On-chip Optical Interconnect Roadmap: Challenges and Critical Directions

Mikhail Haurylau*, Hui Chen†, Jidong Zhang*, Guoqing Chen*, Nicholas A. Nelson*,
David H. Albonesi‡, Eby G. Friedman* and Philippe M. Fauchet*

* Department of Electrical and Computer Engineering
† Department of Physics and Astronomy
University of Rochester, Rochester NY, 14627, USA
‡Department of Electrical and Computer Engineering
Cornell University, Ithaca, NY 14853, USA

***Abstract*— Intrachip optical interconnects can outperform electrical wires but the required parameters for optical components are yet unknown. Here the ITRS is used as a reference point to derive the requirements that optical components must meet.**

I. Introduction

A significant effort has recently been made to develop building blocks for on-chip optical interconnects (OIs), including light sources, waveguides, WDM (wavelength division multiplexing) components, modulators, detectors, etc. While some predictions have already been made [1], [2], there is still no clear performance specifications that optical components need to exceed in order to effectively replace electrical interconnects (EIs).

We are involved in a system-level interdisciplinary project on the intrachip OIs. In the present work we use the International Technology Roadmap for Semiconductors (ITRS) predictions for the EI performance as reference point for the OI requirements. Analysis of such parameters as delay, bandwidth density and power consumption is used to obtain requirements for the individual optical interconnects components. This methodology also provides prospective for OI weaknesses and missing components.

II. Electrical Interconnects Roadmap

In order to estimate EI performance, an RLC interconnect with optimized repeaters and geometry obtained from the global interconnect predictions in the ITRS is examined for different technology nodes [3]. Three degrees of freedom - the wire width and the number and size of the repeaters (or electrical signal amplifiers) – are optimized for minimum signal propagation delay. The delay model for the interconnect is an extension of work described in [4] by including the effects of repeater output capacitance and input signal transition time.

Two of the main parameters for on-chip interconnects are propagation delay and interconnect bandwidth density. For optimal EI width and repeater size and spacing, the propagation delay remains effectively fixed at around 20-25 ps/mm when technology scales from 90 nm (year 2004) to 22 nm (year 2016), whereas the bandwidth density increases due to the smaller wire pitch. Power consumption per unit length is $\sim$ 1 mW/mm and is expected to slowly increase. In the present work, only delay-optimized EIs are considered for comparison with OIs.

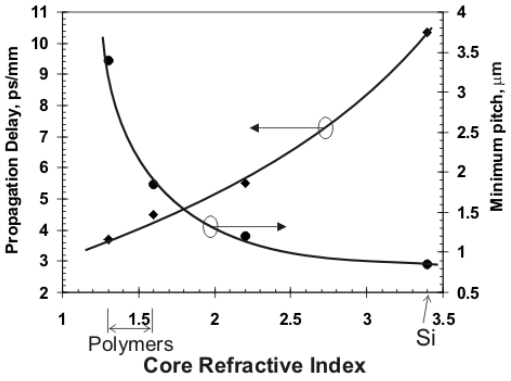

Fig. 1. The tradeoff between the waveguide density and propagation delay per unit length. The graph is plotted for a 10 mm interconnect with a maximum allowed crosstalk of 20%.

III. Optical Interconnects: Configuration and Requirements

A. Monolithic VLSI Technology: Advantages and Limitations

The introduction of OIs into VLSI circuits requires monolithic integration with logic circuits. Microelectronic monolithic fabrication is perhaps one of the most robust and high yield technologies in modern industry, resulting in low cost and ultrahigh levels of device integration. However, the number of materials and processes available for OI fabrication is significantly limited to those technologies compatible with microelectronics.

An important consequence of these limitations is the absence of efficient monolithic on-chip light sources. While a number of exciting scientific achievements have recently been published in the area of optical gain in silicon [5],[6], high speed electrically driven monolithic light sources are far from reality. The most probable device of choice is thus expected to be a silicon-compatible electrooptical modulator, with an external laser light source. In this paper, an OI system is considered that consists of three main parts: an on-chip light modulator for signal switching, a waveguide for the light guidance, and a photodetector as a receiver.

