

Silicon Nanophotonics for Future Multicore Architectures: Opportunities and Challenges

Yi Xu

Macau University of Science and Technology

Sudeep Pasricha

Colorado State University

Editor's notes:

The high integration density, low power consumption, and support for wavelength division multiplexing have made optical interconnects a viable choice to realize intrachip data exchange in multicore processing platforms. This paper presents a tutorial regarding the possibilities and challenges of photonic network-on-chip as the communication backbone for multicore chips.

—Partha Pratim Pande, Washington State University

nanoscale CMOS fabrication technologies. These wires are becoming slower, more power hungry, and less reliable. We are thus at a critical juncture where the power, bandwidth, and latency of communication must scale favorably to meet the needs of process-

■ **THE NEED FOR** high-performance and energy-efficient communication between processing cores has never been more critical. Rapidly increasing application complexity and limited computing power budgets have led to more and more lightweight cores replacing fewer bulky cores in emerging processor chips. The increase in core counts has put more pressure on the communication fabric which must now support many more streams of higher bandwidth data transfers than ever before. A direct consequence of this trend is that chip power and performance are now dominated not by processor cores but by the need to transport data between processors and to memory. Unfortunately, traditional electrical wires that make up the backbone of communication fabrics on computing chips today are facing unprecedented challenges in ultradeep

ing chips in the near future. A failure to respond to this challenge will create a brick wall that will impede advances in all forms of computing.

Silicon nanophotonics represents one of the more promising solutions to overcome the challenge of worsening communication performance with electrical wires. By transferring data using light signaling between cores and memory, orders of magnitude improvement in bandwidth, latency, and energy are possible. Silicon nanophotonics also promises to pair well with existing board-to-board and chip-to-chip photonics offerings that are rapidly being adopted today. Many semiconductor companies have begun investing heavily into silicon nanophotonics. For example, in 2012, IBM prototyped a silicon-integrated nanophotonics chip operating at 25 Gb/s, using a standard 90-nm silicon-on-insulator (SOI) process. In 2013, Intel and Fujitsu demonstrated silicon CMOS-compatible optical PCI-Express interconnects. Emerging silicon photonic foundries, e.g., OpSIS and ePIXfab, are now allowing low-cost fabrication of nanophotonic components on silicon

Digital Object Identifier 10.1109/MDAT.2014.2332153

Date of publication: 20 June 2014; date of current version:

