On-Chip Copper-Based vs. Optical Interconnects: Delay Uncertainty, Latency, Power, and Bandwidth Density Comparative Predictions

Guoqing Chen, Hui Chen*, Mikhail Haurylau, Nicholas A. Nelson,
David H. Albonesi**, Philippe M. Fauchet, and Eby G. Friedman
Department of Electrical and Computer Engineering, University of Rochester, Rochester, New York, 14627
*Department of Physics and Astronomy, University of Rochester, Rochester, New York, 14627
**Department of Electrical and Computer Engineering,Cornell University, Ithaca, New York, 14853

***Abstract*— As CMOS technology is scaled, it has become increasingly difficult for conventional copper interconnect to satisfy different 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-art optical technologies. Based on these predictions, electrical and optical interconnects are compared for delay uncertainty, latency, power, and bandwidth density.**

I. Introduction

In deep submicrometer VLSI technologies, it has become increasingly difficult for conventional copper based electrical interconnect to satisfy the design requirements of delay, power, bandwidth, and delay uncertainty. One promising candidate to solve this problem is optical interconnect. Based on a practical prediction of optical device development, a comprehensive comparison between optical and electrical interconnects is described in this paper for different technology nodes. As compared with [1], more accurate optical device models are adopted. Delay uncertainty is also considered. The paper is organized as follows. In section II, a delay-optimal design of RLC interconnect is presented. In section III, predictions of the performance characteristics of next generation optical devices are made based on current technology trends. In section IV, electrical and optical interconnect are evaluated for different design criteria. Some conclusions are offered in section V.

II. Scaling of Electrical Interconnect

The delay model of an RLC interconnect with repeaters described in [1] is used for electrical interconnect analysis. Three degrees of freedom (the wire width, and the number and size of the repeaters) are explored in the electrical interconnect design process to achieve the minimum delay. The minimum delay per unit length is plotted as a function of wire width in Fig. 1. The interconnect widths are normalized to the minimum wire width W_{min} as predicted by the ITRS [2]. 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 and the minimum delay per unit length is approximately in the range of 20 to 22 ps/mm for all of the technology nodes of interest.

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

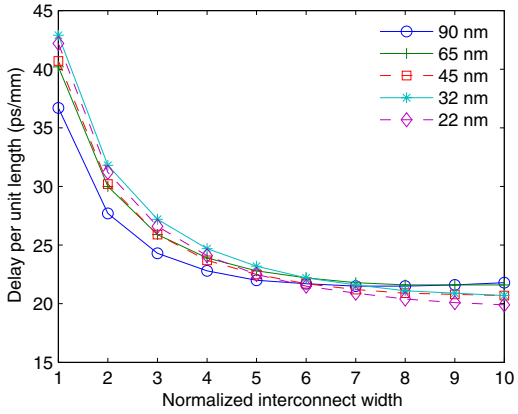

Fig. 1. Minimum delay per unit length as a function of interconnect width.

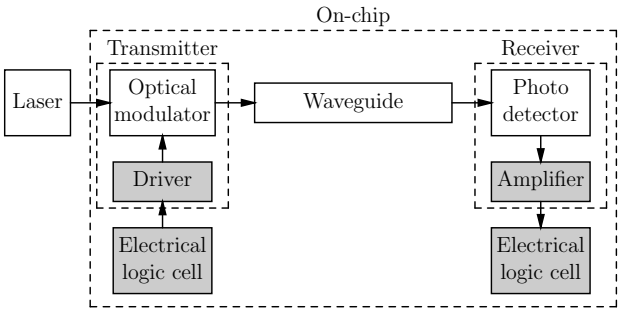

Fig. 2. An on-chip optical interconnect data path.

III. On-Chip Optical Data Path

Introducing optical interconnects into VLSI architectures requires compatibility with CMOS technology. Due to the absence of an efficient silicon-based laser, only those configurations that utilize an external laser as a light source are considered. A diagram of an optical interconnect system is shown in Fig. 2. Considering compatibility with a CMOS technology, a practical solution is a 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. The performance and integration ability of optical devices, however, are expected to be further improved by technology inventions and structural optimization.

A transmitter is composed of an electro-optical modulator and a driver circuit. The design of a fast and cost efficient

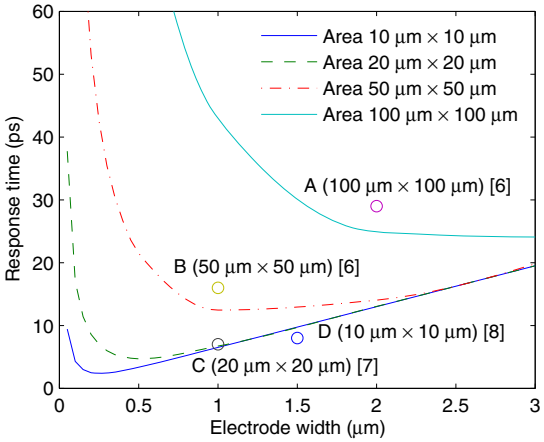

Fig. 3. Detector response time versus electrode width.