*This research was supported by National Science Foundation Grant CCR-0304574.

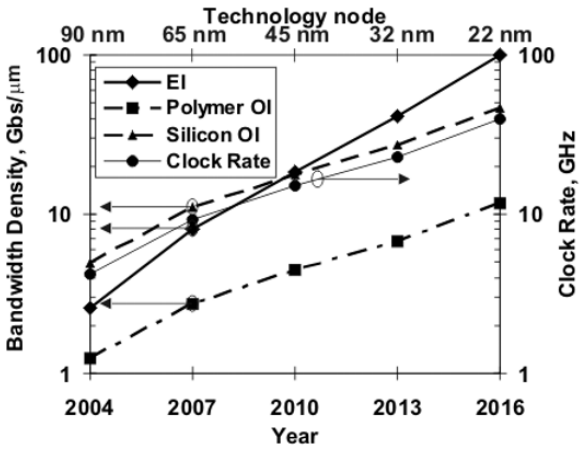

Fig. 2. Comparison of bandwidth density of electrical wires and optical interconnects as a function of year and technology node. Bandwidth density is an important metric, which defines the information throughput of an interconnect through a unit cross section.

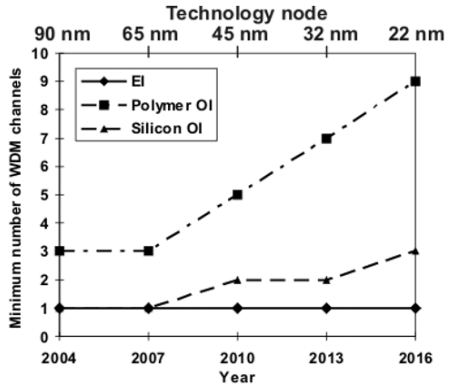

Fig. 3. The number of OI WDM channels required to exceed the EI bandwidth density as a function of year.

B. Optical Waveguides

Optical waveguides can be reliably analyzed as their size and propagation delay depend only on their geometry and refractive index and on the wavelength of light. While there are many material systems available for CMOS-compatible waveguides, a high-index core offers a smaller waveguide pitch, while a low-index core offers a lower propagation delay. This tradeoff is depicted in Fig. 1, where the waveguide delay and minimum pitch are plotted vs. the refractive index of the core. From this graph, the two essential CMOS interconnect requirements can be compared - propagation delay and bandwidth density. Optical waveguides exhibit a smaller propagation delay than electrical wires for any technology node, which clearly demonstrates a significant advantage of optics in delay-limited architectures.

However, due to the large mode diameter, optical waveguides need to be spaced approximately 0.5–3 μm from each other to avoid significant crosstalk. In contrast, a delay-optimized pitch for electrical wires is around 5–7 node sizes, which significantly improves the bandwidth density. A comparison of the bandwidth density for delay-optimized electrical interconnects and optical waveguides is illustrated in Fig. 2. While any increase in optical bandwidth density is solely due to the higher bit rate through the waveguides with a fixed pitch, electrical interconnects can also exploit more efficient repeaters, which allows for smaller wires.

A viable solution to the bandwidth problem is to use the WDM to enhance OI bandwidth density. The number of necessary WDM channels required to match the EI bandwidth density is shown in Fig. 3 for both SOI and polymer waveguides.

C. Transmitter and Receiver: Conversion Cost

In order to utilize the propagation delay advantage, it is necessary to convert the electrical signal into light and back into electrical signals. This requirement has a fixed conversion delay, which is nearly independent of the interconnect length for a given technology. Hence, OIs tend to have a delay advantage in longer connections, when the waveguide propagation delay dominates the overall delay.

To be considered as a candidate for replacing EIs, OIs need to exhibit advantages in both delay and power for critical long distance intrachip interconnections. If the average length of the global interconnects in a target architecture is known, it is possible to extract the conversion cost (*i.e.* delay and power penalty) requirement for OI.