07 October 2014.

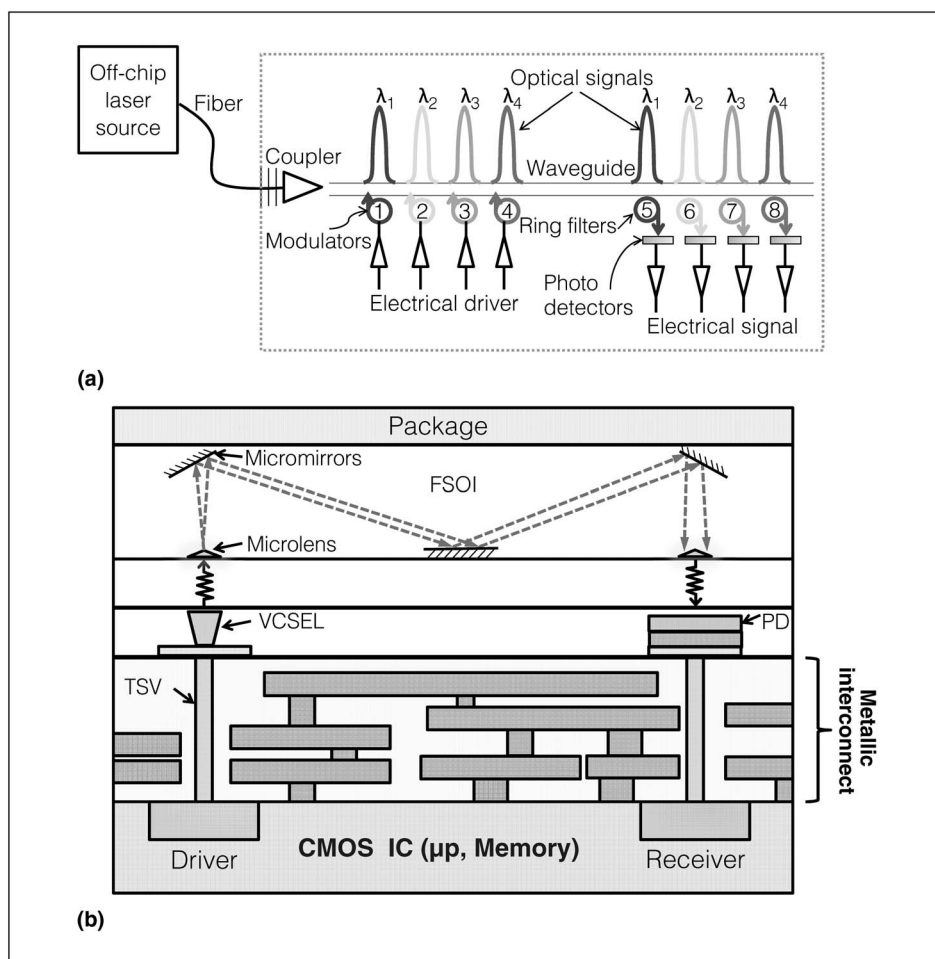


Figure 1. (a) Channel-guided DWDM nanophotonic link. (b) Free-space nanophotonic link.

chips by sharing costs among multiple users with multiproject-wafer shuttle runs. Despite these promising developments, fabrication and integration of silicon nanophotonics with conventional CMOS-based processing and memory elements is still not mature.

In the research community, the challenge of designing silicon nanophotonic communication fabrics is actively being pursued from a variety of different perspectives. New nanophotonic devices (e.g., modulators, photodetectors, waveguides) are emerging as a result of these concerted efforts from industry and academia, as are new on-chip and chip-to-chip network architectures, and tools for rapid design and analysis. Successful solutions will likely adopt and encompass elements from all of these research thrusts. The rest of this article provides an overview of these developments.

Background

Underlying transfer mediums for nanophotonic signals fit into one of two categories: channel guided and free space.

Channel-guided photonic interconnects

A typical on-chip channel-guided photonic network includes an off-chip laser source that provides on-chip light, waveguides that route optical signals, ring modulators that convert electrical signals to optical ones, and ring filters to detect light signals and translate them back into electrical signals. Figure 1a illustrates a dense wavelength division multiplexing (DWDM) photonic link. As light of different wavelengths can be transmitted and modulated in a single waveguide, DWDM technology enables multiple data channels per waveguide, providing high network bandwidth density. At the sender side, electrical signals are converted to laser light by wavelength-selective silicon modulators that absorb and

pass the light for signal “0” and “1,” respectively. The power of the laser should be large enough to sustain all types of light losses such that the detector can receive sufficient optical power. For modulation, microring resonators are preferred over other modulators due to their high modulation speed (10–20 Gb/s), low power (47 fJ/b), and small footprint (μm^2) [2]. The same ring structure can be used as a wavelength-selective filter to extract light out of the waveguide, if the microring is doped with a photodetecting material such as CMOS-compatible germanium. The resonant light will be absorbed by the germanium and converted into an electrical signal.

The waveguide and resonators are critical components in channel-guided nanophotonic architectures. Prototypes of SOI waveguides have been designed with a high-index silicon core ($n = 3.47$)

surrounded by a silicon–oxide bottom cladding ($n = 1.44$) and a low-index top cladding (oxide or air). The high-index contrast between core and cladding confines the light and enables very small bend radius without radiation losses. Recent advances in process technology have brought the propagation losses down to 1–1.3 dB/cm with air cladding [3]. Microring resonators consist of a loop-shaped optical waveguide. The length of the optical path, depending on the effective index and ring size, is exactly a whole number of wavelengths. Therefore, microring resonators have multiple resonances. The spacing between two neighboring resonances is defined as the free spectral range (FSR). A relatively large FSR is preferred for DWDM photonic links to pack in a large number of wavelengths and thereby enhance aggregate link bandwidth.

Free-space photonic interconnects

Channel-guided photonic interconnects rely on microring resonators that have relatively small bandwidth and must achieve very stringent spectral and loss requirements, which translates into extremely fine device geometries and little tolerance for fabrication or thermal variability. An alternative that eliminates the need for microring resonators entirely is to utilize free space instead of waveguides to transfer light [4]. As shown in Figure 1b, in a free-space link, digital data modulate onto a laser, with each modulated light beam emitted by a laser being collimated by a microlens, guided by a series of micromirrors, and then detected by a photodetector (PD), followed by a conversion of the received signals into digital data.

