



Electrical and Optical On-Chip Interconnects in Scaled Microprocessors

Kapil Tiwari ^{1*}, Dr. Rajesh Kumar Rai ², Dr. Reeta Pawar ³

¹⁻³ Madhyanchal Professional University (MPU) - Educational Institute Bhopal, Madhya Pradesh, India

* Corresponding Author: **Kapil Tiwari**

Article Info

ISSN (online): 3049-1215

Volume: 02

Issue: 04

July – August 2025

Received: 10-05-2025

Accepted: 08-06-2025

Published: 04-07-2025

Page No: 116-122

Abstract

The rapid scaling of microprocessor technologies has led to unprecedented advancements in processing speed and transistor density. However, as device dimensions shrink below the nanometre scale, traditional electrical on-chip interconnects face critical limitations including increased signal delay, power consumption, and reduced bandwidth. These challenges have emerged as major bottlenecks in the performance and energy efficiency of modern microprocessors, especially in multi-core and many-core architectures.

To overcome these limitations, optical interconnects have gained significant attention as a high-performance alternative for on-chip and chip-to-chip communication. Optical interconnects offer key advantages such as higher data rates, lower latency, reduced power dissipation, and immunity to electromagnetic interference. The integration of optical components like waveguides, modulators, photodetectors, and silicon photonics within the chip architecture has the potential to revolutionize interconnect design by addressing the shortcomings of conventional electrical wiring. This paper explores the evolution of on-chip interconnects, comparing the performance, scalability, and integration challenges of electrical and optical solutions in scaled microprocessors. It also discusses emerging hybrid interconnect architectures that combine the strengths of both technologies. A comparison is made between electrical and optical interconnects for different design criteria, based on projections about the future of optical devices. Around a tenth of the length of the chip's edge is the crucial dimension beyond which optical connectivity becomes preferable to electrical interconnect at the 22 nm technological node.

DOI: <https://doi.org/10.54660/IJFEI.2025.2.4.116-122>

Keywords: Waveguides, Modulators, Photodetectors

Introduction

As microprocessors continue to evolve under Moore's Law, the relentless scaling of transistor dimensions has led to significant improvements in logic density and computational capabilities. However, this miniaturization has introduced critical challenges, particularly in the realm of on-chip communication. One of the most pressing issues is the performance bottleneck caused by traditional electrical interconnects, which struggle to keep pace with the increasing data transmission demands and shrinking geometries of modern microprocessors.

Electrical interconnects, composed primarily of copper or aluminum lines, suffer from increased resistance-capacitance (RC) delay, signal integrity issues, and power dissipation as feature sizes scale below the 10-nm regime. These challenges impede the overall system performance and limit scalability, making it difficult to meet the bandwidth and latency requirements of high-performance computing systems. To address these limitations, optical interconnects have emerged as a promising alternative. Optical interconnects leverage light to transmit data, offering several advantages over electrical counterparts, including higher bandwidth, lower latency, reduced crosstalk, and better energy efficiency over longer distances. Technologies such as silicon photonics, VCSELs (Vertical Cavity Surface Emitting Lasers), and waveguides are being explored for their integration potential in on-chip and chip-to-chip communication networks.

This transition from purely electrical to hybrid or fully optical interconnect systems marks a paradigm shift in microprocessor design. The integration of optical interconnects on silicon chips presents new opportunities but also introduces complex design, fabrication, and compatibility challenges that must be carefully addressed. As research and development continue in this field, the future of on-chip communication may rely heavily on the successful co-integration of electrical and optical technologies to overcome the limitations of traditional scaling.

Electrical Interconnect

Repeater insertion is commonly employed in submicrometer CMOS technologies to mitigate interconnect delay, shorten transition times, and reduce crosstalk noise. A variety of papers have been published in this domain, detailing design methodologies that meet various design criteria. This section analyses an RLC interconnect with repeaters across various technology nodes, referencing the ITRS. The capacitance and resistance per unit length of the interconnect can be derived directly from the geometrical configurations, assuming that the spacing between adjacent interconnects is equal to the minimum interconnect width. The interconnect inductance is contingent upon the distribution of current return paths, which are challenging to estimate prior to the completion of the circuit's physical design. The signal waveform exhibits low sensitivity to errors in the on-chip inductance, and the magnitude of the on-chip inductance varies slowly as a function of the wire geometry. A fixed value of 0.5 pH/μm is assumed for all technology nodes based on these two characteristics. Figure 1 illustrates a distributed RLC interconnect of length l , which is uniformly segmented into k sections by uniform repeaters. The repeaters are h times larger than a minimum sized repeater, exhibiting an output resistance of R_{tr0}/h , an output capacitance of hC_{d0} , and an input capacitance of hC_{g0} . Here, R_{tr0} , C_{d0} , and C_{g0} represent the output resistance, output capacitance, and input capacitance of a minimum sized repeater, respectively. Repeaters are commonly designed using CMOS inverters. In this analysis, it is assumed that the PMOS transistor has a size