As an example, Fig. 4 shows the requirements for OI conversion for an interconnect length equal to the ITRS projected chip edge length of 17.6 mm for both polymer core and silicon core waveguides. Two types of tradeoffs can be identified. Polymer core waveguides provide a larger conversion delay window but require more aggressive WDM to match the bandwidth density. Silicon core waveguides, however, permit the WDM parameters to be relaxed, but require faster transmitters and receivers.

The conversion penalty consists of two parts - transmitter and receiver. From the aforementioned discussion, to be competitive for intrachip global interconnects, the following requirements should be satisfied.

1. The total delay should be less than 280–370 ps for polymer waveguides and 180–270 ps for silicon waveguides.

2. The total power consumption should be comparable to that of EI ($\sim$ 18 mW) for the chip-length global interconnect.

3. The maximum bandwidth, or bit rate, should exceed the ITRS prediction for the clock rate (Fig. 2, right axis).

4. General CMOS requirements, most significantly technology compatibility and temperature stability, should be satisfied.

There has recently been a number of reports on high speed, low power detectors that already satisfy these global interconnects requirements [7], [8]. In contrast, there are currently no reports of electro-optical modulators suitable for intrachip optical interconnects. Although significant progress in silicon-based modulators has recently been made [9], [10], these modulators do not satisfy all the necessary conditions for replacing the EIs. The main material parameter for the

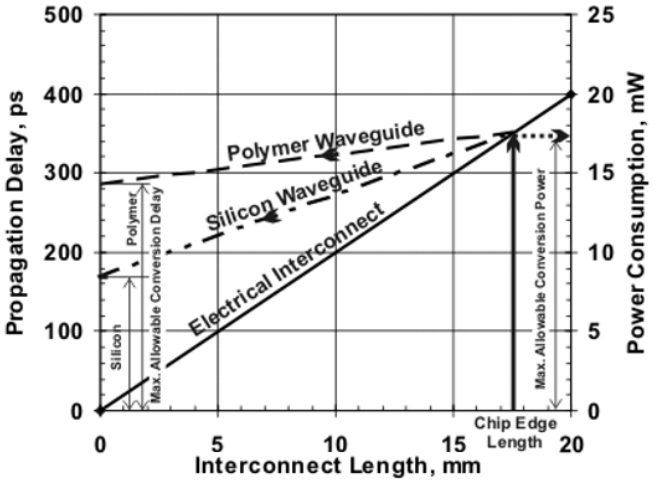

Fig. 4. Maximum conversion cost for an on-chip OI.

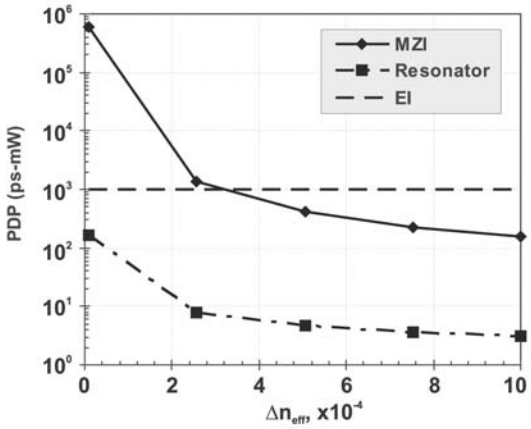

Fig. 5. PDP of transmitter (modulator and receiver circuits) as a function of the effective refractive index change for the 90 nm technology node. The dashed line shows the PDP of a 10 mm EI for comparison.

modulator is an effective refractive index change Δn_{eff}. The higher Δn_{eff}, the more compact the modulator can be, thereby reducing the propagation delay and power consumption. The product of power and delay (PDP) is routinely used in the circuit design process to evaluate system performance. The dependence of the total PDP on Δn_{eff} for both Mach-Zehnder interferometers and microresonators-based transmitter with driver circuits is shown in Fig. 5. The PDP of the delay-optimized EI at the 90 nm technology node is also shown for comparison.

Resonator-based modulators can effectively fold the active device region, thereby significantly reducing the power consumption and the driver delay. Resonant structures should easily exceed EIs in terms of the PDP as shown in Fig. 5. In addition, resonant structures are naturally suited for WDM architectures. However, temperature stability issues remain to be solved before any resonator structure can successfully be used for intrachip applications.

