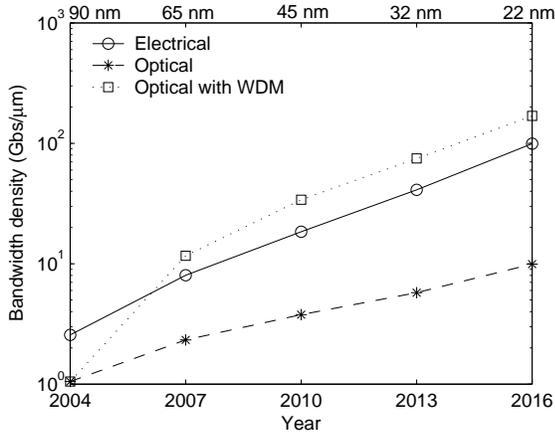


Figure 7: Comparison of bandwidth density of electrical and optical interconnects.

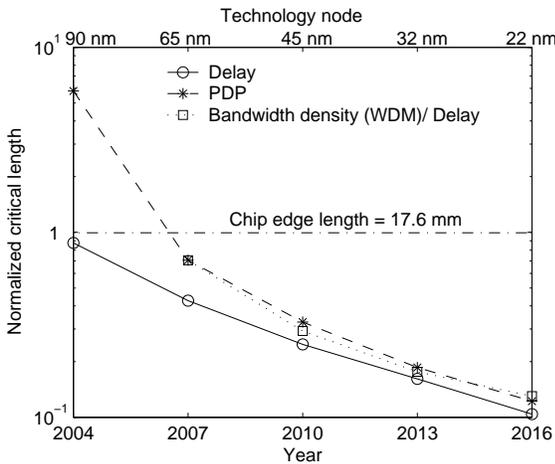


Figure 8: Normalized critical length beyond which optical interconnect is advantageous over electrical interconnect.

(WDM) [11]. The bandwidth density of different interconnects is compared in Fig. 7. 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.

4.4 Discussion

The critical length beyond which optical interconnect overcomes electrical interconnect is plotted in Fig. 8 for different design criteria. The lengths are normalized to the edge of the chip die dimension. As shown in Fig. 8, 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, a whole 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 the 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 more optimistic for optical interconnect. The reason for this optimism is the device models adopted in this approach. Rather than a nitride waveguide [3], a polymer waveguide is used in this analysis, increasing the light speed in the waveguide. Furthermore, a more aggressive WDM scheme is assumed 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.

Note that in this analysis, the optical interconnect design is fixed. The optical interconnect can be improved with respect to a specific criterion by further optimizing the modulator driver and receiver circuits. An additional advantage of optical interconnect is the smaller crosstalk noise as compared with electrical interconnect.

5. 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 interconnect are compared for various design criteria at different technology nodes. Critical lengths beyond which optical interconnect becomes advantageous in order to achieve the minimum delay, PDP, and maximum bandwidth density/delay are presented. With technology scaling, these lengths are well below expected chip die size dimensions.

6. REFERENCES

- [1] J. W. Goodman *et al.*, “Optical Interconnects for VLSI Systems,” *Proceedings of the IEEE*, Vol. 72, No. 7, pp. 850–866, July 1984.
- [2] P. Kapur and K. C. Saraswat, “Comparisons Between Electrical and Optical Interconnects for On-Chip Signaling,” *Proceedings of the IEEE International Interconnect Technology Conference*, pp. 89–91, June 2002.
- [3] M. J. Koblinsky *et al.*, “On-Chip Optical Interconnects,” *Intel Technology Journal*, Vol. 8, No. 2, pp. 129–141, May 2004.
- [4] *The International Technology Roadmap for Semiconductors*, Semiconductor Industry Association, California, 2003.

