6

Large-Scale Integrated Photonics for High-Performance Interconnects

RAYMOND G. BEAUSOLEIL, HP Laboratories

Moore's Law has set great expectations that the performance of information technology will improve exponentially until at least the end of this decade. Although the physics of silicon transistors alone might allow these expectations to be met, the physics of the long metal wires that cross and connect packages almost certainly will not. Global-level interconnects incorporating large-scale integrated photonics fabricated on the same platform as silicon microelectronics hold the promise of revolutionizing computing by enabling parallel many-core and network switch architectures that combine unprecedented performance and ease of use with affordable power consumption.

Over the last decade, remarkable progress has been made in research on low-power silicon photonic devices for interconnect applications, and CMOS-compatible fabrication technologies promise a "Moore's Law for photonics" that could completely change the economics of integrated optics. In this survey, photonic technologies amenable to large-scale CMOS integration are reviewed from the perspective of high-performance interconnects operating over distance scales of 1mm to 100m. An overview of the requirements placed on integrated optical devices by a variety of modern computer applications leads to discussions of active and passive photonic components designed to generate, guide, filter, modulate, and detect light in the telecommunication bands. Critical challenges and prospects for large-scale integration are evaluated with an emphasis on silicon-on-insulator as a platform for photonics.

Categories and Subject Descriptors: B.4.3 [**Input/Output and Data Communications**]: Interconnections (Subsystems)—*Fiber optics; Interfaces; Parallel I/O; Physical structures and topology*; B.7.0 [**Integrated Circuits**]: General; B.7.1 [**Integrated Circuits**]: Types and Design Styles—*Advanced technologies; input/output circuits; VLSI*

General Terms: Performance, Reliability

Additional Key Words and Phrases: Hardware, interconnects, Optical interconnects, DWDM

ACM Reference Format:
Beausoleil, R. G. 2011. Large-scale integrated photonics for high-performance interconnects. ACM J. Emerg. Technol. Comput. Syst. 7, 2, Article 6 (May 2011), 54 pages.
DOI = 10.1145/1970406.1970408 http://doi.acm.org/10.1145/1970406.1970408

1. INTRODUCTION

In 1965, Gordon Moore claimed in an article entitled "Cramming More Components onto Integrated Circuits" that by 1975 the economics of semiconductors could allow as many as 65,000 electronic elements to be incorporated onto a single silicon chip [Moore 1965]. His general observation that the number of components on semiconductor dies had doubled every year since 1959 became known as "Moore's Law," and has remained largely true over the last 50 years. Annual sales for the global semiconductor industry

Author's address: R. G. Beausoleil, Information and Quantum Systems, HP Laboratories, 1501 Page Mill Rd., Palo Alto, CA 94304-1123; email: ray.beausoleil@hp.com.

DOI 10.1145/1970406.1970408 http://doi.acm.org/10.1145/1970406.1970408

have grown to US$300B in 2010 [Gartner 2010], and public awareness of Moore's Law has established expectations that the computing power per unit cost of commercially available information technology will continue to double every 18 months over many years to come.

Until recently, the computer processor industry has been able to concentrate on improving the performance of single computer cores designed to support single-threaded architectures and software applications. As the feature sizes of transistors have decreased over the last decade, however, physical limitations such as leakage current and temperature, as well as the sheer cost of large-scale integration and fabrication, have forced electrical engineers and computer scientists to implement parallel processors consisting of multiple cores. The number of cores is likely to increase exponentially over the next 10–20 years, but the complexity and clock rate of individual cores is unlikely to change significantly [Muller 2005; Colwell 2007]. The result will be highly multi-threaded architectures that must access code and data that may reside in a non-local cache or off-die in main memory, requiring a high-performance interconnect with high bandwidth and low latency to allow efficient programming.

Without a breakthrough in interconnect technology, the performance characteristics of multicore processors are likely to be severely limited by the physical scaling characteristics of wires [Kumar et al. 2005]. As the size of an individual core shrinks, the dimensions of the wires in the core will contract as well, and are unlikely to affect the performance of the processor [Ho et al. 2001; Ho 2003]. However, the wires in the upper metal layers that form the global interconnect will still need to span the die to distribute power and clock signals, and must be used to connect cores to each other and (through the data pins in the processor package) to main memory. The information-carrying capacity of a wire with cross-sectional area A and length L is proportional to A/L^2 [Miller 2000], so reducing the area of wires to increase interconnection density is less effective then simply adding more layers of wires. Longer wires are typically divided into short segments and connected through vias to repeaters, at an increased cost in energy and latency. If the repeaters are designed to maintain bandwidth and minimize latency across the die [Banerjee and Mehrotra 2002], then the corresponding power cost can become an appreciable fraction of the total dissipation of the chip [Owens et al. 2007]. The number of data pins available for off-chip communication is increasing very slowly [ITRS 2009], and as the data rate per pin increases the data modulation frequency is approaching the electrical carrier frequency, creating significant signal integrity issues. All of these limits have made it difficult to reduce the energy consumed per transmitted bit in high-bandwidth on-chip and off-chip global interconnects below 2 picojoules (pJ) [Poulton et al. 2007].

When faced with bandwidth, power, and signal integrity issues on long-haul networks, the telecommunications industry adopted and extended high-capacity fiber-optic technology and wavelength-division multiplexing (WDM) based on lasers operating near 1550nm [Gnauck et al. 2008; Essiambre et al. 2008]. As similar constraints have overtaken large-scale computer performance, active optical cables based on vertical-cavity surface-emitting lasers (VCSELs) have been introduced into data centers and supercomputers [Narayan 2009; Henning and White 2009; Offrein and Pepeljugoski 2009; Pepeljugoski et al. 2010], and board-level intermodular interconnects based on VCSELs have entered the advanced development phase [Doany et al. 2009; Tan et al. 2009a, 2009b].

As interconnect bandwidth requirements scale to 10 terabits/second (Tb/s) and beyond, higher levels of integration will be required to increase bandwidth per pin (or optical connector) at significantly reduced cost. Over the last decade, remarkable progress has been made in research on low-power silicon photonic devices for interconnect applications [Lipson 2005; Jalali and Fathpour 2006; Soref 2006], and complementary

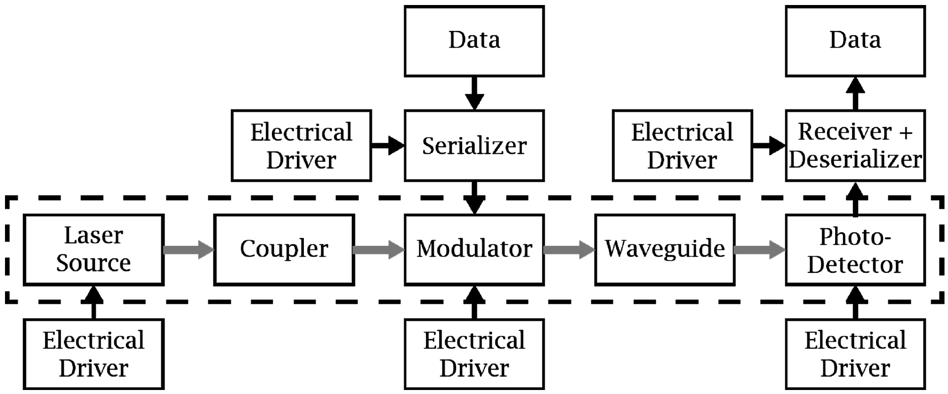

Fig. 1. Simplified schematic of the major electrical and optical components in an on-chip data link. Optical components have been drawn within the dashed rectangle, and are connected to each other by light-gray arrows indicating optical signals; electrical components are connected to each other and to optical components by black arrows indicating electrical signals in wires. An accurate accounting of energy dissipation in this system should include the contributions of all electrical drivers. Not shown are supporting system functions like clocking.

metal-oxide-semiconductor (CMOS) fabrication technologies promise a "Moore's Law for photonics" [Lipson 2009; Liang and Bowers 2009; Yoo 2009] that would probably not arise from telecommunications supply and demand alone. Optical communication signals have carrier frequencies of approximately 200 terahertz (THz), orders of magnitude larger than modulation frequencies likely to be used in short-haul interconnects ($\sim$10GHz), avoiding the signal integrity limitations of electronic interconnects. Propagation losses in optical fibers over data center distance scales are negligible, and in on-chip multimode silicon waveguides they can be reduced to much less than 1 decibel/centimeter (dB/cm) [Fischer et al. 1996], allowing the elimination of repeaters. Therefore, overall network latency is essentially determined by the propagation speed of light in the fiber or waveguide, and network energy consumption becomes independent of distance. On the processor die, waveguides with submicron cross sections can be lithographically defined at pitches of 2–4 microns (μm), and when used in a dense wavelength division multiplexing (DWDM) system based on 64 channels operating at 10 gigabits per second (Gb/s), the bandwidth density of 160–320Gb/s/μm exceeds that of electrical wire interconnects by more than a factor of 100. With these potential breakthroughs in mind, computer architects have begun studying the impact that photonic interconnects could have on core-to-core interconnects [Kirman et al. 2007; Bergman and Carloni 2008; Vantrease et al. 2008; Lee et al. 2010] and processor-to-memory interconnects [Batten et al. 2009], with the principal goal of significantly reducing the degree to which programming models must rely on the locality of code and data.

In this survey, photonic technologies amenable to large-scale CMOS integration are reviewed from the perspective of high-performance interconnects over distance scales of 1mm to 100m, beginning with a brief overview in Section 2 of the requirements placed on integrated optical devices by a variety of modern computer applications. A very simplified schematic diagram of a point-to-point optical link is shown in Figure 1, with the optical components drawn within the dashed rectangle and connected to one another by light-gray arrows. Several passive silicon photonic devices (such as couplers and waveguides) designed to guide and filter light in the telecommunication bands are introduced in Section 3, and the current state of the art in CMOS-compatible planar modulators and photodetectors shown in Figure 1 is presented in Section 4. Another important class of electro-optically active devices that are critical for the operation of a

Fig. 2. The performance of the top supercomputers in May 2010. (Figure reproduced from TOP500 [2010] with permission © 2010 TOP500.org.)

low-cost, low-power optical network is an efficient laser, and current research in CMOS-compatible, III-V, and hybrid light sources is evaluated in Section 5. Finally, prospects for integration are evaluated in Section 6, with an emphasis on silicon-on-insulator (SOI) as a platform for photonics.

2. INTERCONNECT REQUIREMENTS FOR HIGH-PERFORMANCE COMPUTING

2.1. Supercomputers and Data Centers

As shown in Figure 2, top supercomputer performance has tracked Moore's Law quite closely [TOP500 2010], with the June 2010 record holder (the Jaguar XT5 system built by Cray at the National Center for Computational Sciences at the Department of Energy's Oak Ridge National Laboratory [ORNL 2010]) relying on 18,688 compute nodes to demonstrate 1.75 petaflops running the LINPACK benchmark. Each node contains 16GB of DDR-800 memory and two hex-core AMD Opteron 2435 (Istanbul) processors running at a clock speed of 2.6GHz, and is connected to the network through a SeaStar2+ router with a peak bandwidth of nearly 60GB/s, or approximately half a byte per flop [Bland et al. 2009]. Jaguar has a scalable I/O network using almost 5km of Zarlink active optical cables (in lengths up to 60m) [Zarlink 2010] to connect over 3,000 Infiniband ports distributed over the two-story facility. The previous record holder was the Roadrunner supercomputer built by IBM at Los Alamos National Laboratory [LANL 2010; Narayan 2009; Henning and White 2009; Offrein and Pepeljugoski 2009]. Roadrunner also uses active optical cables, incorporating VCSELs, large-area photodetectors, and 92km of multimode optical fiber [Pepeljugoski et al. 2010].

If an exascale supercomputer were to be built using the same number of compute nodes as Jaguar or Roadrunner, then progress in processor technology equivalent to ten Moore's Law generations (or 15 years) will be needed. But the power requirements of such a system could be excessive: if we simply scaled the performance per unit power of the current leader in efficient computing ($\sim$0.8 gigaflops/Watt [Green500 2010]), an exaflop-class machine would require over a gigawatt of continuous power. A major challenge for the large-scale computer industry will be to improve this specification by a factor of 25 to approximately 20 gigaflops/Watt, enabling an exascale system to operate with 50 megawatts of aggregate electrical power. If we assume that approximately 40% of this energy consumption is used by the transistors in the cores and main memory, and another 40–50% by the hard drives and cooling system, this leaves only 10–20% of the system power for the interconnect at all levels. It is clear that optics will be needed to meet this requirement [Benner 2009; Pepeljugoski et al. 2010], and it is likely that silicon photonics will become necessary to lower costs and enable DWDM for increased bandwidth per connector.

While it is difficult to predict the demand for exaflop-class supercomputers in the next decade, market trends already indicate that many data centers with exascale storage and communication capacity will be required during the same time frame [Glick 2008; Astfalk 2009; Morris 2009]. Unlike the highly parallel nature of the processing done by Jaguar or Roadrunner, data center traffic is highly random and will require interconnect bandwidths at least as high as that of a supercomputing system. In 2007, data equivalent to 650,000 copies of the entire contents of the Library of Congress (nominally 10^{13} bytes) were transmitted over the Internet every month, and at that time the five-year-forward compound annual growth rate in traffic was expected to be 46% [Astfalk 2009]. HP has announced long-term plans to instrument the earth with over one trillion sensors, each capable of generating and transmitting data at 1–1000kb/s, requiring an Internet bandwidth about 1,000 times greater than in 2007 [HP 2010]. If estimates are correct that every byte transmitted over the Internet causes $\mathcal{O}(10^6)$ bytes of data communications within at least one data center, the implications for data center interconnects at the end of this decade are dire [Astfalk 2009].

2.2. Computer and Network Architectures

Computer architects have long anticipated the limitations of electrical interconnects as electronic devices shrink [Meindl et al. 2002], and a number of proposals have appeared recently to incorporate optics at deeper levels in integrated circuits. For example, one approach advocated replacing the long global wires of a multicore processor with an optical broadcast bus [Kirman et al. 2007], and another advanced a circuit-switched photonic network that could significantly improve the performance of GPU-like chip architectures [Bergman and Carloni 2008; Lee et al. 2010]. By contrast, Figure 3 shows the physical layout of the "Corona" architecture for a processor with 256 cores divided into 64 groups, targeted for a 16nm process available in 2017 [Vantrease et al. 2008; Beausoleil et al. 2008; Ahn et al. 2009a]. Three of the four stacked physical layers, each fabricated on a separate special-purpose die [Black et al. 2006], are shown: the cores (each with an integrated L1 cache), the L2 caches (shared by each group), and the global photonic interconnect. The silicon photonic layer is an all-to-all DWDM crossbar with fiber-optic connections to DRAM, allowing high-bandwidth and low-latency communication to all memory in the system at an energy cost of approximately 290 femtojoules per bit (fJ/b). The design is enabled by some of the active photonic devices discussed in Section 4, such as silicon microring modulators, each operating at 10Gb/s (or twice the target clock frequency of 5GHz), and integrated silicon-germanium photodetectors. The extremely high degree of parallelism provided by this design is enabled by overprovisioning the active photonic devices by a factor of 64 above those in use at any given time (accounting for half of the energy consumed by the photonic layer [Ahn et al. 2009a]), placing stringent requirements on the quality of the fabrication and integration technologies described in Section 6. High-bandwidth arbitration is provided by an integrated photonic system (arguably a simple, but useful, optical computer) based on the same components used in the interconnect, with a speed that is limited only by the

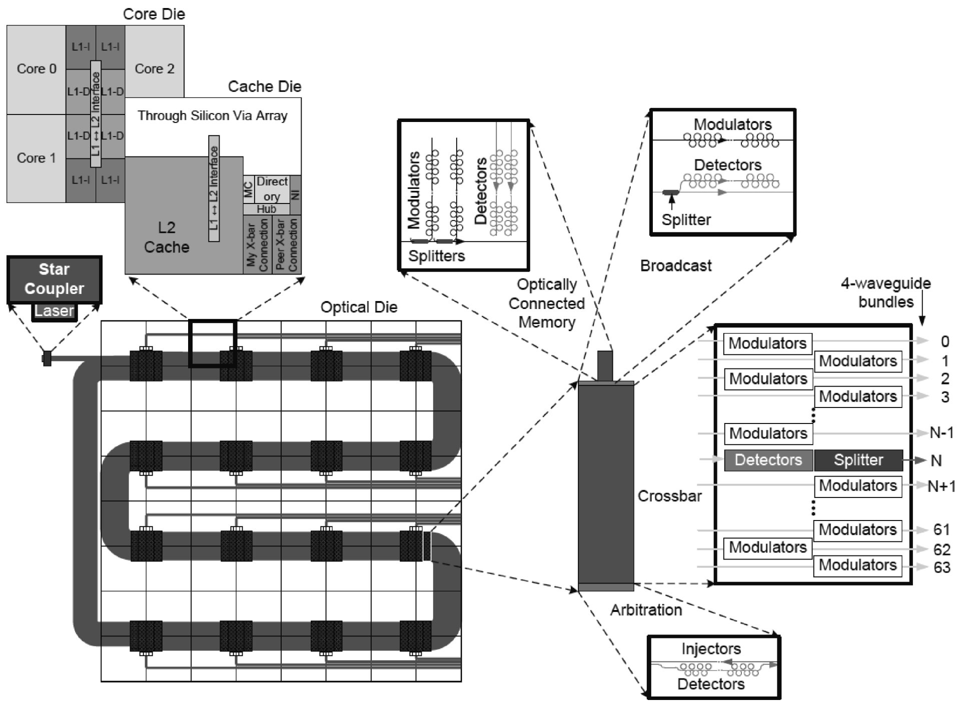

Fig. 3. Physical layout of the Corona architecture for a processor with 256 cores divided into 64 groups. Three of the four stacked physical layers—each integrated on a separate die—are shown: the cores (each with an integrated L1 cache), the L2 caches (shared by each group), and the global photonic interconnect. The L2 caches are connected via the photonics through an analog electronic driver layer (not shown). Fast, high-bandwidth arbitration is provided in the photonic layer by the same types of components used in the interconnect. (Figure reproduced from Vantrease et al. [2008] with permission © 2008 IEEE.)

time required for a light signal to propagate through a silicon waveguide once around the chip [Vantrease et al. 2008, 2009].

An alternative architecture that also provides photonic cross-chip global interconnection between groups of cores, but emphasizes processor-to-main memory communication, is the "local meshes to global switches" (LMGS) topology shown in Figure 4 [Batten et al. 2009]. The optical microring filter matrix supports a processor with 256 cores (divided into 16 groups) and 16 DRAM modules operating at a clock speed of 2.5GHz. This layout is targeted toward a 400mm^2 die fabricated using a 22nm technology, with a poly-crystalline silicon photonic layer implemented on the back of the silicon logic layer. The active and passive devices are under-etched using a novel approach to provide the optical cladding discussed in Section 3; models of this technology show great promise, but as discussed in Section 4 research on the performance of polysilicon active devices is relatively immature compared to that on pure crystalline devices. Performance simulations of the LMGS interconnect architecture predict that photonics enables almost an order of magnitude improvement in bandwidth to DRAM at similar latency and power consumption of an electrical interconnect, with an energy consumption of only 250fJ/b.

Steering the computer industry toward photonic interconnects in and between processors and memory will require a nontrivial collaboration of multinational semiconductor and information technology companies. A nearer-term realization of a photonic application-specific integrated circuit (PhASIC) that could have a substantial impact

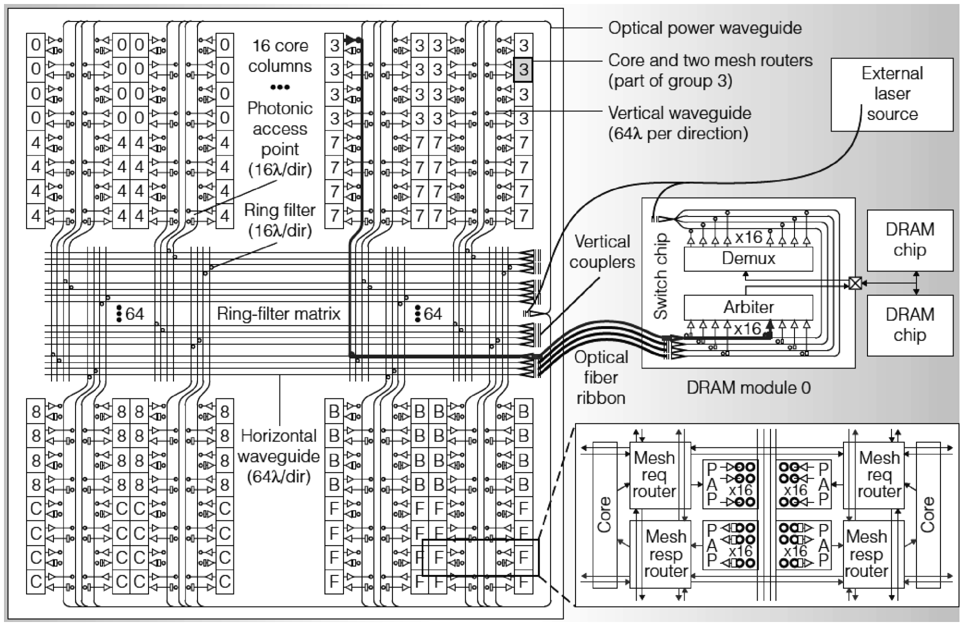

Fig. 4. Schematic of an optical microring filter matrix implementation of a local meshes to global switches topology for a processor with 256 cores (divided into 16 groups indicated on each member core with a hex number) and 16 DRAM modules. The global interconnect channel linking core group 3 to DRAM module 0 is highlighted. (Figure reproduced from Batten et al. [2009] with permission © 2009 IEEE.)

on supercomputer and data center performance may be an optically-enabled network switch. “HyperX” has been proposed as a network architecture (adapted from hypercube and flattened butterfly topologies) that could take advantage of the high-radix switch implementations that large-scale integrated photonics could enable [Ahn et al. 2009b]. HyperX relies on a new adaptive routing algorithm that allows a favorable balance to be found between performance, power, wiring complexity, and fault tolerance. This approach can provide performance levels equivalent to those of a folded-Clos architecture, with fewer switches, and uses photonic components to enable a low-cost packaging strategy for exascale data centers.

2.3. Device and Component Requirements

During the next decade, requirements for lower cost and power dissipation, reduced latency, smaller physical size, and higher levels of integration with mainstream silicon electronics will create inexorable pressure toward adoption of CMOS-compatible photonics. Although it is useful to use the ITRS roadmap [ITRS 2009] to extrapolate interconnect performance requirements over the next ten years [Beausoleil et al. 2008; Miller 2009], the recent leveling-off of on-chip clock speeds around 3GHz [Muller 2005] makes the long-term ITRS projections of on-chip/off-chip clock rates of 14.3/67.5GHz seem highly unlikely to be realized. Furthermore, projections of processor power scaled to the 16nm technology node predict that the total processor, cache, memory controller, and intermediate electronic interconnect of a 10-teraflops multicore CPU should dissipate between 82 and 155 watts (W) [Vantrease et al. 2008]. If we double the half-byte per flop supercomputer bandwidth discussed above to accommodate random data center traffic, then the bidirectional on-chip and off-chip communication bandwidth should be 20 terabytes/second (TB/s). Therefore, an interconnect that consumes no more than

30% of the power used by the cores (consistent with ITRS interconnect projections) should operate with an energy cost of 150–300fJ/b.

This tight constraint on dissipated power discourages the use of serialization and deserialization (SERDES), so at best a particular physical channel may operate at twice the clock frequency (i.e., read or write two bits per clock to/from digital registers), or 10Gb/s if it is assumed that on-chip and off-chip clock rates grow to 5 GHz over the next decade. Therefore, 16,000 point-to-point physical channels will be needed to connect cores to each other and to memory with a total bidirectional bandwidth of 20TB/s, and even more would be needed for an all-to-all crossbar [Ahn et al. 2009a]. In principle, a point-to-point interconnect could be implemented in a tall package using a free-space optical interconnect built with arrays of VCSELs, photodetectors, and MEMS-actuated mirrors [Xue et al. 2008], but (assuming that the packaging issues can be solved using the wafer-bonding techniques described in Section 5) different interconnect technologies would still be required for interconnects to memory and other off-chip resources. Although over the next several years integrated single-channel photonics or multilaser coarse wavelength-division multiplexing (CWDM) could be used to communicate between packages, in the long term the use of DWDM in optical interconnects is inevitable, and would allow the 20 TB/s bandwidth requirement to be met with a single photonic global interconnect layer using 64-wavelength DWDM [Xu et al. 2006; Vantrease et al. 2008; Beausoleil et al. 2008; Ahn et al. 2009a; Manipatruni et al. 2010b].

Based on these observations, in the following sections the current state of the art of CMOS-compatible photonic devices is evaluated, and their future prospects for monolithic integration are assessed from the perspective of the high-bandwidth interconnect requirements developed thus far. Key components shown in Figure 1 include:

Low-loss SOI waveguides. Waveguides fabricated on SOI using standard photolithographic patterning and reactive ion etching techniques have measured losses as low as 0.2dB/cm [Liu et al. 2004b], and there are indications that losses as low as 0.1 dB/cm are within reach [Dong et al. 2010c]. However, as discussed in Section 6, the commercially available SOI wafers used for this purpose today have top silicon layers with poorly controlled thicknesses, and result in fabricated devices that require “trimming” into compliance.

Low-power germanium photodetectors. Germanium has a smaller bandgap than silicon (0.7eV compared to 1.1eV), allowing detection of optical signals in the International Telecommunication Union (ITU) wavelength bands near 1310nm (O-band) or 1550nm (C-band). As described in Section 4, monolithically integrated Ge-on-Si photodetectors offer the possibility of low-capacitance detectors that eliminate the need for power-hungry amplifiers and clock recovery to build a nearly “receiverless” detection scheme [Bhatnagar et al. 2004] for data rates (per channel) of 5Gb/s. At higher data rates, the gain of the transimpedance amplifier can be modest, and the clock signal can be shared by detector arrays in a WDM system.

Resonant modulators. In DWDM systems, the ring resonators described in Section 3 selectively modulate a single wavelength on a given waveguide, and can be moved to an “OFF” state where they are transparent to the data flux in the waveguide. Using methods outlined in Section 4, modulators work by changing the index of silicon using carrier injection, depletion, or accumulation. Published results indicate that the target modulation rate of 10Gb/s with energy consumptions as low as a few femtojoules per bit can be achieved with CMOS technology [Xu et al. 2007; Manipatruni et al. 2007b; Watts et al. 2008a; Zortman et al. 2010a]. These devices will need to be kept resonant with the chosen wavelength by thermal and/or current tuning of their refractive indices. Since the incident power needed to shift the resonance of a typical compact silicon microring

Table I.
Sellmeier coefficients for silicon in Eq. (2), valid over the ranges 20K $\leq T \leq$ 300K and 1.1μm $\leq \lambda \leq$ 5.6μm. Units are K^{-j}.

j	S_{1j}	S_{2j}	S_{3j}
0	1.04907E+01	−1.34661E+03	4.42827E+07
1	−2.08020E−04	2.91664E+01	−1.76213E+06
2	4.21694E−06	−2.78724E−01	−7.61575E+04
3	−5.82298E−09	1.05939E−03	6.78414E+02
4	3.44688E−12	−1.35089E−06	1.03243E+02

j	Λ_{1j}	Λ_{2j}	Λ_{3j}
0	2.99713E−01	−3.51710E+03	1.71400E+06
1	−1.14234E−05	4.23892E+01	−1.44984E+05
2	1.67134E−07	−3.57957E−01	−6.90744E+03
3	−2.51049E−10	1.17504E−03	−3.93699E+01
4	2.32484E−14	−1.13212E−06	2.35770E+01

resonator [Xu et al. 2008] by one linewidth is approximately 50mW [Soljačić et al. 2002; Leuthold et al. 2010], an aggregate optical input power of ~3W is the upper limit for a single-waveguide 64-channel DWDM system. Careful consideration of the number of channels that can fit within a single ring's "free spectral range" is another critical system parameter [Manipatruni et al. 2010b].