IV. Conclusions and Upcoming Challenges

Using the semiconductor technology roadmap as a starting point, the requirements and critical directions are derived for intrachip optical interconnects. The combined transmitter and receiver delay needs to fall into the 280–370 ps window for polymer waveguides and 180–270 ps window for silicon waveguides. Since the bandwidth density is expected to grow for EIs, an increasing number of WDM channels is needed for OIs to exceed the EI performance, up to nine in the case of low-index waveguides and three for high-index waveguides.

Temperature stability is a separate issue that poses a serious challenge for successful OI operation. While the introduction of OIs may help manage the thermal budget in multi-core processor architectures [11], OIs are susceptible to temperature variations. A method of either active or passive optical control similar to that published in [12] is required to maintain stable device operation.

This analysis has identified the primary challenges for intrachip OIs. First, the size, delay and power consumption of silicon-compatible modulators need to be significantly reduced before any state-of-the-art modulator can be considered for on-chip applications. Second, the introduction of WDM requires the development of ultracompact integrated wavelength-selective components and efficient broadband external lasers. Finally, passive or active temperature drift compensation is needed for reliable OIs operation.

References

[1] M. J. Kobrinsky, B. A. Block, J.-F. Zheng, B. C. Barnett, E. Mohammed, M. Reshotko, F. Robertson, S. List, I. Young, and K. Cadien, "On-Chip Optical Interconnects," *Intel Technology Journal*, Vol. 8, No. 2, pp. 129-141, May 2004.

[2] G. Chen, H. Chen, M. Haurylau, N. Nelson, D. Albonesi, P. M. Fauchet, and E. G. Friedman, " Predictions of CMOS Compatible On-Chip Optical Interconnect," *Proc. of the ACM/IEEE International Workshop on System Level Interconnect Prediction*, pp. 13–20, (2005).

[3] http://public.itrs.net/

[4] 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* **8(2),** 195–206 (2000).

[5] J. Ruan, P. M. Fauchet, L. Dal Negro, M. Cazzanelli, and L. Pavesi, "Stimulated Emission in Nanocrystalline Silicon Superlattices," *Appl. Phys. Lett.* **83(26),** 5479–5481 (2003).

[6] O. Boyraz, and B. Jalali, "Demonstration of a Silicon Raman Laser," *Opt. Express* **12(21),** 5269–5273 (2004).

[7] S. J. Koester, J. D. Schaub, G. Dehlinger, J. O. Chu, Q. C. Ouyang, and A. Grill, "High-efficiency, Ge-on-SOI Lateral PIN Photodiodes with 29 GHz Bandwidth," *Proc. of the Device Research Conference*, Notre Dame, Indiana, June 22, (2004).

[8] M. R. Reshotko, D. L. Kencke, and B. Block, "High-speed CMOS Compatible Photodetectors for Optical Interconnects," *Proc. SPIE* **5564,** 146–155 (2004).

[9] Q. F. Xu, B. Schmidt, S. Pradhan, and M. Lipson, "Micrometre-scale Silicon Electro-Optic Modulator," *Nature* **435,** 325–327 (2005).

[10] A. Liu, R. Jones, L. Liao, D. Samara-Rubio, D. Rubin, O. Cohen, R. Nicolaescu, and M. Paniccia, "A High-Speed Silicon Optical Modulator Based on a Metal-Oxide-Semiconductor Capacitor," *Nature* **427,** 615–618 (2004).

[11] N. Nelson G. Briggs, M. Haurylau, G. Chen, H. Chen, D. H. Albonesi, E. G. Friedman, and P. M. Fauchet, "Alleviating Thermal Constraints While Maintaining Performance via Silicon-Based On-Chip Optical Interconnects," *Proc. of the Workshop on Unique Chips and Systems*, pp. 45–52 (2005)

[12] S. M. Weiss, M. Molinari, and P. M. Fauchet, "Temperature Stability for Silicon-Based Photonic Band-Gap Structures," *Appl. Phys. Lett.* **83(10),** 1980–1982 (2003).