Silicon Photonic Microring Links for High-Bandwidth-Density, Low-Power Chip I/O

Silicon photonic microrings have drawn interest in recent years as potential building blocks for high-bandwidth off-chip communication links. The authors analyze a terabit-per-second scale unamplified microring link based on current best-of-class devices. The analysis provides quantitative measures for the achievable energy efficiency and bandwidth density that could be realized within several years. The results highlight key device attributes that require significant advancement to realize sub-pJ/bit scale optical links.

Noam Ophir
Columbia University

Christopher Mineo
David Mountain
Laboratory for Physical Sciences

Keren Bergman
Columbia University

Performance scalability of computing systems built on chip multiprocessor (CMP) multicore architectures is becoming increasingly constrained by limitations in power dissipation, chip packaging, and the data throughput achievable by the interconnection networks. In particular, chip- and package-level I/O bandwidth (to other chips and memory) must scale with compute capabilities to avoid becoming a performance bottleneck while improving energy efficiency (pJ/bit performance) to operate within a microprocessor's power envelope.[1]

Researchers have touted silicon microring modulators as ultra-low-power wavelength-division-multiplexed (WDM)-compatible building blocks for achieving high-aggregate bandwidth I/O links.[2] However, researchers have only partially addressed full-link optical power budget approaches and the resulting aggregate bandwidth limitations.[2,3] (For more information on other approaches, see the "Potential Chip and Package I/O Technologies" sidebar.)

In this article, we analyze the optical power budget of a single unamplified microring-based WDM link assuming current best-of-class reported devices. We provide a power efficiency estimate for end-to-end operation of the link inclusive of electrical chip I/O and optical link operation from the electronic modulator driver to the sense- or limiting-amplifier output (excluding data serialization and deserialization costs). We examine the bound on the achievable single-link (single data-conveying fiber) aggregate bandwidth and highlight several key photonic components that require further research.

The silicon microring: A compact resonator

In the last few decades, advancements in silicon fabrication processes have provided the capability to create silicon structures in

Published by the IEEE Computer Society

0272-1732/13/$31.00 © 2013 IEEE

Potential Chip and Package I/O Technologies

Although the potential I/O density coming out of current chips is largely dictated by the ball-grid-array (BGA) pitch at the chip-to-package interface (about 150 μm), I/O density can be severely constrained by the limitations associated with traversing through the package core, socket, and board-level transmission lines.[1] Two distinct, relatively near-term approaches have recently emerged that bypass the traditional interconnect architecture.

- Multichip module (MCM) packaging designs place multiple chips or memory modules within a single large package to leverage a high-density in-package interconnect. Beyond bonding multiple chips to a shared substrate, recent efforts have highlighted the use of a shared silicon carrier to which the chips are bonded.[2] Using a silicon carrier with a back-end-of-line copper interconnect allows finer BGA pitch (50 μm) and dense wiring between the mounted chips. Such solutions (as well as other, more advanced 3D integration schemes) improve the local I/O bandwidth and power efficiency between the copackaged chips but don't directly address board-level and rack-level interconnect bandwidth.
- Top-of-the-package connector-based designs use high-quality RF transmission lines for out-of-package interconnects instead of traditional FR4 wiring.[3] These approaches enable cable lengths from a few centimeters to greater than 50 cm with high data rates (36 Gbps), but they typically require connector pitches on the order of several hundred micrometers to avoid RF crosstalk.

Along with these approaches, much emphasis has been put into improving the transmit side (Tx) and demux and receive side (Rx) power efficiency and lowering circuitry footprint by using optimized transceiver designs.[4] Such mechanisms are crucial for overcoming electronic channel distortion, in particular at multigigabits per second signaling speeds.[3]

In conjunction with these efforts, researchers have suggested optical communications as a complimentary I/O technology.[5] Vertical-cavity surface-emitting laser (VCSEL)-based active cables have been deployed in large computing systems as modules on the board for board-to-board and rack-to-rack communications. Introducing VCSEL technology into the package level with dense integration schemes must address energy, footprint, and packaging challenges because each fiber can traditionally carry only a single data channel.[5]

Planar-waveguide technology in silicon[6] or III-V's[7] is compelling for optical I/O because of the potential to implement WDM solutions, allowing multiple data channels to be transmitted over a single fiber. Although they increase complexity aspects in terms of fiber-packaging accuracy, thermal stabilization, and on-chip link design, reducing the number of required connectors makes WDM solutions attractive. CMOS compatibility and cost favor silicon photonics over III-V's for close integration with electronics in chip-level and package-level optical I/O.[6]

References

1. F. O'Mahony et al., "The Future of Electrical I/O for Microprocessors," *Proc. Int'l Symp. VLSI Design, Automation and Test,* IEEE CS, 2009, pp. 31-34.
2. J.U. Knickerbocker et al., "3-D Silicon Integration and Silicon Packaging Technology Using Silicon Through-Vias," *IEEE J. Solid State Circuits,* Aug. 2006, pp. 1718-1725.
3. J.E. Jaussi et al., "Multi-Gbit I/O and Interconnect Co-design for Power Efficient Links," *Proc. IEEE 19th Conf. Electrical Performance of Electronic Packaging and Systems,* IEEE CS, 2010, pp. 1-4.
4. B. Kim et al., "A 10-Gb/s Compact Low-Power Serial I/O With DFE-IIR Equalization in 65-nm CMOS," *IEEE J. Solid State Circuits,* Dec. 2009, pp. 3526-3538.
5. M.A. Taubenblatt, "Optical Interconnects for High-Performance Computing," *J. Lightwave Technology,* Feb. 2012, pp. 448-457.
6. C. Gunn, "CMOS Photonics for High-Speed Interconnects," *IEEE Micro,* Mar. 2006, pp. 58-66.
7. R. Nagarajan et al., "InP Photonic Integrated Circuits," *IEEE J. Selected Topics Quantum Electronics,* Sept. 2010, pp. 1113-1125.

standardized CMOS fabrication processes with nanometer-scale accuracy and smooth features. This nano-level engineering of silicon with accuracy much finer than the typical optical communications wavelength (about 1.5 μm) has enabled the emergence of silicon photonics technology—the creation of low-loss optical devices in silicon.[4] With typical silicon waveguide cross-section areas of approximately 0.1 μm^2 and bending radii as small as 1 μm, silicon photonic technology offers significantly smaller-footprint waveguide-based devices compared to existing telecom components.[4]

Significant attention has recently been given to the development of microresonators in the silicon platform. Microresonators form smaller-footprint, higher-power-efficiency devices compared to nonresonant interferometric designs.[4] However, the same resonant property renders microresonators sensitive to temperature and fabrication variation, requiring tuning and stabilization mechanisms.[2] Resonator-based devices are naturally adaptable to WDM communications because of their wavelength-localized operation (each resonator can be used to address only one on-resonance WDM

channel). The spacing between adjacent resonances of a single resonator is measured as the free spectral range (FSR).