- [5] Y. I. Ismail and E. G. Friedman, "Sensitivity of Interconnect Delay to On-Chip Inductance," *Proceedings of the IEEE International Symposium on Circuits and Systems*, pp. 403–406, May 2000.
- [6] V. Adler and Eby G. Friedman, "Repeater Design to Reduce Delay and Power in Resistive Interconnect," *IEEE Transactions on Circuits and System II: Analog and Digital Signal Processing*, Vol. 45, No. 5, pp. 607–616, May 1998.
- [7] P. Kapur, *Scaling Induced Performance Challenges/Limitations of On-Chip Metal Interconnects and Comparisons with Optical Interconnects*, Ph.D. Thesis, Stanford University, Stanford, California, 2002.
- [8] Y. I. Ismail and E. G. Friedman, "Effects of Inductance on the Propagation Delay and Repeater Insertion in VLSI Circuits," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Vol. 8, No. 2, pp. 195–206, April 2000.
- [9] N. H. Mahmoud and Y. I. Ismail, "Accurate Rise Time and Overshoots Estimation in *RLC* Interconnects," *Proceedings of the IEEE International Symposium on Circuits and Systems*, pp. 477–480, May 2003.
- [10] G. Chen and E. G. Friedman, "Low Power Repeaters Driving *RC* Interconnect with Delay and Bandwidth Constraints," *Proceedings of the IEEE International SOC Conference*, pp. 335–339, September 2004.
- [11] G. P. Agrawal, *Fiber-Optic Communication Systems*, Wiley, New York, 1997.
- [12] C. A. Barrios, V. R. Almeida, and M. Lipson, "Low-Power-Consumption Short-Length and High-Modulation-Depth Silicon Electrooptic Modulator," *Journal of Lightwave Technology*, Vol. 21, No. 4, pp. 1089–1098, April 2003.
- [13] R. A. Soref and B. R. Bennett, "Electrooptical Effects in Silicon," *IEEE Journal of Quantum Electronics*, Vol. 23, No. 1, pp. 123–129, January 1987.
- [14] A. Liu *et al.*, "A High-Speed Silicon Optical Modulator Based on a Metal-Oxide-Semiconductor Capacitor," *Nature*, Vol. 427, pp. 615–618, February 2004.
- [15] B. S. Cherkauer and E. G. Friedman, "A Unified Design Methodology for CMOS Tapered Buffers," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Vol. 3, No. 1, pp. 99–111, March 1995.
- [16] N. Hedenstierna and K. O. Jeppson, "CMOS Circuit Speed and Buffer Optimization," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Vol. 6, No. 2, pp. 270–281, March 1987.
- [17] T. Sakurai and A. R. Newton, "Alpha-Power Law MOSFET Model and its Applications to CMOS Inverter Delay and Other Formulas," *IEEE Journal of Solid-State Circuits*, Vol. 25, No. 2, pp. 584–594, April 1990.
- [18] O. Boyraz and B. Jalali, "Demonstration of a Silicon Raman Laser," *Optical Express*, Vol. 12, No. 21, pp. 5269–5273, October 2004.
- [19] Y. A. Vlasov and S. J. McNab, "Losses in Single-Mode Silicon-On-Insulator Strip Waveguides and Bends," *Optical Express*, Vol. 12, No. 8, pp. 1622–1631, April 2004.
- [20] L. Eldada and L. W. Shacklette, "Advances in Polymer Integrated Optics," *IEEE Journal of Selected Topics in Quantum Electronics*, Vol. 6, No. 1, pp. 54–68, January/February 2000.
- [21] S. V. Averine, Y. C. Chan, and Y.L. Lam, "Geometry Optimization of Interdigitated Schottky-Barrier Metal-Semiconductor-Metal Photodiode Structures," *Solid-State Electronics*, Vol. 45, No. 3, pp. 441–446, March 2001.
- [22] J. Oh *et al.*, "Interdigitated Ge p-i-n Photodetectors Fabricated on a Si Substrate Using Graded SiGe Buffer Layers," *IEEE Journal of Quantum Electronics*, Vol. 38, No. 9, pp. 1238–1241, September 2002.
- [23] D. Buca *et al.*, "Metal-Germanium-Metal Ultrafast Infrared Detectors," *Journal of Applied Physics*, Vol. 92, No. 12, pp. 7599–7605, December 2002.
- [24] J. Oh, S. K. Banerjee, and J.C. Campbell, "Metal-Germanium-Metal Photodetectors on Heteroepitaxial Ge-on-Si with Amorphous Ge Schottky Barrier Enhancement Layers," *IEEE Photonics Technology Letters*, Vol. 15, No. 5, pp. 745–747, May 2003.
- [25] L. Pavesi and D. J. Lockwood, *Silicon Photonics*, Springer, New York, 2004.
- [26] S. Y. Chou and Y. Liu, "32 GHz Metal-Semiconductor-Metal Photodetector on Silicon," *Applied Physics Letter*, Vol. 61, No. 15, pp. 1760–1762, October 1992.
- [27] M. A. El-Moursy and E. G. Friedman, "Optimum Wire Sizing of *RLC* Interconnect with Repeaters," *Proceedings of the IEEE Great Lakes Symposium on VLSI*, pp. 27–32, April 2003.
- [28] R. Ho, K. W. Mai, and M. A. Horowitz, "The Future of Wires," *Proceedings of the IEEE*, Vol. 89, No. 4, pp. 490–504, April 2001.