In addition to benefitting from the lack of a need to fabricate and manage microring resonators, a free-space architecture has other benefits over waveguide-based architectures: 1) no waveguide crossing, coupling, or crosstalk losses; 2) the ability to utilize vertical cavity surface emitting lasers (VCSELs) that can be integrated on the die stack and provide light emission without the need for expensive, always-on external laser sources; and 3) higher bandwidth densities than with waveguides for which there exist interwaveguide spacing constraints to avoid crosstalk that limits bandwidth density. Despite these advantages, free-space photonic interconnects introduce their own set of limitations: for example, precise fabrication of micromirrors using micromolding techniques [4] and

efficient alignment of micromirrors requiring precise control is a major challenge. Fabrication of microlenses on top of VCSELs or on the backside of GaAs substrates also requires further research.

Fortunately, a few recent efforts have begun to address the aforementioned challenges. For example, Ouyang et al. [5] propose replacing micromirrors with a single refractive grating layer (RGL) for beam redirection. An RGL is a planar thin-film diffractive optical element with the interesting property that the refractive angle (the extent to which the light direction is changed) depends on the wavelength of the incident light. Thus, controlling the optical path is as simple as selecting the proper wavelength.

Nanophotonic network architecture

Several research efforts have proposed intrachip nanophotonic networks as well as interchip nanophotonic communication solutions. These efforts are discussed below.

Intrachip single-layer nanophotonic networks

In many cases, an optical crossbar is favored as an intrachip network backbone for cache coherence management [6] and data transmission due to its high bandwidth, natural support for broadcast, as well as short and uniform latency that simplify protocol design. We classify previous crossbar designs into two categories: static and dynamic, based on their channel allocation strategy. A channel here is a set of wavelengths used to transfer one data flit. The number of wavelengths per channel depends on the flit size and modulation speed of microrings. Crossbars using static channel allocation include single-write–multiple-read (SWMR) and multiple-write–single-read (MWSR) [7]. The microarchitectural designs of the crossbars are shown in Figure 2a and b, using a radix-4 crossbar as a simple example. I_n/O_n represent the sending/receiving interface of the optical router at node n . The different indices of rings in Figure 2 indicate different optical channels. Each node in SWMR has one dedicated channel to send data and can receive data from all channels. In contrast, MWSR provides each node with a dedicated channel to read data and allows any node to write to the given channel. With exclusive sending channels, SWMR avoids starvation and does not need global arbitration to handle contention, which reduces design complexity and network latency.

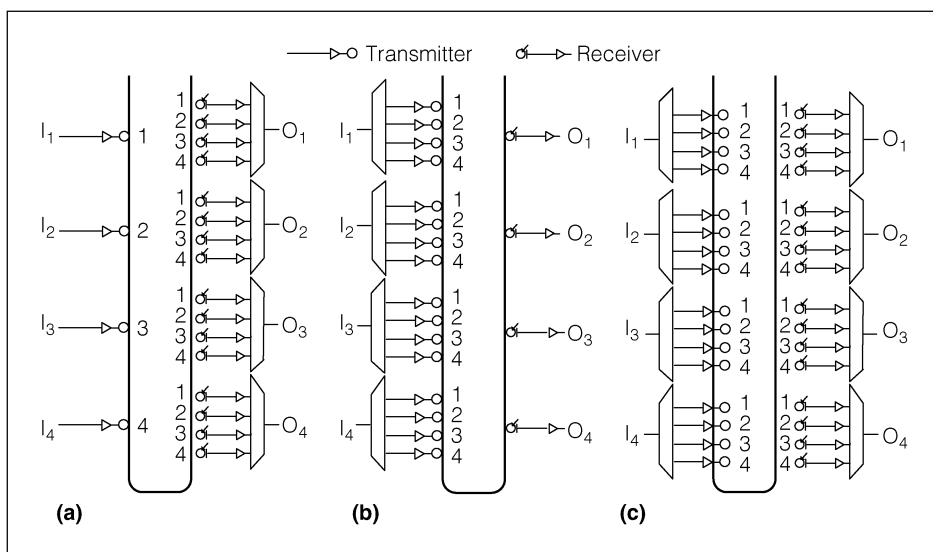


Figure 2. Crossbar microarchitectural designs [1].

When traffic loads on the channels are evenly distributed, SWMR and MWSR can perform well and provide high channel utilization. However, for unbalanced traffic distribution, their dedicated channels will have low utilization and contribute little to the network throughput. Increasing throughput would require overprovisioning of channels and cause notable static power increase. Therefore, the low channel utilization of SWMR and MWSR results in low energy efficiency.