Two designs that have attracted significant interest are the microring[5,6] and the microdisk.[7] In both structures, waves travel around the cavity, which is evanescently coupled to either one or two waveguides. Despite some design and fabrication differences between microrings and microdisks, and a potential performance advantage in microdisks,[7] the essence of their operation is the same, and the performance of best-of-class reported devices is roughly equivalent. For simplicity's sake, we will refer to both device classes as *microrings*.

Electrical control of microrings is implemented either by controlling the cavity's charge carrier concentration[5-7] or through the cavity's local temperature.[8] Both effects facilitate shifting the cavity resonant wavelength through a change in the refractive index. In this analysis, we focus on two microring functionalities needed to implement a point-to-point link.

In the microring modulator, high-speed electrical data signals are modulated onto a continuous-wave (CW) light passing in a waveguide in close proximity to the ring. The wavelength of the channel being modulated is set slightly off-resonance to realize on-off-keyed (OOK) light modulation. Current leading designs for low-power (less than 10 fJ/bit), high-speed (25 Gbps) operation are based on depletion-mode P-N diode structures.[5-7] Multiple microrings coupled with a single waveguide can form a WDM-modulator array, which operates simultaneously on a set of wavelength-parallel channels (each modulator addressing a single wavelength channel).

In the wavelength demultiplexer, microrings can also be used to filter out individual channels from a shared-bus waveguide. Each ring routes on-resonance channels to its drop port, whereas off-resonance channels continue propagating on the bus waveguide. Typically, the drop port leads to a photodetector to convert the modulated signals back to the electrical domain. Because a single ring might not provide sufficient suppression of unwanted channels, multiple rings can be cascaded or even coupled directly to each other to form a higher-order filter to improve suppression.[9]

Microring link design and analysis

For this case study, we assume current best research devices as possible candidates for production with certain improvements and assumptions detailed within the link buildup. We attempt to perform the analysis regardless of the underlying compute system to keep the results general and applicable to multiple integration options while capturing such a microring link's key features and tradeoffs.

Possible integration schemes in the intermediate time frame

Full monolithic integration of silicon photonics within a computing chip offers a significant potential performance boost thanks to close integration with the electronics,[10] but this is only likely to become commercially feasible in the long term (greater than 10 years) because it will require significant modifications and the addition of fabrication steps to existing chip-fabrication processes. In the near term, board-level optical modules currently offer significant flexibility but don't provide a solution for getting the data in and out of the chip package. Therefore, we limit the discussion of feasible intermediate-future systems to integration of silicon photonic technology as separate optical dies within a package. These dies are connected to the processor either through a shared substrate, a silicon carrier (see Figure 1a), or die stacked. Such optical dies would include a limited amount of electronics—as much as required to drive the optical link.

Link design for maximal bandwidth

We optimize the link design for maximal bandwidth density, with power efficiency as a secondary goal. Therefore, the design includes a dense WDM implementation with multiple microring modulators coupled with a shared-bus waveguide for the transmit side (Tx) and multiple ring filters coupled to a shared-bus waveguide for the demux and receive side (Rx) (see Figure 1b). The rings in the Tx and Rx arrays are staggered in size with a gradually increasing radius so as

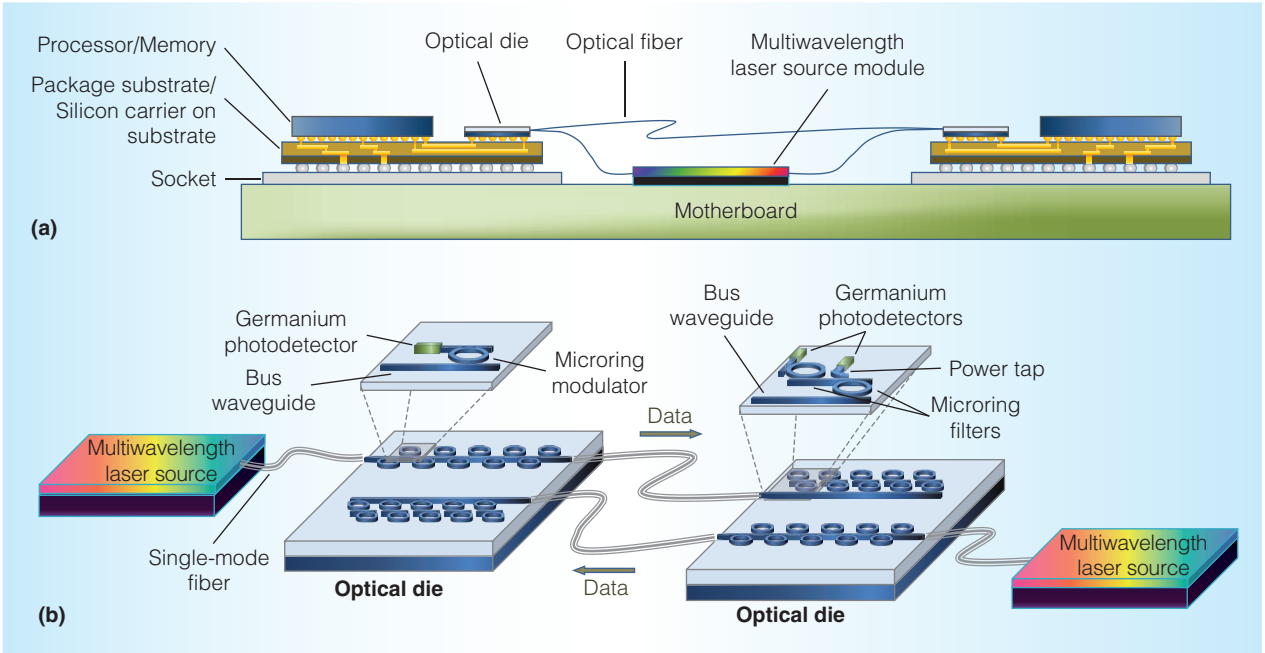

Figure 1. Microring-based point-to-point silicon photonic link concept linking two processor and memory nodes. Depiction of a multidie packaging solution based on direct bonding of both chips on a shared substrate or silicon carrier (a). The optical die would contain the driver and receiver electronics to the extent needed in close proximity to the optical devices. In-package electrical wires transfer data to and from the optical die. Microring-based point-to-point silicon photonic link (b). Two optical dies communicate over fiber. Each die includes a transmit module based on a microring modulator array and a receive module based on a two-stage microring demultiplexing array. Germanium photodetectors provide feedback for thermal stabilization and high-speed signal detection. Laser sources are assumed to be off-chip separate units.

to create a uniform frequency comb from the combined rings' resonances.

For the WDM link design, we first need to set reasonable parameter values to explore in terms of the number of wavelength channels, the channel wavelength spacing, and the per-channel modulation rates.