Multiwavelength lasers. Multichannel light sources for low-cost DWDM systems with precisely controlled and locked frequency intervals are treated in Section 5. One of the possible approaches is the Fabry-Perot comb laser based on quantum dots [Kovsh et al. 2007], which has already been used to demonstrate a bit-error-rate of 10^{-13} at 10Gb/s over ten longitudinal modes [Gubenko et al. 2007; Innolume 2010]. Another possible approach is the mode-locked hybrid Si/III-V evanescent laser [Koch et al. 2007], which uses a silicon-waveguide laser cavity wafer-bonded to a III-V gain region. The laser need only provide 1–2W of total optical power over 64 wavelength channels to supply a 20TB/s network if the detector capacitance is low enough that only 10,000 photons are needed to drive a 1V swing at the detector's output terminal [Vantrease et al. 2008; Ahn et al. 2009a].

3. PASSIVE SILICON PHOTONIC DEVICES: WAVEGUIDES, FILTERS, AND COUPLERS

Silicon is transparent at wavelengths longer than 1.1μm (corresponding to energies below the bandgap of 1.12eV), and has a refractive index with a dependence on wavelength and temperature that can be modeled by a Sellmeier equation as [Frey et al. 2006]

$$n^2(\lambda, T) - 1 = \sum_{i=1}^{3} \frac{s_i(T)\lambda^2}{\lambda^2 - \lambda_i^2}, \tag{1}$$

where $\lambda = c/f$ is the *vacuum* wavelength in microns of an electromagnetic field with optical frequency f, c is the speed of light, T is the temperature in degrees Kelvin,

$$s_i(T) \equiv \sum_{j=0}^{4} S_{ij}\, T^j, \text{ and} \tag{2a}$$

$$\lambda_i(T) \equiv \sum_{j=0}^{4} \Lambda_{ij}\, T^j. \tag{2b}$$

The Sellmeier coefficients S_{ij} and Λ_{ij} for silicon are given in Table I [Frey et al. 2006]. Generally, the form of Equation (1) and Equation (2) are chosen to allow

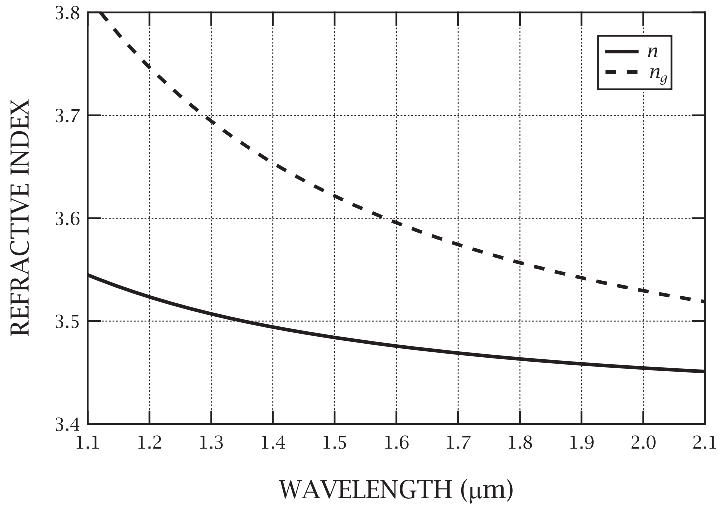

Fig. 5. Plot of the phase and group refractive indices of silicon as a function of vacuum wavelength at a temperature of 300K.

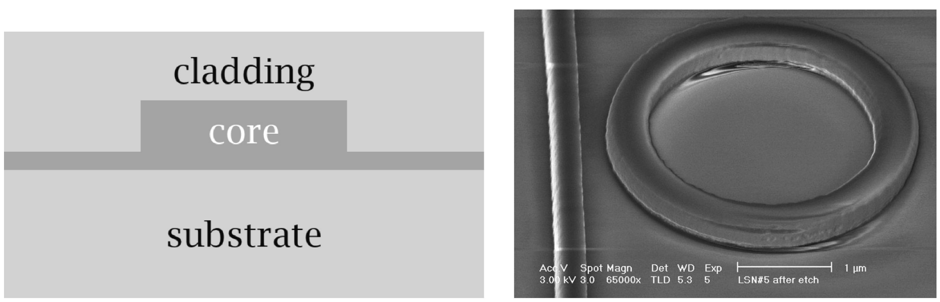

(a) Cross-section of a ridge waveguide (b) Small microring resonator on SOI

Fig. 6. High-index-contrast step-index waveguides for small-footprint large-scale integration. (a) Cross-section of a step-index ridge waveguide. In SOI, the substrate is SiO_2 (on silicon) and has a refractive index of ~1.5 in the telecom band, and the Si top layer has a material index of 3.5. The cladding layer could be air or a passivating material similar to SiO_2. (b) An SEM picture of a silicon wire microring resonator with 1.5μm radius adjacent to a tapered bus waveguide. The field circulating in the ring is coupled to the bus waveguide through evanescent field leakage from the outer surface of the ring. (Figure 6(b) reproduced from Xu et al. [2008] with permission © 2008 OSA.)

an accurate fit to available optical data, and allow computation of the dispersion $(dn(\lambda, T)/d\lambda)$, the group refractive index $(n_g(\lambda, T) = n(\lambda, T) - \lambda\partial n(\lambda, T)/\partial\lambda)$, and the thermooptic coefficient $(dn(\lambda, T)/dT)$ at wavelengths and temperatures where direct measurements are not available. A plot of the refractive index (also as known as the phase index) and the group refractive index of silicon as a function of wavelength and temperature is shown in Figure 5.

3.1. Waveguides

The waveguides most commonly used in large-scale silicon integrated photonics are variations of the dielectric slab waveguide [Hu and Menyuk 2009] that is now commonly treated in textbooks [Okamoto 2006]. As shown in Figure 6(a), the dielectric ridge waveguide has a light-guiding layer with a high-index core that is defined by a ridge atop a pedestal, situated between lower-index substrate and cladding layers

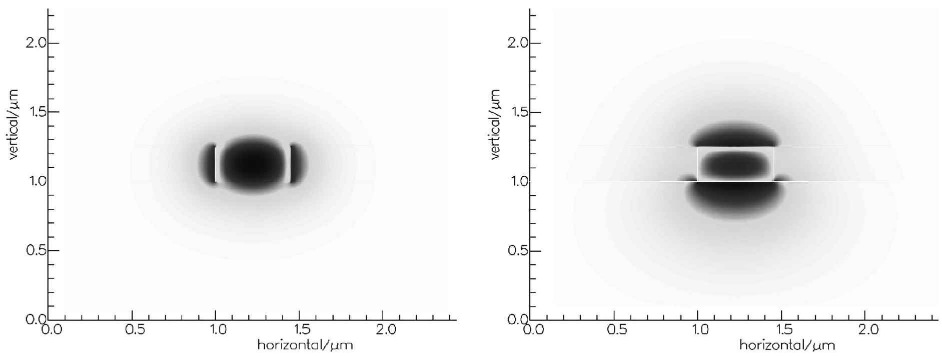

(a) TE mode (b) TM mode

Fig. 7. Simulations of electric field mode profiles for a silicon wire waveguide with a width of 450nm and a height of 250nm, with an SiO_2 substrate and no cladding (air). In the case of the TE mode, the electric field is polarized horizontally, and the effective refractive index is $n = 2.385$, for the TM mode, the magnetic field is polarized horizontally, and $n = 1.772$. (Courtesy of Di Liang.)

[Lipson 2005]. When the thickness of the pedestal is reduced to zero, the resulting structure is known as a "wire" waveguide (or a "channel" waveguide when the refractive index of the substrate and cladding are the same) as shown in Figure 6(b). In SOI, the substrate is SiO_2 (on silicon) and has a material refractive index of $\sim$1.5 in the telecom band, and the Si top layer has a material index of $\sim$3.5. The cladding layer could be air or a passivating material similar to SiO_2. This high index contrast between the materials in the core and the substrate/cladding creates a strong confinement of the propagating field within the silicon, with exponentially decaying evanescent fields having extents of $\sim$100–300nm at the index steps surrounding the core. Physically, this process can be understood in terms of either total internal reflection in a ray-tracing model or energy concentration in high-index regions in an electromagnetic boundary-value problem.

Propagating modes in optical waveguides are characterized by their transverse field profiles and the polarization (orientation) of their electric and magnetic fields. A transverse electric (TE) mode has an electric field vector aligned primarily perpendicular to the z (propagation) axis, while (depending on the waveguide composition and geometry) the magnetic field may have a component oriented along the z-axis. By contrast, a transverse magnetic (TM) mode has a magnetic field vector aligned perpendicular to the z-axis. (Generally, there are no propagating modes in purely dielectric ridge or wire waveguide that are exactly TE or TM, so the corresponding field polarizations are often designated as "quasi-TE" and "quasi-TM," respectively [Mashanovich et al. 2008].) Transverse core dimensions of 4–8μm enable multi-transverse-mode fields with relatively low propagating losses of $\sim$0.1dB/cm [Fischer et al. 1996], while the wire waveguides shown in Figure 7 that are nominally 500nm wide by 250nm tall operate in a single transverse mode with typical propagation losses of 1dB/cm [Gnan et al. 2008] when fabricated using e-beam lithography (EBL). The phase velocity of light in these waveguides is c/n, where n is an *effective* refractive index that generally has a value between the material indices of the core and cladding; the corresponding group velocity is c/n_g, where n_g is an effective group refractive index that typically has a value that is $2\times$ that of the phase index. Recent demonstrations of large-core waveguides with lateral dimensions of $\sim$1.5μm fabricated in SOI using standard photolithography and etching techniques have demonstrated average propagation losses as low as 0.2dB/cm for quasi-TE modes [Liu et al. 2004b], and there are indications that 0.1dB losses in shallow-ridge waveguides with widths of 2μm are within reach [Dong et al. 2010c].

A number of technologies have been proposed and demonstrated to complete the set of optical components amenable to large-scale integration on-chip with CMOS compatible processes. Polarization rotators are commonly used in fiber-optic networks, and an offset dual-core silicon channel waveguide structure with a rotation element with a width of $\sim 2\mu$m and a length of 35μm has been demonstrated with a rotation angle of 72° and a polarization extinction ratio of 11dB [Fukuda et al. 2008], and a design based on coupling through an intermediate multimode waveguide with similar dimensions has been proposed that promises 90° rotation with an extinction ratio as high as 17dB [Yue et al. 2009]. Although none of the architectures described in Section 2 require purely passive crossed waveguides, as optical circuit element density grows, waveguide intersections in the silicon plane may become necessary. SOI waveguides with crossings showing losses of -0.16dB and cross-talk of -40dB have been demonstrated, based on a design incorporating parabolically broadened waveguides to reduce lateral refractive index contrast while maintaining high lateral mode confinement [Bogaerts et al. 2007]. Models show that similar results could be obtained with more conventional planar patterning and etching [Li et al. 2009]. However, an intriguing new approach to CMOS-compatible waveguides relies on refractive index engineering using subwavelength gratings rather than ridge or wire waveguides. This approach has been used to demonstrate index-guided waveguides fabricated in a single etch step with propagation losses as low as 2.1dB/cm with a low group index of $n_g \approx 1.5$ over the entire telecom C-band [Bock et al. 2010a]. When crossed waveguides are fabricated using the same approach, the loss per crossing was measured to be as low as -0.02dB with crosstalk less than -40dB [Bock et al. 2010b]. Other compact planar photonic devices such as input couplers and multiplexor circuits have also been implemented with performance characteristics comparable to those of ridge and wire waveguides [Cheben et al. 2010].

One of the principal contributions to the propagation and scattering losses of single-mode high-index-contrast (HIC) channel and ridge waveguides is the roughness of the sidewalls induced by patterning and etching processes, causing the observed loss rates of $\sim$1dB/cm that are two orders of magnitude higher than those of low-index-contrast waveguides such as optical fiber. This optical loss is not simply incoherent scattering from random surface perturbations into the substrate or cladding; rather, a nontrivial fraction of the scattered power has been found to couple to backward-propagating modes in HIC waveguides [Morichetti et al. 2010a, Morichetti et al. 2010b]. This backscattering can be extensive even in single-mode waveguides only a few hundred micrometers in length, creating a possible barrier to some applications of high-density photonic integrated circuits. A model of the scattering process is shown in Figure 8(a) for a waveguide of length L_w and nominal width w with rough sidewall surfaces. The complex backscattering amplitude $h_R(z)$ for a waveguide segment of length dz depends on the average loss coefficient $\alpha \cong (\pi/\lambda_B)(\partial n/\partial w)\delta w$, where λ_B is the Bragg wavelength corresponding to a simple sinusoidal corrugation in the waveguide sidewalls having amplitude δw, and n is the effective index of the waveguide. In the small perturbation regime, the backscattered power is proportional to $|\alpha|^2$, and as shown in Figure 8(b) can reach nontrivial values for TE polarizations over relatively short distances. These measurements were performed on waveguides fabricated using EBL [Gnan et al. 2008], and should be repeated for guides fabricated using high-resolution CMOS photolithography and etching technologies.

3.2. Filters

A complete toolbox of silicon-based optical components that enable low-cost, compact, and/or WDM interconnect applications should include passive multiplexors and demultiplexors, splitters/couplers, and filters. These generally fall into two classes: nonresonant devices, which are well studied and understood and tend to have comparatively

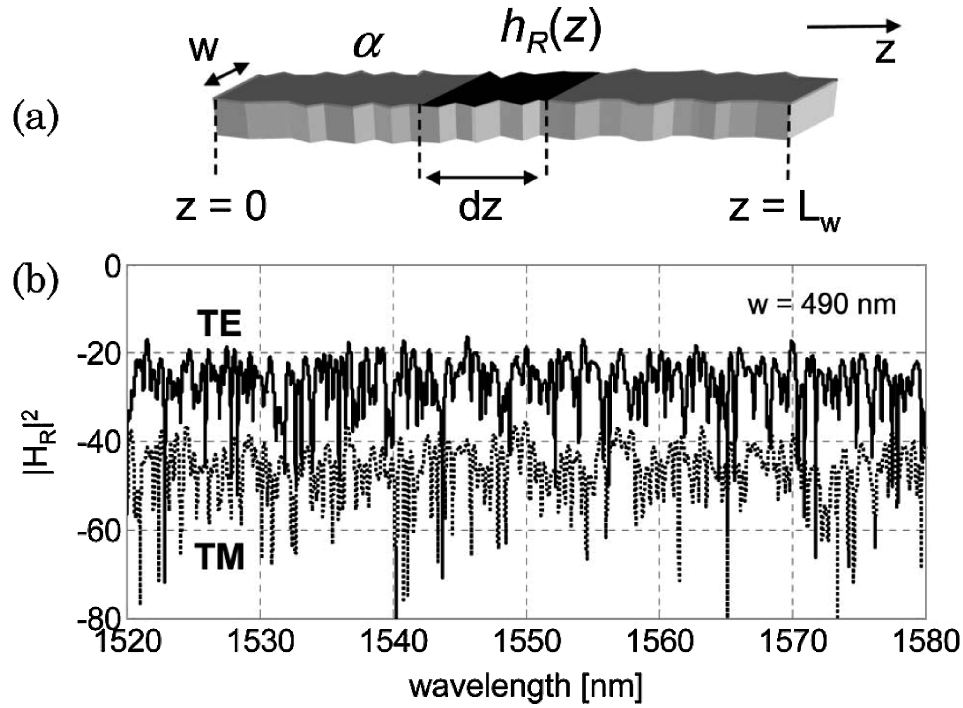

Fig. 8. Backscattering induced by sidewall roughness in high-index-contrast SOI channel waveguides. (a) Model of a waveguide of length L_w and nominal width w with rough sidewall surfaces. The complex backscattering amplitude $h_R(z)$ for a waveguide segment of length dz depends on the average loss coefficient α. (b) Measurements of the backscattered power spectral density for an SOI waveguide with L_w = 1mm and w = 490nm for TE (solid curve) and TM (dotted curve) input polarization. (Figure reproduced from Morichetti et al. [2010b] with permission © 2010 OSA.)

large footprints in silicon photonic integrated circuits; and resonant devices (such as microring cavities) which can be quite compact, but with sensitivities to environmental perturbations that present particular challenges for integration.

3.2.1. Nonresonant Couplers and Filters. The simplest nonresonant coupler, the 1×2 Y-junction beamsplitter, can be cascaded to distribute optical power over many waveguides with low loss, minimal sensitivity to wavelength, and even power distribution over the output ports [Tao et al. 2008]. SOI-based multimode interference couplers also divide power from a single input port to multiple output ports by relying on spatial interference in an intermediary slab region [Lipson 2005; Jiao et al. 2010]. Wavelength-division multiplexing applications may require that different wavelength channels be distributed into different waveguides, requiring spectral filters such as an arrayed waveguide or echelle grating [Bogaerts et al. 2010].

Perhaps the most commonly used spectral filter in CMOS-compatible WDM photonics is the silicon-wire Mach-Zehnder interferometer shown schematically in Figure 9. The narrowband input field E_{in} enters the device through a single-mode wire waveguide, is split into two equal parts by a Y-junction or multimode interference coupler, and then is recombined using the same structure to produce the output field E_{out}. The field in the upper waveguide propagates through a section of length L containing an integrated resistive heater [Liu et al. 2010a] which changes the local temperature by ΔT, causing a change in the effective refractive index $\Delta n = (dn/dT)\Delta T$, where the thermooptic

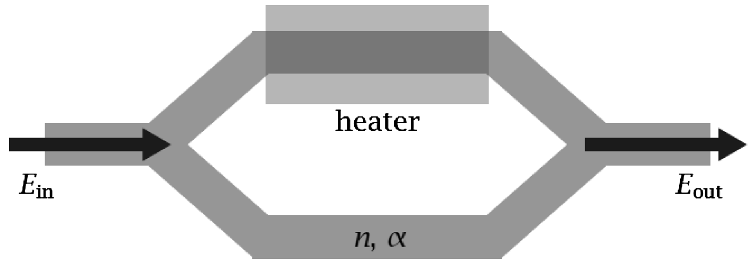

Fig. 9. Schematic of a silicon Mach-Zehnder interferometer used as a spectral filter. An input field is split by a 3-dB Y-junction or multimode interference coupler, and the mode in the upper arm is phase-shifted by an integrated resistive heater. When the beams are recombined in another 3-dB coupler, the magnitude and phase of the output field depends on the applied phase shift.

coefficient $dn/dT \approx 2 \times 10^{-4}$ K^{-1} for highly-confined modes in silicon [Frey et al. 2006]. In this case, the output field is given by

$$E_{\text{out}}(\phi) = \frac{1}{2}(1 + e^{i\phi})E_{\text{in}}, \tag{3}$$

where $\phi = 2\pi \Delta n f L/c$ and f is the optical (carrier) frequency of both fields. The output power of the interferometer is proportional to

$$|E_{\text{out}}(\phi)|^2 = \cos^2\left(\frac{\phi}{2}\right)|E_{\text{in}}|^2 . \tag{4}$$

Therefore, if ΔT is chosen so that $\phi = 2m\pi$ for some integer m (equivalent to the condition $\Delta nL = m\lambda$) the fields in both arms constructively interfere at the output coupler, and $E_{\text{out}} = E_{\text{in}}$. However, if $\phi = (2m+1)\pi$, the fields destructively interfere, and $E_{\text{out}} = 0$. There are a number of variations on this simple design; for example, the waveguides in the arms can be fabricated with different lengths to optimize filter performance at a particular wavelength for $\Delta T = 0$, but special care is then required to reduce the sensitivity of the entire device to ambient temperature [Uenuma and Moooka 2009].

3.2.2. Resonant Couplers and Filters. A particularly relevant example of a resonant filter is the two-port microring resonator shown in Figure 10(a) consisting of two bus waveguides and a ring waveguide with an effective refractive index n and power absorbtion coefficient α. The mean radius of the ring, measured at the center of the circular waveguide, is a, and the perimeter is $p = 2\pi a$. Light is weakly coupled between the curved waveguides and the nearby straight bus waveguides through the short-range evanescent fields leaking from both structures. The magnitude and phase of these couplings are determined by the dimensions of the waveguides and the separation between them, and are represented by the (ideally) broadband power reflectance $R_{1,2}$ and transmittance $T_{1,2}$ at ports 1 and 2. (For the time being, we neglect backscattering in all waveguides, but we include the effects of dissipative optical loss in the couplers by allowing $R_{1,2} + T_{1,2} \leq 1$.) The microring resonator shown in Figure 6(b) [Xu et al. 2008] is a single-port device with $R_2 = 1$ and $T_2 = 0$. A narrowband input field E_{in} produces an output field E_{out} at port 1 and a cross-coupled field E_{xc} at port 2, given by

$$E_{\text{out}}(\phi) = \frac{1}{\sqrt{R_1}} \frac{R_1 - (R_1 + T_1)\Gamma e^{i\phi}}{1 - \Gamma e^{i\phi}} E_{\text{in}}, \tag{5a}$$

$$E_{\text{xc}}(\phi) = \frac{\sqrt{T_1 T_2 e^{-\alpha p/2}}\, e^{i\phi/2}}{1 - \Gamma e^{i\phi}} E_{\text{in}}, \tag{5b}$$

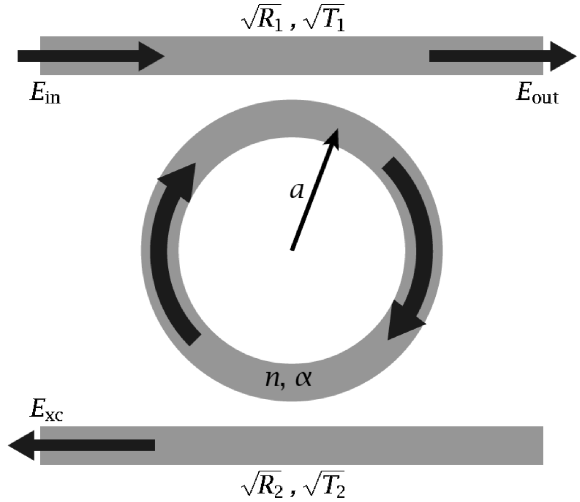

(a) Schematic of a microring resonator

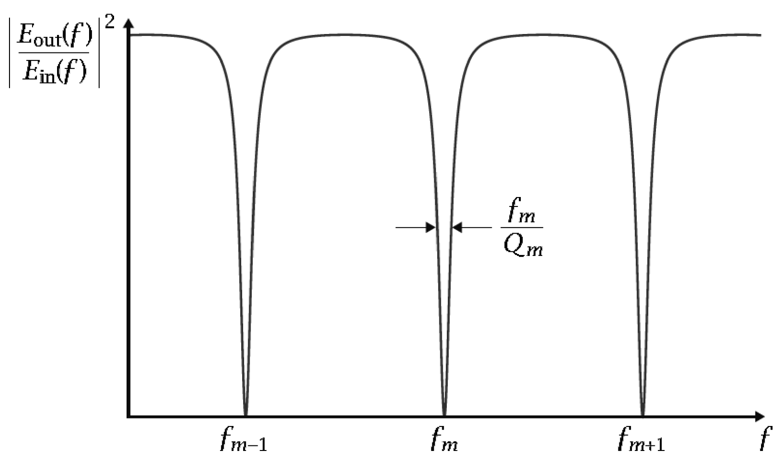

(b) Output spectrum of a microring resonator

Fig. 10. Two-port microring resonator. (a) Schematic of a two-port microring resonator consisting of two bus waveguides and a ring waveguide with an effective refractive index n and power absorbtion coefficient α. The radius of the ring—measured at the center of the curved waveguide—is a, and the perimeter is $p = 2\pi a$. The broadband power reflectance R and transmittance T at ports 1 and 2 are determined by the dimensions of the waveguides and the separation between them. An input field E_{in} produces an output field E_{out} at port 1 and a cross-coupled field E_{xc} at port 2. (b) Plot of $|E_{\text{out}}(f)/E_{\text{in}}(f)|^2$ as a function of frequency, showing axial modes with resonant frequencies $f_m = mc/np$. The width of each resonance is determined by the corresponding value of Q_m, and the separation between adjacent resonance frequencies is approximately $c/n_g p$, where n_g is the average effective group refractive index in that frequency range.

where $\phi \equiv 2\pi npf/c$ and $\Gamma^2 \equiv R_1 R_2 e^{-\alpha p}$ is the fraction of the circulating power remaining in the ring after one round trip.

When $f = f_m \equiv mc/np$ for some positive integer m, $\phi = 2m\pi$, and the microring *resonates* at this frequency: the output field reaches a minimum value, while the cross-coupled field reaches a maximum value. Since the frequency f and physical wavelength $\tilde{\lambda} = \lambda/n$ of an electromagnetic field propagating through a material with effective refractive index n satisfy the relation $\tilde{\lambda} f = c/n$, this condition may also be expressed in terms of wavelength as $m\tilde{\lambda}_m = p$; in other words, a resonance is obtained whenever an integer number of physical wavelengths fits exactly around the ring perimeter. The value of the frequency f_m will shift when the mode index m, perimeter p, or the refractive index n change. If these variations are small enough that we can ignore second-order terms in the Taylor expansion, then

$$\frac{\delta f_m}{f_m} = \frac{\delta m}{m} - \frac{\delta p}{p} - \frac{\delta n}{n}. \tag{6}$$

Now the effective refractive index $n \equiv n(f, T, q)$ may vary due to changes in the frequency itself, the temperature T, and (as discussed in Section 4.1.1) the local carrier density q, according to

$$\delta n = \frac{\partial n}{\partial f}\,\delta f + \frac{\partial n}{\partial T}\,\delta T + \frac{\partial n}{\partial q}\,\delta q, \tag{7}$$

so that the fractional shift in the resonance frequency of longitudinal mode m becomes

$$\frac{\delta f_m}{f_m} = \frac{n}{n_g}\left[\frac{\delta m}{m} - \frac{\delta p}{p}\right] - \frac{1}{n_g}\left[\frac{\partial n}{\partial T}\,\delta T + \frac{\partial n}{\partial q}\,\delta q\right], \tag{8}$$

where $n_g = n + f\,\partial n/\partial f$ is the effective group refractive index introduced earlier as a function of the vacuum wavelength. Similar considerations show that the corresponding fractional change in the *physical* wavelength is $\delta\tilde{\lambda}/\tilde{\lambda} = -(n_g/n)\,\delta f/f$. As an example of the utility of Equation (8), in the case where $\delta m = \pm 1$, $\delta p = 0$, $\delta T = 0$, and $\delta q = 0$, the corresponding frequency separation between adjacent minima in $|E_{\text{out}}|^2$ is known as the *free spectral range* and is given by $\Delta f_{\text{FSR}} = c/n_g p$. In general, careful simulations of the transverse and longitudinal eigenmodes of microring resonators will also depend on the waveguide width and height, and on the refractive indices of the substrate and cladding.