Dynamic channel allocation design, e.g., multiple-write-multiple-read (MWMR) [8] shown in Figure 2c, can improve channel utilization and network throughput with channel sharing. Each network node can write to or read from any channel via more transmitters/receivers and MUXes than in SWMR and MWSR. Thus, under uneven traffic distribution, the nodes with high injection rate can utilize multiple channels to improve channel usage. However, full channel sharing in MWMR requires more microring resonators than SWMR or MWSR. Most of the time, the majority of microring resonators are idle as only M out of $N \times M$ (N is the crossbar radix, M is the number of channels) transmitters are used simultaneously. But idle microrings still consume significant trimming power (see the Design Constraints section) and cause more light losses in the waveguide. Hence, full sharing architectures also have low energy efficiency because of a large number of microrings.

A global crossbar topology can provide contention-free communication, high performance, and low

design complexity. But the large amount of microrings becomes a serious issue for high node counts, requiring high laser and trimming power. Many recent works build switch-based nanophotonic network topologies such as clos, mesh [9]–[11], etc., to reduce the number of microrings. For example, Shacham et al. [11] proposed a circuit-switched nanophotonic network that transmits large messages through the nanophotonic mesh network and delivers small messages using electrical wires to improve overall energy efficiency. One of the challenges for such switch-based optical

networks is high loss due to waveguide crossings within optical switches. The design constraint on the input power per waveguide limits the maximum number of wavelengths transmitted in a waveguide. Therefore, light passing through more switches requires higher power, which results in less wavelengths and lower network bandwidth.

Another aspect to consider when designing photonic network architectures is the transmission quality. Datta et al. [32] proposed multiple-segment bus-based topologies and error-correction codes for high-speed intercluster interconnection [31] and demonstrated better bit-error-rate (BER) performance within the same limit on launched power per wavelength, compared to 2-D torus and MWSR architectures.

Intrachip multilayer nanophotonic networks

Most intrachip nanophotonic networks are envisioned to be implemented with a dedicated photonic layer stacked on top of conventional CMOS metal and logic layers. To support higher bandwidths in the future, single-layer topologies will need to be modified by increasing the number of waveguides and microring resonators in the photonic layer. But this may not be practically feasible due to limitations imposed by physical design concerns, such as waveguide spacing requirements and higher waveguide crossing losses.

For these reasons, recent efforts have focused on multilayer intrachip nanophotonic network

implementations. Such architectures use multiple nanophotonic layers, much like the multiple metal layers in contemporary CMOS chip designs. For example, the OPAL architecture [11] stacks multiple layers of nanophotonic rings with no optical waveguide crossover points, coupled to an electrical mesh topology. OPAL also allows reconfiguration of traffic between the electrical and photonic layers to improve transfer bandwidths, latency, and power dissipation over single-layer topologies. The R-3PO architecture [13] is another architecture that uses multiple layers with no optical waveguide crossover points. R-3PO integrates a crossbar topology, as well as support for dynamically reallocating bandwidth from underutilized to overutilized links. Such multilayer photonic network architectures are likely to be more practical to fabricate, as they simplify intralayer physical design over single-layer alternatives.

Interchip communication solutions

Realizing the full potential of silicon nanophotonics will require a tight integration of intrachip photonic networks with interchip photonic links. Perhaps the most significant impact that photonics can have beyond the chip level is to allow high-bandwidth and low-latency transfers between processing chips and main memory, to alleviate main-memory bottlenecks.

Early work [14], [15] made the case for connecting topologically local but physically remote main-memory racks to CPU racks via ultrahigh-bandwidth parallel optical interconnects, to enable flexible memory architecture options using low-cost commodity memory technology. Various CPU-memory photonic network topologies were explored in [12]. Koka et al. concluded that point-to-point networks provide superior performance and energy over switched photonic networks, due to the prohibitive loss characteristics of additional photonic devices required for switched photonic networks.

Continuing the trend of integrating silicon nanophotonics into commodity components such as processors and off-chip interconnects, a few efforts have proposed a tighter use of photonics in memory modules. Corona [16] briefly considered the system level implications of optically connected 3-D memory stacks. In [17], the PIDRAM chip was proposed with aggressive photonic penetration into the DRAM chip. The architecture modified DRAM banks and

enhanced their interconnections with photonic links to reduce intrachip energy. In [18], a less invasive 3-D stacked memory architecture was presented that contains an interface die to handle the conversion between optics and electronics, with the remaining layers containing commodity DRAM bank/chip arrangements.

Nanophotonic network protocols

Designing hardware cache coherence protocols for future multicore computing is challenging. Traditional topologies, e.g., mesh, face well-acknowledged challenges in performance scalability. For example, core-to-core communications require up to 64 hops for a 1024-core mesh network, making it prohibitively expensive to support indirection in directory-based coherence protocols.