Because each ring modulator or filter should address only a single channel, the working spectrum is limited by the system's smallest FSR (corresponding to the largest ring). Therefore, the rings should be as small as possible to provide the largest possible available FSR. For a silicon ring with a sufficiently high quality factor (Q) of about 10,000, this radius is calculated at about 1.7 μm.[3,11] We can add increasingly larger rings to the array until the largest ring resonances start overlapping the smallest ring's resonance. This limitation dictates approximately 50 nm of operational spectrum with an optimized choice of ring radii ranging from about 1.8 to 1.9 μm, assuming operation at wavelengths around the telecom band of 1.55 μm.[3] Because most silicon photonic research to date has focused primarily on the 1.55 μm region, we based our analysis on parameters reported for devices in these wavelengths unless stated otherwise.

The traditional notion of increasing per-channel signaling rate to increase the aggregate bandwidth doesn't necessarily hold in dense WDM unamplified links, primarily because of the reduction in receiver sensitivity at higher data rates (requiring an increase in laser power). Operating at high data rates also results in additional SerDes and transceiver electronics power dissipation to an extent that might make high data rates less attractive. However, low channel rates could be irrelevant for the packaging options considered in this analysis because we need to get the data off the processor chip (to the optical die) through a limited amount of pins in the chip BGA. For the same pin-count limitation reasons, SerDes circuits

are emerging for electrical chip I/O regardless of the photonic components.[12,13] Furthermore, to maintain a high aggregate link bandwidth with low channel rates, one would also need to implement an aggressive WDM channel spacing, which would likely be limited by achievable Q ring values and would also enhance sensitivity to thermal fluctuations.

As depletion-mode microring devices have been reported at up to 25 Gbps operation with CMOS-compatible voltage levels,[5-7] we consider two primary working points—at 12.5- and 25-Gbps channel modulation rates. For 12.5-Gbps operation, we assume a ring-modulator Q about 12,000 to be feasible (cavity loss limited), whereas for 25 Gbps, we set the Q at 8,000 (photon-lifetime limited).

The minimal channel wavelength spacing depends on the optical power constraints because the losses and power penalties are a function of the channel spacing. Therefore, we start by walking through the analysis for the different channel rates at equivalent aggregate bandwidths set to 1.55 terabits per second (Tbps). We will revisit the question of maximal achievable channel density later.

After accounting for the available operational spectrum (about 50 nm) and possible modulation rates (12.5 and 25 Gbps), we can define two preliminary work points for analysis: design A, 12.5-Gbps channels at 50-GHz channel spacing (124 channels), and design B, 25-Gbps channels at 100-GHz channel spacing (62 channels).

Fabrication tolerances and thermal variation effect on microring array

Although resonators are useful for low-power, low-footprint WDM operation, their resonant wavelength is highly sensitive to dimension and temperature variations. Hence, we must pay special attention to the tuning mechanisms to correct for the static fabrication-related deviations from design as well as dynamic thermally induced fluctuations.

Recent studies have identified the primary causes of fabrication variation in microring arrays to result from the silicon-on-insulator (SOI) thickness unevenness (within a single wafer level as well as the wafer-to-wafer level) and fabrication-processes variability.[14,15] However, these studies found that the SOI thickness and fabrication processes are fairly uniform on the single-die scale. Because each modulator occupies a small footprint, it's feasible to create a full array on a single die, which in turn implies that the relative resonance location of the rings with respect to each other can be fairly well controlled at up to 100 GHz deviation of each ring from the designed spacing.[15] Furthermore, given a design that fully spans the spectral range with equidistant resonances, the cyclical nature of the rings' resonances ensures that a single ring resonance will fall within half the channel spacing of the channel location for each wavelength channel. Therefore, such an array's fabrication variability is translated into a postfabrication mapping problem of the wavelength channels to the appropriate physical rings along with some thermal fine tuning.

Although the fabrication variation is static, the temperature variations within a processing node's package can dynamically fluctuate over tens of degrees Celsius. Because silicon has a strong thermo-optic effect, ring resonances shift by as much as 10 GHz/°C.[14] Assuming a 40°C range of temperatures realistically experienced within a package, this can result by as much as 400 GHz dynamic shifts of the microrings' resonances. Therefore, a dynamic solution must be included to either adjust for the ring resonance drift or alternatively stabilize the ring temperature locally to keep its resonant wavelength fixed (using an integrated heater). Thermal crosstalk between the array elements can be avoided by forming trenches between devices.[16]

The fabrication-related post-channel-mapping residual thermal tuning must be able to shift each ring to up to the greater value of half the channel spacing and the random shifts from an equidistant frequency comb. For the dense WDM channel spacing we're considering, the random shifts from the comb (100 GHz) dominate. However, because the random thermal variations in these settings are much larger (400 GHz), we can approximate the upper bound on the amount of thermal compensation required to account for about 400-GHz resonance shifts.

Modulator array

A WDM modulator array (Figure 1b) can be constructed of multiple electro-optically active microrings coupled to a shared bus waveguide, each imprinting on-off key modulation on a single wavelength channel. Such a design avoids the need for WDM multiplexers and demultiplexers, greatly reducing footprint as well as insertion losses.

Optical properties. Because the link would ideally operate in a CMOS environment, we limit high-speed signals to operate at CMOS-compatible voltage levels—that is, 1 V_{pp}. Researchers have demonstrated extremely low-power (10 fJ/bit) microring modulators with these driving conditions.[5-7] Given a driving voltage limitation of 1 V_{pp} (and therefore, a given amount of resonance shifting), we now must determine the channel's optimal relative wavelength positioning with respect to the resonance. A fundamental tradeoff exists here: a small wavelength detuning from the resonance results in a high extinction ratio (ER) but along with a high insertion loss (IL).[6] For a device demonstrated with Q about 10,000, the optimal working point at 12.5 Gbps corresponded to a 4-dB IL and 8-dB ER.[6] Similarly, for a Q about 8,000, the optimal theoretical working point at 25 Gbps would correspond to similar values of 3.5 dB IL and best-case 6.45 dB ER.[5] We compute the IL related to the OOK power extinction in the ring for these ERs assuming a 50-percent 0-to-1 ratio. The corresponding ER power penalty (compared to a commercial-grade modulator) is computed for a shot-noise limited model.[17]

Beyond the power loss resulting from each modulator in a stand-alone configuration, integrating these devices into a dense wavelength array results in additional losses and intermodulation crosstalk from adjacent channels' modulators resonance tails.[18] We compute the IL (Table 1) similarly to the method suggested by Sherwood-Droz et al.[3] while assuming a 42-pm (5.3-GHz) dynamic shift of each modulator.[8]

Thermal stabilization. Although the modulation bias voltage control is a built-in heating mechanism,[18] researchers view the inclusion of bias tees in a CMOS environment as difficult. DeRose et al. have demonstrated integrated heaters with 4.4 μW/GHz heating efficiency.[8] Cunningham et al. have shown potential for improved heating efficiency (1.6 μW/GHz) at the cost of increased fabrication complexity by etching the underlying silicon substrate.[19]

Researchers have demonstrated feedback mechanisms for controlling the modulators' heaters based on average power sensing[20] or BER measurements.[21] Integrating these mechanisms in a microring WDM array would most likely require the inclusion of a drop port on the microring modulators, which would feed into a photodetector. Because the optical power reaching the modulators is relatively high, the photodetector would require only a small fraction of the light to be coupled to a drop port. Therefore, it should be possible to attain the designed Q factors even with the inclusion of the drop port. A locally integrated feedback thermistor can augment or replace these mechanisms as well.[8] Assuming 1-mW power consumption of the feedback mechanism (estimated for the type of feedback implemented elsewhere[18]) and the need to stabilize a maximum resonance shift of 400 GHz (with an average tuning of only half that shift), heating stabilization dissipates approximately 1.9 mW per modulator.