Given values of α, R_2, and any dissipative losses in port 1, if we choose the dimensions of the waveguides in port 1 so that $R_1 = (R_1 + T_1)\Gamma$, then the ring is *critically coupled*, with

$$E_{\text{out}}(\phi) = \sqrt{R_1}\,\frac{1 - e^{i\phi}}{1 - \Gamma e^{i\phi}}\,E_{\text{in}}. \tag{9}$$

In this ideal case, whenever the frequency $f = f_m$, the output power vanishes, and as shown in Figure 10(b), at frequencies midway between f_m and f_{m+1} the output power achieves a maximum value proportional to $|E_{\text{out}}|^2 = |E_{\text{in}}|^2 4R_1/(1+\Gamma)^2$. At a frequency $f = f_m + \Delta f$, where $\Delta f \ll \Delta f_{\text{FSR}}$, Equation (9) may be approximated by

$$E_{\text{out}}(\Delta f) = \sqrt{\frac{R_1}{\Gamma}}\,\frac{-i\Delta f\ \Delta f_m/2 + [2\sqrt{\Gamma}/(1+\Gamma)]\Delta f^2}{\Delta f^2 + (\Delta f_m/2)^2}\,E_{\text{in}}(f), \tag{10}$$

where $\Delta f_m \equiv f_m/Q_m$, and

$$Q_m \equiv m\pi\,\frac{\sqrt{\Gamma}}{1 - \Gamma}. \tag{11}$$

As shown in Figure 10(b), the denominator in Eq. (10) characterizes a Lorentzian resonance with a full-width at half-maximum of Δf_m.

Both the precise values of the resonance frequencies f_m and the quality factors Q_m of the ring are determined by the details of the ring geometry, the core material, and the sidewall roughness of the ring waveguide [Lipson 2005; Xu et al. 2008]. Although it is clear that both intrinsic absorption by the core material (generally negligible in silicon at telecommunication wavelengths) and the scattering by the core sidewall

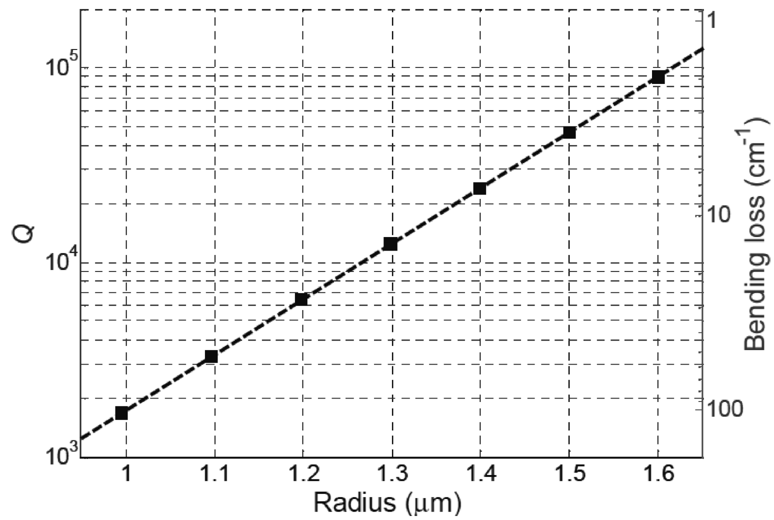

Fig. 11. Model of the intrinsic optical quality factor Q and waveguide bending loss versus the radius for a quasi-TE mode of an uncoupled microring resonator at 1.55μm computed using a 3D finite-difference time-domain simulation. Here the SOI waveguide has a rectangular cross-section with dimensions 450nm $\times$ 250nm. (Figure reproduced from Xu et al. [2008] with permission © 2008 OSA.)

roughness described above causes loss that depends on the ring circumference, there is a fundamental upper limit placed on Q by the curvature of the ring itself. As shown in Figure 7, because the boundaries of a straight channel or ridge waveguide are dielectric rather then metal in nature, the corresponding transverse mode profile has evanescent wings that extend beyond the high index core into the lower index cladding. In the case of the straight waveguide, these evanescent fields have phase profiles linked to the propagating modes within the guide, traveling at a speed c/n. In a curved waveguide, the evanescent fields on the outside of the core must travel faster than the propagating field in the core, and beyond a critical distance the required propagation speed for the phase front exceeds c/n_{cladding}, coupling the evanescent fields to modes propagating away from the ring [Lipson 2005]. As shown in Figure 11, this bending loss creates a fundamental upper limit to the intrinsic (uncoupled) Q of a microring resonator [Xu et al. 2008]. A completely uncoupled ring (i.e., with $R_1 = R_2 = 1$ and $\alpha p \ll 1$) has an intrinsic $Q_m \approx 2m\pi/\alpha p$, and the same ring when critically coupled through a single port as in Figure 6(b) with $R_1 = e^{-\alpha p}$ has a loaded quality factor that is smaller by a factor of 2.

Microring resonators can function as photonic integrated circuit elements in a variety of configurations, particularly in the case where WDM is used to encode data on a comb of multiple independent regularly-spaced frequency channels [Lipson 2005; Dong et al. 2010b; Manipatruni et al. 2010b]. In Figure 12(a), a ring tuned off-resonance allows all data channels to propagate through the bus waveguide, while in (b) a ring with a resonance frequency aligned with that of a particular data channel diverts light from that channel into the ring, where it is scattered by the rough, curved surface of the ring. As described in Section 4.1.3, electrically activating this tuning at high speed allows the microring to be used as an electro optic modulator [Xu et al. 2005; Chen et al. 2009; Lee et al. 2010]. In Figure 12(c), a ring (which need not be critically coupled) with two bus waveguides can be used to divert one (resonant) channel from one waveguide to the other, and Figure 12(d) shows that incorporating a photodetection element (such as germanium) in a ring allows it to absorb diverted light and act as

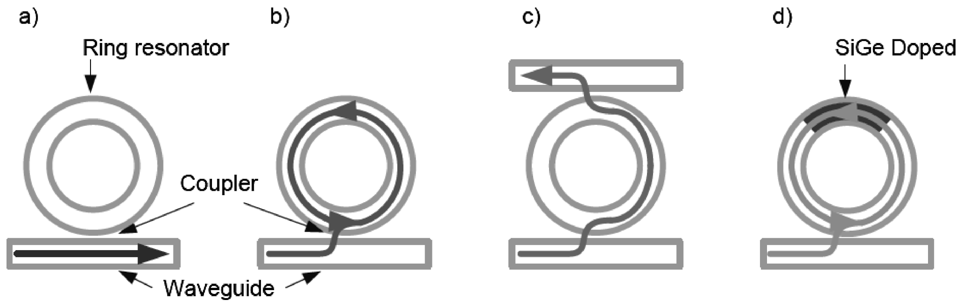

Fig. 12. Microring resonators as photonic integrated circuit elements. (a) A ring tuned off-resonance allows all data channels to propagate through the bus waveguide. (b) A ring with a resonance frequency aligned with that of a particular data channel diverts light from that channel into the ring, where it is scattered. (c) A ring coupled to two bus waveguides can be used to divert one (resonant) channel from one waveguide to the other. (d) Incorporating a photodetection element (such as germanium) in a ring allows it to absorb diverted light and act as part of a receiver. (Figure reproduced from Vantrease et al. [2008] with permission © 2008 IEEE.)

part of a receiver. Recent research has proposed that microrings can be used to provide the optical isolation that is critical to minimize the effects of coherent backscattering discussed below [Yu and Fan 2009]. Resonant structures similar to microrings, such as microdisks [Watts et al. 2008a; Zortman et al. 2010a; DeRose et al. 2010], offer similar opportunities for integrated photonics, but many of these devices have resonant higher-order transverse modes that are problematic for WDM architectures.

As bidirectional bandwidth across a single CMOS die scales to 10 TB/s and beyond [Vantrease et al. 2008], the increased circuit density enabled by compact resonant filters, modulators, and switches will become a significant driver for large-scale integrated photonics. A microring filter with a smaller circumference has a larger free spectral range, allowing a larger number of microrings resonant with different, equally-spaced data channels to be attached in a small area to a common bus waveguide carrying a higher aggregate data bandwidth [Manipatruni et al. 2010]. In the particular case of the microring modulator discussed in Section 4.1.3, the power consumed by the device is proportional to the capacitance (and therefore the circumference) and inversely proportional to the Q of the ring. The electrical capacitance of devices demonstrated in the laboratory has been reduced to a few femtofarads [Chen and Lipson 2009], lowering their energy consumption to $\sim$10fJ/b [Watts et al. 2008a; Zortman et al. 2010a]. After an SOI microring resonator with a diameter of only 4 μm was incorporated into an add/drop filter and demonstrated a $Q = 6730$ [Nawrocka et al. 2006], further reductions in device footprint required careful attention to spatial and spectral mode-matching to minimize the coupling between the fundamental transverse mode of the ring's curved waveguide and higher-order propagating modes in the straight bus waveguide [Xu et al. 2008]. By tapering the bus waveguide to a width of $\sim$290nm and increasing the gap between the the straight waveguide and the $\sim$440nm ring waveguide to $\sim$340nm, a microring resonator fabricated using EBL with a diameter of 3μm was demonstrated with a coupled $Q \approx 9000$.

In the ITU wavelength O-band and C-band, the thermooptic effect in silicon causes the refractive index of tightly confined modes to shift with temperature by approximately $dn/dT \approx 2 \times 10^{-4}\ \mathrm{K}^{-1}$. Using Equation 9, this corresponds to a temperature-dependent wavelength shift (TDWS) of a microring resonator operating near 1310nm of $\delta\lambda_m \approx$ 75pm/K—or $\delta f_m \approx$ 13GHz/K—independent of the radius of the ring. Therefore, a local temperature change of only 1K is sufficient to switch a microring coupled to a bus waveguide with a Q chosen to provide a FWHM $\Delta f_m = 20$GHz from an "off" state

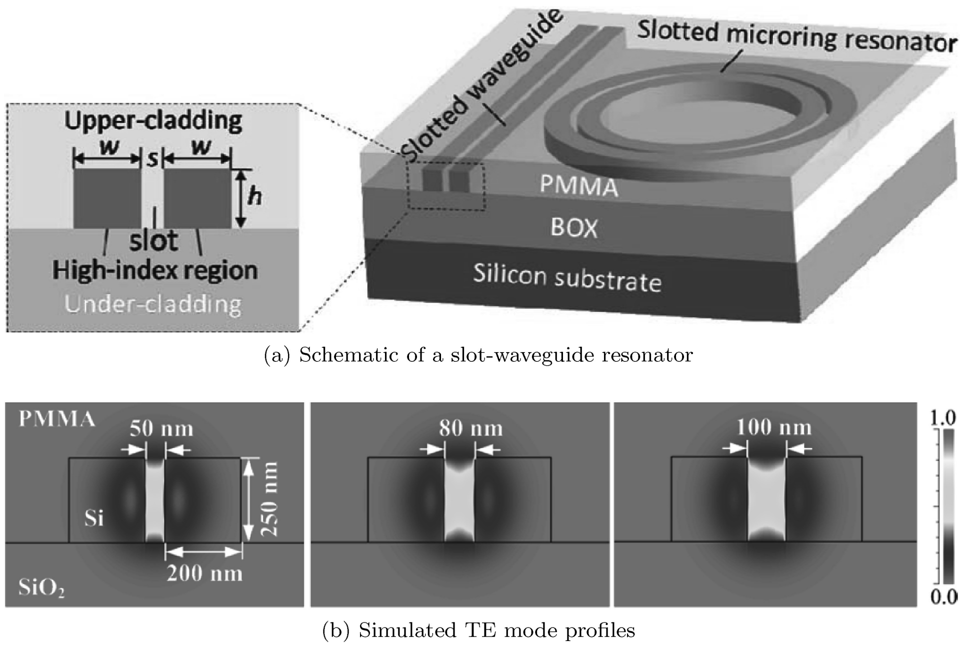

Fig. 13. Models of athermal coupled microring resonators based on slot waveguides incorporating polymers. (a) Schematic of a slotted-waveguide microring resonator with an inset showing a cross-section of the waveguide. (b) Simulated electric field profiles for TE modes in waveguides with $w = 200$nm and $h = 250$nm and varying slot width s, showing that a substantial fraction of the mode energy is carried by the slot. (Figure reproduced from Zhou et al. [2009a] with permission © 2009 by the authors.)

with $|E_{\text{out}}(f_m)/E_{\text{out}}(f_m)|^2 \rightarrow 0$ to a state where the output power is 60% of the input. Since temperature fluctuations within single-die CPU packages can exceed 40K, it is clear that the use of resonant devices in information technology applications requires techniques that mitigate thermooptic effects. One way to reduce the impact of temperature shifts on a microring resonator is to couple it to another optical system (such as a Mach-Zehnder interferometer [Guha et al. 2010]) that has an effective optical path length nL designed to cancel changes in that of the ring, at the expense of an increase in the aggregate size of the device. Alternatively, the device footprint can be retained by changing the structure of the ring and bus waveguides to include a high-confinement low-index slot waveguide [Lipson 2005] that can be filled with a material (such as a polymer) that has a value of $dn/dT < 0$ [Lee et al. 2008; Zhou et al. 2009a]. Figure 13(a) presents a schematic of a slotted-waveguide microring resonator with an inset showing a cross-section of the waveguide, while Figure 13(b) illustrates simulated electric field profiles for TE modes in waveguides with $w = 200$nm and $h = 250$nm and varying slot width s, showing that a substantial fraction of the mode energy is carried by the slot. Careful choice of the polymer and design of the slots and the coupling junction between the ring and bus waveguides results in a cancelation of the TDWS [Lee et al. 2008], but the high field intensity at the slot sidewalls shown in Figure 13(b) causes higher levels of surface scattering and therefore lower cavity Q [Zhou et al. 2009a]. Recently, an intermediate approach that reduces both the level of confinement and the effective index of polymer-cladded silicon wire waveguides has shown a reduction of the TDWS to 0.5pm/K for TM modes [Raghunathan et al. 2010]. The reduced confinement allows a greater fraction of the transverse mode profile to extend into the polymer, but increases the footprint of the waveguides and resonators. One design constraint common to all of these approaches is the need to tailor the waveguide dimensions to provide the same confinement factor for each wavelength, complicating the design of DWDM systems.

In an integrated photonic DWDM system, the frequency channels propagate along a common bus waveguide, interacting with devices such as modulators, filters, and detectors resonant with particular channels. For example, as described in Section 4.1.3, a microring resonator designed to act as a modulator for frequency channel m will create a resonant loss at frequency f_m according to Equation (5a) with a magnitude that depends on the values of R_1, T_1, Γ, and on the value of the refractive index of the ring. (However, recent measurements of coherent backscattering by microrings resonant with an incident field show that a significant fraction of the measured round-trip loss in cavities fabricated using EBL can be caused by counterpropagating fields caused by sidewall roughness of the ring waveguide [Morichetti et al. 2010]. This effect will need to be studied for the case of critically coupled compact rings built using photolithography and CMOS materials processing.) In addition, there will be a quasi-resonant loss induced by microrings designed to operate at channels $m - 1$ and $m + 1$. Consider a critically-coupled high-Q ring with $R_1 \approx \Gamma \to 1$ and $\Delta f_m = 20$GHz. If the channel spacing is 80GHz, then by Equation (10) the total power loss in channel m due to the two rings resonant with the adjacent channels is $2 \times 10 \log_{10} |E_{\text{out}}(f)/E_{\text{in}}(f)|^2 \approx -0.13$dB. Rings more than a few channels away from m cause nonresonant losses that are not predicted very accurately by Equation (10), because they arise from interactions of the evanescent fields of the bus waveguide and curved waveguides comprising the rings. If surface roughness is neglected, finite-difference time-domain models of the compact microring cavity shown in Figure 6(b) estimate that the nonresonant scattering loss along the bus waveguide due to the presence of a nearby curved waveguide is less than -0.001dB per ring for both TE and TM modes.

As both physical component and channel densities increase to accommodate the bandwidths demanded by Moore's Law, the Lorentzian lineshape shown in Figure 10(b) may not provide the out-of-band rejection ratios needed for adequate bit-error rates. Instead, resonant filters incorporating multiple coupled microrings can provide higher-order filter lineshapes without the complexity typical of thin-film filters [Little et al. 2004; Xiao et al. 2007]. Purely passive, low-loss thermally tuned filters incorporating up to 11 microrings have been fabricated using chemical vapor deposition in Hydex with an index contrast of 17% (corresponding to a core index of 1.70 and a cladding index of 1.45) [Little et al. 2004], and a third-order silicon microring add-drop filter fabricated in SOI has been demonstrated with low loss, a passband of 125GHz, and out-of-band signal rejection of 40dB [Xiao et al. 2007]. More complex resonator configurations, such as microrings embedded within compact racetracks [Zhang et al. 2008a], allow frequency responses tailored to particular applications.

4. ACTIVE SILICON PHOTONIC DEVICES: MODULATORS AND PHOTODETECTORS

4.1. Integrated CMOS-Compatible Modulators

The promise of large-scale integrated photonic interconnect links based on CMOS-compatible devices has inspired an explosion of research and development on silicon-based electro-optic modulators. A recent review of this progress by Reed et al. [2010] is quite complete and includes a table of device performance results.

4.1.1. The Physics of Silicon-Based Switches and Modulators. Optoelectronic modulators generally encode data on propagating light channels by changing the amplitude, phase, or polarization of a continuous optical field through variations in the complex propagation vector of that field. In telecommunications systems, most modulators are based on semiconductor linear and nonlinear optical mechanisms such as the Franz-Keldysh, Pockels, and Kerr effects. However, in single-crystal silicon, the magnitude of these effects at the telecommunication wavelengths near 1310nm and 1550nm are too small to allow energy-efficient operation of high-performance modulators [Soref and Bennett 1987]. The most promising candidates for silicon optical modulators rely on the free-carrier plasma-dispersion effect to change the silicon refractive index by varying electron and hole densities within the optical waveguide, or the quantum-confined Stark effect to vary the refractive index [Reed et al. 2010]. The thermo-optic effect in silicon is appreciable, but it is too slow to be used effectively for high-bandwidth modulation [Lipson 2005].

Soref and Bennett used a numerical Kramers-Kronig approach to analyze absorption data in pure crystalline silicon over widely varying electron and hole densities and computed the corresponding perturbations in the refractive index and absorption coefficients at telecommunications wavelengths [Soref and Bennett 1987]. For example, at an optical wavelength $\lambda \cong 1310$nm, their results can be summarized by the formulas [Gardes et al. 2008; Reed et al. 2010]

$$\Delta n_{\mathrm{Si}}(1310\mathrm{nm}) = -6.2 \times 10^{-22}\mathrm{cm}^3\, N_e - 6.0 \times 10^{-18}\mathrm{cm}^{2.4}\, N_h^{0.8}, \tag{12a}$$

$$\Delta \alpha_{\mathrm{Si}}(1310\mathrm{nm}) = 6.0 \times 10^{-18}\mathrm{cm}^2\, N_e + 4.0 \times 10^{-18}\mathrm{cm}^2\, N_h, \tag{12b}$$

and near $\lambda = 1550$nm,

$$\Delta n_{\mathrm{Si}}(1550\mathrm{nm}) = -8.8 \times 10^{-22}\ \mathrm{cm}^3\, N_e - 8.5 \times 10^{-18}\ \mathrm{cm}^{2.4}\, N_h^{0.8}, \tag{13a}$$

$$\Delta \alpha_{\mathrm{Si}}(1550\mathrm{nm}) = 8.5 \times 10^{-18}\ \mathrm{cm}^2\, N_e + 6.0 \times 10^{-18}\ \mathrm{cm}^2\, N_h. \tag{13b}$$

Here Δn_{Si} is the change in the (dimensionless) refractive index, $\Delta \alpha_{\mathrm{Si}}$ is the absorption coefficient in cm^{-1}, N_e is the electron density in cm^{-3}, and N_h is the density of holes in cm^{-3}. (Note that Soref and Bennett found that the dependence of the refractive index on the hole density is not linear, in contrast to predictions of simple free carrier or Drude theory, but instead is characterized by the exponent 0.8.) In principle, moderate changes in carrier density should provide significant refractive index modulation, but in practice the absorption becomes a significant limit through device insertion loss. For example, at a carrier density of $10^{17}\mathrm{cm}^3$, we find that $\Delta n(1310\mathrm{nm}) \approx -3.0 \times 10^{-4}$ and $\Delta\alpha(1310\mathrm{nm}) \approx 1.0\mathrm{cm}^{-1}$, and $\Delta n(1550\mathrm{nm}) \approx -4.3 \times 10^{-4}$ and $\Delta\alpha(1550\mathrm{nm}) \approx 1.45\mathrm{cm}^{-1}$.

Modulators based on refractive index perturbation through the free-carrier plasma-dispersion effect generally operate in one of three modes:

Carrier injection. [Xu et al. 2005; Reed et al. 2010] As shown in Figure 14(a), the waveguide region defining the optical mode is confined within the intrinsic region of a p-i-n diode to avoid optical absorption losses in the heavily doped p-type and n-type regions. When the junction is forward-biased, carriers can be injected into the waveguide through the doped regions in contact with metal wires, where they modify the refractive index of the intrinsic silicon. Modulators based on injection mode operation generally require low bias voltages (< 1 V) and allow large extinction ratios—the ratio of light intensities representing bit values 1 and 0 in an on-off keying modulation scheme—but are generally limited in speed by the long minority carrier lifetime of silicon ($\sim$ 1ns for devices with few-micron-scale lateral dimensions). This limitation can be partially mitigated either by using drive signal pre-emphasis [Xu et al. 2007] or through the significantly smaller device volumes available from compact microring resonators [Manipatruni et al. 2010].

Carrier depletion. If the waveguide in the carrier injection device is lightly doped as shown in Figure 14(b), the resulting p-n diode can be operated in reverse bias to deplete carriers from a central region that has a width that depends on the applied voltage [Reed et al. 2010]. If the intrinsic capacitance and resistance of the device is low enough, then the minority carrier lifetime is no longer a significant limitation

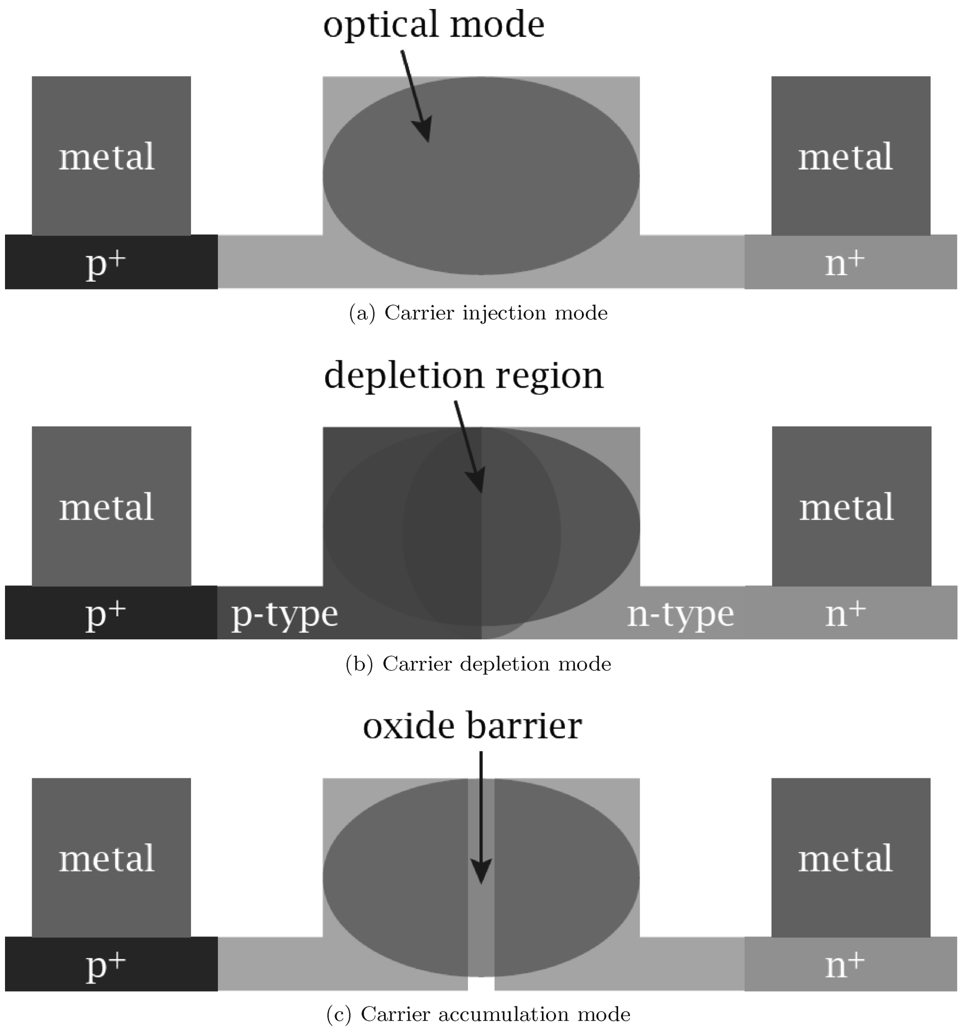

Fig. 14. Three commonly used carrier modes in silicon modulators. (a) In carrier injection, the heavily p-doped and n-doped regions are separated by the intrinsic region containing the optical waveguide to form a p-i-n junction. (b) In carrier depletion, lightly doped p-type and n-type regions contact in the waveguide region to form a p-n junction. (c) In carrier accumulation (the MOS effect), a capacitative structure is formed by placing a thin oxide barrier in the center of the intrinsic region. (After Reed et al. [2001].)

[Liu et al. 2007; Liao et al. 2007]. The dopant concentration in the waveguide must be low enough to avoid significant absorption of the optical mode, thereby limiting the maximum change in refractive index to the value obtained when the available carriers are fully depleted. Although forward-bias operation [Spector et al. 2010] and more complicated diode structures are available [Rasigade et al. 2010] in depletion mode, p-n structures are commonly used to obtain low-power operation at moderate modulation depths [Zheng et al. 2010; Watts et al. 2010; Zortman et al. 2010a].

Carrier accumulation. If the device structure of the injection-mode device is modified to include a thin oxide barrier at the center of the waveguide, then the metal-oxide-semiconductor (MOS) configuration shown in Figure 14(c) can be used to accumulate carriers within the optical mode and thereby modify the local refractive index [Liu et al.

2004a; Barrios and Lipson 2004]. The capacitative structure enforced by the barrier allows the device to operate at speeds that depend primarily on its resistance and capacitance rather than the minority carrier lifetime [Passaro and Dell'Olio 2008], and a device with a length less than 0.5mm recently demonstrated by Lightwire delivers a 9dB extinction ratio at a data rate of 10Gb/s [D'Andrea 2009].

A promising alternative to refractive index modulation in pure silicon based on the free-carrier plasma-dispersion effect is the electric-field-induced change in absorption observed in waveguide-integrated silicon-germanium devices. For example, an electro-absorption (EA) modulator based on the Franz-Keldysh effect [Liu et al. 2008] has demonstrated a 3dB-bandwidth of 1.2GHz and an energy consumption of 50fJ/b at 1540nm. A similar effect that becomes significant in thin quantum wells is the quantum-confined Stark effect (QCSE) [Kuo et al. 2005; Roth et al. 2007]. At room temperature, QCSE absorption in germanium quantum wells grown on silicon has been observed to be as large as that in III-V materials [Kuo et al. 2005; Rong et al. 2010], and EA at 10Gb/s has been demonstrated in 30μm waveguides [Rong et al. 2010]. Although the insertion loss of QCSE devices is likely to remain quite high ($\sim$3–4dB) compared to that of electro-refractive modulators fabricated from pure silicon, the short interaction length that SiGe QCSE enables justifies vigorous further research.

4.1.2. Nonresonant Mach-Zehnder Modulators. The Mach-Zehnder interferometer structure discussed in Section 3.2.1 has been used by many groups to demonstrate high-speed modulation in pure silicon, primarily due to the relative insensitivity of the passive performance characteristics to temperature. Although some MZI modulators have been implemented using charge accumulation [D'Andrea 2009] and charge injection [Green et al. 2007], most recent dramatic progress has been made using carrier depletion [Liu et al. 2004a; Liao et al. 2007; Liu et al. 2007]. One of the most important figures of merit of an MZI modulator is $V_\pi L$, the product of the potential difference between the two interferometer arms needed to produce a phase shift of π at the output port and the length of the arms. Over the past year, several groups have reported depletion-mode devices [Feng et al. 2010b; Liu et al. 2010a; Watts et al. 2010] with $V_\pi L \approx$ 1V-cm for devices with lengths of 0.5–1.0mm operating at bandwidths of approximately 10Gb/s with extinction ratios of 4–7dB. The large sizes of these devices results in an energy expenditure of approximately 5–10pJ/b, but depending on the application this is competitive with electronics at distances greater than about a meter, and could enable high-bandwidth energy-efficient point-to-point links using CWDM [Liu et al. 2010a].