Recent efforts have embraced emerging nanophotonic technologies for supporting cache coherence protocols as nanophotonic networks provide high bandwidth density, low propagation delay, and low power consumption for remote communications. For snoopy-based coherence traffic, Corona [16] and Optical bus [3] leveraged a single broadcast bus with an arbitration scheme, while Spectrum [19] employed antennas to form multiple broadcast channels. To maintain the ordering in unordered network, atomic coherence [6] serializes the transactions with mutexes (optical tokens) to reduce additional racing transitions and ensure correctness in a directory protocol, where fast propagation speed of the optical network mitigates the performance penalty of serialization. Only the node which holds the mutex of the required address is allowed to perform the coherence activity, while the competing coherence requests need to retry after the mutex is released. Token-based arbitration schemes have also been proposed to solve the contentions occurring in optical channels, e.g., MWMR [7]. The implementation of the token can be one-bit-wide light pulse, the presence of which represents the availability of the optical channel.

Challenges

While nanophotonic interconnects provide many advantages, there are also fundamental challenges in integration, fabrication, testing, and architecture design to provide high-performance, energy-efficient, low-cost, and reliable on-chip communication.

Design constraints

High static power consumption. In conventional electrical networks, dynamic power, which depends on the activities of routers and channels, typically dominates the total network power. But for nanophotonic networks, static power surpasses dynamic power and dominates total network power. The static power of a nanophotonic network is mainly composed of laser power and microring trimming power. The laser power is determined by the total number of wavelengths, the conversion efficiency from electrons to photons of the laser, and all types of transmission losses, including both on-resonance microrings and scattering losses from off-resonance microrings. Hence, laser power increases with the total number of microrings. The resonance wavelength of a microring also drifts with temperature variation. Such drift can be corrected, or trimmed, via either heating or carrier injection. Both methods consume power. Recent studies have shown that laser and trimming power together contribute over 60% of the total on-chip network power. For example, the Corona nanophotonic network in 17-nm technology from HP is estimated to consume ~ 26 W in trimming microrings, out of ~ 48 W of the total network power [20]. Hence, it is unwise to increase the throughput of a nanophotonic network by increasing the number of channels (and microrings) because idle channels still consume significant static power. Instead, nanophotonic networks that achieve high throughput via improvements in channel utilization need to be designed to improve energy efficiency.

Peak bandwidth. The peak bandwidth of a nanophotonic network is determined by the number of waveguides, the number of wavelengths supported in one waveguide, and modulation rate. The number of waveguide fibers is limited by the perimeter of the chip/macrochip, the size of fiber pitch, and spacing between fibers [12]. The number of wavelengths transferred in a single waveguide depends on the bandwidth of spectrum, channel crosstalk [21], the limited optical power that single waveguides can carry, and light losses caused by the devices [12]. As mentioned earlier, the available spectral bandwidth is around one FSR, which is inversely proportional to the microring radius. Although rings with small size can yield larger FSR, they have larger bending losses and thus suffer more

channel crosstalk. Hence, arbitrarily large FSR or spectrum bandwidth is infeasible. The launched power per waveguide should be no less than the sum of aggregated light losses along the propagation path and the power to switch the photodetector at the desired rate. Meanwhile, the maximum launched power is limited by the nonlinear effects, which induce free carrier absorption (FCA) and cause the light loss to increase with the input power [12]. Thus, the upper bound on launched power results in the limited number of wavelengths supported in each waveguide. Unless these technology constraints are overcome, the peak bandwidth of a nanophotonic network will remain limited.

Architecture designs

At the architecture level, there are several directions that will need to be pursued. New channel-guided or free-space intrachip nanophotonic network topologies are needed that can minimize losses and reduce laser power requirements; are able to coexist with process, voltage, and thermal (PVT) variations; and are resilient to the effects of soft errors and device aging phenomena. Interchip photonic network topologies must cope with the increasing memory bandwidth requirements from emerging parallel and multiapplication workloads, while also minimizing latency and access energy. New protocols for fast arbitration, effective data flow control, memory coherence and consistency, and quality-of-service management must complement these new topologies within photonic networks. A better unification of the off-chip and on-chip photonic data transfer paths will also be essential in eliminating bottlenecks at the I/O interface. Finally, the penetration of photonics into accelerators (e.g., GPUs [22]) and integration with new memory hierarchies based on nonvolatile memories (e.g., PCRAM [23]) will also require rearchitecting nanophotonic networks to meet the data transfer needs for these unique components.

Intrachip variations

Intrachip thermal and process variations represent two of the most fundamental challenges in integrating nanophotonic devices in emerging multicore chip architectures.