Modulator driver electronics. A modulator driver with some low-power amplification stages might be required to create an optimal driving signal for a 1-V_{pp} modulator. Researchers have demonstrated low-power modulator driver circuits for modulators with higher voltage swings with 0.14 pJ/bit efficiency at 10 Gbps.[22] We assume that the same numbers can be reasonably applied to 12.5-Gbps operation and that 0.3 pJ/bit is achievable for 25 Gbps (quadruple the amount of power dissipated but also double the bit rate).

Wavelength demultiplexer and receiver array

We implement the Rx side using microring filters for spectrally demultiplexing the WDM channels followed by high-speed receivers for optical-to-electrical (O/E) conversion based on germanium photodetectors. Microrings are suitable for a

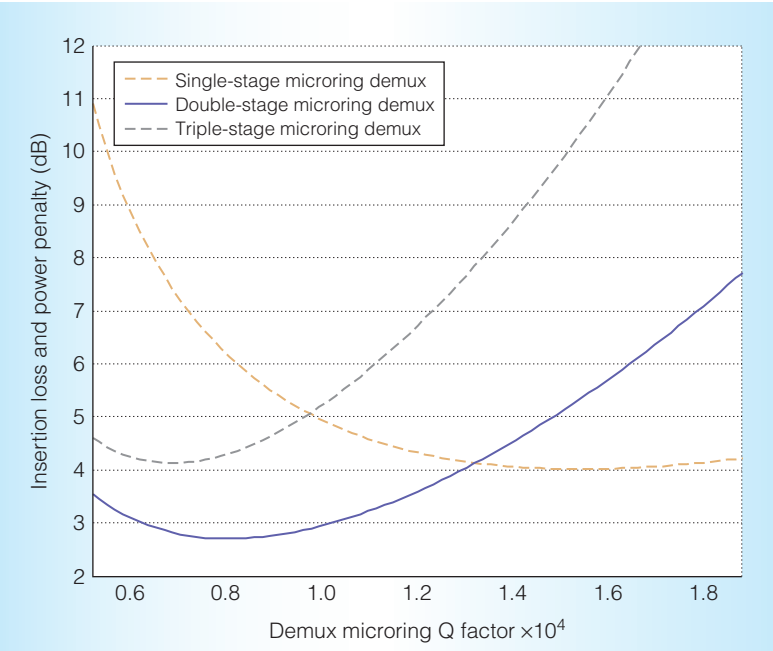

Figure 2. Design considerations of microring-based wavelength demultiplexer at design point A (12.5 Gbps channels, 50 GHz spacing). A double-stage microring demultiplexer (with the appropriate choice of an optimal Q factor) provides sufficient interchannel crosstalk suppression while keeping sideband truncation and insertion losses reasonable.

dense high-channel count WDM demux. Designs based on Echelle gratings or arrayed waveguide gratings (AWG) occupy a large silicon footprint and are likely to be limited to no more than about 50 ports.[23]

For simplicity of analysis, we consider first-order ring filters rather than higher-order filters. Germanium PIN photodetectors or avalanche photodetectors (APD) are likely candidates for CMOS-compatible integration with silicon. The choice of microrings implies that polarization control is required either using polarization-maintaining (PM) fiber or regular single-mode fiber with some polarization correction mechanism.

Germanium photodetectors. Waveguide-integrated photodetectors operating at up to 40 Gbps have had record performances in recent years.[22,24,25] Although some of these reported values are relevant for detection at short wavelengths, we assume that sufficient advancements can be made to enable these sensitivities to apply to 1,550 nm. From Assefa et al.,[24] we interpolate a sensitivity of −16 dBm at 12.5 Gbps and assume a sensitivity of −12 dBm at 25 Gbps. Closer integration with electronics (with lower parasitic capacitance) and improved APD designs provide potential for sensitivity improvements. Therefore, we assume a uniform 4-dB improvement to be possible at all bit rates within a few years, which implies a sensitivity of −20 dBm at 12.5 Gbps and −16 dBm at 25 dBm.

Multistage microring filtering. For dense WDM demultiplexing, we consider having more than a single ring filter per channel. The IL dropping through a ring filter is 0.5 dB[26] and a power tap (if needed) for each stage's thermal stabilization is assumed to result in an additional 0.5 dB loss. Because APDs seem a likely candidate for photodetection, we consider shot-noise-limited power-penalty models.[17,27]

The choice of Q for the ring-filtering stages trades off interchannel crosstalk suppression with signal degradation due to sideband attenuation (see Figure 2). Some extra IL is also accumulated on each channel as power is lost to neighboring filters (similar to the modulator array loss). However, the loss of power into erroneous channels is also accompanied by the reduction of those channels as crosstalk contributors. Because the penalty from crosstalk outweighs the extra IL, we consider the worst-case crosstalk no-extra-IL scenario. As Figure 2 shows, two-stage filtering achieves superior performance compared to single- and triple-stage designs. The optimal Q factor is 8,000 for design A and 4,000 for design B, both resulting in about 1.2-dB total power penalty (crosstalk and sideband truncation) and 1.5-dB IL.

Thermal stabilization determines the demultiplexer array's power consumption. The first stage requires 1.9 mW for both heater and feedback circuit, whereas the second stage dissipates only 0.9 mW because the feedback can be based on photodetection by the receiver.

Receiver electronics. A primary source of power dissipation in such a link occurs at the receiver circuitry (amplification, equalization, clock, and data recover). Zheng et al. have demonstrated power-efficient circuitry

at 10 Gbps with 0.4 pJ/bit based on a transimpedance amplifier (TIA) and sense-amp (SA) design.[22] We assume that this power efficiency can equally apply for the 12.5 Gbps bit rate. Although Assefa et al. have demonstrated much higher power circuitry for 25 Gbps using a TIA and limited amplifiers (LA) with several pJ/bit power consumption,[24] we assume the need for improved power efficiency will allow 25 Gbps to be developed to operate at 1 pJ/bit. High-speed equalization circuits at the receiver, which were not included in this analysis, might also help improve receiver sensitivity by negating tight-filtering penalties (both optical and electrical),[28] but such circuits typically dissipate significant power and could come at an unacceptable power efficiency cost.

Waveguides and fibers

Within the system proposed, light is guided by an optical fiber between the Tx and Rx module and by silicon-photonic bus waveguides within the photonic modules.