that is double that of the NMOS transistor. The interconnect delay model extends the findings presented in, incorporating the effects of repeater output capacitance and input slew. The output capacitance of the repeater is considered to be equivalent to the input gate capacitance.




Fig 1: Repeater insertion in an RLC interconnect.

The sensitivity of the delay model to this assumption is relatively low. By including the repeater output capacitance, the variable ζ used to characterize inductance effects becomes

Where
 $RT = kR_{tr0}/(hRl)$ and
 $CT = hkC_{g0}/(Cl)$.

The delay of a single stage interconnect assuming a step input signal can be obtained by curve fitting,

Where
 $\omega_n = k/Ll(Cl + C_{g0}hk)$

With technology scaling, interconnect resistance increases, therefore, ζ is normally greater than 0.5. In this range, the signal transition time at the far end of an interconnect exhibits a linear dependence on ζ . The impact of input transition time on delay and far end transition time is addressed in a manner analogous to that of an RC interconnect.

The interconnect power models utilised in this analysis are identical to the models outlined in The design of the electrical interconnect examines three degrees of freedom: wire width, as well as the quantity and dimensions of the repeaters. Multiple combinations are analysed to identify the most effective design in relation to a defined criterion.

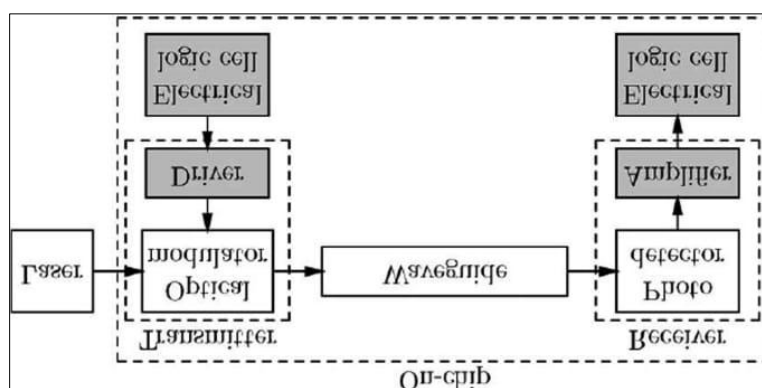


Fig 2: An on-chip optical interconnect data path.

On-Chip Optical Data Link

An on-chip optical data link refers to a communication system embedded within a microprocessor or integrated circuit (IC) that transmits data using optical signals (light) rather than traditional electrical signals. It forms a critical part of optical interconnect systems, enabling high-speed, energy-efficient data transfer between different functional blocks or cores on a chip. The use of optical interconnects into VLSI

systems necessitates compatibility with CMOS technology. This requirement substantially restricts the selection of materials and procedures available for the fabrication of optical components. A major challenge in optical interconnects is the lack of an efficient silicon-based laser that can be monolithically integrated. Only configurations employing an external laser are taken into account. Figure 2 illustrates a diagram of an optical connectivity system. The

system comprises three principal optical components: an off-chip laser, an optical modulator, and an optical detector. This paper assumes low-refractive index polymer waveguides with an effective index of 1.4.