Thermal variations. These variations change the refractive index of nanophotonic components, e.g.,

photonic microring resonators, such that those components fail to resonate designated wavelengths in the waveguide. Studies have reported that a microring's resonance wavelength typically drifts by ~ 0.102 nm/°C [24], while intrachip temperatures can fluctuate well beyond 30 °C. New device-level innovations and architecture-level mechanisms are needed to efficiently regulate temperature variations and ensure correct operation.

Process variations. These variations refer to variations in critical physical dimensions, e.g., thickness of silicon, width of waveguide, caused by lithography imperfection and etch nonuniformity of devices. Process variations (PVs) thus directly affect the resonant wavelengths of a microring [24], which is a critical nanophotonic component used as a modulator, a filter, or a switching element. For example, as much as ~ 4.79 nm of PV drift within a wafer has been observed in a demonstration of a photonic platform leveraging a state-of-the-art CMOS foundry infrastructure [25].

PV drifts and thermal variations can cause a microring to resonate at a completely different wavelength that is several channels away [26]. As a result, drifted microrings cannot be used for communication since they will create erroneous signals. A few solutions to this challenge have been recently proposed. A dynamic regulation method was proposed in which adjustments to chip temperature are used to compensate chip-wise PV drifts [27]. For example, if the PV drift of microrings in a chip region is toward blue, then the regulator would heat up, i.e., red shift, the region via dynamic voltage/frequency scaling. Nitta et al. [28] proposed using error detection/correction codes to tackle variation-induced microring faults. More complicated approaches such as Mintrim [26] that employ redundancy have also been shown to be effective in tolerating variations. New runtime solutions that are computationally inexpensive and can compensate for PVT variations will be essential in realizing predictable nanophotonic network designs.

CAD challenges

The successful adoption of silicon nanophotonics will also rely to a large extent on a robust ecosystem of CAD tools. Such tools will be indispensable to design, optimize, and test multibillion transistor 3-D chips with disparate technologies such as silicon

nanophotonics. At the physical level, floorplanners will require optical signals to be identified early to accommodate optical components that are larger than their electrical counterparts. Chip routing and layout tools will likely require new routing rules to determine allowable signal crossings, signal turns, crosstalk estimation, and wavelength-specific characterizations. Power analysis tools will need to integrate models for trimming control as well as photonic transmitters and receivers that will consume significant static power, unlike their electrical equivalents.

Such CAD tools are slowly emerging as the possibility of realizing silicon nanophotonic edges closer to reality. For instance, Hendry et al. [29] propose VANDAL, a CAD tool for placing functional photonic devices, modifying their parameters, and routing waveguides between devices to create complex network structures. CAD frameworks at the system level that allow exploration and analysis of nanophotonic network architectures will also become increasingly essential for multiobjective optimization. As an example, a recently proposed system-level CAD framework [30] uses CAD optimization algorithms to balance power, bandwidth, and latency goals while customizing various facets of a hybrid nanophotonic ring/electrical mesh network architecture, such as DWDM degree, waveguide count, photonic uplinks/downlinks, and modulation frequency.

Device-level challenges

Innovation at the device level will also be essential to realizing practical silicon nanophotonic architectures. As an example, consider microring resonator devices. At present, there are two types of low-level techniques that can restore the resonance frequency of microrings. The first type is post-fabrication physical trimming, where high-energy particles such as ultraviolet light or an electron beam is used to adjust the refractive index of microrings to achieve resonance correction. Given that the number of intrachip microrings is on the order of thousands to millions, it is unclear if such physical trimming is practical for volume production. In addition, physical trimming may create degradation of the quality factor “ Q ” of a microring, and faster aging in trimmed devices. The second type and most widely used technique for restoring the resonance frequency is power trimming, in which heating or current injection into a microring is used to correct

its resonance wavelength. The former causes the wavelength to shift toward the red end and the latter toward the blue end of the resonance spectrum. Although trimming could address drifts introduced by both PV and thermal variations, it can result in significant power consumption. In addition, current injection has very limited correction range, as it would generate thermal runaway beyond the trimming range [24].

New device-level components and compensation techniques that overcome the drawback of existing solutions will be vital in realizing error-free nanophotonic architectures. Such developments are already underway, e.g., recently designed microring resonators clad with polymethyl methacrylate possess very high thermal stability. Such athermal resonators do not require extra tuning power but give rise to new challenges, e.g., correcting PV in athermal resonators is problematic. Much research is thus still needed, not just into microring resonator-based modulators, couplers, filters, and detectors, but also cost-effective off-chip and on-chip laser sources and photonic transfer media.