Single-mode fiber and connectors. We assume fibers are used to guide light from the off-chip laser source to the Tx module and to connect the Tx and Rx modules. Because most silicon photonic components are single mode and operate at a single polarization, the fibers must be single mode as well as have some control of polarization to couple the light into the optical dies on the correct polarization. Because a standard single-mode fiber has extremely low propagation losses (less than 0.2 dB/km), the propagation losses are not a concern. Dispersion is unlikely to play a major role as well at distances of up to a few kilometers at these modulation rates.[29]

Single-mode-fiber connectors, however, can contribute significant loss (0.1 to 0.5 dB) and require accurate connecting, which might be challenging in a data center or high-performance computer. We will account for only 0.5 dB loss total of fiber connector and propagation loss, but this number can change drastically depending on the realistic settings within the computing facility. We can achieve polarization control by using PM fibers or some polarization rotation mechanism of the fiber (for instance, with feedback from an on-chip photodetector and some circuitry to optimize the polarization).

Overall link performance will be mostly agnostic to the actual link distance—ranging from a few centimeters (on board links) to a kilometer (across a data center or high-performance computer) and will mostly depend on the number and quality of fiber connectors required. The resulting system-wide equal-power-cost connectivity enabled by optical links fundamentally differs from electrical interconnects' connectivity costs at link distances greater than a few meters.

On-chip waveguides. Researchers can implement on-chip bus waveguides using silicon (1.7 dB/cm)[30] or silicon nitride (0.1 dB/cm).[31] Although silicon nitride provides the potential for extremely low-loss propagation, special tapers are required to couple the light between the silicon nitride sections and silicon sections (coupling to the silicon microrings), adding extra loss. Because the bus waveguides can be laid down to occupy only a few millimeters' length on the dies, crystalline silicon waveguides result in overall lower loss. Using an overestimated 3,600-μm^2 footprint for each microring (including contacts and thermal stabilization), and laying microrings on both sides of the waveguide, we arrive to a total length of 5 mm on the Tx side and 4.5 mm on the Rx side for design A (4 mm and 3.5 mm, respectively, for design B). Because the waveguides are relatively short, and the overall powers coupled into the waveguide less than 125 mW, nonlinear loss remains insignificant.

Waveguide-to-fiber couplers. In this design, we propose using best-case 0.5-dB IL edge couplers with inverse tapers[32,33] as opposed to much higher-loss and bandwidth-limited grating couplers. However, all such couplers are extremely sensitive to alignment displacements (1 dB for even 1-μm initial misalignment). Therefore, we use 1 dB/facet as a more representative value for each coupler loss.

Link power budget and power efficiency

After accounting for all the optical loss and power penalty factors (see Table 1), we

Table 1. Link insertion loss and power penalty.

Module/Power-Budget Parameter	Component	Design A	Design B
Tx module	Two edge couplers	2 dB	2 dB
	Waveguide	0.85-dB IL	0.7-dB IL
	Modulator array	4-dB IL (ring modulator)	3.5-dB IL (ring modulator)
		1-dB IL (array induced)	0.6-dB IL (array induced)
		2.4 dB IL (OOK modulation)	2.1-dB IL (OOK modulation)
		2-dB penalty (limited ER)	2.8-dB penalty (limited ER)
		0-dB penalty (intermodulation)	0-dB penalty (intermodulation)
Optical fiber	Fiber connectors	0.5-dB IL	0.5-dB IL
Rx module	One edge coupler	1-dB IL	1-dB IL
	Waveguide	0.085-dB IL (just till first ring)	0.085-dB IL (just till first ring)
	Demux array	1.5-dB IL (two rings and power tap)	1.5-dB (two rings and power tap)
		1.2-dB penalty (filtering and crosstalk)	1.2-dB penalty (filtering and crosstalk)
Jitter	—	2-dB penalty	3-dB penalty
Receiver sensitivity	—	–20 dBm	–16 dBm
Required laser power per channel	—	–1.4 dBm (0.7 mW)	2.9 dBm (1.9 mW)

can construct the following total IL and power penalty for the each module under each design.

We also add estimated additional power penalties of 2 dB at 12.5 Gbps and 3 dB at 25 Gbps to account for jitter and temporal closing of the eye diagram.

Once we account for the receiver sensitivity, power penalties, and IL, we can compute the required laser power at the link's input. To achieve error-free transmission through the link, we need 0.7 mW per channel for design A and 1.9 mW per channel for design B. The difference between the two scenarios is mainly due to the 4-dB difference in receiver sensitivity.

Power efficiency. As Table 2 shows, and assuming full link utilization, the aggregate power dissipation in the optical dies equals 0.9 pJ/bit in design A and 1.5 pJ/bit in design B. We estimate electrical data transmission to and from the optical dies to consume 1 pJ/bit for design A and 2 pJ/bit for design B, assuming electronic transceivers will continue to scale in power efficiency with new CMOS technology nodes.[34,35] Depending on the availability of high wall-plug efficiency lasers, the overall power consumption can be dominated just by the laser power, therefore driving optical design to put even more emphasis on low-loss structures and high-sensitivity detectors. Partial link utilization (noncontinuous data transmission) would significantly degrade the power efficiency because the lasers are assumed to be constantly on for wavelength and power stability.

Open issues, design tradeoffs, and alternatives

Although Tables 1 and 2 summarize the power budget and wall-plug efficiency of the microring link designs we worked through for the case analysis, it's important to inspect the link's physical scaling limitations as well as touch on some important design tradeoffs and concerns for implementation in computing systems.

Maximal aggregate bandwidth

Using a dense WDM grid results in increased insertion losses as well as power penalties. We can quantify this tradeoff assuming a two-stage design (not necessarily optimal for large channel spacing). As the

Table 2. Link power efficiency by component.

Module/Functionality	Operation	Design A	Design B
Tx module	Modulation	0.01 pJ/bit	0.01 pJ/bit
	Modulator driver	0.10 pJ/bit	0.30 pJ/bit
	Thermal stabilization	0.15 pJ/bit	0.08 pJ/bit
Rx module	Thermal stabilization	0.22 pJ/bit	0.11 pJ/bit
	Receiver PD and electronics	0.40 pJ/bit	1.00 pJ/bit
Total on-die dissipation	—	0.90 pJ/bit	1.50 pJ/bit
Laser source	(Considering different possible wall-plug efficiencies)	0.70 mW implies 5.6 pJ/bit at 1% efficiency, 0.56 pJ/bit at 10% efficiency	1.80 mW implies 7.8 pJ/bit at 1% efficiency, 0.78 pJ/bit at 10% efficiency
Electronic data transmission to/from die	—	1 pJ/bit	2 pJ/bit
Total link power efficiency	—	7.5 pJ/b at 1% laser efficiency, 2.5 pJ/b at 10% laser efficiency	11.3 pJ/b at 1% laser efficiency, 4.3 pJ/b at 10% laser efficiency

channel spacing is reduced, higher losses and penalties are incurred, requiring higher per-channel powers (Figure 3a). Furthermore, as channel spacing is reduced along with a channel count increase, the aggregate laser power is further increased. At a certain point, nonlinear waveguide losses[36] and optical instabilities in the microring modulators[37] will start appearing. We set an upper limit of 21 dBm (125 mW) aggregate power launched into the waveguides to avoid nonlinear loss, which in turn determines how much aggregate bandwidth can be achieved through the link. This limit is 1.8 Tbps using 12.5-Gbps channels and 1.56 Tbps using 25-Gbps channels (Figure 3b).