A. Transmitter

The transmitter consists of an electro-optical modulator along with a driver circuit. The development of a silicon-compatible electro-optical modulator that is both fast and cost-efficient presents significant challenges in the pursuit of on-chip optical interconnects. The modulator executes the conversion between electrical and optical signals through a two-step process. Initially, specific optical properties of the medium, such as the refractive index or absorption coefficient, are altered by the electrical signals. Subsequently, the optical signals undergo modulation, which can occur in either amplitude or phase, through the alteration of optical properties. Unstrained bulk crystalline silicon does not demonstrate a linear Pockels effect, and the changes in refractive index resulting from the Kerr effect are minimal. One of the limited mechanisms available for altering the refractive index in pure silicon is the free carrier plasma dispersion effect. Various methods are available for the generation of free carrier plasma. The displacement of carriers occurs at a faster rate compared to injection schemes, as it does not involve any slow carrier recombination processes. Liu *et al.* demonstrated the first MOS capacitor electro-optical modulator utilising the carrier displacement effect, achieving operation at frequencies exceeding one gigahertz. Through design optimisation and advancements in technology, including the thinning of the gate oxide and the implementation of an epitaxial over-growth technique, the modulator's bandwidth is projected to increase to between 30 GHz and 40 GHz, while the delay is anticipated to be reduced to less than 20 ps by the year 2016.

The device structure employed is a Mach-Zehnder interferometer, resulting in a substantial footprint of 10 mm in length. This configuration leads to elevated capacitance, consequently causing increased delay and power consumption in the driver circuits. Simulations and preliminary experiments conducted by Barrios *et al.* indicate that an alternative modulator topology is feasible. An optical micro-cavity can significantly reduce the modulator size to approximately 10 μm to 30 μm , while preserving the same operating principle and speed. Considering these factors, it is anticipated that the capacitance of the modulator structure will decrease to below 10 pF in the coming years.




Fig 3: Circuit model of an optical receiver.

A series of tapered inverters is utilised to drive the modulator. When the output capacitance of the inverter matches the input gate capacitance, the ideal size ratio for two adjacent inverters is determined to be 3.6. The initial stage employs a minimum sized inverter. The calculation for the number of stages is represented by the formula

$$N = \log(CM) / \log(3.6)$$

With CM denoting the modulator.

B. Receiver

The receiver comprises two components: a photodetector that transforms light into electricity, and a receiving circuit that amplifies the analogue electrical signal to a digital voltage level. A simplified equivalent circuit model is depicted in Fig. 3. The simultaneous operation of the optical modulator and detector at the same wavelength in each optical connection creates a conflict regarding the needs of the optical material. Unlike a modulator, which necessitates minimal optical absorption, the functioning of a detector is predicated on the absorption of light. To ensure compatibility with CMOS technology, an effective approach is to utilise light at a wavelength of 1.5 μm in conjunction with a SiGe or Ge photodetector.

This work examines interdigitated SiGe p-i-n and Metal-Semiconductor-Metal (MSM) detectors, attributed to its rapid response and satisfactory quantum efficiency. The signal rising time (response time) of the detector is represented as $T_r = T_{tr} + T_2$, where T_{tr} denotes the duration necessary for photo-generation.

Carriers will drift towards the electrical contact, while TRC represents the RC response time of the detector. The 3 dB bandwidth of a detector is defined as $\Delta f_{dec} = 0.35/T_r$. The delay of the photo-detector, according to a one pole approximation, is expressed as $\tau_{dec} = 0.315T_r$, where τ_{dec} represents the delay and T_r denotes the rise time. In 2002, a Ge p-i-n detector with an interdigitated design was fabricated on a Si substrate, achieving a 3 dB bandwidth of 3.8 GHz at a wavelength of 1.3 μm . Multiple additional studies have been conducted on SiGe detectors, demonstrating comparable performance levels. The bandwidth and delay characteristics of the majority of these detectors are constrained by the carrier transit time, which can be enhanced via device optimisation techniques. The performance trend of future detectors is projected based on a model proposed by Averine *et al.* The anticipated response time is projected to decrease substantially from tens of picoseconds to a few picoseconds in the near future. The decrease can be attributed to the generally bulky nature of current detectors, which necessitates a longer transit time for carriers. Consequently, emphasis has been placed on the development of smaller detectors. Upon achieving efficient coupling between the waveguides and the detectors, a significant reduction in the size of the detectors is anticipated, which will greatly decrease the response time. This trend is anticipated to decelerate and ultimately reach saturation as a result of inherent constraints in material properties.




Fig 4: Minimum delay per unit length as a function of interconnect width.