SILICON NANOPHOTONICS IS widely considered as a promising alternative to electrical wires that are scaling poorly in deep submicrometer technologies. Much recent work has been focused on the areas of silicon nanophotonic device prototyping and new network architectures and protocols. But there exist several challenges related to stringent design constraints, need for more efficient network architectures, need for coping with PVT variations, lack of CAD tools, and need for more resilient nanophotonic devices. These challenges represent an obstacle toward realizing reliable and robust silicon nanophotonic communication and mass-volume production. However, the rapid pace of advances in each of these problem areas, as can be observed from recent works attempting to overcome many of these limitations, bodes well for the adoption of silicon nanophotonics into emerging multicore architectures, in the not-too-distant future. ■

■ References

- [1] C. Batterm, "Designing nanophotonic interconnection networks," in *Workshop Interaction Between Nanophoton. Devices Syst.*, 2010.
- [2] Q. Xu, S. Manipatruni, B. Schmidt, J. Shaky, and M. Lipson, "12.5 gbit/s carrier-injection-based silicon micro-ring silicon modulators," *Opt. Exp.*, pp. 430–436, 2007.
- [3] N. Kirman et al., "Leveraging optical technology in future bus-based chip multiprocessors," in *Proc. 39th Annu. IEEE/ACM Int. Symp. Microarchitect.*, 2006, pp. 492–503.
- [4] J. Xue et al., "An intra-chip free-space optical interconnect," in *Proc. 37th Annu. Int. Symp. Comput. Architect.*, 2010, pp. 94–105.
- [5] J. Ouyang, C. Yang, D. Niu, Y. Xie, and Z. Liu, "F2BFLY: An on-chip free-space optical network with wavelength-switching," in *Proc. Int. Conf. Supercomput.*, 2011, pp. 348–358.
- [6] D. Vantrease, M. H. Lipasti, and N. L. Binkert, "Atomic coherence: Leveraging nanophotonics to build race-free cache coherence protocols," in *Proc. IEEE 17th Int. Symp. High Performance Comput. Architect.*, 2011, pp. 132–143.
- [7] Y. Pan et al., "Firefly: Illuminating future network-on-chip with nanophotonics," in *Proc. 36th Annu. Int. Symp. Comput. Architect.*, 2009, pp. 429–440.
- [8] Y. Pan, J. Kim, and G. Memik, "FlexiShare: Channel sharing for an energy-efficient nanophotonic crossbar," in *Proc. IEEE 16th Int. Symp. High Performance Comput. Architect.*, 2010, DOI: 10.1109/HPCA.2010.5416626.
- [9] M. J. Cianchetti, J. C. Kerekes, and D. H. Albonese, "Phastlane: A rapid transit optical routing network," in *Proc. 36th Annu. Int. Symp. Comput. Architect.*, 2009, pp. 441–450.
- [10] S. Pasricha and S. Bahirat, "OPAL: A multi-layer hybrid photonic NoC for 3D ICs," in *Proc. 16th Asia South Pacific Design Autom. Conf.*, 2011, pp. 345–350.
- [11] A. Shacham, K. Bergman, and L. Carloni, "Photonic networks-on-chip for future generations of chip multiprocessors," *IEEE Trans. Comput.*, vol. 57, no. 9, pp. 1246–1260, Sep. 2008.
- [12] P. Koka et al., "A micro-architectural analysis of switched photonic multi-chip interconnects," in *Proc. 39th Annu. Int. Symp. Comput. Architect.*, 2012, pp. 153–164.
- [13] R. Morris, A. K. Kodi, and A. Louri, "Dynamic reconfiguration of 3D photonic on-chip interconnects for maximizing performance and improving fault tolerance," in *Proc. 45th Annu. IEEE/ACM Int. Symp. Microarchitect.*, 2012, pp. 282–293.
- [14] C. Batterm et al., "Building manycore processor-to-DRAM networks with monolithic silicon photonics," in *Proc. 16th IEEE Symp. High Performance Interconnects*, 2008, pp. 21–30.