This result shows that operating at 25-Gbps channel rates rather than 12.5-Gbps channel rates not only comes at a significant cost to power efficiency (Table 2) but is also detrimental to bandwidth scalability. However, in eventual systems, the choice of channel rate might be dictated by packaging limitations rather than optimal optical performance.

Because the bandwidth limit is set at this point by the aggregate power in the silicon waveguide, it is also important to consider alternative bus waveguide configurations, such as Si_3N_4 (silicon nitride), which have lower loss, a higher damage threshold, and lower effective nonlinearity. However, the challenges there lie in integrating the active devices with the waveguide in a low-loss manner as well as the more complex fabrication processes.

Sensitivity to misalignment and additional losses

As Figure 3b shows, any additional loss (due to misalignment or higher number of optical fiber connectors) would directly impact the achievable aggregate bandwidth and the link's power efficiency. Even a 3-dB additional loss would immediately result in about a 30 percent reduction of link bandwidth. As such, it is not unlikely that optical amplification will be required to enable high-throughput operation with higher-complexity (and loss) optical interconnects.

Laser efficiency

Commercial wavelength-stabilized sources used in WDM telecom systems typically operate at about 1 percent wall-plug efficiency because of the significant portion of power spent on cooling and temperature stabilization of each laser cavity as well as the IL associated with multiplexing multiple signals into a single fiber. Higher efficiency relaxed-wavelength-stabilization concepts allowing for a finite amount of wavelength drift aren't suitable for dense WDM grids. Therefore, developing telecom-grade laser sources with high wall-plug efficiency is key for creating low-power optical WDM interconnects.

Multiwavelength lasers that use a single cavity to generate multiple laser channels simultaneously[38,39] are highly attractive

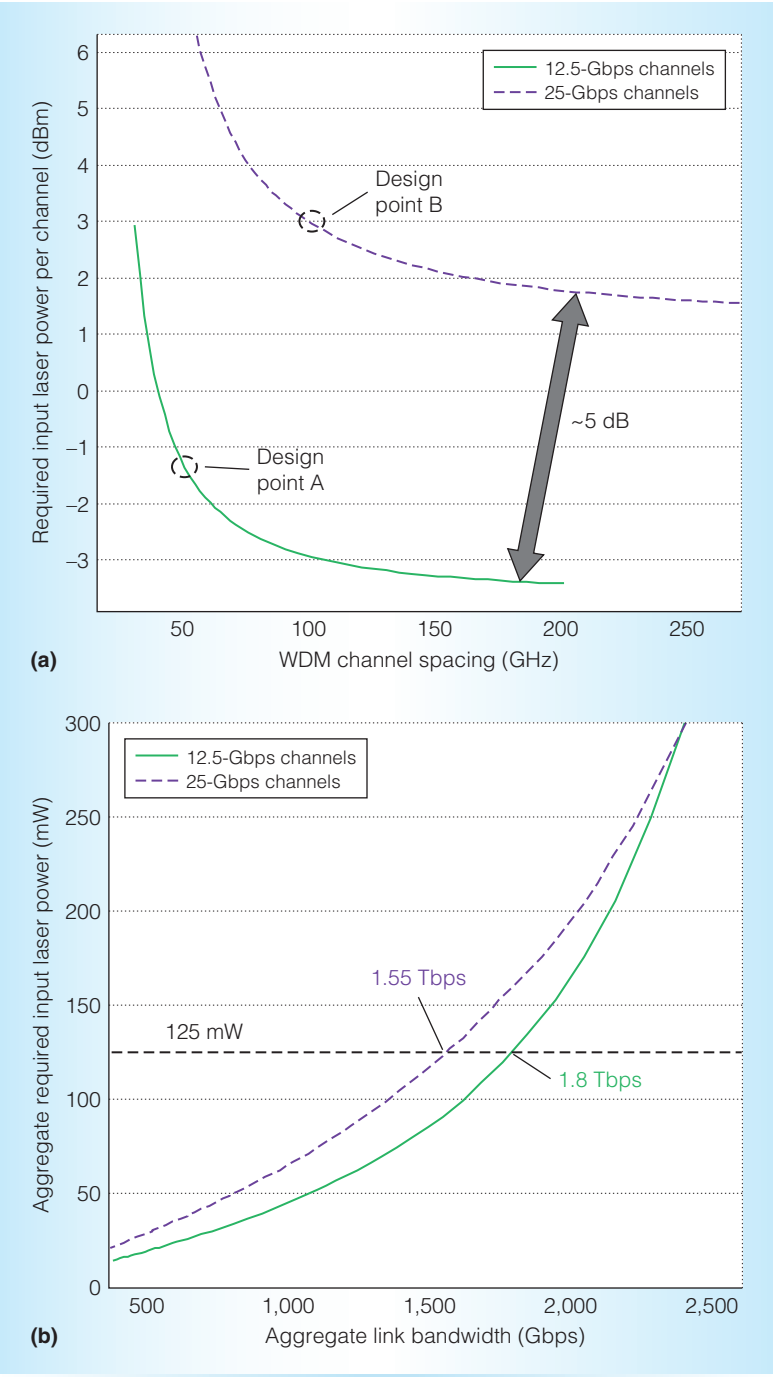

Figure 3. Bandwidth scaling limitations of a WDM microring link. Required optical power per channel as a function of channel spacing (a). The smaller the channel spacing, the larger the overall insertion losses and power penalties. The dominant factor differentiating the two channel rates is the reduced receiver sensitivity at the higher data rate. Aggregate input laser power as function of link bandwidth (b). Assuming a 125-mW launched power upper limit (100 mW in the waveguide), the aggregate bandwidth is limited to 1.55 Tbps (62 channels) for a 25-Gbps modulation rate and 1.8 Tbps (144 channels) for a 12.5-Gbps modulation rate.

because only a single cavity must be stabilized. Such systems are starting to be commercially available and are a subject of ongoing research at both 1.3 and 1.55 µm. Such quantum dot fabry-perot lasers show great promise for low relative-intensity noise (RIN) operation,[38,39] but further scaling of the number of lasing lines and output power is required, as well as improved channel power uniformity.

A different possible solution could be based on using parametric oscillators[40] to generate multiple wavelengths. In these solutions, only the seed laser and the parametric cavity must be thermally stabilized. However, such parametric solutions still require significant improvements to provide suitable stable sources.