4.1.3. Resonant Microring and Microdisk Modulators and Switches. Systems and applications requiring bandwidths approaching 1Tb/s per optomechanical connector (e.g., a single fiber optical core affixed to a chip or board) will require DWDM and therefore resonant filters such as the microrings or microdisks discussed in Section 3.2.2. The first compact silicon microring modulator capable of operating at 1.5Gb/s is shown schematically in Figure 15(a) [Xu et al. 2005]. The device had a radius R = 6μm and was based on a p-i-n diode operated in carrier injection mode using a forward bias voltage of $V_F \approx$ 4V. DC transmission measurements of the microring at three different bias voltages are plotted in Figure 15(b), with an inset showing the transmission of the probe wavelength of 1,573.9nm as a function of bias. An extinction ratio as high as 15dB was found at this wavelength for a bias voltage change of only 0.3V. Subsequently, the same group demonstrated a WDM link using four microring modulators operated at 4Gb/s each [Xu et al. 2006], and a single modulator with bandwidths of 12.5 and 18Gb/s [Xu et al. 2007; Manipatruni et al. 2007b], leading to a demonstration of a four-ring 50Gb/s transmitter [Manipatruni et al. 2010b]. At these modulation speeds,

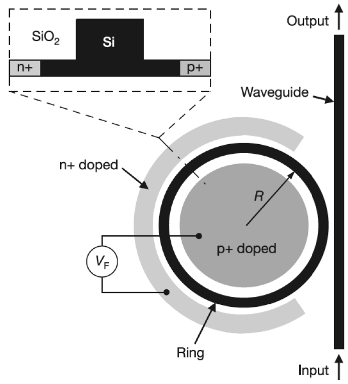

(a) Schematic of the p-i-n microring

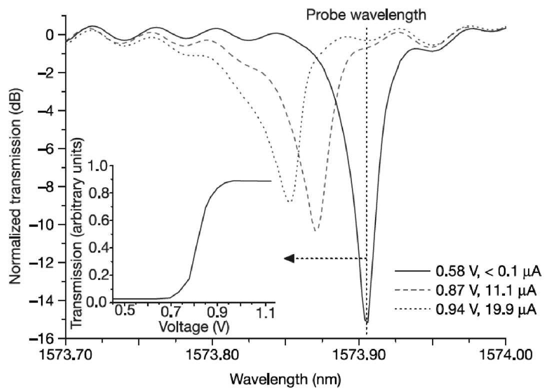

(b) DC characteristics of the active cavity

Fig. 15. The first compact silicon microring modulator capable of operating at 1.5Gb/s. (a) A schematic diagram of the microring device with $R = 6\mu$m based on a p-i-n diode. The modulator was operated in charge injection mode using a forward bias voltage of $V_F \approx 4$V. (b) DC transmission measurements of the microring at three different bias voltages, and an inset showing the transmission of the probe wavelength of 1,573.9nm as a function of bias. An extinction ratio as high as 15dB was found at this wavelength for a bias voltage change of only 0.3V. (Figure reproduced from Xu et al. [2005] with permission © 2005 NPG).

the problem of the long minority carrier lifetime in pure silicon is often compensated by using signal pre-emphasis to overdrive the leading edges of square waves to load the junction with carriers. The small size (and correspondingly low capacitance and resistance) of these resonant devices allows them to be operated at the low voltages and low driving power required by nanoelectronic technologies beyond 22nm CMOS [Chen et al. 2009; Lee et al. 2010]. A 2.5μm-radius modulator has been driven at 1Gb/s using only 150mV peak-to-peak in carrier injection mode [Manipatruni et al. 2010]. A similar device has been operated at 3Gb/s while consuming 120fJ/b [Chen et al. 2009], and an injection-mode ring with a 2.5μm radius has shown an 18dB DC extinction ratio and a 6Gb/s modulation rate while dissipating only 45fJ/b. Silicon microdisk modulators [Zhou and Poon 2006]—which unlike microrings typically have higher-order resonant modes—formed as vertical p-n junctions and operated in carrier depletion mode at about 10Gb/s with moderate extinction ratios of approximately 4dB have been demonstrated that consume only 85fJ/b with a drive voltage of 3.5V [Watts et al. 2008a] and only 3fJ/b at 1V [Zortman et al. 2010a]. Silicon microring modulators operated in depletion mode [Dong et al. 2009] recently have been demonstrated at 12.5Gb/s with a 1V swing and an 8dB modulation depth (with a 3.5dB insertion loss) that dissipate only 10fJ/b [Dong et al. 2010a]. The aggregate power consumption of a compact resonant modulator or switch [Ng et al. 2008; Watts et al. 2008b] will increase due to other factors, including integrated temperature control [Manipatruni et al. 2008; DeRose et al. 2010; Dong et al. 2010a] and drive electronics [Zheng et al. 2010]. In principle, significantly higher modulation frequencies—even equal to the free spectral range of a microring resonator—at similar fractional power dissipation could be obtained by perturbing the coupling between the bus waveguide and the ring waveguide, with the drive electronics becoming the chief performance limiter [Sacher and Poon 2008]. Alternatively, a hybrid device based on a p-i-n-i-p structure in silicon has been proposed [Manipatruni et al. 2007a] that would have characteristic carrier injection and extraction times of approximately 10 ps, allowing a micron-scale resonator to operate at 40Gb/s with 12dB extinction ratio while dissipating several fJ/b. Finally, some progress has been made recently with polysilicon waveguides [Zhu et al. 2009] and microring modulators [Preston et al. 2009], with (respectively) linear losses that are ~7dB/cm higher than that of the best single-crystal waveguides and power consumptions that are ~100× that of the best demonstrated silicon modulators.

4.1.4. Resonant Hybrid Silicon/Polymer Modulators. An entirely different approach to integrated optical modulators relies on a purely electro-optic (EO) effect in polymer claddings on microring resonators fabricated in silicon nitride on silicon oxide [Block et al. 2008]. EO polymers are intriguing for large-scale integrated photonic interconnects because many have electro-optic coefficients greater than 300pm/V (10× that of lithium niobate), potentially very high modulation frequencies of more than 100GHz, and a dielectric constant of 2.5–4 that is small enough to allow fabrication of low-capacitance DWDM devices. The EO polymer cladding can be spin-cast onto existing photonic devices after wafers emerge from the processing line, and the subsequent polymer poling and annealing can occur at temperatures much lower than the 450°C thermal process limit. Initial measurements of the frequency response of hybrid microring devices with $R = 21\mu$m as a function of drive voltage and frequency showed a 1dB modulation depth using a 2.7V drive with 50% duty cycle at 10GHz [Block et al. 2008], but more recent experiments have demonstrated 8dB modulation depth at 6V and 20Gb/s transmission performance [Young et al. 2010]. Almost purely capacitive modulation at these size scales offers the possibility of low-power, high-density integration at the 45-nm technology node and beyond [Palaniappan and Palermo 2010].

4.1.5. Multi-Element Modulators and Switches. Recently, multielement structures have been studied as a means to overcome some of the temporal limitations of single-element resonant structure, such as a frequency passband with a relatively narrow Lorentzian lineshape, and a modulation bandwidth that is limited by the long minority carrier lifetime in silicon. For example, a silicon microring has been used in conjunction with a Mach-Zehnder modulator to provide bandwidth equalization [Gill et al. 2010]. For DWDM applications, dual-microring modulators and switches offer the potential for much higher bandwidths with superior filter performance. High-speed control of the Q of a silicon microcavity comprised of two separated microrings coupled by silicon waveguides has been demonstrated with a switching time of 100ps [Manipatruni et al. 2008]. A spectrally hitless switch based on two microring resonators (each with an independently driven p-i-n diode) has demonstrated a 60GHz bandwidth with a footprint of only 600μm^2 [Lira et al. 2009]. Models of a compact silicon coupled-ring modulator structure where only one ring is actively driven and over-coupled to a waveguide show that the Vernier effect allows 20GHz modulation of the first ring to provide 60Gb/s modulation, with a notch filter profile that is higher-order than that of a single ring. As simulated, the combined system exhibits a smaller resonance shift and lower driving electrical power than a single-ring configuration [Li et al. 2008]. Similarly, simulation of the silicon dual-microring modulator shown in Figure 16 show that such a structure should allow 40Gb/s operation with a modulation depth of 10dB and a factor of three lower operating power than a single ring of the same dimension [Xu 2009]. The layout of the silicon rings and waveguides used in this design is shown in Figure 16(a) with an overlay of doped electrodes and driver circuits, and a cross-section of one ring shown in the inset. Figure 16(b) illustrates an equivalent circuit model of the modulator, showing the point A representing the n^+-doped regions connected by wires, and the diodes D_1 and D_2 representing the p-i-n junctions across the two rings. The effective contact resistance R_C between the metal wires and the doped silicon should be much smaller than the driver resistance R_S. The device is designed to operate in carrier-injection mode, and the diodes D_1 and D_2 are biased with opposite signs depending on the polarity of the driving voltage V_A. As the driving voltage changes sign, carriers are driven from one microring to the other at a rate determined by the value of R_C. Note that the single-ring wavelength shift shown in Figure 15(b) occurs in opposite spectral

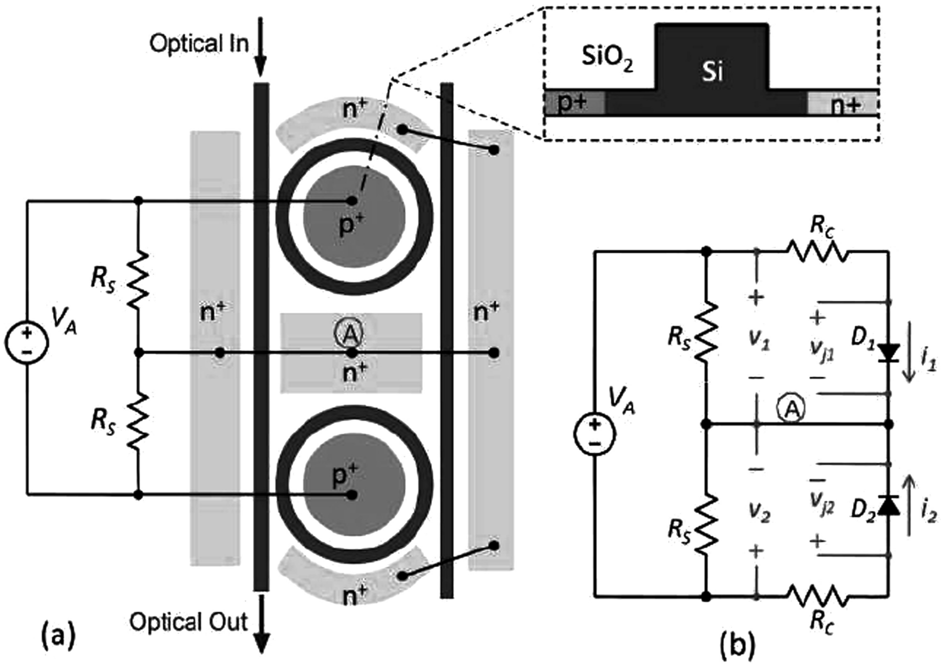

Fig. 16. Schematic diagram of a silicon dual-microring modulator. (a) Layout of the silicon rings and waveguides with an overlay of doped electrodes and driver circuits. A crosssection of one ring is shown in the inset. (b) An equivalent circuit model of the modulator, showing the point A representing the n^{+}-doped regions connected by wires, and the diodes D_1 and D_2 representing the p-i-n junctions across the two rings. The effective contact resistance R_C between the metal wires and the doped silicon should be much smaller than the driver resistance R_S. (Figure reproduced from Xu [2009] with permission © 2009 OSA.)

directions in the dual-ring configuration, and the small quantity of charge dissipated in the resistor R_S allows the total operating energy of the modulator (including both the rings and the drivers) to be as low as 300fJ/b.

4.1.6. Signal Processing for Alternative Modulation Schemes. The previous section concentrated primarily on critically-coupled microring modulators using simple on-off-keying protocols, but other configurations and modulation schemes are under active investigation. For example, as shown in Figure 17(a), single-ring modulators with integrated MOS electrodes can be divided into three types according to their operating regimes in the optical domain: Type I (single-waveguide under-coupled), Type II (single-waveguide critically-coupled), and Type III (dual-waveguide). For each type, the phase transition across a resonance has been plotted as a function of frequency. In Figure 17(b), the signal quality (Q) as a function of increasing carrier transit time across the MOS capacitor is shown, demonstrating that Type II modulators exhibit higher tolerance to longer transit times [Zhang et al. 2009; Zhang 2010a]. More complex multiring designs, such as the "embedded-ring" modulator shown schematically in Figure 17(c), can be operated in a highly nonlinear optical phase regime to enable high-speed modulation at lower voltage with higher power efficiency. Near the resonance, the plot of the simulated signal Q shown in Figure 17(d) demonstrates insensitivity to variations in the value of the coupling between the rings near the optimum [Zhang 2008], a critical requirement for manufacturability. Finally, a schematic of a microring-based NRZ-DPSK modulation and demodulation schemes is shown in Figure 17(e). During modulation, a continuous-wave channel experiences a phase shift of π across the

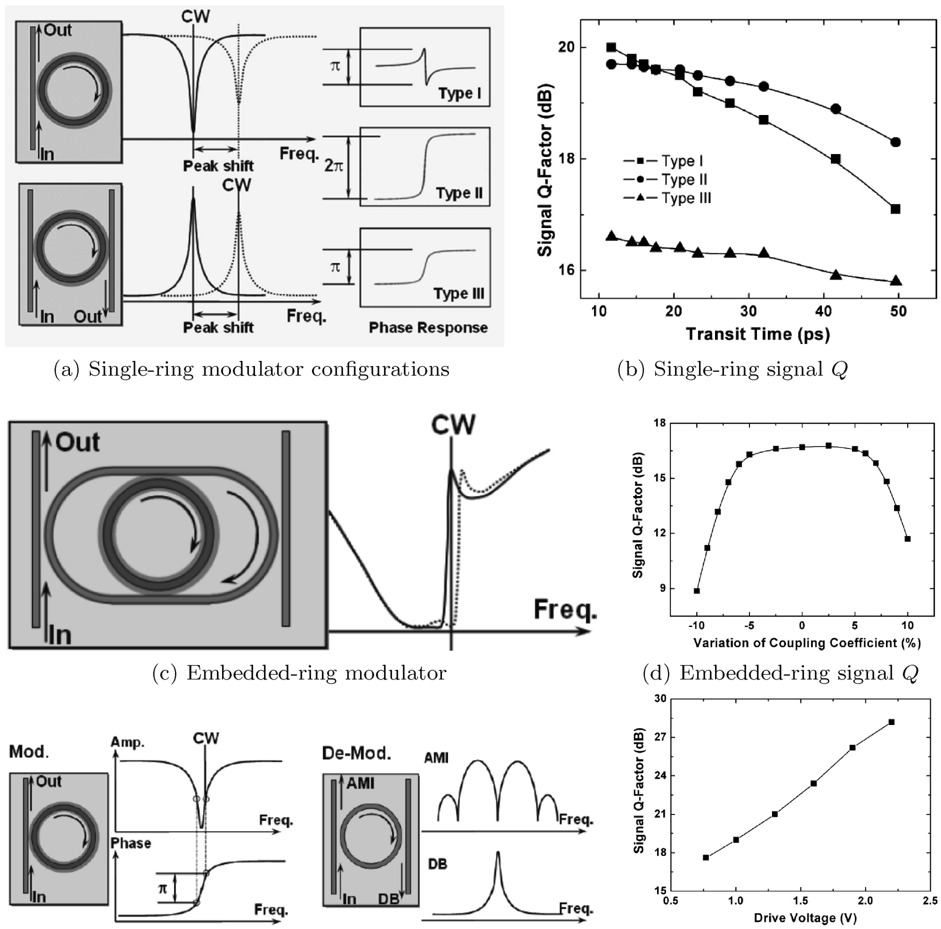

(a) Single-ring modulator configurations (b) Single-ring signal Q

(c) Embedded-ring modulator (d) Embedded-ring signal Q

(e) NRZ-DPSK modulation and demodulation (f) NRZ-DPSK signal Q

Fig. 17. Signal processing for alternative modulation schemes. (a) Single-ring modulators are divided into three types according to their operating regimes in the optical domain: Type I (single-waveguide under-coupled), Type II (single-waveguide over-coupled), and Type III (dual-waveguide). (b) Simulations show that Type II modulators exhibit higher tolerance to longer transit times. (c) A schematic of an embedded-ring modulator operated in a highly nonlinear optical regime to enable high speed modulation. (d) Near the resonance, the signal Q is insensitive to the precise value of the coupling between the rings. (e) Microring-based NRZ-DPSK modulation and demodulation schemes. (f) Dramatic improvements in signal Q are provided by modest increases in drive voltage. (Figure reproduced from Zhang et al. [2009] with permission © 2009 by the authors.)

optical resonance, and a double-waveguide microring filter enables demodulation of both duobinary and alternate-mark inversion components and thus provides balanced detection with 50% less power consumption that on-off-keying. As shown in Figure 17(f), signal Q improved dramatically when the ring modulator was driven until the phase shift became greater than π [Zhang et al. 2008a, 2008b].

4.2. CMOS-Compatible Ge-on-Si Photodetectors

Photodetectors at all wavelengths operate by absorbing incident light, and then creating charge carriers that can be accumulated and measured by electronic circuits. At the standard telecommunication wavelengths proposed for short-haul photonic interconnects, silicon is transparent and therefore not useful for the active element of a photodetector. But in recent years the epitaxial integration of germanium (which has an appreciable absorption at both 1310nm and 1550nm) with silicon waveguides and CMOS current and voltage read-out circuits has led to several device structures and topographies that are promising for high-bandwidth interconnects. An excellent review of the historical development and current status of high-performance Ge-on-Si photodetectors can be found in Michel et al. [2010].

4.2.1. Model of Photodetection. We begin by building a simple model of the bit-error-ratio as a function of the parameters that are likely to be achieved beyond 22nm CMOS. The detector will sense a current $I(b)$ when photons representing a bit value b are incident on the diode, and there will be a noise variance $\sigma(b)$ associated with each of these values that corresponds with the width of the signal current distribution. If we choose a threshold current approximately given by $I_{\text{th}} = [\sigma(0)I(1)+\sigma(1)I(0)]/[\sigma(0)+\sigma(1)]$, then we can define the bit value "0" for $I < I_{\text{th}}$ and the bit value "1" for $I > I_{\text{th}}$ [Ramaswami et al. 2009]. Given finite values of the noise variances, we will misidentify the bit value at a rate that is determined by the overlap of the distributions of $I(0)$ and $I(1)$. For our definition of I_{th}, this bit-error-rate is

$$\text{BER} = K\left[\frac{I(1) - I(0)}{\sigma(0) + \sigma(1)}\right], \tag{14}$$

where

$$K(z) = \frac{1}{2}\,\text{erfc}\left(\frac{z}{\sqrt{2}}\right) \tag{15}$$

is the overlap integral. The current produced by a photodetector with a responsivity A and gain multiplier M is given by

$$I(b) = AMP(b), \tag{16}$$

where the optical power carried by a single "1" bit incident on the photodetector is

$$P(1) = N_\gamma \frac{hc}{\lambda} B. \tag{17}$$

Here we have represented the average optical power indicating the bit value "1" as N_γ photons per bit multiplied by the energy per photon (given by the product of Planck's constant h and the speed of light c divided by the vacuum wavelength of the light in the detection region) and the average number of "1" bits arriving per unit time as the bandwidth B. We specify the optical power carried by photons representing the bit value "0" in terms of the extinction ratio $x = P(0)/P(1)$.

If we assume that we will not require either optical preamplification or significant transimpedance amplification (discussed below), and we approximate the electrical bandwidth of the receiver as B, then the noise variance for bit value b is given by

$$\sigma^2(b) = \left[2qMFI(b) + 2q(I_{\text{ds}} + I_{\text{db}}M^2F) + \frac{4k_BT}{R_L}\right]B. \tag{18}$$

A model of the excess noise factor F is given by McIntyre [1966] and generally depends on the dimensions of the depletion layer and the ionization probabilities of the electrons and holes. In the case where those probabilities are equal, $F = M$; p-i-n photodiodes can be modeled by taking $M = 1$, giving $F = 1$. The first term on the right-hand side of Equation (18) represents the "shot noise" arising from fluctuations in the number of photoelectrons created in the detection region [Ramaswami et al. 2009], while the second term describes the contributions from dark current arising in the surface (I_{ds})

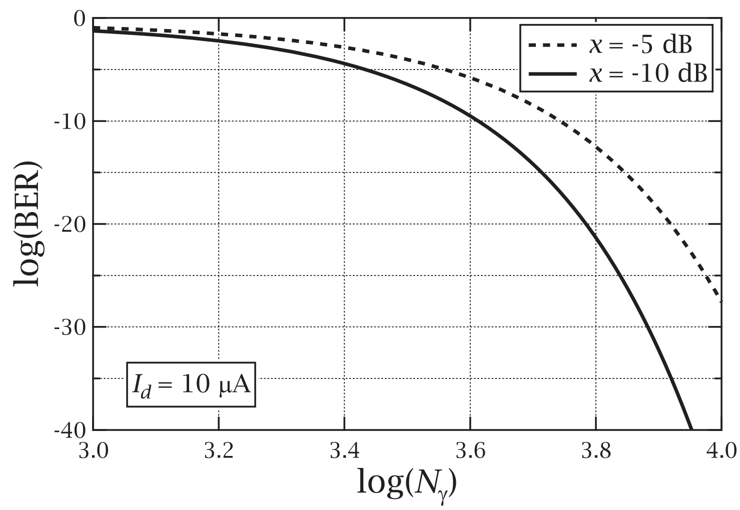

(a) Photodetector with large I_d

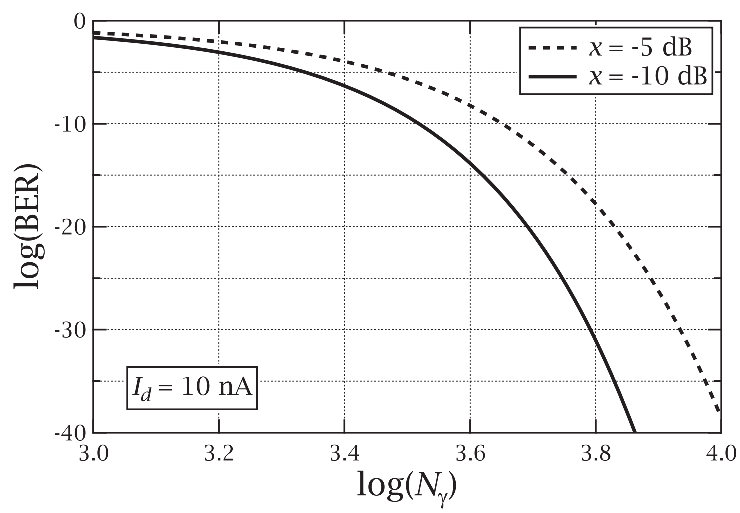

(b) Photodetector with small I_d

Fig. 18. Bit error rates given by Equation (14) as a function of photon number representing a bit value of "1" incident on a photodetector with responsivity $A = 0.7$A/W and gain multiplier $M = 1$. We have assumed a data rate commensurate with a bandwidth $B = 10$GHz, and a load resistance $R_L = 3$kΩ consistent with device sizes in 22nm CMOS. The detector operates in silicon-germanium at 1310nm, and we have compared BERs for two different modulation extinction ratios. (a) When $x = 10$dB and the dark current is 10μA, a BER of 10^{-20} can be obtained with 6000 absorbed photons. (b) When $x = 10$dB and the dark current is 10nA, a BER of 10^{-20} can be obtained with about 5000 absorbed photons.

and the gain region (I_{db}) of the photodetector. The final term arises from thermal noise (also known as "Johnson-Nyquist" noise) caused by thermal fluctuations in the energy of the charge carriers in the detector circuit; here k_B is Boltzmann's constant, T is the temperature, and R_L is the load resistance of the circuit.

Figure 18 illustrates the bit error rates given by Equation (14) as a function of photon number representing a bit value of "1" incident on a p-i-n photodetector with responsivity $A = 0.7$A/W. We have assumed a data rate commensurate with a bandwidth $B = 10$GHz, and a load resistance $R_L = 3\text{k}\Omega$ consistent with device sizes in 22nm CMOS. The detector operates in silicon-germanium at 1310nm, and we have compared BERs for modulation extinction ratios of 5 and 10dB. In Figure 18(a) we consider the case of a detector with a high total dark current of 10μA, and Figure 18(b) we treat the low-dark current case by assuming that $I_d = 10$nA. In both cases, for the high-extinction-ratio case we can obtain a BER of 10^{-20} with 5000–6000 absorbed photons, corresponding to an absorbed power for $b = 1$ of about 10μW (–20dBm). With an extinction ratio of 5dB, we require approximately 8000 absorbed photons per bit to obtain a BER of 10^{-20}.

In typical point-to-point links, without impedance matching the receiver bandwidth would be limited by the RC time constant determined by the load resistance R_L (chosen as large as possible to minimize thermal noise) and the capacitance C of the detector element. Instead, transimpedance amplifiers (TIA) allow high-bandwidth operation of the detection system by reducing the apparent load resistance at the receiver output by a factor of $G + 1$, where G is the gain of the amplifier [Ramaswami et al. 2009]. However, if the detection element has a size typical of the dimensions of a high-index-contrast SOI waveguide, then the natural capacitance of silicon ($\sim$ 200aF/μm [Miller 2009]) allows detectors to be built with capacitances that are already well-matched to the 10fF values typical of registers in an L2 cache beyond 22nm CMOS [ITRS 2009]. Since $1/(2\pi \times 3\ \text{k}\Omega \times 10\ \text{fF}) \approx 5$GHz, in moderate-frequency DWDM links we should be able to eliminate the TIA altogether, and design a "receiverless" detection circuit that allows the charge accumulated (and perhaps amplified) from the absorbing element to produce a voltage swing commensurate with the registers [Bhatnagar et al. 2004]. However, if modulation frequencies increase to 10Gb/s and beyond, it is unlikely that a completely receiverless design could be used, because even if the capacitance of the detector could be reduced significantly, some degree of impedance-matching would still be needed with the digital silicon registers.

4.2.2. Epitaxial Ge-on-Si Waveguide Photodetectors. Achieving high bandwidth and responsivity is challenging in "free-space" Ge-on-Si photodetectors used to measure normally-incident light, because thicker Ge volume is needed to increase the fraction of photons absorbed, and smaller Ge volume is needed to reduce the characteristic response time of the absorber and thereby increase the bandwidth. One way to decouple the competing factors in this volume trade-off is to integrate the Ge active region with a silicon waveguide [Michel et al. 2010], allowing the length of the waveguide to establish the absorption depth of the Ge, while the small lateral dimensions of the waveguide allow rapid carrier collection. Furthermore, the relatively small area of a high-index-contrast ridge or wire waveguide limits the dark current at a given operating bias voltage, since the dark current density tends to be independent of device size.