Comparison between electrical and optical interconnects

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

A. Delay

Minimising latency along an RLC interconnect can be achieved by determining the ideal number and size of repeaters [7]. Increasing the wire width further decreases this minimum delay [20]. Figure 4 shows the smallest delay per unit length that can be achieved for various wire widths. The ITRS predicts a minimal wire width W_{min} , which is used to normalise the connector widths. Fig. 4 shows that scaling has a minimal impact on the latency of repeater interconnects, which is in line with the findings in [21]. When the wire width is more than $3W_{min}$, the rate at which the delay decreases with increasing width slows down. For all technology nodes, the smallest delay per unit length that can be achieved is around 20 ps/mm. See Table I for the distribution of delays for an optical data route of 1 cm. The transmitter's delay is substantially greater than the receiver's. Electrical interconnect is predicted to run slower than optical interconnect by 2007 (the 65 nm technology node).

B. Power

For electrical interconnects, power must be assessed according to certain design criteria, including delay and bandwidth. A minimally sized wire devoid of repeaters expends the least power.

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

Year	2004	2007	2010	2016	2020
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
Reciver amplifier	37.5	16.9	10.4	6.9	4.0
Total optical	28.3	162.0	113.6	90.2	74.7
Electrical	200.0	200.0	200.0	200.0	200.0

Table 2: Power consumption (mW) in an optical data path

Year	2004	2007	2010	2016	2020
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	177.9	18.7	8.8	6.3	5.3

This configuration is impractical for global interconnect because of considerable delay and limited bandwidth. The power-delay product (PDP) serves as a significant design criterion. Each wire size allows for the determination of a local minimum Power Delay Product (PDP) through the adjustment of repeater size and quantity. Simulations indicate that the global minimum power-delay product (PDP) can be attained with a wire size ranging from $4W_{min}$ to $5W_{min}$ across various technology nodes.

The power consumption of the optical interconnect is largely independent of its length, as the length is sufficiently short for optical power loss in the waveguide to be negligible. This paper evaluates only the electrical power for the optical data path, as presented in Table II. The transmitter's power consumption exceeds that of the receiver, contradicting the assumption presented in This difference arises because the modulator considered in this analysis is compatible with CMOS technology. The modulator's size and capacitance are substantial, necessitating a correspondingly large driver circuit. There is a notable reduction in power from the 90 nm technology node to the 65 nm technology node, indicating anticipated advancements in modulator structures from Mach-Zehnder to micro-cavity designs. Figure 5 presents a comparison of the PDP of electrical interconnects with that of optical interconnects for an interconnect length of 1 cm. Observe the intersection between the 65 nm and 45 nm technology nodes.

Bandwidth density

The maximum bit rate for a single interconnect is considered to be equivalent to the clock rate, with one bit transmitted per clock period. Appropriate design enables the attainment of

this bandwidth in both electrical and optical interconnects. The bandwidth density is solely determined by the interconnect pitch. The dimensions of the waveguide must exceed those of the optical mode. Due to this limitation, the waveguide pitch is set at $4\ \mu\text{m}$, significantly exceeding the electrical interconnect pitch, which results in a reduced bandwidth density. This limitation can be mitigated through the implementation of wavelength division multiplexing

(WDM). Figure 6 presents a comparison of the bandwidth density across various interconnects. The channel number in a waveguide is considered to be one at the 90 nm technology node, with an increase of four for each subsequent technology node.

Figure 7 illustrates the critical length at which optical interconnects surpass electrical interconnects, based on various design criteria.