Polarization control and polarization-insensitive devices

Because microrings are polarization sensitive, a system built with such devices must incorporate polarization control. For our analysis, we assumed such a mechanism was implemented either by using PM fiber or by an off-chip polarization controller operated with feedback from the die to optimize the polarization state of the light coupled into the die. If it's possible to implement this polarization control for all signals simultaneously, the additional power cost of the feedback and control mechanisms would be amortized over the aggregate link bandwidth and therefore would have little impact on the power efficiency. An alternative set of partial solutions includes the following:

- On-chip laser sources directly coupled to the waveguides. Integrating III-V materials on silicon is challenging and requires low-temperature laser stabilization,[41] as well as the addition of a wavelength-multiplexing structure.
- Wavelength-demultiplexing using an AWG or Echelle grating, which can potentially be less polarization sensitive than rings. This might limit the number of channels used because of port-count limitations.
- Implementing a polarization diversity scheme to receive both polarizations.[42]

This solution comes at the cost of doubling the Rx size and power consumption.

A combination of such solutions can be jointly implemented if PM fiber is deemed unfeasible, though at the cost of increasing the insertion loss and reducing the link aggregate bandwidth.

Operational wavelength

Choosing to operate at 1.3 μm could provide improved performance primarily because of higher-efficiency lasing and higher-sensitivity detectors. However, there might be disadvantages with regard to propagation losses and larger sensitivity to misalignment at shorter wavelengths. External factors such as the convenience of R&D at 1,550 nm, availability of Erbium-doped fiber amplifiers (EDFAs), and compatibility with telecom infrastructure might be the eventual deciding factors for wavelength band choice.

As this case study showed, microring-based silicon photonics should be able to provide bandwidth densities as high as 1.8 Tbps on a single optical link occupying approximately 1.5-mm^2 net silicon and requiring only three optical connectors to the die. The potential for multifiber packaging to the optical die as well as the inclusion of multiple links within a single in-package module clearly poises the technology to provide sufficient off-chip bandwidth for processor-memory and processor-processor transactions.

Our analysis shows a high dependence of the overall power efficiency on laser wall-plug efficiency, driving the need for further development of higher-sensitivity receivers as well as more efficient multiwavelength laser sources. Furthermore, because laser power plays a key role in the overall efficiency, the efficiency is inversely proportional to the channel spacing and hence the link aggregate bandwidth.

At the same time, silicon photonic photodetectors will have to continue improving to realize sensitivities sufficient to meet the power budget constraints of unamplified links. Further improvements will also be required to enable a transition from point-to-point links to higher-complexity and consequently higher-loss interconnects that pass through system networks. The potential for such dramatic performance improvements in germanium photodetectors lies in the close integration with the electronic circuitry as well as continued focus on avalanche-based receivers, which have only recently started to emerge in the silicon photonic world.

In addition, these system-scale interconnects will further drive the need for chip-scale, energy-efficient optical amplification. Individual component performance improvements must at the same time be accompanied with continued emphasis on dense device integration and testability along with optical packaging. Fiber packaging with micrometer-level precision is one of the key challenges the industry will have to tackle going forward to enable low-loss single-mode die-to-fiber interfaces.

Although nonresonant commercial silicon photonics communication modules already exist for high-performance computing and datacenter environments, silicon resonator-based devices provide the best potential for high-bandwidth, high-density, low-power I/O that can viably be inserted within the chip packaging solution and overcome the I/O bottleneck. MICRO

References

1. F. O'Mahony et al., ''The Future of Electrical I/O for Microprocessors,'' *Proc. Int'l Symp. VLSI Design, Automation and Test,* IEEE CS, 2009, pp. 31-34.
2. M. Georgas et al., ''Addressing Link-Level Design Tradeoffs for Integrated Photonic Interconnects,'' *Proc. Custom Integrated Circuits Conf.,* IEEE CS, 2011, pp. 1-8.
3. N. Sherwood-Droz et al., ''Device Guidelines for WDM Interconnects Using Silicon Microring Resonators,'' *Proc. Workshop on the Interaction Between Nanophotonic Devices and Systems,* 2010; http://www.csl.cornell.edu/winds2010/abstracts/sherwood-winds2010.pdf.
4. M. Lipson, ''Guiding, Modulating, and Emitting Light on Silicon—Challenges and Opportunities,'' *J. Lightwave Technology,* Dec 2005, pp. 4222-4238.

5. G. Li et al., ''25Gb/s 1V-Driving CMOS Ring Modulator with Integrated Thermal Tuning,'' *Optics Express,* Oct. 2011, pp. 20,435-20,443.
6. P. Dong et al., ''High-Speed and Compact Silicon Modulator Based on a Racetrack Resonator with a 1 V Drive Voltage,'' *Optics Letters,* Oct. 2010, pp. 3246-3248.
7. M.R. Watts et al., ''Vertical Junction Silicon Microdisk Modulators and Switches,'' *Optics Express,* Oct. 2011, pp. 21,989-22,003.
8. C.T. DeRose et al., ''Silicon Microring Modulator with Integrated Heater and Temperature Sensor for Thermal Control,'' *Proc. Conf. Lasers and Electro-Optics,* Optical Soc. of America, 2010, no. CThJ3.
9. M.A. Popovic et al., ''Multistage High-Order Microring-Resonator Add-Drop Filters,'' *Optics Letters,* Sept. 2006, pp. 2571-2573.
10. J. Orcutt et al., ''Open Foundry Platform for High-Performance Electronic-Photonic Integration,'' *Optics Express,* May 2012, pp. 12,222-12,232.
11. Q. Xu, D. Fattal, and R.G. Beausoleil, ''Silicon Microring Resonators with 1.5-μm Radius,'' *Optics Express,* Mar. 2008, pp. 4309-4315.
12. J.E. Jaussi et al., ''Multi-Gbit I/O and Interconnect Co-design for Power Efficient Links,'' *Proc. IEEE 19th Conf. Electrical Performance of Electronic Packaging and Systems,* IEEE CS, 2010, pp. 1-4.
13. B. Kim et al., ''A 10-Gb/s Compact Low-Power Serial I/O With DFE-IIR Equalization in 65-nm CMOS,'' *IEEE J. Solid State Circuits,* Dec. 2009, pp. 3526-3538.
14. W.A. Zortman et al., ''Silicon Photonics Manufacturing,'' *Optics Express,* Oct. 2010, pp. 23,598-23,607.
15. A.V. Krishnamoorthy et al., ''Exploiting CMOS Manufacturing to Reduce Tuning Requirements for Resonant Optical Devices,'' *IEEE Photonics J.,* June 2011, pp. 567-579.
16. P. Dong et al., ''Low Power and Compact Reconfigurable Silicon Multiplexing Devices,'' *Proc. Integrated Photonics Research, Silicon and Nanophotonics,* Optical Soc. of America, 2010, no. IWG5.
17. R. Ramaswami, K.N. Sivarajan, and G.H. Sasaki, *Optical Networks,* Morgan Kaufman, 2010.
18. K. Padmaraju et al., ''Intermodulation Crosstalk from Silicon Microring Modulators in Wavelength-Parallel Photonic Networks-on-Chip,'' *Proc. Photonic Soc. Ann. Meeting,* Optical Soc. of America, 2010, no. ThB5.
19. J.E. Cunningham et al., ''Highly-Efficient Thermally-Tuned Resonant Optical Filters,'' *Optics Express,* Aug. 2010, pp. 19,055-19,063.
20. K. Padmaraju et al., ''Dynamic Stabilization of a Microring Modulator Under Thermal Perturbation,'' *Proc. Optical Fiber Communications Conf.,* Optical Soc. of America, 2012, no. OW4F.2.
21. A.L. Lentine et al., ''Active Wavelength Control of Silicon Microphotonic Resonant Modulators,'' *Proc. IEEE Optical Interconnects Conf.,* IEEE CS, 2012, no. TuC5.
22. X. Zheng et al., ''Ultra-Efficient 10Gb/s Hybrid Integrated Silicon Photonic Transmitter and Receiver,'' *Optics Express,* Mar. 2011, pp. 5172-5186.
23. L. Chen et al., ''Monolithically Integrated 40-Wavelength Demultiplexer and Photodetector Array on Silicon,'' *IEEE Photonics Technology Letters,* July 2011, pp. 869-871.
24. S. Assefa et al., ''20Gbps Receiver Based on Germanium Photodetector Hybrid-Integrated with 90nm CMOS Amplifier,'' *Proc. Conf. Lasers and Electro-Optics,* Optical Soc. of America, 2011, no. PDPB11.
25. S. Assefa, F. Xia, and Y.A. Vlasov, ''Reinventing Germanium Avalanche Photodetector for Nanophotonic On-Chip Optical Interconnects,'' *Nature,* vol. 464, 4 Mar. 2010, pp. 80-84.
26. A. Biberman et al., ''Demonstration of All-Optical Multi-Wavelength Message Routing for Silicon Photonic Networks,'' *Proc. Optical Fiber Communications Conf.,* Optical Soc. of America, 2008, no. OTuF6.
27. J.L. Rebola and A.V.T. Cartaxo, ''Power Penalty Assessment in Optically Preamplified Receivers with Arbitrary Optical Filtering and Signal-Dependent Noise Dominance,'' *J. Lightwave Technology,* Mar. 2002, pp. 401-408.
28. A. Rylyakov, C. Schow, and J. Kash, ''A New Ultra-High Sensitivity, Low-Power Optical Receiver Based on a Decision-Feedback Equalizer,'' *Proc. Optical Fiber Communications Conf.,* Optical Soc. of America, 2011, no. OThP3.
29. A. Biberman et al., ''First Demonstration of Long-Haul Transmission Using Silicon