The first demonstration of high-performance monolithically waveguide-integrated Ge photodetector used silicon nitride and silicon oxynitride channel waveguides fabricated above a Ge p-i-n junction [Ahn et al. 2007]. Even though the principal absorption mechanism of this device was evanescent, it exhibited a responsivity of 1.1A/W at 1550nm, with a bandwidth over 7GHz at 1V reverse bias. Soon thereafter, similar results were demonstrated using an SOI rib waveguide butt-coupled to a Ge metal-semiconductor-metal (MSM) photodetector [Vivien et al. 2007], and a vertical n-i-p photodiode formed by epitaxial Ge grown on silicon rib waveguides in SOI [Yin et al. 2007]. Recent work on butt-coupled n-i-p devices has increased the bandwidth to 42GHz with a 4V reverse bias at 1550nm, and it appears that the bias voltage will drop significantly for operation at 1310nm [Vivien et al. 2009]. A vertical Ge-on-SOI p-i-n photodiode has been demonstrated using evanescent coupling with multimode waveguides [Feng et al. 2010a], and a lateral p-i-n structure using a cost-effective 200nm Ge membrane evanescently coupled to an SOI rib waveguide has shown moderate internal responsivity at 1550nm with an 18GHz bandwidth and 1V reverse bias [Wang et al. 2008]. As interconnect bandwidth requirements grow beyond 22nm CMOS, device sizes and capacitances must shrink and further integration of multiple-device components will be necessary. Ge-on-Si MSM photodetectors with capacitances of only 2.4fF have been integrated into a four-channel WDM system using passive microring drop filters, and can be operated at 15Gb/s per channel [Chen and Lipson 2009]. A similar photodetector has been monolithically integrated into a proof-of-concept on-chip single-channel optical data link operating at 3Gb/s with a 0.5V modulator voltage swing and 1.0V detector bias [Chen et al. 2009]. Excluding the drive electronics, the total energy consumed by this link has been estimated to be $\sim$120fJ/b. An 8-channel receiver based on microring drop filters and vertical p-i-n Ge-on-Si photodiodes with responsivities of $\sim$1A/W has demonstrated an aggregate data rate of 160Gb/s [Fang et al. 2010]. The BER of each channel when driven by –20dBm in the bus waveguide was about 10^{-11} at 10Gb/s, and should decrease further when the passive coupling through the drop filters has been optimized.

Improvements in detector sensitivity of 5–10dB are possible using photoelectron gain multiplication in an avalanche photodiode (APD) configuration [Michel et al. 2010]. In this structure, absorption of photons create a photocurrent which is amplified by cascading electron impact ionizations in a high-electric-field layer. A monolithic free-space Ge-on-Si APD operating at 1300nm with a gain-bandwidth product of 340GHz has been demonstrated with a sensitivity of –28dBm at a bandwidth of 10Gb/s and a bias voltage of 25V [Kang et al. 2009]. This device relied on photoelectron generation in the Ge layer, followed by amplification in the silicon layer, and was $\sim 2\mu$m thick, thereby requiring a bias voltage of 25V. By moving to a monolithically integrated structure with Ge deposited on a SiON insulating layer on an SOI ridge waveguide, an APD has been demonstrated that achieves avalanche gain of 10dB at 30GHz with a bias voltage of only 1.5V [Assefa et al. 2010]. In this case, careful nanophotonic and nanoelectronic engineering shaped the optical and DC electromagnetic field profiles, allowing the gain to occur within the Ge layer with an overall reduction in the amplification noise level of 70%.

5. LIGHT SOURCES FOR LARGE-SCALE INTEGRATED DWDM

Although some progress has been made in fabricating and packaging light-emitting diodes for high-speed short-haul data transmission [Chen et al. 1999; Fattal et al. 2008], the workhorse light source of telecommunication networks remains the laser. A laser is a regenerative oscillator comprised of three fundamental components: a gain medium (or amplifier) that can store energy that is transferable to an optical field through stimulated emission; a pump mechanism—typically electrical or optical—that can inject energy into the amplifier; and a feedback system that recirculates the optical field through the amplifier. Only two years after the demonstration of the first laser fifty years ago, and long before silicon transistors were integrated into information processing devices, laser oscillation was demonstrated for the first time in III-V materials. Since that time, extensive research and development has brought to market many single-wavelength sources for directly modulated, CWDM, and DWDM networks, particularly vertical cavity surface-emitting lasers (VCSELs) and distributed feedback (DFB) lasers with output powers typically in the range of 1–10mW and total electrical-to-optical conversion efficiencies of 10–30%. However, future computer networks linking nodes over distances from 100m down to 1mm will require more highly integrated multiwavelength sources capable of providing equivalent performance per channel over many frequency channels, at a cost consistent with CMOS information

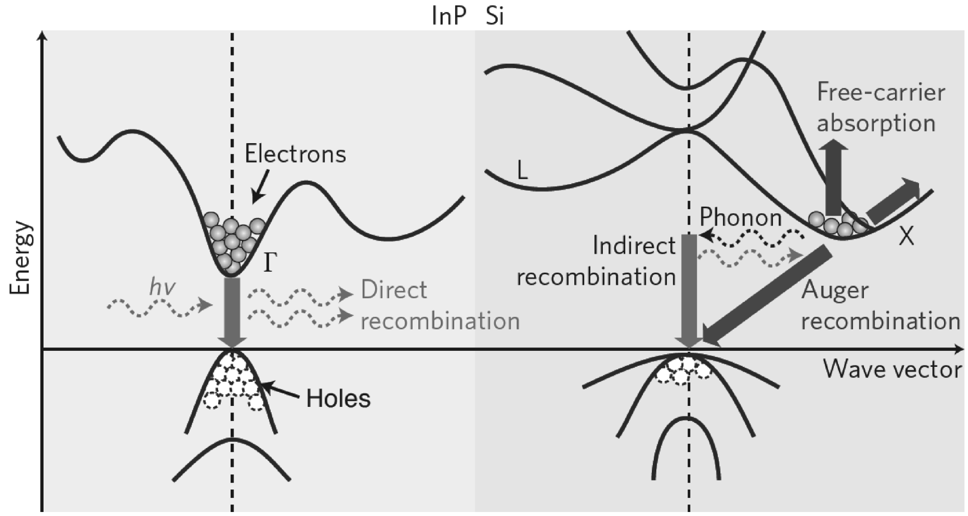

Fig. 19. The energy band structure of a direct-bandgap material such as indium phosphide (left) differs significantly from that of an indirect-bandgap material like silicon (right). In InP, the probability that electron-hole recombination results in the emission of a photon is high, but in Si light generation by indirect recombination is dominated by other physical processes such as Auger recombination and free carrier absorption. Here the vertical axis represents carrier energy, while the horizontal axis represents momentum. (Figure reproduced from Liang and Bowers [2010] with permission © 2010 NPG.)

technologies. Furthermore, research in CMOS-compatible optical isolators [Yu and Fan 2009] will be critical in enabling stable operation of lasers that are copackaged with dies incorporating the active devices described in Section 4.

5.1. Silicon Light Sources

In principle, silicon itself would be an ideal gain medium for large-scale integrated lasers, but as shown in Figure 19, the physics of light generation processes in silicon are highly unfavorable. In commonly used direct-bandgap III-V materials such as indium phosphide, the plot of allowed carrier energies as a function of momentum shows that the minimum energies of the conduction and valence bands occur at the same momentum. Therefore, injected carriers can recombine efficiently to generate optical photons because simultaneous conservation of energy and momentum are possible. However, in an indirect bandgap material like silicon, the energy extrema of the conduction and valence bands are significantly different, rendering optical emission highly unlikely. Instead, indirect recombination is dominated by other physical processes such as Auger recombination and free carrier absorption [Liang and Bowers 2010], which couple electron-hole pairs to phonons and have far shorter (nonradiative) lifetimes than the (radiative) lifetime associated with photon emission.

5.1.1. Silicon Raman lasers. One approach to the momentum conservation problem for optical transitions in silicon shown by Figure 19(b) is to use the inelastic scattering of photons by phonons at optical frequencies, a process known as Raman scattering. When electromagnetic radiation interacts with a material system with many available energy levels, the most likely result is the elastic Rayleigh scattering responsible for phenomena like the blue sky: the scattered photon has essentially the same energy as the incident photon. However, with a probability of only 10^{-6}–10^{-9}, a photon can interact with an optical phonon and change the electronic or mechanical (e.g., rotational or vibrational) state of the material. When the incident photon creates a phonon (Stokes scattering), the emitted photon has lower energy; when the incident photon annihilates a phonon (anti-Stokes scattering), the emitted photon has higher energy. Laser amplification in silicon based on stimulated Raman scattering can be obtained by driving the material with an optical field tuned to a frequency below the bandgap of Si (which suppresses Auger and free carrier recombination), allowing an optical field tuned to the Stokes transition to realize gain [Jalali and Fathpour 2006]. Initial measurements demonstrated stimulated Raman gain in high-index-contrast SOI waveguides as large as 0.25dB at a wavelength of 1542nm [Claps et al. 2002] but the continuous-wave power of 1.6W at the pump wavelength of 1427nm was high enough to create free carriers through the nonlinear optical process of two-photon absorption (TPA) [Leuthold et al. 2010]. The subsequent demonstration of a pulsed silicon Raman laser used timing and delay techniques to allow carriers generated by TPA to recombine prior to recirculation of the laser pulse [Boyraz and Jalali 2004]. Soon thereafter, continuous-wave gain [Jones et al. 2005] and laser oscillation [Rong et al. 2005] were demonstrated in SOI by placing the laser amplifier in a reverse-biased p-i-n diode structure to extract the free carriers created by TPA [Liu et al. 2004b]. The characteristics and performance of the most efficient continuous-wave silicon Raman laser demonstrated to date are summarized in Figure 20 [Rong et al. 2007]. In Figure 20(a), a schematic of the laser resonator shows that the incident Raman pump power is coupled into the resonator, and the corresponding initial pump power $I_p(0)$ is depleted by the amplifier until it reaches the value $I_p(L)$ after one full trip around the high-Q SOI racetrack resonator of length L. The directional coupler shown in a cross-sectional view in Figure 20(b) was designed to maximize the value of $I_p(0)$ at the pump wavelength of 1550nm, but to maintain a small output coupling (with correspondingly small coupling loss and higher gain saturation) at the Raman laser signal wavelength of 1686nm. The output power of the laser as a function of the coupled input power is plotted in Figure 20(c), showing a laser threshold of 25mW and an internal optical conversion efficiency of about 17%. Despite this impressive technical achievement, the aggregate operating efficiency of this laser is substantially reduced due to the electrical-to-optical conversion efficiency of the pump laser, the coupling efficiency of the pump field to the laser cavity, and the power required to operate the large p-i-n junction used to remove the TPA free carriers from the resonator.

5.1.2. Epitaxial Ge-on-Si Lasers. One approach to a CMOS-compatible light source that could mitigate the large indirect bandgap momentum mismatch in silicon is the epitaxial growth of germanium on silicon. Although Ge is also an indirect bandgap material, the momentum misalignment between the valence and conduction band is much smaller, significantly increasing the probability that radiative recombination could occur. The differences between the lattice constants and thermal expansion coefficients mismatch of Si and Ge is substantial, but the tensile strain of a Ge layer epitaxially grown on a Si substrate actually reduces the band mismatch and increases the optical gain available in the material system [Sun et al. 2010; Liang and Bowers 2010]. The first demonstration of a Ge-on-Si laser operating at room temperature used pulsed optical pumping of a Ge waveguide resonator with dimensions $1.6\mu\text{m} \times 0.5\mu\text{m} \times 4.8\text{mm}$ grown on silicon [Liu et al. 2010b]. Laser threshold was reached in the Ge material when the system was pumped with 1.5ns Q-switched laser pulses carrying 6μJ at a wavelength of 1064nm, with a peak pump intensity of 300kW/cm^2 at 50μJ per pulse. Studies of electrically pumped geometries are under way.

5.1.3. Parametric Amplification Using Four-Wave-Mixing in Silicon. In addition to Raman scattering, silicon exhibits other third-order nonlinear optical phenomena such as self-phase modulation, cross-phase modulation, third-harmonic generation, and four-wave mixing (FWM) [Leuthold et al. 2010]. (Nonlinear optical effects are characterized by an

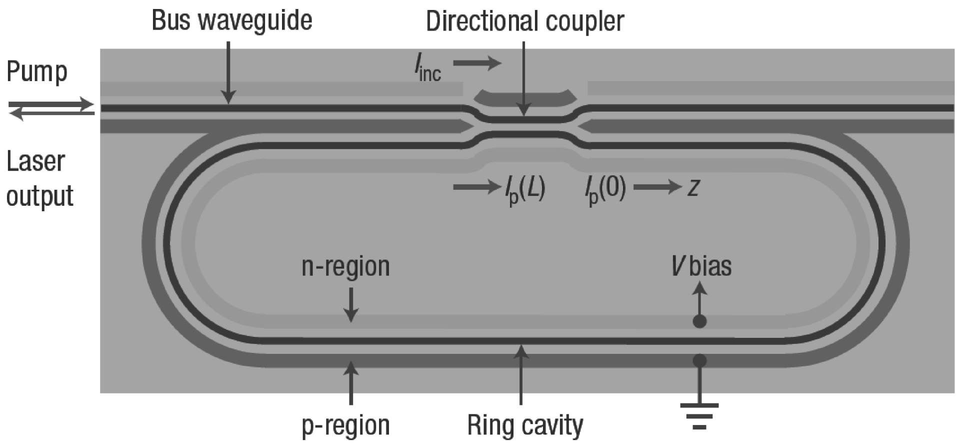

(a) p-i-n device schematic

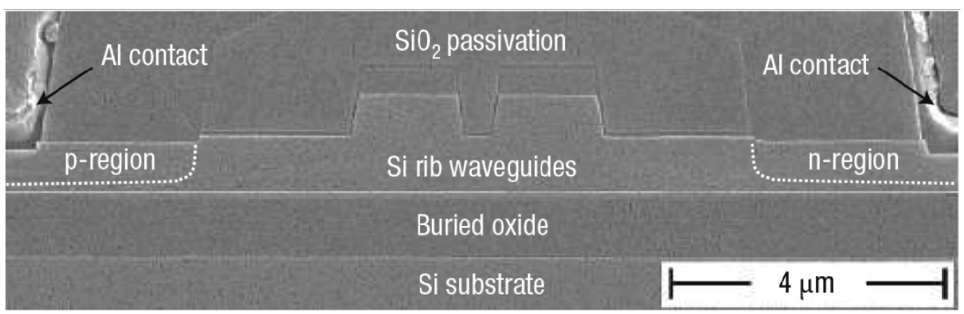

(b) SEM of the p-i-n region

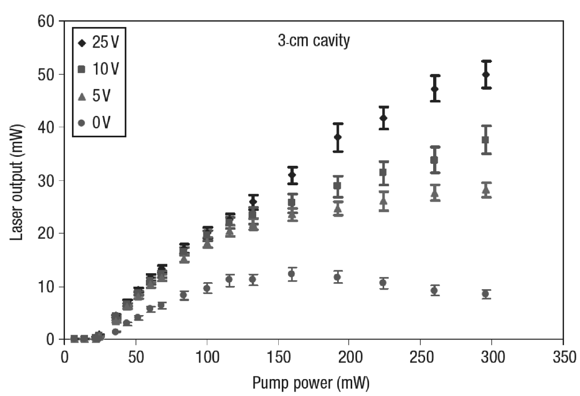

(c) Output power

Fig. 20. Continuous-wave silicon Raman laser based on a p-i-n junction design. (a) Schematic of the laser cavity with the incident Raman pump. Note that $I_p(L)$ is the value of the coupled input Raman pump power $I_p(0)$ after traveling once around the resonator. (b) An SEM of the laser resonator in the region of the input/output coupler. (c) Laser output power for a 3-cm cavity as a function of coupled input pump power. (Figure reproduced from Rong et al. [2007] with permission © 2007 NPG.)

"order" equal to the number of input optical fields that interact to produce an output field.) Recently, degenerate FWM has been used to demonstrate broadband parametric gain in silicon waveguides with dimensions of roughly 300nm $\times$ 550nm $\times$ 6.4mm [Foster et al. 2006]. In this configuration, two of the four interacting photons are absorbed from the pump laser operating near 1525nm and produce two photons at nearby "signal" and "idler" wavelengths, allowing on-off gain to be demonstrated over a wavelength band of 28nm, and wavelength conversion in the range 1511–1,591nm with peak conversion efficiencies of +5.2dB. Similar nonlinear optical techniques allowed the subsequent demonstration of a silicon-chip-based ultrafast optical oscilloscope [Foster et al. 2008] and ultrafast waveform compression using a time-domain telescope [Foster et al. 2009], and ensuing work showed electrical control of the phase mismatch in silicon FWM wavelength converters [Tsia et al. 2008]. Subsequently, a monolithically integrated CMOS-compatible frequency-comb source based on FWM was fabricated using a silicon-nitride microring resonator coupled to a SiN ridge waveguide [Levy et al. 2010]. When pumped with 50mW of input optical power tuned to the resonance of the microring near 1557.8nm, optical parametric oscillation in the ring generated more than 100 new wavelengths simultaneously in the adjacent telecom band. Although the absolute wavelengths of the channels depends on temperature, at high pump powers one significant advantage of this source (shared by the frequency-comb sources discussed below) is that the frequency *spacing* between adjacent channels is fixed at a value determined by the resonator geometry, and to first order is not affected by the temperature of the system. A more recent demonstration of a FWM device based on high-index glass microrings and waveguides fabricated within SiO_2 showed promising conversion efficiency, with a threshold of 54mW and a differential slope efficiency of 7.4% for a single oscillating mode [Razzari et al. 2010]. At higher pump powers, cascaded FWM resulted in a comb of wavelengths, and at a pump power of 101mW approximately 9% of the pump power was converted into oscillating output modes. One disadvantage of all sources based on third-order optical effects is that they tend to be inefficient at the overall system level, relying on external optical pump lasers to operate.

5.2. Hybrid III-V/Si Laser Structures

The problem of low intrinsic optical gain in silicon has led researchers toward laser sources that couple the high single-pass direct-bandgap gain of III-V materials with the CMOS compatibility of high-index-contrast silicon waveguides on SOI wafers.

5.2.1. SOA and Si External Cavity Lasers. The most direct approach to a CMOS-compatible laser source for integrated optical interconnects is independent fabrication and optimization of a semiconductor optical amplifier (SOA) and a silicon resonator that defines the longitudinal and transverse mode structure of the light emitted by the device. Recently, a tunable laser was demonstrated using an SOA and an external resonator incorporating two silicon microring cavities connected through a multimode interference coupler to form a loop-back reflecting resonator [Chu et al. 2009]. The device has a footprint less than $0.32mm^2$, and can cover the C telecommunication bands for DWDM applications. The dual microring filter allowed a tuning range of 38nm with a total power consumption of 26mW, an $\sim$85% reduction over devices that use external cavities relying on SiON waveguides.

5.2.2. Heterogeneously Integrated III-V/Si Lasers. A promising approach to the problem of fabricating compact, highly efficient laser sources for chip-scale data networks is the heterogeneous integration of InP membranes onto an SOI substrate with an embedded silicon ridge waveguide [Van Campenhout et al. 2007; Van Campenhout et al. 2008; Liu et al. 2010]. The CMOS-compatible substrate of the laser shown in Figure 21

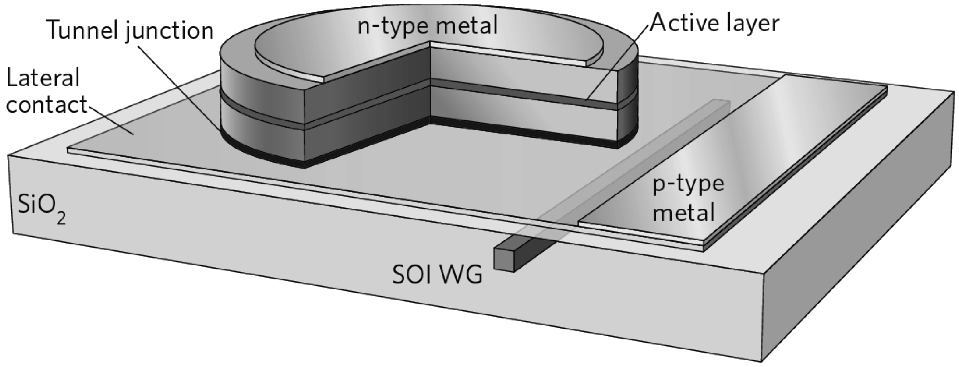

Fig. 21. Schematic of a heterogeneously integrated III-V microdisk laser evanescently outcoupled through a silicon ridge waveguide buried in oxide. The InP-based microdisk has a diameter of 7.5μm and a thickness of 1μm. Gain in the active region is provided electrically by current injection. The device structure incorporates a tunnel junction to efficiently contact the p-side of the p-n junction. (Figure reproduced from Liang and Bowers [2010] with permission © 2010 NPG.)

was fabricated in a 200mm SOI wafer, with a 550nm $\times$ 220nm waveguide etched in silicon above a 1μm buried-oxide layer and then covered with a 750nm SiO_2 layer. The InP-based epilayer structure was grown by molecular beam epitaxy (MBE) on a two-inch InP wafer, incorporating an active layer consisting of three InAsP quantum wells [Van Campenhout et al. 2007] and a 300nm sacrificial InGaAs etch-stop layer for subsequent substrate removal. After dicing the InP wafer, the dies were molecularly bonded to the SOI waveguide wafer. The microdisk structures were patterned to have 7.5μm diameters using EBL and then etched to a thickness of 1μm, leaving the 100nm bottom contact layer shown in Figure 21. After deposition of the electrodes, laser gain is provided by injecting carriers into the active region of the disk. This structure is capable of room-temperature continuous-wave operation at a wavelength of 1.6μm with a threshold current as low as 0.5mA, and has a thermally limited output power of about 10μW for an input current of 1.5mA [Van Campenhout et al. 2007]. Four similar devices have been integrated onto the same bus waveguide to provide multiwavelength output suitable for CWDM, although coupling of the gain in the active region to higher-order modes in the relatively thick disks resulted in unequalized output at the operating wavelengths [Van Campenhout et al. 2008].

5.2.3. Wafer-Scale III-V/Si Hybrid Evanescent Lasers. Another approach to hybridized lasers that heterogeneously integrate III-V laser gain and silicon waveguides uses a laser resonator and output-coupling waveguide defined completely in the CMOS-compatible SOI layer, but relies on evanescent coupling between the passive laser cavity modes and a wafer-bonded III-V layer to provide gain [Liang and Bowers 2010]. The hybrid propagating electromagnetic field mode is confined primarily in the Si waveguide, but the dimensions of the waveguide are reduced until approximately 3–8% of the power in the mode is carried along through the III-V material. In initial demonstrations of this concept, Fabry-Perot laser resonators were defined in the top silicon layer of SOI prior to bonding with an AlGaInAs-based wafer, and then after substrate removal the membranes containing the laser amplifiers were processed using standard techniques used for III-V lasers [Park et al. 2005; Fang et al. 2006]. This research has been extended to CWDM sources that include other linear cavity topographies that rely on more sophisticated wavelength-selective feedback methods that do not require the same level of facet dicing and polishing, such as distributed feedback (DFB), distributed Bragg reflector (DBR), and sampled-grating DBR [Fang et al. 2009; Alduino

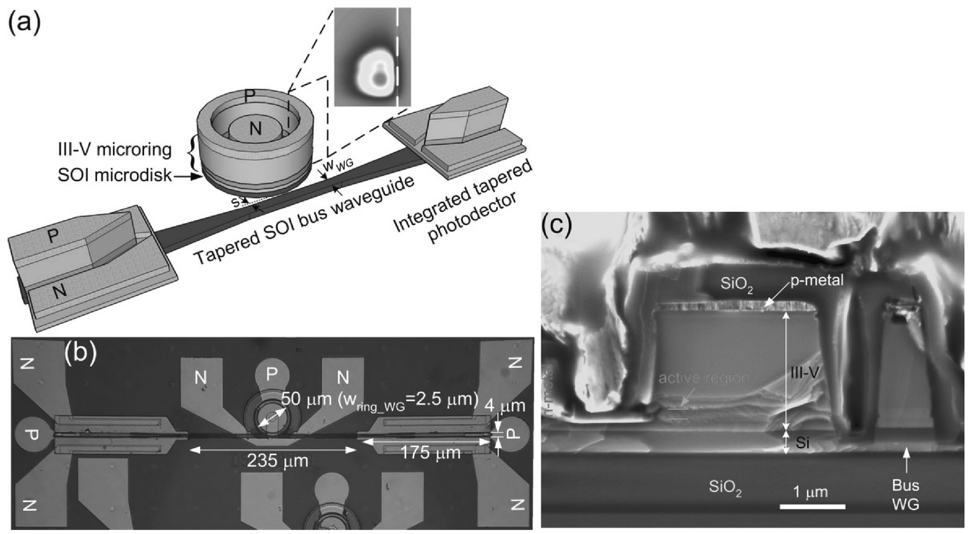

Fig. 22. A hybrid integrated SOI/III-V microring laser. (a) A schematic of the laser structure showing the SOI ring waveguide defining the laser resonator, the III-V gain region, and the SOI bus waveguide that transmits power to tapered photodetectors. A transverse mode profile computed using the beam propagation method is shown in the inset. (b) A microscope image of the integrated devices. (c) An SEM cross-sectional image of the structure near the coupler between the laser resonator and the bus waveguide. (Figure reproduced from Liang et al. [2009b] with permission © 2009 OSA.)

et al. 2010]. An alternate approach relied on a Si on-chip "racetrack" resonator laser that in principle could be one of several such devices evanescently coupled to a common bus waveguide for CWDM [Fang et al. 2007]. This laser demonstrated continuous-wave output at 1590nm with a threshold current of 175mA, an output power of 29mW, and a maximum operating temperature of 60C. Further reductions in the threshold current require that the round-trip length of the racetrack cavity (which is greater than 2mm) be decreased significantly. For example, the hybrid microring laser structure shown in Figure 22 has a diameter of only 50μm—as well as very low optical and electrical losses—which allowed the device to be operated at a wavelength near 1530nm with a threshold current as low as 5.4mA at 10°C [Liang et al. 2009b]. If high-temperature performance can be demonstrated for this device class, then ring diameters below 10μm should enable threshold currents below 1mA, and the correspondingly small capacitance of the device should allow direct on-off-keying modulation rates above 10Gb/s [Liang and Bowers 2010]. As in the case of microdisk lasers [Van Campenhout et al. 2008], several of these microring lasers could be coupled to the same waveguide bus, enabling highly integrated CWDM data transmission at low energy expenditures per bit. Initial demonstrations of mode-locked hybrid evanescent lasers based on Fabry-Perot cavities and generating 4ps pulses at 40GHz have been demonstrated with both passive and active/passive locking techniques [Koch et al. 2007], an important first step toward compact DWDM laser sources. More complex external cavity quantum-well structures have demonstrated pulse durations as short as 60fs [Quarterman et al. 2009], but it isn't yet clear that these devices are candidates for large-scale integration.

5.3. Frequency-Comb Sources for DWDM

Although the potential for the large-scale integration of heterogeneous III-V/Si lasers is very promising, the case of ultra-high-bandwidth DWDM interconnects requires that we consider optical "power supplies" with channel counts approaching 64 (or even 128) and channel spacings of 40–120GHz, and optical output powers of multiple tens of milliwatts per channel so that we can link many network nodes to each other without repeaters. Aggregating individual DFB lasers commonly used in the telecommunications industry for computer network applications is probably infeasible: not only is the cost prohibitive, but also the channel wavelength spacings would vary significantly with time. The extensive monitoring and control required to maintain high data rate parallel links would also dramatically increase both system complexity and energy consumption per received bit. Instead, as in the case of the optical parametric oscillator source described above, laser systems that provide a comb of frequencies with fixed channel spacing would allow the absolute wavelength of only one or two of the channels to be monitored, and a modification of the temperature of the package (for example) would tune the entire comb.