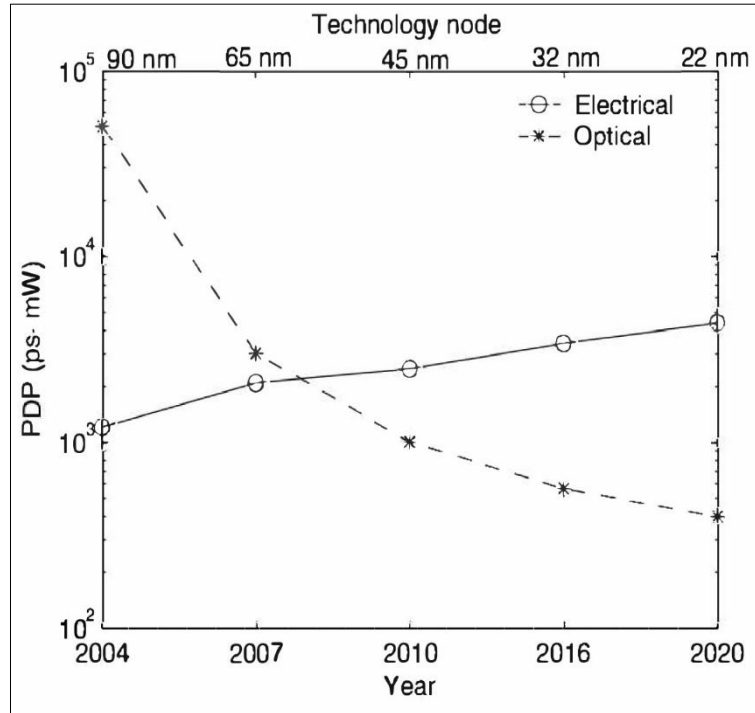


Fig 5: Comparison of the PDP of electrical and optical interconnects with a length of 1 cm.

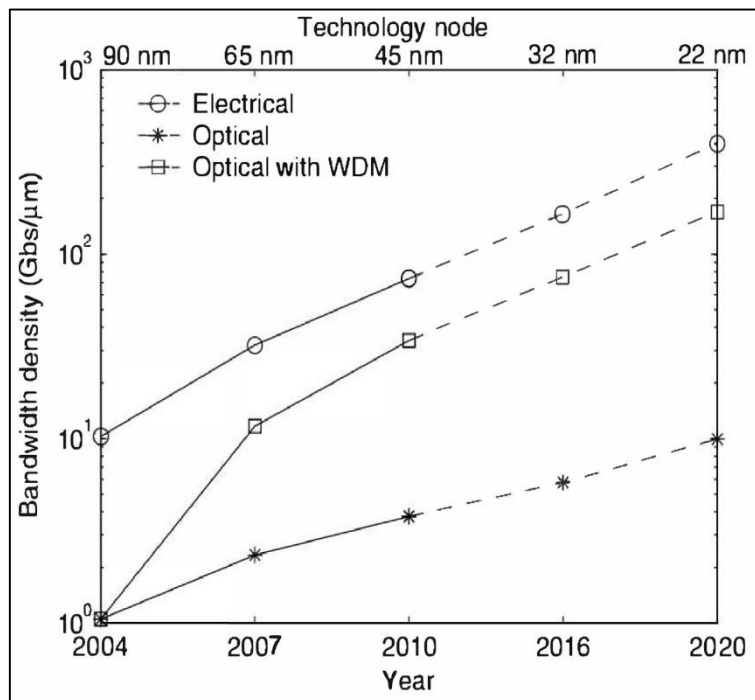


Fig 6: Comparison of bandwidth density of electrical and optical interconnects.

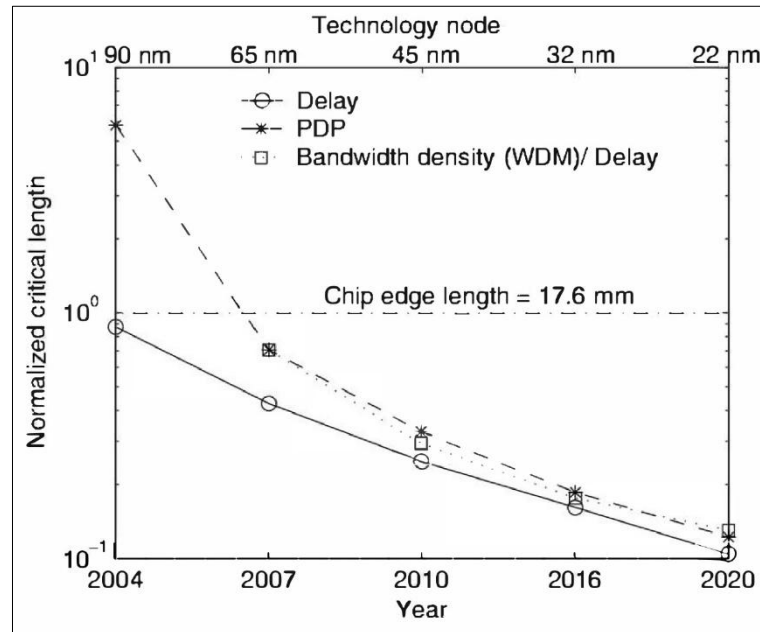


Fig 7: Normalized critical length beyond which optical interconnect is advantageous over electrical interconnect.