Microring Modulators,'' *Optics Express,* July 2010, pp. 15,544-15,552.
30. F. Xia, L. Sekaric, and Y.A. Vlasov, ''Ultra-compact Optical Buffers on a Silicon Chip,'' *Nature Photonics,* vol. 1, Jan. 2007, pp. 65-71.
31. A. Gondarenko, J.S. Levy, and M. Lipson, ''High Confinement Micron-Scale Silicon Nitride High Q Ring Resonator,'' *Optics Express,* July 2009, pp. 11,366-11,370.
32. S. McNab,, N. Moll, and Y.A. Vlasov, ''Ultra-Low Loss Photonic Integrated Circuit with Membrane-Type Photonic Crystal Waveguides,'' *Optics Express,* Nov. 2003, pp. 2927-2939.
33. M. Pu et al., ''Ultra-Low-Loss Nano-Taper Coupler for Silicon-on-Insulator Ridge Waveguide,'' *Proc. 36th European Conf. Optical Communication,* IEEE CS, 2010, no. Tu5C6.
34. K. Fukuda et al., ''A 12.3mW 12.5Gb/s Complete Transceiver in 65nm CMOS,'' *Proc. IEEE Int'l Solid-State Circuits Conf.,* IEEE CS, 2010, no. 20.5.
35. J.E. Proesel and T.O. Dickson, ''A 20-Gb/s, 0.66-pJ/bit Serial Receiver with 2-Stage Continuous-Time Linear Equalizer and 1-Tap Decision Feedback Equalizer in 45nm SOI CMOS,'' *Proc. IEEE Symp. VLSI Circuits,* IEEE CS, 2011, no. 20.1.
36. B.G. Lee et al., ''Demonstration of Broadband Wavelength Conversion at 40 Gb/s in Silicon Waveguides,'' *IEEE Photonics Technology Letters,* Feb. 2009, pp. 182-184.
37. Q. Li et al., ''Experimental Characterization of the Optical-Power Upper Bound in a Silicon Microring Modulator,'' *Proc. IEEE Optical Interconnects Conf.,* IEEE CS, 2012, no. TuB5.
38. G.L. Wojcik et al., ''A Single Comb Laser Source for Short Reach WDM Interconnects,'' *Proc. SPIE,* vol. 7230, 2009, pp. 72,300M-1-72,300M-12.
39. Z.G. Lu et al., ''Low Noise InAs/InP Quantum Dot C-Band Monolithic Multiwavelength Lasers for WDM-PONs,'' *Proc. Optical Fiber Communications Conf.,* Optical Soc. of America, 2009, no. JWA27.
40. M.A. Foster et al., ''Silicon-Based Monolithic Optical Frequency Comb Source,'' *Optics Express,* July 2011, pp. 14,233-14,239.
41. S.R. Jain et al., ''Integrated Hybrid Silicon DFB Laser-EAM Array Using Quantum Well Intermixing,'' *Optics Express,* June 2011, pp. 13,692-13,699.
42. L. Chen, C.R. Doerr, and Y.K. Chen, ''Compact Polarization Rotator on Silicon for Polarization-Diversified Circuits,'' *Optics Letters,* Feb. 2011, pp. 469-471.

Noam Ophir is a PhD candidate in the Lightwave Research Laboratory at Columbia University. His research focuses on silicon photonics for nonlinear processing and optical interconnects. Ophir has an MS in electrical engineering from Tel Aviv University.

Christopher Mineo is a computer architecture researcher at the Laboratory for Physical Sciences. His research interests include network-on-chip optimization, low-power circuit design, architectures for massively parallel systems, and optical interchip communication for high-performance computing. Mineo has a PhD in computer engineering from North Carolina State University, Raleigh.

David Mountain is a senior technical director at the Laboratory for Physical Sciences. His research interests include power efficiency, chip I/O, system-level interconnects, file system I/O, productivity, and resilience. Mountain has an MS in electrical engineering from the University of Maryland, College Park. He is a senior member of IEEE.

Keren Bergman is the Charles Batchelor Professor and chair of electrical engineering at Columbia University, where she also directs the Lightwave Research Laboratory. Her research interests include optical interconnection networks for advanced computing systems, data centers, optically interconnected memory, and chip multiprocessor nanophotonic networks-on-chip. Bergman has a PhD in electrical engineering from the Massachusetts Institute of Technology. She is a fellow of IEEE and the OSA.

Direct questions and comments about this article to Noam Ophir, 1300 Seeley W. Mudd, 500 West 120th St., New York, NY 10027; ophir@ee.columbia.edu.