5.3.1. Quantum-Dot Frequency-Comb Lasers. The Fabry-Perot (FP) resonator configuration, a linear cavity of length L formed by two mirrors, is ubiquitous in the laser industry, and (given a sufficiently broadband gain medium) is capable of producing a large number of discrete wavelength channels separated from their nearest neighbors by the FSR frequency $\Delta f_{\mathrm{FSR}} = c/2n_g L$. For III-V laser amplifier materials with $n_g \approx 3$, a laser with length $L = 0.5$mm has $\Delta f_{\mathrm{FSR}} \approx 100$GHz, compatible with the ITU O-band or C-band. Since temperature and other systematic effects cause only small fractional changes in $n_g L$, the FP comb spacing remains fixed even when the absolute wavelengths of the constituent channels shift significantly, allowing all channels to be stabilized and tracked simultaneously.

In contrast with the quantum-well gain materials described previously, quantum dot laser amplifiers tend to provide broadband gain profiles due to significant variation in the size of the dots. The gain spectrum of an individual quantum dot is typically smaller than the FSR of a short FP cavity and the inhomogeneity of the size distribution of many dots grown using MBE can provide gain over a wavelength range of 40–60nm [Kovsh et al. 2007], sufficient to support 50–100FP modes with 100GHz channel spacing. This relatively large frequency separation also allows spectral hole burning in the gain medium, which is an order of magnitude larger in quantum dots than quantum wells: dots that interact strongly with a particular longitudinal resonator mode generally interact weakly with other modes. This ensures that the relative intensity noise of each mode will be small enough to achieve low bit error rates (BER) under high speed modulation of each channel [Gubenko et al. 2007]. The laser structure reported in Wojcik et al. [2009] is grown by MBE on a GaAs substrate and incorporates 10 stacked planes of InAs/InGaAs quantum dots in a $3\mu\mathrm{m} \times 440$nm waveguide structure fabricated using standard photolithography and etching techniques. The FP cavity is defined by cleaved facets, and the output of a typical device is shown in Figure 23. Here the frequency channel spacing is 43GHz, and each of the 13 channels within 3dB of the most powerful channel is carrying at least 5mW. The relative intensity noise shown in the figure for one channel over the range 100kHz to 10GHz is characteristic of a single-channel device, indicating that the gain coupling between adjacent channels is negligibly small.

5.3.2. Externally Modulated Frequency Comb Sources. Another method that can be used to generate a fixed comb of frequency channels is the external modulation of conventional III-V or fiber-optic laser sources. For example, the system shown schematically in Figure 24(a) relies on a dual-electrode Mach-Zehnder (amplitude and phase) modulator (DEMZM) to generate a broad and stable optical comb from a stand-alone continuous-wave laser source [Zhou et al. 2009a]. Experiments using this approach have demonstrated the output spectrum plotted in Figure 24(b) with 175 optical modes spanning over 3.5THz at 1550nm. The pulses emerging from the DEMZM are strongly frequency

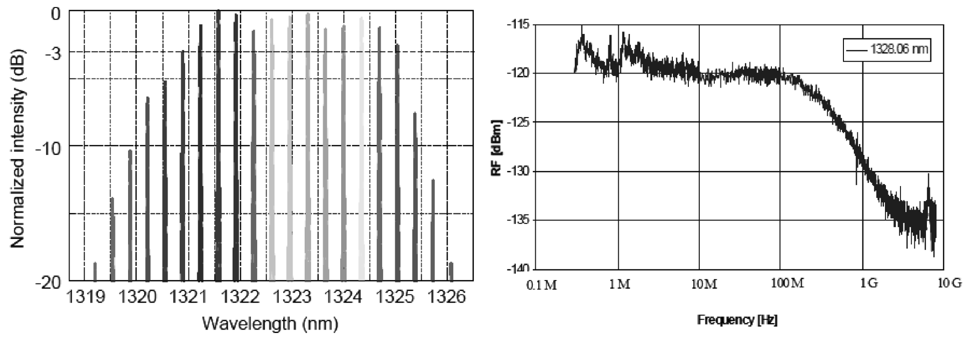

Fig. 23. Electrically-pumped inhomogeneous quantum-dot gain regions can provide laser amplification over a broad spectrum, allowing a Fabry-Perot laser to generate a comb of frequencies with a channel spacing that depends only on the geometry of the resonator (left). The relative intensity noise of a given channel (right) shows that the coupling between adjacent channels is negligibly small. (Figure reproduced from Wojcik et al. [2009] with permission © 2009 SPIE.)

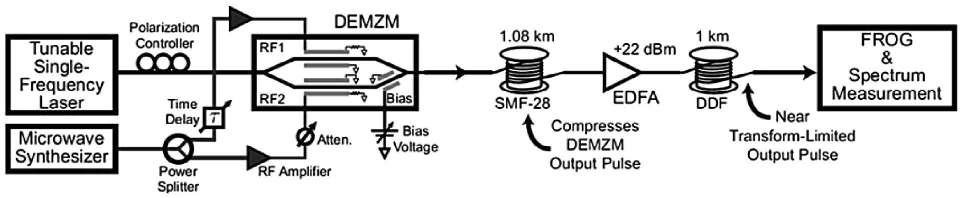

(a) Schematic of the frequency comb generator

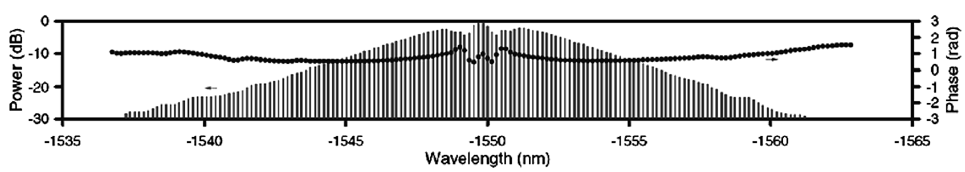

(b) Comb channel power and phase

Fig. 24. Broadband frequency-comb generation for "optical power supply" of a DWDM interconnect system. (a) Schematic of a frequency-comb generation system comprised of a dual-electrode Mach-Zehnder modulator and dispersion-decreasing fiber for pulse chirp reduction. (b) Plot of a comb of 175 frequency channels near 1550nm spanning a 3.5 THz optical spectrum with nearly flat spectral phase. (Figure reproduced from Zhou et al. [2009] with permission © 2009 by the authors.)

chirped, so dispersion engineering in both conventional optical fiber and dispersion-decreasing fiber was used to regularize the phase profile of the comb. This approach has the advantage that the laser source can be chosen primarily for efficiency, output power, and spectral stability, and the active and dispersive elements could be monolithically integrated [Zhou et al. 2009a; Osgood et al. 2009]. However, the mode spacing is limited by the modulation frequency of the DEMZM, and the laser output (after experiencing the gain provided by the optical amplifier located between the two lengths of fiber) must be large enough that ~1W power consumption of the DEMZM does not significantly reduce the efficiency of the system.

6. PHOTONIC INTEGRATED CIRCUITS: INTEGRATION AND PACKAGING

Given that long-haul telecommunication has been the driver for much of the progress made in the field of photonics over the last thirty years, the most common platform for photonic integrated circuits (PICs) has been indium phosphide (InP) [Razavi 2002;

Nagarajan et al. 2007]. Over the next decade, as bandwidth requirements for each optomechanical connector grow to 1Tb/s and beyond, distances demanding high-bandwidth interconnects shrink from meters to centimeters to millimeters, and the cost per optical component falls from US$100–1000 to US$0.001–1, silicon as a platform for PICs has become a strong contender [Gunn 2006; Liang and Bowers 2009; Hochberg and Baehr-Jones 2010]. Moving from InP to Si PICs could inspire a "Moore's Law" for photonics, and provide enormous leverage for dramatic improvements in integrated optical technologies: annual sales for the global semiconductor industry have grown to US$300B in 2010 [Gartner 2010], and historically investments in capital equipment and R&D have been 10–15% and 15–20% of sales, respectively.

6.1. Integration: SOI as a PIC Platform

As discussed in Section 3, SOI is a promising PIC platform because it provides a 1–3μm-thick low-index oxide substrate below a thin silicon top layer, providing an ideal high-index-contrast environment for strongly confined optical waveguides. Much of the progress in SOI PICs over the past decade has been made on individual devices or partially integrated optical systems built using EBL [Chen et al. 2009], but there has been substantial recent progress in moving photonics to low-cost commodity CMOS-compatible processes. A collaboration of MIT, BAE Systems, Cornell, and Lucent has demonstrated monolithically-integrated tunable optical filters, germanium electro-absorption modulators, and germanium p-i-n diode photodetectors using silicon wire waveguides within a full 150nm CMOS process flow [Beals et al. 2008]. Luxtera demonstrated a 40Gb/s transceiver fabricated using a 130nm CMOS process in SOI [Gunn 2006; Pinguet et al. 2008], made possible by high-speed operation of electronic driver circuits built in SOI as well as photonic devices with submicron feature sizes enabled by high-volume deep-UV photolithography. Intel demonstrated a WDM transmitter using rib waveguides on SOI [Liu et al. 2010a] that has eight channels capable of operating at 25Gb/s each, with integrated modulators and DEMUX/MUX based on Mach-Zehnder interferometers (MZIs). The modulators were reversed-biased p-n diodes operated in carrier depletion mode (discussed in Section 4), and the DEMUX and MUX used a cascaded asymmetric MZI design that allowed the entire device to have single fiber input and output. A 5Gb/s single-channel transmitter based on an integrated microring modulator flip-chip bonded to an electrical driver circuit was fabricated using a 130nm SOI CMOS process by Sun and Luxtera [Zheng et al. 2010], who showed that the entire device dissipated 400fJ/b at CMOS drive voltages.

The incorporation of germanium in silicon PICs complicates prospects for monolithic CMOS integration using SOI [Michel et al. 2010]. Previous demonstrations have relied on a "front-end-of-line" (FEOL) approach, where after transistor fabrication SOI wafers are shuttled from a silicon foundry (which usually does not have the capacity to grow pure germanium) to a Ge epitaxy facility to create photodetectors in the same layer as the Si electronics [Gunn 2006; Beals et al. 2008; Pinguet et al. 2008; Liow et al. 2010; Michel et al. 2010]. Although in the near term this approach is likely to remain dominant for CMOS integration, in the long term a "back-end-of-line" (BEOL) process, where photodetectors in amorphous global interconnect layers could be connected to Si transistors through metal vias, would be preferable [Michel et al. 2010]. Over the next decade, as wafer thinning and bonding technologies reaches mainstream semiconductor products, a 3D die-stacking approach [Black et al. 2006; Vantrease et al. 2008; Ahn et al. 2009a] that connected separate electronic and photonic thinned wafers using through-silicon vias [ITRS 2009] would allow separate optimizations of large-scale integrated photonics and electronics (even those based on entirely different substrates).

It remains to be seen whether the SiGe lasers discussed in Section 5.1.2 can operate with multi-milliwatt output powers and $\sim$25% overall efficiency necessary for adoption in high-bandwidth computer interconnects, but their prospects for wafer-scale integration are similar to those of epitaxial Ge photodetectors on SOI. The III-V frequency comb sources reviewed in Section 5.3 are designed to be used as stand-alone packages, and will need to reach economies of scale (e.g., cost per unit power per channel) that are unprecedented for the III-V semiconductor industry that currently provides lasers for telecommunications. The hybrid III-V/Si lasers described in Section 5.2 are approaching the complexity and level of integration of the most advanced III-V photonic devices [Liang and Bowers 2009]. III-V epitaxial transfer of wafers as large as 150mm in diameter [Liang et al. 2009a] and III-V die attachment to larger SOI wafers [Fedeli et al. 2008; Liu et al. 2010c] indicate the potential of this approach for high-volume manufacturability. Intel's recent demonstration of a CWDM 50-Gb/s link—incorporating hybrid lasers, low-voltage MZMs, passive silicon MUX and DEMUX, and SiGe photodetectors—is a tour de force of integration performed in standard fabrication facilities [Alduino et al. 2010]. Despite this progress, it remains difficult to predict the long-term reliability, performance, device uniformity, and cost of the hybrid approach [Liang and Bowers 2010].

6.2. Trimming of Fabrication Imperfections in DWDM Systems

The feasibility of silicon as a platform for photonic integrated circuits requires that devices are fabricated uniformly at both the die and wafer scales [Hochberg and Baehr-Jones 2010]. Given spatial variations in feature sizes that are typical of CMOS materials and fabrication processes, the impact on the performance characteristics of photonic devices is likely to be much greater than that of electronic devices [Selvaraja et al. 2010], particularly in the case of resonant WDM systems [Peng et al. 2010]. For example, the resonance frequencies of microring resonators depend on the dimensions and effective refractive index of the waveguide forming the ring, and the simple discretization of identical microrings at different locations within the mask reticle can cause significant variation in their passband locations. Variations in etch depths across a wafer that would have little impact on the operation of a transistor could cause a significant shift in the resonance wavelengths of microrings formed from ridge waveguides. Batch-to-batch changes in the thickness of the top silicon of an SOI wafer—as well as perturbations in this thickness across a single wafer—can cause problems even if the etching is perfectly uniform.

To gain an appreciation for the magnitude of the problem for DWDM systems, consider a silicon wire waveguide operating at 1550nm that is 500nm wide and 250nm high on SOI. The effective group index of this waveguide is approximately $n_g \approx 4$, and the sensitivity of this index to small changes in the height of the waveguide is $dn/dh \approx 3.3 \times 10^{-3}\text{nm}^{-1}$ [Selvaraja et al. 2010]. From the discussion in Section 3.2.2, the shift in resonance m of a microring resonator formed from this wire waveguide is $\Delta\lambda_m = (\lambda_m/n_g)(dn/dh)\Delta h$ for a change in height Δh, or equivalently

$$\frac{\Delta f_m}{\Delta h} = \frac{c}{n_g \lambda_m}\frac{dn}{dh} \approx 160\text{GHz/nm}. \tag{19}$$

The worst-case 3-σ batch-to-batch plus intra-wafer edge-to-edge variation in the thickness of the top silicon layer of an 8-inch nominally 250nm Thick Unibond SOI wafer available from Soitec [Soitec 2010] is $\pm 5\%$, or ± 12.5nm, corresponding to a variation in microring resonance frequency of $\Delta f_m \approx \pm 2$THz. Recent measurements of microring resonators operating at 1310nm and fabricated using 248nm photolithography are consistent with a 3-σ edge-to-edge wafer thickness nonuniformity of 6nm peak-to-peak [Peng et al. 2010].

Although as design rules for future lithographic technology nodes are adopted the problem may not be as severe for large-footprint nonresonant silicon photonic devices [Selvaraja et al. 2010], it seems clear that some degree of "trimming" will be needed to bring resonant devices in DWDM systems into compliance with their design [Zortman et al. 2010b]. One approach may be to use negative-thermooptic polymers as a cladding material and/or in slot waveguides to athermalize the microrings, as described in Section 3.2.2. Alternatively, as part of postprocessing a PIC wafer, the characteristics of individual devices could be diagnosed [Cooper et al. 2010] and then corrected by exposing the polymer to UV light [Zhou et al. 2009b]. The most straightforward approach would use a slow dynamic process to tune a monitored microring such as (for example) a liquid crystal cladding [Cort et al. 2009] or an integrated resistive heater [Watts et al. 2009; Dong et al. 2010b; DeRose et al. 2010; Dong et al. 2010d; Zortman et al. 2010b]. Integrated heaters have been demonstrated with tuning ranges as large as 32nm and response times as short as 1μs, at power cost of 4.4μW/GHz. Since $dn/dT > 0$ for silicon, heating a microring will redshift (i.e., decrease) the resonance frequency of a ring resonator. Alternatively, from Equation (12a), injecting carriers into a microring will blueshift the ring. Using the DC shift data provided in [Xu et al. 2005] as a guide, the blue shift caused by charge injection for a 5-μm ring biased at 0.8V would dissipate approximately 125μW/nm (or 1μW/GHz) at 1550nm [Ahn et al. 2009a]. It is clear that these corrections represent a nontrivial contribution to the total system power budget if tighter control in the uniformity of fabricated rings can't be achieved.

6.3. Packaging: Fiber-to-Waveguide Couplers

One of the most critical technologies that must be made cost effective before silicon photonics can reach volume applications is the packaging of an interface between the small transverse mode of a high-index-contrast wire or rib waveguide (with typical area 0.10–$0.15\mu m^2$) and the larger transverse mode profile of a low-index-contrast waveguide such as an optical fiber (with typical area $50\mu m^2$). A practical solution—which may require an optomechanical mount with either active [Zhang et al. 2010b] or passive alignment [Zimmermann et al. 2008] in addition to an optical mode convertor—should compete with the cost of a typical data pin (~US$0.01–0.02 per pin) on a standard electronic IC package. In high-bandwidth applications requiring WDM, the optical connector/convertor cost may be amortized over the total number of data channels, resulting in a potentially significant reduction in packaging costs.

The adiabatically tapered spot-size convertor shown in Figure 25(a) first couples the high-index-contrast wire waveguide mode to a larger low-index-contrast waveguide before launching the mode into a laterally aligned butt-coupled (and possibly lensed) optical fiber [Almeida et al. 2003; Gunn 2006; Tsuchizawa et al. 2005]. The device in the figure has a silicon taper with a length of at least 200μm and a tip that is typically less than 100nm wide. The polymer waveguide core has $3\mu m \times 3\mu m$ transverse extent and a 3% index contrast. Coupling losses of ~0.5dB have been reported for tapered fiber with a mode size of 4.3μm, but this loss increases to 2.5dB for standard single-mode fiber with a 9μm core [Tsuchizawa et al. 2005] (compared to an insertion loss of 20dB in the absence of a spot-size converter). In addition, the lateral alignment tolerance for these couplers is generally less than 1μm for an additional loss of 1dB [Almeida et al. 2003; Zimmermann et al. 2008].

Recent work on vertical mode conversion using high-index-contrast subwavelength gratings in the configuration shown in Figure 25(b) has improved matching between wire waveguide and standard single-mode fiber with relaxed lateral alignment requirements [Van Laere et al. 2007a; Van Laere et al. 2007b; Roelkens et al. 2008; Zimmermann et al. 2008; Tang et al. 2010; Vermeulen et al. 2010]. Because of the high index, the grating can be relatively short (~25 periods) while providing a relatively

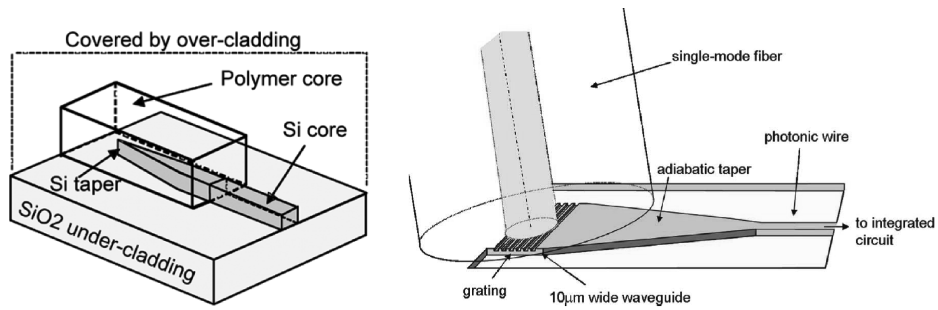

(a) Spot-size converter (b) Experimental reflected mode profiles

Fig. 25. Mode convertors coupling optical fibers to silicon wire waveguides. (a) A schematic of an adiabatically tapered spot-size converter in SOI. (b) Schematic of an adiabatically tapered subwavelength-grating single-mode fiber-to-Si wire waveguide coupler in SOI. (Tsuchizawa et al. [2005] and Van Laere et al. [2007b], respectively, with permission © 2005 IEEE.)

large wavelength band where coupling is efficient. The grating is designed to operate with light incident at an angle of about 10° to minimize backreflection into the fiber. A subwavelength grating coupler fabricated in SOI using a CMOS pilot line with 193nm DUV photolithography has been demonstrated with a peak coupling efficiency of about 70% [Vermeulen et al. 2010] (similar to that reported in Pinguet et al. [2008]), a 3dB coupling bandwidth of 80nm, and a 1dB misalignment tolerance of 2μm. The primary advantage of this approach over that of the end-coupled adiabatic taper is that light can be coupled onto the chip at any location that is convenient, which allows both wafer-scale testing and packaging layout optimization.

Although a mechanically fixed connection between the cores of a single-mode waveguide and a single-mode fiber is the most direct solution to the packaging of an optical data "pin," in some applications it may be economically more feasible to couple spatially separated reference planes using a free-space optical interposer. Traditionally this has been done using costly compound glass or plastic lenses with volumes on the order of a few cubic millimeters, but recently a compact optic with strong focusing power has been demonstrated that can be fabricated using planar CMOS fabrication technology [Fattal et al. 2010]. Following earlier work on high-index contrast subwavelength gratings (SWGs) for VCSELs [Huang et al. 2007, 2008], nonperiodic dielectric structures [Gerke and Piestun 2010] have been shown to shape optical phasefronts to steer or focus incident light fields. Figure 26(a) is a microscope image of a mirror with a diameter of 150μm fabricated in 450nm of amorphous silicon on quartz. The mirror is formed in a single etch step from a grating with a fixed period of 670nm, with grooves that vary in width as a function of position. As shown in Figure 26(b), flat silicon cylindrical and spherical mirrors imaged an incident collimated beam with a focal length of 20mm, and more complicated groove patterns can focus light to diffraction-limited spots in both mirror and lens configurations [Fattal et al. 2010].

7. CONCLUSION

Research on individual silicon photonic components has been so productive and rapid over the last decade that the primary uncertainty in predicting system-level performance has become the trade-offs which must be made to enable wafer-scale integration. If SOI is ultimately the preferred platform for integration, then both batch-to-batch and intra-wafer edge-to-edge variation of commercially available wafers must be more tightly controlled. Uniformity of devices after lithography and etching—particularly resonant structures such as microring filters and modulators—will determine the

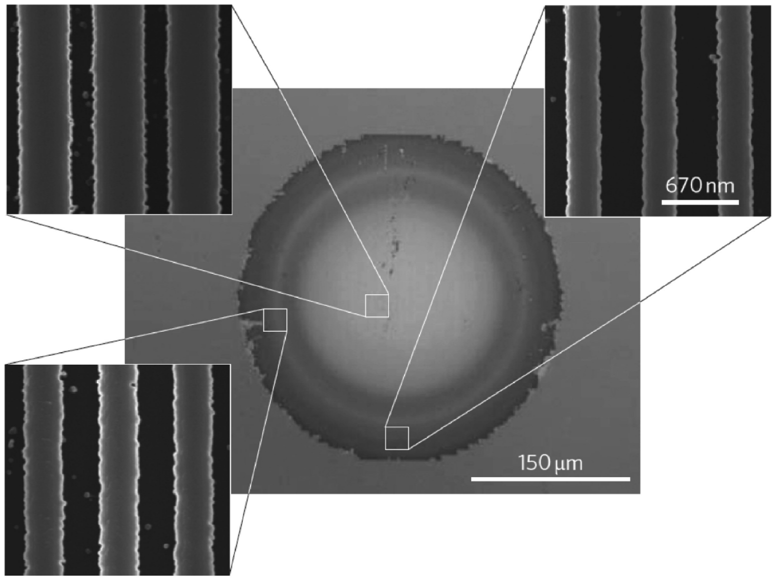

(a) Microscope image of a SWG mirror

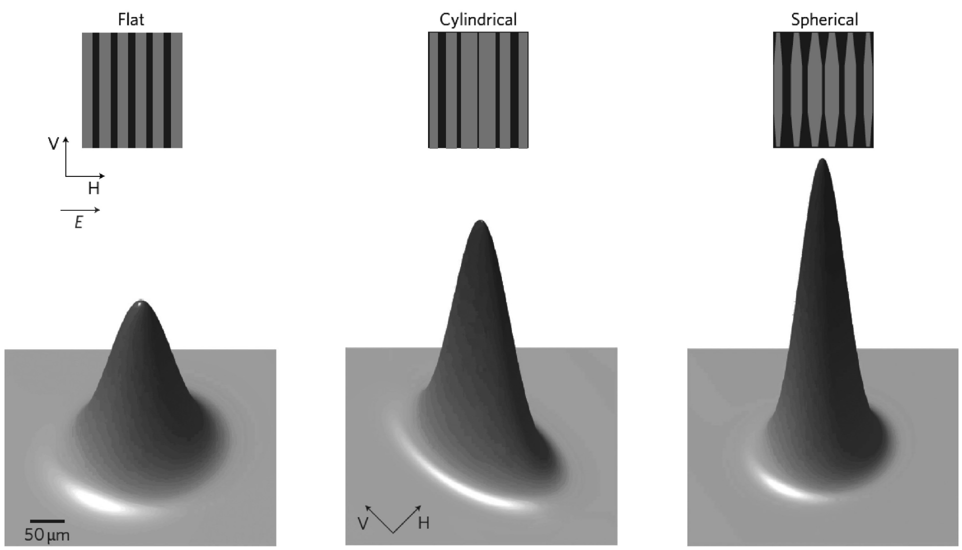

(b) Experimental reflected mode profiles

Fig. 26. Planar nonperiodic subwavelength grating reflector with focusing power. (a) Optical microscope picture of a fabricated spherical SWG mirror, with SEM insets showing the groove width in particular locations. (b) Schematic of the groove distributions flat (left), cylindrical (middle) and spherical (right) SWG mirrors, designed for electric field polarizations oriented perpendicular to the grooves, with measured beam profiles at the foci for the three mirrors. (Figure reproduced from Fattal et al. [2010] with permission © 2010 NPG.)

extent to which static post-processing (which adds cost and complexity) or dynamic trimming (which contributes to the energy consumption of the interconnect) are necessary. Once these effects are well understood, then other trade-offs can be brought forward for consideration. For example, a low-cost multiwavelength continuous-wave light source with a fixed channel spacing is vital as an “optical power supply” in a DWDM system. There are several promising candidates, either under active research or commercially available in prototype form, but until a route to a low-cost high-volume product can be identified it is difficult to specify the details of the on-chip DWDM components.

Therefore, although a strong case can be made using raw numbers and benchmark models that future integrated circuits must inevitably rely on large-scale integrated photonics for global-level interconnects, it is not yet clear exactly how this radically new technology will reach the market. To limit costs, supercomputers and data centers tend to rely on commodity components, which are unlikely to incorporate photonics without a sustained collaboration of multinational semiconductor and information technology companies and government funding agencies. It is possible that a nearer-term realization of a monolithically-integrated PhASIC—such as an optically-enabled network switch—could have a substantial impact on high-performance computing. In either case, the Communications Technology Roadmap effort led by the MIT Microphotonics Center [MIT 2010] should be encouraged and recognized as a critical step towards bringing industry and academia together to establish a Moore’s Law for photonics.

ACKNOWLEDGMENTS

This work has benefited greatly from discussions and collaboration with J. Ahn, G. Astfalk, N. Binkert, J. E. Bowers, A. Davis, D. Fattal, M. Fiorentino, N. P. Jouppi, D. Liang, M. McLaren, T. Morris, N. Muralimanohar, C. M. Santori, R. S. Schreiber, M. R. T. Tan, D. Vantrease, S. Y. Wang, R. S. Williams, A. E. Willner, Q. Xu, and L. Zhang. Particular thanks are due to three anonymous referees who read the entire manuscript and made many constructive comments that significantly improved the content and presentation of this material.

REFERENCES

AHN, D., YIN HONG, C., LIU, J., GIZIEWICZ, W., BEALS, M., KIMERLING, L. C., MICHEL, J., CHEN, J., AND KÄRTNER, F. X. 2007. High performance, waveguide integrated Ge photodetectors. *Opt. Express 15,* 7, 3916–3921.

AHN, J., FIORENTINO, M., BEAUSOLEIL, R. G., BINKERT, N., DAVIS, A., FATTAL, D., JOUPPI, N. P., MCLAREN, M., SANTORI, C. M., SCHREIBER, R. S., SPILLANE, S. M., VANTREASE, D., AND XU, Q. 2009a. Devices and architectures for photonic chip-scale integration. *Appl. Phys. A 95*, 989–997.