This paper employs a fixed optical interconnect design. The optical interconnect can be enhanced regarding a specific criterion through further circuit optimisation. Optical interconnects offer the benefit of reduced crosstalk noise in comparison to electrical interconnects.

Conclusions

For various technological nodes, we offer a thorough comparison of electrical and optical on-chip connectivity based on a performance metric projection for future CMOS compatible optical devices. Delays, PDPs, and bandwidth density/delay all have critical lengths beyond which optical connectivity starts to pay off. Due to technological scaling, their lengths are significantly less than the dimension of the semiconductor die size.

References

- Goodman JW, Leonberger FJ, Kung SY, Athale RA. Optical interconnects for VLSI systems. *Proceedings of the IEEE*. 1984;72(7):850-866.
- Kapur P, Saraswat KC. Comparisons between electrical and optical interconnects for on-chip signaling. *Proceedings of the IEEE International Interconnect Technology Conference*. 2002:89-91.
- Kobrinjsky MJ, Block BA, Zheng JF, Barnett BC, Mohammed E, Reshotko M, *et al.* On-chip optical interconnects. *Intel Technology Journal*. 2004;8(2):129-141.
- Semiconductor Industry Association. *The International Technology Roadmap for Semiconductors*. San Jose, CA: Semiconductor Industry Association; 2003.
- Ismail YI, Friedman EG. Sensitivity of interconnect delay to on-chip inductance. *Proceedings of the IEEE International Symposium on Circuits and Systems*. 2000:403-406.
- Adler V, Friedman EG. Repeater design to reduce delay and power in resistive interconnect. *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*. 1998;45(5):606-616.
- Ismail YI, Friedman EG. Effects of inductance on the propagation delay and repeater insertion in VLSI circuits. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*. 2000;8(2):195-206.
- Chen G, Friedman EG. Low power repeaters driving RC interconnect with delay and bandwidth constraints. *Proceedings of the IEEE International SOC Conference*. 2004:335-339.
- Eldada L, Shacklette LW. Advances in polymer integrated optics. *IEEE Journal of Selected Topics in Quantum Electronics*. 2000;6(1):54-68.
- Barrios CA, Almeida VR, Lipson M. Low-power-consumption short-length and high-modulation-depth silicon electrooptic modulator. *Journal of Lightwave Technology*. 2003;21(4):1089-1098.
- Liu A, Jones R, Liao L, Samara-Rubio D, Rubin D, Cohen O, *et al.* A high-speed silicon optical modulator based on a metal-oxide-semiconductor capacitor. *Nature*. 2004;427:615-618.
- Cherkauer BS, Friedman EG. A unified design methodology for CMOS tapered buffers. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*. 1995;3(1):99-111.
- Hedenstierna N, Jeppson KO. CMOS circuit speed and buffer optimization. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*. 1987;6(2):270-281.
- Sakurai T, Newton AR. Alpha-power law MOSFET model and its applications to CMOS inverter delay and other formulas. *IEEE Journal of Solid-State Circuits*. 1990;25(2):584-594.
- Averine SV, Chan YC, Lam Y. Geometry optimization of interdigitated Schottky-barrier metal-semiconductor-metal photodiode structures. *Solid-State Electronics*. 2001;45(3):441-446.
- Oh J, Csutak SM, Campbell JC. Interdigitated Ge p-i-n photodetectors fabricated on a Si substrate using graded SiGe buffer layers. *IEEE Journal of Quantum Electronics*. 2002;38(9):1238-1241.
- Oh J, Banerjee SK, Campbell JC. Metal-germanium-metal photodetectors on heteroepitaxial Ge-on-Si with amorphous Ge Schottky barrier enhancement layers. *IEEE Photonics Technology Letters*. 2003;15(5):745-

- 747.
18. Pavesi L, Lockwood DJ. Silicon Photonics. New York, NY: Springer; 2004.
 19. Kapur P. Scaling induced performance challenges/limitations of on-chip metal interconnects and comparisons with optical interconnects [Ph.D. dissertation]. Stanford, CA: Stanford University; 2002.
 20. El-Moursy MA, Friedman EG. Optimum wire sizing of RLC interconnect with repeaters. Proceedings of the IEEE Great Lakes Symposium on VLSI. 2003:27-32.
 21. Ho R, Mai KW, Horowitz MA. The future of wires. Proceedings of the IEEE. 2001;89(4):490-504.