CMOS compatible electro-optical modulator is one of the most challenging tasks on the path towards realizing on-chip optical interconnects. In this paper, a predictive modulator model [3] is used that combines the advantages of the structures used in [4] and [5]. To optimize the performance of a modulator, a comprehensive closed-form model [3] is used to determine a proper tradeoff among all physical parameters of a MOS modulator. A series of tapered inverters is used to drive the modulator.

For a specific operating wavelength of $1.5\,\mu m$, low-refractive index strip polymer waveguides are assumed with a core cross section of $1.5\,\mu$m$\times 1.5\,\mu$m. The core index and cladding index are 1.6 and 1.1, respectively. The mode effective index can be determined as 1.48.

The receiver has two components: a photo-detector and an amplifier. 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 trend in the performance of future detectors is projected based on a model proposed by Averine *et al.* [6]. In Fig. 3, the MSM detector response time as a function of electrode width is plotted for different detector sizes and is compared with experimental results [6], [7], [8]. The detector response time is expected in the near future to drop significantly, from tens of picoseconds to a few picoseconds. The amplifier is designed to satisfy bandwidth and noise constraints [9].

IV. Electrical vs. Optical Interconnects

In this section, different criteria used in the design of the two interconnect systems described in sections II and III are compared, including delay uncertainty, latency, power dissipation, and bandwidth density. The interconnect length is 10 mm.

Delay uncertainty is caused by geometric process variations and environmental changes. 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. The parameters are extracted from [2], [10], [11]. The delay and 3σ value for different parts of a 1 cm optical data path are listed in Table I. A comparison of the standard deviation of the delays of the electrical and optical interconnect is shown in Fig. 4. The delay uncertainty of the optical interconnect 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.

TABLE I

Delay and 3σ value of a 1 cm optical data path.

Tech. node	90 nm		65 nm		45 nm		32 nm		22 nm	
	Delay (ps)	3σ (%)	Delay (ps)	3σ (%)	Delay (ps)	3σ (%)	Delay (ps)	3σ (%)	Delay (ps)	3σ (%)
Mod. 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
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

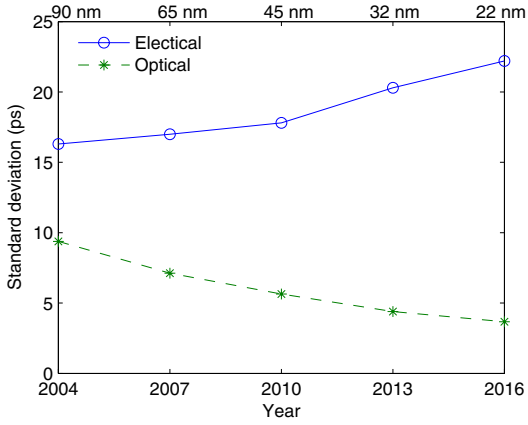

Fig. 4. Comparison of standard deviation of delays of electrical and optical interconnects.

In order for the data to be correctly latched at the receiving register, certain setup and hold constraints should be satisfied. In this paper, the timing budget assigned to T_{setup} and T_{hold} is assumed to be 20% of the clock period, *i.e.*, the delay uncertainty cannot exceed 80% of the clock period. If this requirement is not satisfied, pipeline registers are inserted such that the timing requirements of each stage are satisfied. The actual delay of the interconnect considering delay uncertainty is $T_{total} = m(T_{max} + T_{setup} + T_{C-Q})$, where m is the number of register stages, T_{C-Q} is the clock-to-data delay, and T_{max} is the maximum delay of each stage. $T_{setup}+T_{C-Q}$ is also assumed to be 20% of the clock period. Since no register-like device can be inserted into an optical data path, the delay uncertainty provides an upper bound on the optical channel bandwidth. As listed in Table II, 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 higher performance of the electrical circuits in the modulator driver and receiver amplifier.

TABLE II

DELAY (ps) OF ELECTRICAL AND OPTICAL INTERCONNECTS.

Technology node	90 nm	65 nm	45 nm	32 nm	22 nm
Electrical	311.9	313.2	291.3	312.0	317.8
Optical	238.9	173.3	145.4	127.7	114.9

TABLE III

POWER (mW) OF OPTICAL AND ELECTRICAL INTERCONNECTS.