AHN, J. H., BINKERT, N., DAVIS, A., MCLAREN, M., AND SCHREIBER, R. S. 2009b. HyperX: topology, routing, and packaging of efficient large-scale networks. In *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis (SC)*. ACM, New York, NY, 1–11.

ALDUINO, A., LIAO, L., JONES, R., MORSE, M., KIM, B., LO, W.-Z., BASAK, J., KOCH, B., LIU, H.-F., RONG, H., SYSAK, M., KRAUSE, C., SABA, R., LAZAR, D., HORWITZ, L., BAR, R., LITSKI, S., LIU, A., SULLIVAN, K., DOSUNMU, O., NA, N., YIN, T., HAUBENSACK, F., WEI HSIEH, I., HECK, J., BEATTY, R., PARK, H., BOVINGTON, J., LEE, S., NGUYEN, H., AU, H., NGUYEN, K., MERANI, P., HAKAMI, M., AND PANICCIA, M. 2010. Demonstration of a high speed 4-channel integrated silicon photonics WDM link with hybrid silicon lasers. In *Proceedings of the OSA Topical Meeting on Integrated Photonics Research (IPR)*. Optical Society of America, PDIWI5.

ALMEIDA, V. R., PANEPUCCI, R. R., AND LIPSON, M. 2003. Nanotaper for compact mode conversion. *Opt. Lett. 28,* 15, 1302–1304.

ASSEFA, S., XIA, F., AND VLASOV, Y. A. 2010. Reinventing germanium avalanche photodetector for nanophotonic on-chip optical interconnects. *Nature 464,* 7285, 80–84.

ASTFALK, G. 2009. Why optical data communications and why now? *Appl. Phys. A 95*, 933–940.

BANERJEE, K. AND MEHROTRA, A. 2002. A power-optimal repeater insertion methodology for global interconnects in nanometer designs. *IEEE Trans. Electron. Dev. 49,* 11, 2001–2007.

BARRIOS, C. A. AND LIPSON, M. 2004. Modeling and analysis of high-speed electro-optic modulation in high confinement silicon waveguides using metal-oxide-semiconductor configuration. *J. Appl. Phys. 96,* 11, 6008–6015.

BATTEN, C., JOSHI, A., ORCUTT, J., KHILO, A., MOSS, B., HOLZWARTH, C., POPOVIC, M., LI, H., SMITH, H., HOYT, J., KARTNER, F., RAM, R., STOJANOVIC, V., AND ASANOVIC, K. 2009. Building many-core processor-to-DRAM networks with monolithic CMOS silicon photonics. *IEEE Micro 29,* 4, 8–21.

BEALS, M., MICHEL, J., LIU, J. F., AHN, D. H., SPARACIN, D., SUN, R., HONG, C. Y., KIMERLING, L. C., POMERENE, A., CAROTHERS, D., BEATTIE, J., KOPA, A., APSEL, A., RASRAS, M. S., GILL, D. M., PATEL, S. S., TU, K. Y., CHEN, Y. K., AND WHITE, A. E. 2008. Process flow innovations for photonic device integration in CMOS. *Proc. Soc. Photo-Opt. Instrum. Eng. 6898,* 1, 689804.

BEAUSOLEIL, R. G., AHN, J., BINKERT, N., DAVIS, A., FATTAL, D., FIORENTINO, M., JOUPPI, N. P., MCLAREN, M., SANTORI, C. M., SCHREIBER, R. S., SPILLANE, S. M., VANTREASE, D., AND XU, Q. 2008. A nanophotonic interconnect for high-performance many-core computation. *IEEE LEOS Newslett. 22,* 3, 15–22.

BEAUSOLEIL, R. G., KUEKES, P. J., SNIDER, G. S., WANG, S.-Y., AND WILLIAMS, R. S. 2008. Nanoelectronic and nanophotonic interconnect (Invited Paper). *Proc. IEEE 96,* 2, 230–247.

BENNER, A. 2009. Cost-effective optics: Enabling the exascale roadmap. In *Proceedings of the 17th IEEE Symposium on High Performance Interconnects (HOTI)*. 133–137.

BERGMAN, K. AND CARLONI, L. 2008. Power efficient photonic networks on-chip. *Proc. Soc. Photo-Opt. Instrum. Eng. 6898*, 689813.

BHATNAGAR, A., DEBAES, C., THIENPONT, H., AND MILLER, D. A. B. 2004. Receiverless detection schemes for optical clock distribution. *Proc. Soc. Photo-Opt. Instrum. Eng. 5359*, 352–359.

BLACK, B., ANNAVARAM, M., BREKELBAUM, N., DEVALE, J., JIANG, L., LOH, G. H., MCCAULEY, D., MORROW, P., NELSON, D. W., PANTUSO, D., REED, P., RUPLEY, J., SHANKAR, S., SHEN, J., AND WEBB, C. 2006. Die stacking (3D) microarchitecture. In *Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO-39)*. 469–479.

BLAND, A. S., KENDALL, R. A., KOTHE, D. B., ROGERS, J. H., AND SHIPMAN, G. M. 2009. Jaguar: The worlds most powerful computer. In *Proceedings of the Cray User Group*.

BLOCK, B. A., YOUNKIN, T. R., DAVIDS, P. S., RESHOTKO, M. R., CHANG, P., POLISHAK, B. M., HUANG, S., LUO, J., AND JEN, A. K. Y. 2008. Electro-optic polymer cladding ring resonator modulators. *Opt. Express 16,* 22, 18326–18333.

BOCK, P. J., CHEBEN, P., SCHMID, J. H., LAPOINTE, J., DELÂGE, A., JANZ, S., AERS, G. C., XU, D.-X., DENSMORE, A., AND HALL, T. J. 2010a. Subwavelength grating periodic structures in silicon-on-insulator: a new type of microphotonic waveguide. *Opt. Express 18,* 19, 20251–20262.

BOCK, P. J., CHEBEN, P., SCHMID, J. H., LAPOINTE, J., DELÂGE, A., XU, D.-X., JANZ, S., DENSMORE, A., AND HALL, T. J. 2010b. Subwavelength grating crossings for silicon wire waveguides. *Opt. Express 18,* 15, 16146–16155.

BOGAERTS, W., DUMON, P., THOURHOUT, D. V., AND BAETS, R. 2007. Low-loss, low-cross-talk crossings for silicon-on-insulator nanophotonic waveguides. *Opt. Lett. 32,* 19, 2801–2803.

BOGAERTS, W., SELVARAJA, S., DUMON, P., BROUCKAERT, J., DE VOS, K., VAN THOURHOUT, D., AND BAETS, R. 2010. Silicon-on-insulator spectral filters fabricated with CMOS technology. *IEEE J. Select. Topics Quantum Electron. 16,* 1, 33–44.

BOYRAZ, O. AND JALALI, B. 2004. Demonstration of a silicon Raman laser. *Opt. Express 12,* 21, 5269–5273.

CHEBEN, P., BOCK, P. J., SCHMID, J. H., LAPOINTE, J., JANZ, S., XU, D.-X., DENSMORE, A., DELÂGE, A., LAMONTAGNE, B., AND HALL, T. J. 2010. Refractive index engineering with subwavelength gratings for efficient microphotonic couplers and planar waveguide multiplexers. *Opt. Lett. 35,* 15, 2526–2528.

CHEN, C., HARGIS, M., WOODALL, J., MELLOCH, M., REYNOLDS, J., YABLONOVITCH, E., AND WANG, W. 1999. GHz bandwidth GaAs light-emitting diodes. *Appl. Phys. Lett. 74*, 3140.

CHEN, L. AND LIPSON, M. 2009. Ultra-low capacitance and high speed germanium photodetectors on silicon. *Opt. Express 17,* 10, 7901–7906.

CHEN, L., PRESTON, K., MANIPATRUNI, S., AND LIPSON, M. 2009. Integrated GHz silicon photonic interconnect with micrometer-scale modulators and detectors. *Opt. Express 17,* 17, 15248–15256.

CHU, T., FUJIOKA, N., AND ISHIZAKA, M. 2009. Compact, lower-power-consumption wavelength tunable laser fabricated with silicon photonic-wire waveguide micro-ring resonators. *Opt. Express 17,* 16, 14063–14068.

CLAPS, R., DIMITROPOULOS, D., HAN, Y., AND JALALI, B. 2002. Observation of Raman emission in silicon waveguides at 1.54μm. *Opt. Express 10,* 22, 1305–1313.

COLWELL, R. 2007. Computer Architecture Futures 2007. In *Proceedings of the Federated Computing Research Conference*.

COOPER, M. L., GUPTA, G., PARK, J. S., SCHNEIDER, M. A., DIVLIANSKY, I. B., AND MOOKHERJEA, S. 2010. Quantitative infrared imaging of silicon-on-insulator microring resonators. *Opt. Lett. 35,* 5, 784–786.

CORT, W. D., BEECKMAN, J., JAMES, R., FERNÁNDEZ, F. A., BAETS, R., AND NEYTS, K. 2009. Tuning of silicon-on-insulator ring resonators with liquid crystal cladding using the longitudinal field component. *Opt. Lett. 34,* 13, 2054–2056.

D'ANDREA, D. 2009. CMOS photonics today & tomorrow. In *Proceedings of the Optical Fiber Communication Conference (OFC)*. Optical Society of America.

DeRose, C. T., Watts, M. R., Trotter, D. C., Luck, D. L., Nielson, G. N., and Young, R. W. 2010. Silicon microring modulator with integrated heater and temperature sensor for thermal control. In *Proceedings of the Conference on Lasers and Electro-Optics (CLEO)*.

Doany, F., Schow, C., Baks, C., Kuchta, D., Pepeljugoski, P., Schares, L., Budd, R., Libsch, F., Dangel, R., Horst, F., Offrein, B., and Kash, J. 2009. 160 Gb/s bidirectional polymer-waveguide board-level optical interconnects using CMOS-based transceivers. *IEEE Trans. Adv. Pack. 32,* 2, 345–359.

Dong, P., Liao, S., Feng, D., Liang, H., Zheng, D., Shafiiha, R., Kung, C.-C., Qian, W., Li, G., Zheng, X., Krishnamoorthy, A. V., and Asghari, M. 2009. Low vpp, ultralow-energy, compact, high-speed silicon electro-optic modulator. *Opt. Express 17,* 25, 22484–22490.

Dong, P., Liao, S., Liang, H., Qian, W., Wang, X., Shafiiha, R., Feng, D., Li, G., Zheng, X., Krishnamoorthy, A. V., and Asghari, M. 2010a. High-speed and compact silicon modulator based on a racetrack resonator with a 1 V drive voltage. *Opt. Lett. 35,* 19, 3246–3248.

Dong, P., Qian, W., Liang, H., Shafiiha, R., Feng, N.-N., Feng, D., Zheng, X., Krishnamoorthy, A. V., and Asghari, M. 2010b. Low power and compact reconfigurable multiplexing devices based on silicon microring resonators. *Opt. Express 18,* 10, 9852–9858.

Dong, P., Qian, W., Liao, S., Liang, H., Kung, C.-C., Feng, N.-N., Shafiiha, R., Fong, J., Feng, D., Krishnamoorthy, A. V., and Asghari, M. 2010c. Low loss shallow-ridge silicon waveguides. *Opt. Express 18,* 14, 14474–14479.

Dong, P., Shafiiha, R., Liao, S., Liang, H., Feng, N.-N., Feng, D., Li, G., Zheng, X., Krishnamoorthy, A. V., and Asghari, M. 2010d. Wavelength-tunable silicon microring modulator. *Opt. Express 18,* 11, 10941–10946.

Essiambre, R.-J., Foschini, G. J., Kramer, G., and Winzer, P. J. 2008. Capacity limits of information transport in fiber-optic networks. *Phys. Rev. Lett. 101,* 16, 163901.

Fang, A., Sysak, M., Koch, B., Jones, R., Lively, E., Kuo, Y.-H., Liang, D., Raday, O., and Bowers, J. 2009. Single-wavelength silicon evanescent lasers. *IEEE J. Select. Topics Quantum Electron. 15,* 3, 535–544.

Fang, A. W., Jones, R., Park, H., Cohen, O., Raday, O., Paniccia, M. J., and Bowers, J. E. 2007. Integrated AlGaInAs-silicon evanescent race track laser and photodetector. *Opt. Express 15,* 5, 2315–2322.

Fang, A. W., Park, H., Cohen, O., Jones, R., Paniccia, M. J., and Bowers, J. E. 2006. Electrically pumped hybrid AlGaInAs-silicon evanescent laser. *Opt. Express 14,* 20, 9203–9210.

Fang, Q., Phang, Y. T., Tan, C. W., Liow, T.-Y., Yu, M. B., Lo, G. Q., and Kwong, D. L. 2010. Multi-channel silicon photonic receiver based on ring-resonators. *Opt. Express 18,* 13, 13510–13515.

Fattal, D., Fiorentino, M., Tan, M., Houng, D., Wang, S. Y., and Beausoleil, R. G. 2008. Design of an efficient light-emitting diode with 10 GHz modulation bandwidth. *Appl. Phys. Lett. 93,* 24, 243501.

Fattal, D., Li, J., Peng, Z., Fiorentino, M., and Beausoleil, R. G. 2010. Flat dielectric grating reflectors with focusing abilities. *Nature Photon. 4,* 7, 466–470.

Fedeli, J. M., Cioccio, L. D., Marris-Morini, D., Vivien, L., Orobtchouk, R., Rojo-Romeo, P., Seassal, C., and Mandorlo, F. 2008. Development of silicon photonics devices using microelectronic tools for the integration on top of a CMOS wafer. *Adv. Opt. Tech.* 412518.

Feng, N.-N., Dong, P., Zheng, D., Liao, S., Liang, H., Shafiiha, R., Feng, D., Li, G., Cunningham, J. E., Krishnamoorthy, A. V., and Asghari, M. 2010a. Vertical p-i-n germanium photodetector with high external responsivity integrated with large core Si waveguides. *Opt. Express 18,* 1, 96–101.

Feng, N.-N., Liao, S., Feng, D., Dong, P., Zheng, D., Liang, H., Shafiiha, R., Li, G., Cunningham, J. E., Krishnamoorthy, A. V., and Asghari, M. 2010b. High speed carrier-depletion modulators with 1.4 V-cm $V_\pi L$ integrated on 0.25μm silicon-on-insulator waveguides. *Opt. Express 18,* 8, 7994–7999.

Fischer, U., Zinke, T., Kropp, J.-R., Arndt, F., and Petermann, K. 1996. 0.1 dB/cm waveguide losses in single-mode SOI rib waveguides. *IEEE Photon. Technol. Lett. 8,* 5, 647–648.

Foster, M. A., Salem, R., Geraghty, D. F., Turner-Foster, A. C., Lipson, M., and Gaeta, A. L. 2008. Silicon-chip-based ultrafast optical oscilloscope. *Nature 456,* 7218, 81–84.

Foster, M. A., Salem, R., Okawachi, Y., Turner-Foster, A. C., Lipson, M., and Gaeta, A. L. 2009. Ultrafast waveform compression using a time-domain telescope. *Nature Photon. 3,* 10, 581–585.

Foster, M. A., Turner, A. C., Sharping, J. E., Schmidt, B. S., Lipson, M., and Gaeta, A. L. 2006. Broad-band optical parametric gain on a silicon photonic chip. *Nature 441,* 7096, 960–963.

Frey, B. J., Leviton, D. B., and Madison, T. J. 2006. Temperature-dependent refractive index of silicon and germanium. *Proc. Soc. Photo-Opt. Instrum. Eng. 6273,* 62732J, 1–10.

Fukuda, H., Yamada, K., Tsuchizawa, T., Watanabe, T., Shinojima, H., and ichi Itabashi, S. 2008. Polarization rotator based on silicon wire waveguides. *Opt. Express 16,* 4, 2628–2635.

Gardes, F. Y., Reed, G. T., Mashanovich, G. Z., and Png, C. E. 2008. Optical modulators in silicon photonic circuits. In *Silicon Photonics: The State of the Art,* G. T. Reed (ed.), Wiley. 95–146.

GARTNER. 2010. Press Release: “Gartner Says Worldwide Semiconductor Revenue to Grow 31.5 Percent in 2010”. http://www.gartner.com/it/page.jsp?id=1430013.

GERKE, T. D. AND PIESTUN, R. 2010. Aperiodic volume optics. *Nature Photon. 4,* 3, 188–193.

GILL, D., PATEL, S., RASRAS, M., TU, K.-Y., WHITE, A., CHEN, Y.-K., POMERENE, A., CAROTHERS, D., KAMOCSAI, R., HILL, C., AND BEATTIE, J. 2010. CMOS-compatible Si-ring-assisted Mach-Zehnder interferometer with internal bandwidth equalization. *IEEE J. Select. Topics Quantum Electron. 16,* 1, 45–52.

GLICK, M. 2008. Optical interconnects in next generation data centers: An end to end view. In *Proceedings of the 16th IEEE Symposium on High Performance Interconnects (HOTI)*. 178–181.

GNAN, M., THORNS, S., MACINTYRE, D., DE LA RUE, R., AND SOREL, M. 2008. Fabrication of low-loss photonic wires in silicon-on-insulator using hydrogen silsesquioxane electron-beam resist. *Electron. Lett. 44,* 2, 115–116.

GNAUCK, A. H., TKACH, R. W., CHRAPLYVY, A. R., AND LI, T. 2008. High-capacity optical transmission systems. *J. Lightwave Technol. 26,* 9, 1032–1045.

GREEN, W. M., ROOKS, M. J., SEKARIC, L., AND VLASOV, Y. A. 2007. Ultra-compact, low RF power, 10 Gb/s silicon Mach-Zehnder modulator. *Opt. Express 15,* 25, 17106–17113.

GREEN500. 2010. The Green500 List: Environmentally Responsible Supercomputing. http://www.green500.org.

GUBENKO, A., KRESTNIKOV, I., LIVSHTIS, D., MIKHRIN, S., KOVSH, A., WEST, L., BORNHOLDT, C., GROTE, N., AND ZHUKOV, A. 2007. Error-free 10 Gbit/s transmission using individual Fabry-Perot modes of low-noise quantum-dot laser. *Electron. Lett. 43,* 25, 1430–1431.

GUHA, B., KYOTOKU, B. B. C., AND LIPSON, M. 2010. CMOS-compatible athermal silicon microring resonators. *Opt. Express 18,* 4, 3487–3493.

GUNN, C. 2006. CMOS photonics for high-speed interconnects. *Micro, IEEE 26,* 2, 58–66.

HENNING, P. AND WHITE, A. 2009. Trailblazing with Roadrunner. *Comput. Sci. Engin. 11,* 4, 91–95.

HO, R. 2003. On chip wires: Scaling and efficiency. Ph.D. thesis, Stanford University.

HO, R., MAI, K., AND HOROWITZ, M. 2001. The future of wires. *Proc. IEEE 89,* 4, 490–504.

HOCHBERG, M. AND BAEHR-JONES, T. 2010. Towards fabless silicon photonics. *Nature Photon. 4,* 8, 492–494.

HP. 2010. Press Release: “Shell to use CeNSE for clearer picture of oil and gas reservoirs”. http://www.hpl.hp.com/news/2009/oct-dec/cense.html.

HU, J. AND MENYUK, C. R. 2009. Understanding leaky modes: slab waveguide revisited. *Adv. Opt. Photon. 1,* 1, 58–106.

HUANG, M. C., ZHOU, Y., AND CHANG-HASNAIN, C. J. 2007. A surface-emitting laser incorporating a high-index-contrast subwavelength grating. *Nature Photon. 1,* 2, 119–122.

HUANG, M. C. Y., ZHOU, Y., AND CHANG-HASNAIN, C. J. 2008. Single mode high-contrast subwavelength grating vertical cavity surface emitting lasers. *Appl. Phys. Lett. 92,* 17, 171108.

INNOLUME. 2010. Innolume GmbH. http://www.innolume.com.

ITRS. 2009. International Technology Roadmap for Semiconductors. http://www.itrs.net.

JALALI, B. AND FATHPOUR, S. 2006. Silicon photonics. *J. Lightwave Technol. 24,* 12, 4600–4615.

JIAO, Y., SHI, Y., DAI, D., AND HE, S. 2010. Accurate and efficient simulation for silicon-nanowire-based multimode interference couplers with a 3D finite-element mode-propagation analysis. *J. Opt. Soc. Amer. B 27,* 9, 1813–1818.

JONES, R., RONG, H., LIU, A., FANG, A., PANICCIA, M., HAK, D., AND COHEN, O. 2005. Net continuous wave optical gain in a low loss silicon-on-insulator waveguide by stimulated Raman scattering. *Opt. Express 13,* 2, 519–525.

KANG, Y., LIU, H.-D., MORSE, M., PANICCIA, M. J., ZADKA, M., LITSKI, S., SARID, G., PAUCHARD, A., KUO, Y.-H., CHEN, H.-W., ZAOUI, W. S., BOWERS, J. E., BELING, A., MCINTOSH, D. C., ZHENG, X., AND CAMPBELL, J. C. 2009. Monolithic germanium/silicon avalanche photodiodes with 340 GHz gain-bandwidth product. *Nature Photon. 3,* 1, 59–63.

KIRMAN, N., KIRMAN, M., DOKANIA, R., MARTINEZ, J., APSEL, A., WATKINS, M., AND ALBONESI, D. 2007. On-chip optical technology in future bus-based multicore designs. *IEEE Micro 27,* 1, 56–66.

KOCH, B. R., FANG, A. W., COHEN, O., AND BOWERS, J. E. 2007. Mode-locked silicon evanescent lasers. *Opt. Express 15,* 18, 11225–11233.

KOVSH, A., KRESTNIKOV, I., LIVSHITS, D., MIKHRIN, S., WEIMERT, J., AND ZHUKOV, A. 2007. Quantum dot laser with 75nm broad spectrum of emission. *Opt. Lett. 32,* 7, 793–795.

KUMAR, R., ZYUBAN, V., AND TULLSEN, D. 2005. Interconnections in multi-core architectures: Understanding mechanisms, overheads and scaling. In *Proceedings of the 32nd International Symposium on Computer Architecture*. 408–419.

KUO, Y.-H., LEE, Y. K., GE, Y., REN, S., ROTH, J. E., KAMINS, T. I., MILLER, D. A. B., AND HARRIS, J. S. 2005. Strong quantum-confined Stark effect in germanium quantum-well structures on silicon. *Nature 437,* 7063, 1334–1336.

LANL. 2010. High-Performance Computing: Raodrunner. See http://www.lanl.gov/roadrunner.

LEE, B., BIBERMAN, A., CHAN, J., AND BERGMAN, K. 2010. High-performance modulators and switches for silicon photonic networks-on-chip. *IEEE J. Select. Topics Quantum Electron. 16,* 1, 6–22.

LEE, J.-M., KIM, D.-J., KIM, G.-H., KWON, O.-K., KIM, K.-J., AND KIM, G. 2008. Controlling temperature dependence of silicon waveguide using slot structure. *Opt. Express 16,* 3, 1645–1652.

LEUTHOLD, J., KOOS, C., AND FREUDE, W. 2010. Nonlinear silicon photonics. *Nature Photon. 4,* 8, 535–544.

LEVY, J. S., GONDARENKO, A., FOSTER, M. A., TURNER-FOSTER, A. C., GAETA, A. L., AND LIPSON, M. 2010. CMOS-compatible multiple-wavelength oscillator for on-chip optical interconnects. *Nature Photon. 4,* 1, 37–40.

LI, J., FATTAL, D. A., AND BEAUSOLEIL, R. G. 2009. Crosstalk-free design for the intersection of two dielectric waveguides. *Opt. Express 17,* 9, 7717–7724.

LI, Y., ZHANG, L., SONG, M., ZHANG, B., YANG, J.-Y., BEAUSOLEIL, R. G., WILLNER, A. E., AND DAPKUS, P. D. 2008. Coupled-ring-resonator-based silicon modulator for enhanced performance. *Opt. Express 16,* 17, 13342–13348.

LIANG, D. AND BOWERS, J. 2009. Photonic integration: Si or InP substrates? *Electron. Lett. 45,* 12, 578–581.

LIANG, D. AND BOWERS, J. E. 2010. Recent progress in lasers on silicon. *Nature Photon. 4,* 8, 511–517.

LIANG, D., BOWERS, J. E., OAKLEY, D. C., NAPOLEONE, A., CHAPMAN, D. C., CHEN, C.-L., JUODAWLKIS, P. W., AND RADAY, O. 2009a. High-quality 150mm InP-to-silicon epitaxial transfer for silicon photonic integrated circuits. *Electrochem. Solid. Lett. 12,* 4, H101–H104.

LIANG, D., FIORENTINO, M., OKUMURA, T., CHANG, H.-H., SPENCER, D. T., KUO, Y.-H., FANG, A. W., DAI, D., BEAUSOLEIL, R. G., AND BOWERS, J. E. 2009b. Electrically-pumped compact hybrid silicon microring lasers for optical interconnects. *Opt. Express 17,* 22, 20355–20364.

LIAO, L., LIU, A., BASAK, J., NGUYEN, H., PANICCIA, M., RUBIN, D., CHETRIT, Y., COHEN, R., AND IZHAKY, N. 2007. 40 Gbit/s silicon optical modulator for highspeed applications. *Electron. Lett. 43,* 22.

LIOW, T.-Y., ANG, K.-W., FANG, Q., SONG, J.-F., XIONG, Y.-Z., YU, M.-B., LO, G.-Q., AND KWONG, D.-L. 2010. Silicon modulators and germanium photodetectors on SOI: Monolithic integration, compatibility, and performance optimization. *IEEE J. Select. Topics Quantum Electron. 16,* 1, 307–315.

LIPSON, M. 2005. Guiding, modulating, and emitting light on silicon—challenges and opportunities. *23,* 12, 4222–4238.

LIPSON, M. 2009. Silicon photonics: the optical spice rack. *Electron. Lett. 45,* 12, 576–578.

LIRA, H. L. R., MANIPATRUNI, S., AND LIPSON, M. 2009. Broadband hitless silicon electro-optic switch for on-chip optical networks. *Opt. Express 17,* 25, 22271–22280.

LITTLE, B., CHU, S., ABSIL, P., HRYNIEWICZ, J., JOHNSON, F., SEIFERTH, F., GILL, D., VAN, V., KING, O., AND TRAKALO, M. 2004. Very high-order microring resonator filters for WDM applications. *IEEE Photon. Technol. Lett. 16,* 10, 2263–2265.

LIU, A., JONES, R., LIAO, L., SAMARA-RUBIO, D., RUBIN, D., COHEN, O., NICOLAESCU, R., AND PANICCIA, M. J. 2004a. A high-speed silicon optical modulator based on a metal-oxide-semiconductor capacitor. *Nature 427*, 615–618.

LIU, A., LIAO, L., CHETRIT, Y., BASAK, J., NGUYEN, H., RUBIN, D., AND PANICCIA, M. 2010a. Wavelength division multiplexing based photonic integrated circuits on silicon-on-insulator platform. *IEEE J. Select. Topics Quantum Electron. 16,* 1, 23–32.

LIU, A., LIAO, L., RUBIN, D., NGUYEN, H., CIFTCIOGLU, B., CHETRIT, Y., IZHAKY, N., AND PANICCIA, M. 2007. High-speed optical modulation based on carrier depletion in a silicon waveguide. *Opt. Express 15,* 2, 660–668.

LIU, A., RONG, H., PANICCIA, M., COHEN, O., AND HAK, D. 2004b. Net optical gain in a low loss silicon-on-insulator waveguide by stimulated Raman scattering. *Opt. Express 12,* 18, 4261–4268.