- [15] Y. Katayama and A. Okazaki, "Optical interconnect opportunities for future server memory systems," in *Proc. IEEE 13th Int. Symp. High Performance Comput. Architect.*, 2007, pp. 46–50.
- [16] D. Vantrease et al., "Corona: System implications of emerging nanophotonic technology," in *Proc. 35th Annu. Int. Symp. Comput. Architect.*, 2008, pp. 153–164.
- [17] S. Beamer et al., "Re-architecting DRAM memory systems with monolithically integrated silicon photonics," in *Proc. 37th Annu. Int. Symp. Comput. Architect.*, 2010, pp. 129–140.
- [18] A. N. Udipi, N. Muralimanohar, R. Balasubramonian, A. Davis, and N. P. Jouppi, "Combining memory and a controller with photonics through 3D-stacking to enable scalable and energy-efficient systems," in *Proc. Annu. Int. Symp. Comput. Architect.*, 2011, pp. 425–436.
- [19] Z. Li et al., "Spectrum: A hybrid nanophotonic-electric on-chip network," in *Proc. 46th ACM/IEEE Design Autom. Conf.*, 2009, pp. 575–580.
- [20] J. Ahn et al., "Devices and architectures for photonic chip-scale integration," *Appl. Phys. Rev.*, vol. 95, no. 4, pp. 989–997, 2009.
- [21] N. Sherwood-Droz, K. Preston, J. S. Levy, and M. Lipson, "Device guidelines for WDM interconnects using silicon microring resonators," in *Workshop Interaction Between Nanophoton. Devices Syst.*, 2010.
- [22] N. Goswami, Z. Li, A. Verma, R. Shankar, and T. Li, "Integrating nanophotonics in GPU microarchitecture," in *Proc. 21st Int. Conf. Parallel Architect. Compilation Tech.*, 2012, pp. 425–426.
- [23] Z. Li, R. Zhou, and T. Li, "Exploring high-performance and energy proportional interface for phase change memory systems," in *Proc. IEEE 19th Int. Symp. High Performance Comput. Architect.*, 2013, pp. 210–221.
- [24] C. Nitta, M. Farrens, and V. Akella, "Addressing system-level trimming issues in on-chip nanophotonic networks," in *Proc. IEEE 17th Int. Symp. High Performance Comput. Architect.*, 2011, pp. 122–131.
- [25] J. S. Orcutt et al., "Nanophotonic integration in state-of-the-art CMOS foundries," *Opt. Exp.*, vol. 19, no. 3, pp. 2335–2346, 2011.
- [26] Y. Xu, J. Yang, and R. Melhem, "Tolerating process variations in nanophotonic on-chip networks," in *Proc. 39th Annu. Int. Symp. Comput. Architect.*, 2012, pp. 142–152.
- [27] Z. Li et al., "Reliability modeling and management of nanophotonic on-chip networks," *IEEE Trans. Very Large Scale Integr. Syst.*, vol. 20, no. 1, pp. 98–111, Jan. 2012.
- [28] C. Nitta, M. K. Farrens, and V. Akella, "Resilient microring resonator based photonic networks," in *Proc. 44th Annu. IEEE/ACM Int. Symp. Microarchitect.*, 2011, pp. 95–104.
- [29] G. Hendry, J. Chan, L. Carloni, and K. Bergman, "VANDAL: A tool for the design specification of nanophotonic networks," in *Proc. Design Autom. Test Eur. Conf. Exhibit.*, 2011, DOI: 10.1109/DATE.2011.5763133.
- [30] S. Bahirat and S. Pasricha, "A particle swarm optimization approach for synthesizing application-specific hybrid photonic networks-on-chip," in *Proc. 13th Int. Symp. Quality Electron. Design* 2012, pp. 78–83.
- [31] A. Biberman et al., "First experimental bit-error rate validation of 12.5-Gb/s silicon modulator enabling photonic networks-on-chip," in *Proc. IEEE/OSA Opt. Fiber Conf.*, 2010.
- [32] I. Datta, D. Datta, and P. P. Pande, "Design methodology for optical interconnect topologies in NoCs with BER and transmit power constraints," *IEEE/OSA J. Lightw. Technol.*, vol. 32, no. 1, pp. 163–175, Jan. 1, 2014.

Yi Xu is an Assistant Professor in the Information Technology Department, Macau University of Science and Technology, Macau, China. Her research interests include the network-on-chip designs for 2-D/3-D chips, thermal modeling, and nanophotonic network architecture. Xu has a PhD in electrical and computer engineering from the University of Pittsburgh, Pittsburgh, PA, USA (2012). She is a Member of the IEEE and the Association for Computing Machinery (ACM).

Sudeep Pasricha is an Associate Professor in the Electrical and Computer Engineering Department, Colorado State University, Fort Collins, CO, USA. His research interests include embedded, mobile, and high-performance computing. Pasricha has a PhD in computer science from the University of California Irvine, Irvine, CA, USA (2008). He is a Senior Member of the IEEE and the Association for Computing Machinery (ACM).

■ Direct questions and comments about this article to Yi Xu, Macau University of Science and Technology, Macau 999078, China; yi.xu.2000@gmail.com.