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

The electrical interconnect power models used in this analysis are the same as those models used in [1]. The power of the registers can be estimated by scaling a typical master-slave D flip-flop. The power due to the registers is negligible as compared to the power of the interconnects. For optical interconnect, only the electrical power is evaluated. The power dissipated by the electrical and optical interconnect is compared in Table III. In optical interconnect, the power consumed by the transmitter dominates the power of the receiver. Both the electrical and optical interconnect power increases due to higher clock frequencies and greater leakage current.

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. From section II, the optimal interconnect width is $7W_{min}$, corresponding to a pitch of $8W_{min}$. Requiring the waveguide size to be larger than the optical mode size, the waveguide pitch is assumed to be 4 μm. Single wavelength optical interconnects are not beneficial if high bandwidth density is desired. The bandwidth of optical interconnects, however, can be significantly improved by introducing wavelength division multiplexing (WDM). The bandwidth density of different interconnects is compared in Fig. 5. 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.

The critical length beyond which optical interconnect overcomes electrical interconnect is plotted in Fig. 6 for different design criteria. The lengths are normalized to the edge of the chip die dimension. As shown in Fig. 6, the critical length is approximately one tenth of the chip edge length at the 22 nm technology node.

V. 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 are presented. These lengths are well below expected chip die size dimensions.

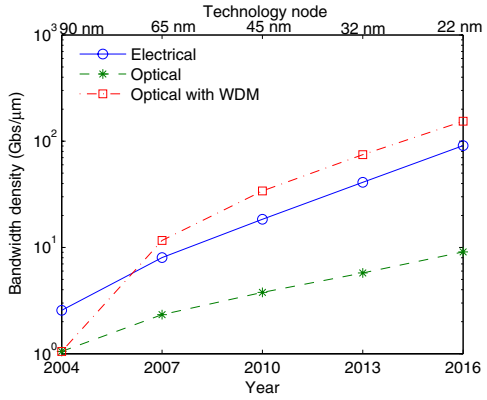

Fig. 5. Bandwidth density of electrical and optical interconnects.

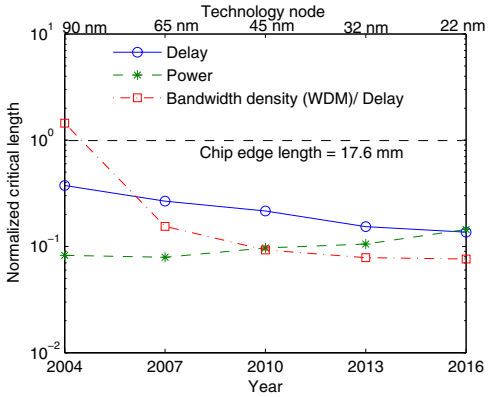

Fig. 6. Normalized critical length beyond which optical interconnect is advantageous over electrical interconnect.

REFERENCES

[1] G. Chen *et al.*, "Predictions of CMOS Compatible On-Chip Optical Interconnect," *Proceedings of the ACM International Workshop on System Level Interconnect Prediction*, pp. 13–20, April 2005.

[2] *The International Technology Roadmap for Semiconductors*, Semiconductor Industry Association, California, 2003.

[3] M. Haurylau *et al.*, "Closed-Form Model of a Capacitor-Based Electro-Optical Modulator," *IEEE Journal of Lightwave Technology*, 2006 (to be published).

[4] Q. Xu *et al.*, "Micrometer-Scale Silicon Electro-Optic Modulator," *Nature*, Vol. 435, pp. 325–327, May 2005.

[5] L. Lia *et al.*, "High Speed Silicon Mach-Zehnder Modulator," *Optics Express*, Vol. 13, No. 8, pp. 3129–3135, April 2005.

[6] 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.

[7] 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.

[8] D. Buca *et al.*, "Metal-Germanium-Metal Ultrafast Infrared Detectors," *Journal of Applied Physics*, Vol. 92, No. 12, pp. 7599–7605, December 2002.

[9] 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.

[10] Y. Cao *et al.*, "Design Sensitivities to Variability: Extrapolation and Assessments in Nanometer VLSI," *Proceedings of the IEEE ASIC/SoC Conference*, pp. 411–415, September 2002.

[11] V. Venkatraman and W. Burleson, "Robust Multi-Level Current-Mode On-Chip Interconnect Signaling in the Presence of Process Variations," *Proceedings of the IEEE International Symposium on Quality Electronic Design*, pp. 522–527, March 2005.