LIU, J., BEALS, M., POMERENE, A., BERNARDIS, S., SUN, R., CHENG, J., KIMERLING, L. C., AND MICHEL, J. 2008. Waveguide-integrated, ultralow-energy GeSi electro-absorption modulators. *Nature Photon. 2,* 7, 433–437.

LIU, J., SUN, X., CAMACHO-AGUILERA, R., KIMERLING, L. C., AND MICHEL, J. 2010b. Ge-on-Si laser operating at room temperature. *Opt. Lett. 35,* 5, 679–681.

LIU, L., KUMAR, R., HUYBRECHTS, K., SPUESENS, T., ROELKENS, G., GELUK, E.-J., DE VRIES, T., REGRENY, P., VAN THOURHOUT, D., BAETS, R., AND MORTHIER, G. 2010c. An ultra-small, low-power, all-optical flip-flop memory on a silicon chip. *Nature Photon. 4,* 3, 182–187.

MANIPATRUNI, S., CHEN, L., AND LIPSON, M. 2010b. Ultra high bandwidth wdm using silicon microring modulators. *Opt. Express 18,* 16, 16858–16867.

MANIPATRUNI, S., DOKANIA, R. K., SCHMIDT, B., SHERWOOD-DROZ, N., POITRAS, C. B., APSEL, A. B., AND LIPSON, M. 2008. Wide temperature range operation of micrometer-scale silicon electro-optic modulators. *Opt. Lett. 33,* 19, 2185–2187.

MANIPATRUNI, S., POITRAS, C. B., XU, Q., AND LIPSON, M. 2008. High-speed electro-optic control of the optical quality factor of a silicon microcavity. *Opt. Lett. 33,* 15, 1644–1646.

MANIPATRUNI, S., PRESTON, K., CHEN, L., AND LIPSON, M. 2010. Ultra-low voltage, ultra-small mode volume silicon microring modulator. *Opt. Express 18,* 17, 18235–18242.

MANIPATRUNI, S., XU, Q., AND LIPSON, M. 2007a. Pinip based high-speed high-extinction ratio micron-size silicon electrooptic modulator. *Opt. Express 15,* 20, 13035–13042.

MANIPATRUNI, S., XU, Q., SCHMIDT, B., SHAKYA, J., AND LIPSON, M. 2007b. High speed carrier injection 18 Gb/s silicon micro-ring electro-optic modulator. In *Proceedings of the IEEE Lasers and Electro-Optics Society (LEOS) 20th Annual Meeting*. 537–538.

MASHANOVICH, G. Z., REED, G. T., TIMOTIJEVIC, B. D., AND CHAN, S. P. 2008. Silicon photonic waveguides. In *Silicon Photonics: The State of the Art*, G. T. Reed (ed.). Wiley, 15–46.

MCINTYRE, R. 1966. Multiplication noise in uniform avalanche diodes. *IEEE Trans. Electron Devices 13,* 1, 164–168.

MEINDL, J. D., DAVIS, J. A., ZARKESH-HA, P., PATEL, C. S., MARTIN, K. P., AND KOHL, P. A. 2002. Interconnect opportunities for gigascale integration. *IBM J. Res. Dev. 46*, 245.

MICHEL, J., LIU, J., AND KIMERLING, L. C. 2010. High-performance Ge-on-Si photodetectors. *Nature Photon. 4,* 8, 527–534.

MILLER, D. 2000. Rationale and challenges for optical interconnects to electronic chips. *Proc. IEEE 88,* 6, 728–749.

MILLER, D. 2009. Device requirements for optical interconnects to silicon chips. *Proc. IEEE 97,* 7, 1166–1185.

MIT. 2010. MIT Microphotonics Center: Communications Technology Roadmap. See http://mphotonics.mit.edu.

MOORE, G. E. 1965. Cramming more components onto integrated circuits. *Electron. 38*, 114–117.

MORICHETTI, F., CANCIAMILLA, A., FERRARI, C., TORREGIANI, M., MELLONI, A., AND MARTINELLI, M. 2010a. Roughness induced backscattering in optical silicon waveguides. *Phys. Rev. Lett. 104,* 3, 033902.

MORICHETTI, F., CANCIAMILLA, A., MARTINELLI, M., SAMARELLI, A., RUE, R. M. D. L., SOREL, M., AND MELLONI, A. 2010. Coherent backscattering in optical microring resonators. *Appl. Phys. Lett. 96,* 8, 081112.

MORICHETTI, F., CANCIAMILLA, A., AND MELLONI, A. 2010b. Statistics of backscattering in optical waveguides. *Opt. Lett. 35,* 11, 1777–1779.

MORRIS, T. 2009. Breaking free of electrical constraints. *Appl. Phys. A 95*, 941–944.

MULLER, D. A. 2005. A sound barrier for silicon? *Nature Mater. 4,* 9 (Sept.), 645–647.

NAGARAJAN, R., KATO, M., PLEUMEEKERS, J., EVANS, P., LAMBERT, D., CHEN, A., DOMINIC, V., MATHUR, A., CHAVARKAR, P., MISSEY, M., DENTAI, A., HURTT, S., BÄCK, J., MUTHIAH, R., MURTHY, S., SALVATORE, R., JOYNER, C., ROSSI, J., SCHNEIDER, R., ZIARI, M., TSAI, H.-S., BOSTAK, J., KAUFFMAN, M., PENNYPACKER, S., BUTRIE, T., REFFLE, M., MEHUYS, D., MITCHELL, M., NILSSON, A., GRUBB, S., KISH, F., AND WELCH, D. 2007. Large-scale photonic integrated circuits for long-haul transmission and switching. *J. Opt. Netw. 6,* 2, 102–111.

NARAYAN, S. 2009. Supercomputers: past, present and the future. *Crossroads 15,* 4, 7–10.

NAWROCKA, M. S., LIU, T., WANG, X., AND PANEPUCCI, R. R. 2006. Tunable silicon microring resonator with wide free spectral range. *Appl. Phys. Lett. 89,* 7, 071110.

NG, H.-Y., WANG, M. R., LI, D., WANG, X., MARTINEZ, J., PANEPUCCI, R. R., AND PATHAK, K. 2008. 4×4 wavelength-reconfigurable photonic switch based on thermally tuned silicon microring resonators. *Opt. Eng. 47*, 044601.

OFFREIN, B. AND PEPELJUGOSKI, P. 2009. Optics in supercomputers. In *Proceedings of the 35th European Conference on Optical Communication (ECOC)*. 1–2.

OKAMOTO, K. 2006. *Fundamentals of Optical Waveguides* 2nd Ed. Academic Press.

ORNL. 2010. National Center for Computational Sciences: Jaguar. http://www.nccs.gov/computing-resources/jaguar.

OSGOOD, R. M., PANOIU, N. C., DADAP, J. I., LIU, X., CHEN, X., HSIEH, I.-W., DULKEITH, E., GREEN, W. M., AND VLASOV, Y. A. 2009. Engineering nonlinearities in nanoscale optical systems: physics and applications in dispersion-engineered silicon nanophotonic wires. *Adv. Opt. Photon. 1,* 1, 162–235.

OWENS, J., DALLY, W., HO, R., JAYASIMHA, D., KECKLER, S., AND PEH, L.-S. 2007. Research challenges for on-chip interconnection networks. *IEEE Micro 27,* 5, 96–108.

PALANIAPPAN, A. AND PALERMO, S. 2010. Power efficiency comparisons of interchip optical interconnect architectures. *IEEE Trans. Circuits Syst. II, Exp. Briefs 57,* 5, 343–347.

PARK, H., FANG, A., KODAMA, S., AND BOWERS, J. 2005. Hybrid silicon evanescent laser fabricated with a silicon waveguide and III-V offset quantum wells. *Opt. Express 13,* 23, 9460–9464.

PASSARO, V. AND DELL'OLIO, F. 2008. Scaling and optimization of MOS optical modulators in nanometer SOI waveguides. *IEEE Trans. Nanotechnology 7,* 4, 401–408.

PENG, Z., FATTAL, D., FIORENTINO, M., AND BEAUSOLEIL, R. G. 2010. Fabrication variations in SOI microrings for DWDM networks. In *Proceedings of the IEEE Photonics Society 7th International Conference on Group IV Photonics (GFP)*. P1.19.

PEPELJUGOSKI, P. K., KASH, J. A., DOANY, F., KUCHTA, D. M., SCHARES, L., SCHOW, C., TAUBENBLATT, M., OFFREIN, B. J., AND BENNER, A. 2010. Low power and high density optical interconnects for future supercomputers. In *Proceedings of the Optical Fiber Communication Conference (OFC)*. Optical Society of America.

PINGUET, T., ANALUI, B., MASINI, G., SADAGOPAN, V., AND GLOECKNER, S. 2008. 40-Gbps monolithically integrated transceivers in CMOS photonics. *Proc. Soc. Photo-Opt. Instrum. Eng. 6898*, 689805.

POULTON, J., PALMER, R., FULLER, A., GREER, T., EYLES, J., DALLY, W., AND HOROWITZ, M. 2007. A 14-mW 6.25-Gb/s transceiver in 90-nm CMOS. *IEEE J. Solid-State Circuits 42,* 12, 2745–2757.

PRESTON, K., MANIPATRUNI, S., GONDARENKO, A., POITRAS, C. B., AND LIPSON, M. 2009. Deposited silicon high-speed integratedelectro-optic modulator. *Opt. Express 17,* 7, 5118–5124.

QUARTERMAN, A. H., WILCOX, K. G., APOSTOLOPOULOS, V., MIHOUBI, Z., ELSMERE, S. P., FARRER, I., RITCHIE, D. A., AND TROPPER, A. 2009. A passively mode-locked external-cavity semiconductor laser emitting 60-fs pulses. *Nature Photon. 3,* 12, 729–731.

RAGHUNATHAN, V., YE, W. N., HU, J., IZUHARA, T., MICHEL, J., AND KIMERLING, L. 2010. Athermal operation of silicon waveguides: spectral, second order and footprint dependencies. *Opt. Express 18,* 17, 17631–17639.

RAMASWAMI, R., SIVARAJAN, K., AND SASAKI, G. 2009. *Optical Networks: A Practical Perspective*. Morgan Kaufmann Publishers.

RASIGADE, G., MARRIS-MORINI, D., VIVIEN, L., AND CASSAN, E. 2010. Performance evolutions of carrier depletion silicon optical modulators: From p-n to p-i-p-i-n diodes. *IEEE J. Select. Topics Quantum Electron. 16,* 1, 179–184.

RAZAVI, B. 2002. *Design of Integrated Circuits for Optical Communications*. McGraw-Hill, New York.

RAZZARI, L., DUCHESNE, D., FERRERA, M., MORANDOTTI, R., CHU, S., LITTLE, B. E., AND MOSS, D. J. 2010. CMOS-compatible integrated optical hyper-parametric oscillator. *Nature Photon. 4,* 1, 41–45.

REED, G. T., MASHANOVICH, G., GARDES, F. Y., AND THOMSON, D. J. 2010. Silicon optical modulators. *Nature Photon. 4,* 8, 518–526.

ROELKENS, G., VERMEULEN, D., THOURHOUT, D. V., BAETS, R., BRISION, S., LYAN, P., GAUTIER, P., AND FÉDÉLI, J.-M. 2008. High efficiency diffractive grating couplers for interfacing a single mode optical fiber with a nanophotonic silicon-on-insulator waveguide circuit. *Appl. Phys. Lett. 92,* 13, 131101.

RONG, H., JONES, R., LIU, A., COHEN, O., HAK, D., FANG, A., AND PANICCIA, M. 2005. A continuous-wave Raman silicon laser. *Nature 433,* 7027, 725–728.

RONG, H., XU, S., KUO, Y.-H., SIH, V., COHEN, O., RADAY, O., AND PANICCIA, M. 2007. Low-threshold continuous-wave Raman silicon laser. *Nature Photon. 1,* 4, 232–237.

RONG, Y., GE, Y., HUO, Y., FIORENTINO, M., TAN, M., KAMINS, T., OCHALSKI, T., HUYET, G., AND HARRIS, J. 2010. Quantum-confined Stark effect in Ge/SiGe quantum wells on Si. *IEEE J. Select. Topics Quantum Electron. 16,* 1, 85–92.

ROTH, J. E., FIDANER, O., SCHAEVITZ, R. K., KUO, Y.-H., KAMINS, T. I., HARRIS, J. S., AND MILLER, D. A. B. 2007. Optical modulator on silicon employing germanium quantum wells. *Opt. Express 15*, 5851–5859.

SACHER, W. D. AND POON, J. K. S. 2008. Dynamics of microring resonator modulators. *Opt. Express 16,* 20, 15741–15753.

SELVARAJA, S., BOGAERTS, W., DUMON, P., VAN THOURHOUT, D., AND BAETS, R. 2010. Subnanometer linewidth uniformity in silicon nanophotonic waveguide devices using CMOS fabrication technology. *IEEE J. Select. Topics Quantum Electron. 16,* 1, 316–324.

SOITEC. 2010. SOI Products. http://www.soitec.com/en/products/soi-products.php.

SOLJAČIĆ, M., IBANESCU, M., JOHNSON, S. G., FINK, Y., AND JOANNOPOULOS, J. D. 2002. Optimal bistable switching in nonlinear photonic crystals. *Phys. Rev. E 66,* 5, 055601.

SOREF, R. 2006. The past, present, and future of silicon photonics. *IEEE J. Select. Topics Quantum Electron. 12*, 1678–1687.

SOREF, R. AND BENNETT, B. 1987. Electrooptical effects in silicon. *IEEE J. Quantum Electron. 23,* 1, 123–129.

SPECTOR, S., SORACE, C., GEIS, M., GREIN, M., YOON, J., LYSZCZARZ, T., IPPEN, E., AND KARTNER, F. 2010. Operation and optimization of silicon-diode-based optical modulators. *IEEE J. Select. Topics Quantum Electron. 16,* 1, 165–172.

SUN, X., LIU, J., KIMERLING, L., AND MICHEL, J. 2010. Toward a germanium laser for integrated silicon photonics. *IEEE J. Select. Topics Quantum Electron. 16,* 1, 124–131.

TAN, M., ROSENBERG, P., YEO, J.-S., MCLAREN, M., MATHAI, S., MORRIS, T., KUO, H. P., STRAZNICKY, J., JOUPPI, N., AND WANG, S.-Y. 2009a. A high-speed optical multidrop bus for computer interconnections. *IEEE Micro 29,* 4, 62–73.

TAN, M., ROSENBERG, P., YEO, J.-S., MCLAREN, M., MATHAI, S., MORRIS, T., KUO, H. P., STRAZNICKY, J., JOUPPI, N., AND WANG, S.-Y. 2009b. A high-speed optical multidrop bus for computer interconnections. *Appl. Phys. A 95*, 945–953.

TANG, Y., WANG, Z., WOSINSKI, L., WESTERGREN, U., AND HE, S. 2010. Highly efficient nonuniform grating coupler for silicon-on-insulator nanophotonic circuits. *Opt. Lett. 35,* 8, 1290–1292.

TAO, S. H., FANG, Q., SONG, J. F., YU, M. B., LO, G. Q., AND KWONG, D. L. 2008. Cascade wide-angle Y-junction 1×16 optical power splitter based on silicon wire waveguides on silicon-on-insulator. *Opt. Express 16,* 26, 21456–21461.

TOP500. 2010. TOP500 supercomputing sites. http://www.top500.org.

TSIA, K. K., FATHPOUR, S., AND JALALI, B. 2008. Electrical control of parametric processes in silicon waveguides. *Opt. Express 16,* 13, 9838–9843.

TSUCHIZAWA, T., YAMADA, K., FUKUDA, H., WATANABE, T., ICHI TAKAHASHI, J., TAKAHASHI, M., SHOJI, T., TAMECHIKA, E., ITABASHI, S., AND MORITA, H. 2005. Microphotonics devices based on silicon microfabrication technology. *IEEE J. Select. Topics Quantum Electron. 11,* 1, 232–240.

UENUMA, M. AND MOOOKA, T. 2009. Temperature-independent silicon waveguide optical filter. *Opt. Lett. 34,* 5, 599–601.

VAN CAMPENHOUT, J., LIU, L., ROJO ROMEO, P., VAN THOURHOUT, D., SEASSAL, C., REGRENY, P., CIOCCIO, L., FEDELI, J., AND BAETS, R. 2008. A compact SOI-integrated multiwavelength laser source based on cascaded InP microdisks. *IEEE Photon. Technol. Lett. 20,* 16, 1345–1347.

VAN CAMPENHOUT, J., ROMEO, P. R., REGRENY, P., SEASSAL, C., THOURHOUT, D. V., VERSTUYFT, S., CIOCCIO, L. D., FEDELI, J.-M., LAGAHE, C., AND BAETS, R. 2007. Electrically pumped InP-based microdisk lasers integrated with a nanophotonic silicon-on-insulator waveguide circuit. *Opt. Express 15,* 11, 6744–6749.

VAN LAERE, F., CLAES, T., SCHRAUWEN, J., SCHEERLINCK, S., BOGAERTS, W., TAILLAERT, D., O'FAOLAIN, L., VAN THOURHOUT, D., AND BAETS, R. 2007a. Compact focusing grating couplers for silicon-on-insulator integrated circuits. *IEEE Photon. Technol. Lett. 19,* 23, 1919–1921.

VAN LAERE, F., ROELKENS, G., AYRE, M., SCHRAUWEN, J., TAILLAERT, D., VAN THOURHOUT, D., KRAUSS, T., AND BAETS, R. 2007b. Compact and highly efficient grating couplers between optical fiber and nanophotonic waveguides. *J. Lightwave Technol. 25,* 1, 151–156.

VANTREASE, D., BINKERT, N., SCHREIBER, R., AND LIPASTI, M. H. 2009. Light speed arbitration and flow control for nanophotonic interconnects. In *Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture*. ACM, 304–315.

VANTREASE, D., SCHREIBER, R., MONCHIERO, M., MCLAREN, M., JOUPPI, N. P., FIORENTINO, M., DAVIS, A., BINKERT, N., BEAUSOLEIL, R. G., AND AHN, J. 2008. Corona: System implications of emerging nanophotonic technology. In *Proceedings of the 35^{th} International Symposium on Computer Architecture*. 153–164.

VERMEULEN, D., SELVARAJA, S., VERHEYEN, P., LEPAGE, G., BOGAERTS, W., ABSIL, P., THOURHOUT, D. V., AND ROELKENS, G. 2010. High-efficiency fiber-to-chip grating couplers realized using an advanced CMOS-compatible silicon-on-insulator platform. *Opt. Express 18,* 17, 18278–18283.

VIVIEN, L., OSMOND, J., FÉDÉLI, J.-M., MARRIS-MORINI, D., CROZAT, P., DAMLENCOURT, J.-F., CASSAN, E., LECUNFF, Y., AND LAVAL, S. 2009. 42 GHz p.i.n germanium photodetector integrated in a silicon-on-insulator waveguide. *Opt. Express 17,* 8, 6252–6257.

VIVIEN, L., ROUVIÈRE, M., FÉDÉLI, J.-M., MARRIS-MORINI, D., DAMLENCOURT, J. F., MANGENEY, J., CROZAT, P., MELHAOUI, L. E., CASSAN, E., ROUX, X. L., PASCAL, D., AND LAVAL, S. 2007. High speed and high responsivity germanium photodetector integrated in a silicon-on-insulator microwaveguide. *Opt. Express 15,* 15, 9843–9848.

WANG, J., LOH, W., CHUA, K., ZANG, H., XIONG, Y., TAN, S., YU, M., LEE, S., LO, G., AND KWONG, D. 2008. Low-voltage high-speed (18 GHz/V) evanescent-coupled thin-film-Ge lateral PIN photodetectors integrated on Si waveguide. *IEEE Photo. Technol. Lett. 20,* 17, 1485–1487.

WATTS, M., TROTTER, D., YOUNG, R., AND LENTINE, A. 2008a. Ultralow power silicon microdisk modulators and switches. In *Proceedings of the IEEE 5^{th} International Conference on Group IV Photonics (GFP)*. 4–6.

WATTS, M., ZORTMAN, W., TROTTER, D., NIELSON, G., LUCK, D., AND YOUNG, R. 2009. Adiabatic resonant microrings (ARMs) with directly integrated thermal microphotonics. In *Proceedings of the Conference on Lasers and Electro-Optics (CLEO)*. CPDB10.

WATTS, M., ZORTMAN, W., TROTTER, D., YOUNG, R., AND LENTINE, A. 2010. Low-voltage, compact, depletion-mode, silicon Mach-Zehnder modulator. *IEEE J. Select. Topics Quantum Electron. 16,* 1, 159–164.

WATTS, M. R., TROTTER, D. C., AND YOUNG, R. W. 2008b. Maximally confined high-speed second-order silicon microdisk switches. In *Proceedings of the Optical Fiber Communication Conference (OFC)*. Optical Society of America, PDP14.

WOJCIK, G. L., YIN, D., KOVSH, A. R., GUBENKO, A. E., KRESTNIKOV, I. L., MIKHRIN, S. S., LIVSHITS, D. A., FATTAL, D. A., FIORENTINO, M., AND BEAUSOLEIL, R. G. 2009. A single comb laser source for short reach WDM interconnects. *Proc. Soc. Photo-Opt. Instrum. Eng. 7230*, 723021.

XIAO, S., KHAN, M. H., SHEN, H., AND QI, M. 2007. A highly compact third-order silicon microring add-drop filter with a very large free spectral range, a flat passband and a low delay dispersion. *Opt. Express 15,* 22, 14765–14771.

XU, Q. 2009. Silicon dual-ring modulator. *Opt. Express 17,* 23, 20783–20793.

XU, Q., FATTAL, D., AND BEAUSOLEIL, R. G. 2008. Silicon microring resonators with 1.5-μm radius. *Opt. Express 16,* 6, 4309–4315.

XU, Q., MANIPATRUNI, S., SCHMIDT, B., SHAKYA, J., AND LIPSON, M. 2007. 12.5 Gbit/s carrier-injection-based silicon micro-ring silicon modulators. *Opt. Express 15,* 2, 430–436.

XU, Q., SCHMIDT, B., PRADHAN, S., AND LIPSON, M. 2005. Micrometre-scale silicon electro-optic modulator. *Nature 435,* 7040, 325–327.

XU, Q., SCHMIDT, B., SHAKYA, J., AND LIPSON, M. 2006. Cascaded silicon micro-ring modulators for WDM optical interconnection. *Opt. Express 14,* 20, 9431–9435.

XUE, J., GARG, A., CIFTCIOGLU, B., WANG, S., HU, J., SAVIDIS, I., JAIN, M., HUANG, M., WU, H., FRIEDMAN, E. G., WICKS, G. W., AND MOORE, D. 2008. An intra-chip free-space optical interconnect. In *Proceedings of the 3rd Workshop on Chip Multiprocessor Memory Systems and Interconnects*.

YIN, T., COHEN, R., MORSE, M. M., SARID, G., CHETRIT, Y., RUBIN, D., AND PANICCIA, M. J. 2007. 31 GHz Ge n-i-p waveguide photodetectors on silicon-on-insulator substrate. *Opt. Express 15,* 21, 13965–13971.

YOO, S. 2009. Future prospects of silicon photonics in next generation communication and computing systems. *Electron. Lett. 45,* 12, 584–588.

YOUNG, I., MOHAMMED, E., LIAO, J., KERN, A., PALERMO, S., BLOCK, B., RESHOTKO, M., AND CHANG, P. 2010. Optical i/o technology for tera-scale computing. *IEEE J. Solid-State Circuits 45,* 1, 235–248.

YU, Z. AND FAN, S. 2009. Complete optical isolation created by indirect interband photonic transitions. *Nature Photon. 3,* 2, 91–94.

YUE, Y., ZHANG, L., SONG, M., BEAUSOLEIL, R. G., AND WILLNER, A. E. 2009. Higher-order-mode assisted silicon-on-insulator 90 degree polarization rotator. *Opt. Express 17,* 23, 20694–20699.

ZARLINK. 2010. Zarlink Semiconductor. http://www.zarlink.com.

ZHANG, L., LI, Y., SONG, M., BEAUSOLEIL, R. G., AND WILLNER, A. E. 2008a. Data quality dependencies in microring-based DPSK transmitter and receiver. *Opt. Express 16,* 8, 5739–5745.

ZHANG, L., LI, Y., SONG, M., YANG, J.-Y., BEAUSOLEIL, R. G., AND WILLNER, A. E. 2009. Silicon microring-based signal modulation for chip-scale optical interconnection. *Appl. Phys. A 95*, 1089–1100.

ZHANG, L., LI, Y., YANG, J.-Y., SONG, M., BEAUSOLEIL, R. G., AND WILLNER, A. E. 2010a. Silicon-based microring resonator modulators for intensity modulation. *IEEE J. Select. Topics Quantum Electron. 16,* 1, 149–158.

ZHANG, L., SONG, M., WU, T., ZOU, L., BEAUSOLEIL, R. G., AND WILLNER, A. E. 2008a. Embedded ring resonators for microphotonic applications. *Opt. Lett. 33,* 17, 1978–1980.

ZHANG, L., YANG, J.-Y., LI, Y., SONG, M., BEAUSOLEIL, R. G., AND WILLNER, A. E. 2008b. Monolithic modulator and demodulator of differential quadrature phase-shift keying signals based on silicon microrings. *Opt. Lett. 33,* 13, 1428–1430.

ZHANG, Q. X., DU, Y., TAN, C. W., ZHANG, J., YU, M. B., YEOH, W. G., LO, G.-Q., AND KWONG, D.-L. 2010b. A silicon platform with MEMS active alignment function and its potential application in Si-photonics packaging. *IEEE J. Select. Topics Quantum Electron. 16,* 1, 267–275.

ZHENG, X., LEXAU, J., LUO, Y., THACKER, H., PINGUET, T., MEKIS, A., LI, G., SHI, J., AMBERG, P., PINCKNEY, N., RAJ, K., HO, R., CUNNINGHAM, J. E., AND KRISHNAMOORTHY, A. V. 2010. Ultra-low-energy all-CMOS modulator integrated with driver. *Opt. Express 18,* 3, 3059–3070.

ZHOU, L., KASHIWAGI, K., OKAMOTO, K., SCOTT, R. P., FONTAINE, N. K., DING, D., AKELLA, V., AND YOO, S. J. B. 2009a. Towards athermal optically-interconnected computing system using slotted silicon microring resonators and RF-photonic comb generation. *Appl. Phys. A 95*, 1101–1109.

ZHOU, L., OKAMOTO, K., AND YOO, S. 2009b. Athermalizing and trimming of slotted silicon microring resonators with UV-sensitive PMMA upper-cladding. *IEEE Photon. Technol. Lett. 21,* 17, 1175–1177.

ZHOU, L. AND POON, A. W. 2006. Silicon electro-optic modulators using p-i-n diodes embedded 10-micron-diameter microdisk resonators. *Opt. Express 14,* 15, 6851–6857.

ZHU, S., FANG, Q., YU, M. B., LO, G. Q., AND KWONG, D. L. 2009. Propagation losses in undoped and n-doped polycrystalline silicon wire waveguides. *Opt. Express 17,* 23, 20891–20899.

ZIMMERMANN, L., TEKIN, T., SCHROEDER, H., DUMON, P., AND BOGAERTS, W. 2008. How to bring nanophotonics to application—silicon photonics packaging. *IEEE LEOS Newsl. 22,* 6, 4–14.

ZORTMAN, W. A., TROTTER, D. C., AND WATTS, M. R. 2010a. Silicon photonics manufacturing. *Opt. Express 18,* 23, 23598–23607.

ZORTMAN, W. A., WATTS, M. R., TROTTER, D. C., YOUNG, R. W., AND LENTINE, A. L. 2010b. Low-power high-speed silicon microdisk modulators. In *Proceedings of the Conference on Lasers and Electro-Optics (CLEO)*.

Received September 2010; revised January 2011; accepted January 2011