Silicon photonics for compact, energy-efficient interconnects [Invited]

T. Barwicz, H. Byun, F. Gan, C. W. Holzwarth, M. A. Popović, P. T. Rakich, M. R. Watts, E. P. Ippen, F. X. Kärtner, and H. I. Smith

Research Laboratory of Electronics, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, Massachusetts 02139, USA

J. S. Orcutt, R. J. Ram, and V. Stojanovic

Research Laboratory of Electronics and Microsystems Technology Laboratories, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, Massachusetts 02139, USA

O. O. Olubuyide and J. L. Hoyt

Microsystems Technology Laboratories, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, Massachusetts 02139, USA

S. Spector, M. Geis, M. Grein, T. Lyszczarz, and J. U. Yoon

Lincoln Laboratories, Massachusetts Institute of Technology, 244 Wood Street, Lexington, Massachusetts 02420, USA

Received September 6, 2006; accepted October 1, 2006; published December 22, 2006 (Doc. ID 74779)

The goal of the research program that we describe is to break the emerging performance wall in microprocessor development arising from limited bandwidth and density of on-chip interconnects and chip-to-chip (processor-to-memory) electrical interfaces. Complementary metal-oxide semiconductor compatible photonic devices provide an infrastructure for deployment of a range of integrated photonic networks, which will replace state-of-the-art electrical interconnects, providing significant gains at the system level. Scaling of wavelength-division-multiplexing (WDM) architectures using high-index-contrast (HIC) waveguides offers one path to realizing the energy efficiency and density requirements of high data rate links. HIC microring-resonator filters are well suited to support add–drop nodes in dense WDM photonic networks with high aggregate data rates because they support high Q's and, due to their traveling-wave character, naturally support physically separated input and drop ports. A novel reconfigurable, "hitless" switch is presented that does not perturb the express channels either before, during, or after reconfiguration. In addition, multigigahertz operation of low-power, Mach–Zehnder silicon modulators as well as germanium-on-silicon photodiodes are presented.

OCIS codes: 250.3140, 250.7360, 230.5750, 230.5170.

1. Introduction

Since 2002, single processor performance has increased far more slowly than it did over the preceding 20 years, owing to increasing power consumption, wire delay, and memory latency, as well as diminishing returns from further exploitation of instruction-level parallelism (ILP). In 2006, single processor performance is three times slower than what historic trends would have predicted in 2002. Modern processor architectures thus use multiple processor cores on a die, coupled in various ways to make use of data and thread-level parallelism. Regardless of application domain, these modern chips will be primarily constrained by the latency, bandwidth, and capacity of the external memory system. With multiple processor cores, and critical connection to an external memory system, the question becomes how to design an efficient on-chip network to facilitate interaction among cores and the external memory.

State-of-the art on-chip busses [1] and crossbars [2] exclusively use inverter-based repeaters, although they are known for their increased power consumption and inherent trade-off between bandwidth and latency [3]. Other electrical on-chip communica-

tion techniques (such as equalized point-to-point links [4,5]) are likely to be used in the future due to increased energy efficiency and short latency. These techniques essentially decouple the latency-bandwidth trade-off at the expense of the trading of the signal-to-noise ratio (SNR) with power, data rate, and latency. In this aspect, these links are much closer to photonic links, which are by default designed with some SNR constraints. For a photonic on-chip link, the energy cost per bit would have to be better than 100 fJ/bit to exceed the efficiency of the equalized on-chip electrical link at the 45 nm process node.

Although becoming similar to electrical links from the SNR point of view, optical links demand a different type of link protocol due to the type of resource (network) sharing. While electrical links use time-division multiple access (TDMA), simple collision detection, or just rely on density, optical links often use wavelength-division multiplexing (WDM) as an additional degree of freedom. A variety of WDM network schemes exist [6] for flexible interconnections among a large number of nodes. To prevent complexity explosion for the large number of nodes, most schemes use either tunable transmitters or tunable receivers or both, and rely on sensing schemes or fixed slot allocation for collision avoidance. Existing on-chip WDM network proposals also involve tunable cross-point resonators [7,8], which unfortunately do not scale well with the number of ports and have low energy efficiency and slow tuning. It is worth emphasizing that unlike the on-chip electrical networks, which are packet switched, for photonic interconnect we consider only circuit-switched multiple-access networks making use of photonic density and avoiding the complicated electronic control and buffering associated with packet-switched routing.

Optical links described above may consist of link-to-link connections at fixed wavelengths as well as fully tunable switchable networks. The key optical and optoelectronic devices include fixed and/or tunable filters that can access or add a given wavelength from or to an optical waveguide, ultralow power electro-optic modulators, and high-sensitivity low-noise detectors to approach fundamental limits in power dissipated per Gbits/s transmission speed.

2. High-index contrast filter bank

Microring resonators were first proposed as integrated optical wavelength filters by Marcatili in 1969 [9]. Applications that envision complex electronic–photonic integrated circuit (EPIC) functionality such as the proposed photonic interconnect call for dense optical integration and require microrings with low loss, a large free spectral range (FSR), and a drive to wavelength-scale resonator sizes. The latter two of these three key requirements invariably call for implementation in high-index contrast (HIC), which has presented many challenges since HIC microrings were first investigated [10–13]. A reconfigurable interconnect requires the ability to accurately control the resonant frequency of each microring and to reduce cross talk between the individual filters in the WDM filter bank.

The microring resonator filter design used for fabrication of the WDM filter banks is very similar to the design discussed in [14]. By utilizing this design with the HIC materials of silicon-rich silicon nitride ($n=2.2$@1550 nm) and silicon dioxide ($n=1.455$@1550 nm), a very wide FSR of 20 nm was realized. Also, the HIC allows for a bending radius as small as 8 μm, allowing for the possibility for the creation of very densely integrated photonic systems. The filter design was optimized to achieve the objective of a 3 dB bandwidth of 50 GHz and less than 30 dB cross talk for 150 GHz spaced channels. These parameters result in a design where the critical feature size of the gap between the bus and ring waveguides is 160 nm [see Fig. 1(a)].

In order to fabricate microring–resonator filters that meet the design requirements, direct write scanning electron beam lithography (SEBL) was used owing to its combination of high-resolution and high-dimensional control. The basic fabrication process used is similar to that described in [15]. The silicon wafer was initially iodized to grow a 3 μm thick undercladding layer of SiO_2 to prevent power leakage to the substrate. Next, a 400 nm layer of silicon-rich silicon nitride (SiN_x) was deposited by low-pressure chemical vapor deposition (LPCVD) to serve as the high-index core material. A Raith 150 SEBL system was used at 30 keV to write the pattern into a 200 nm resist layer of poly-(methyl methacrylate) (PMMA). A layer of Aquasave, a conducting polymer manufactured by Mitsubishi Rayon, was spun on top of the PMMA to prevent

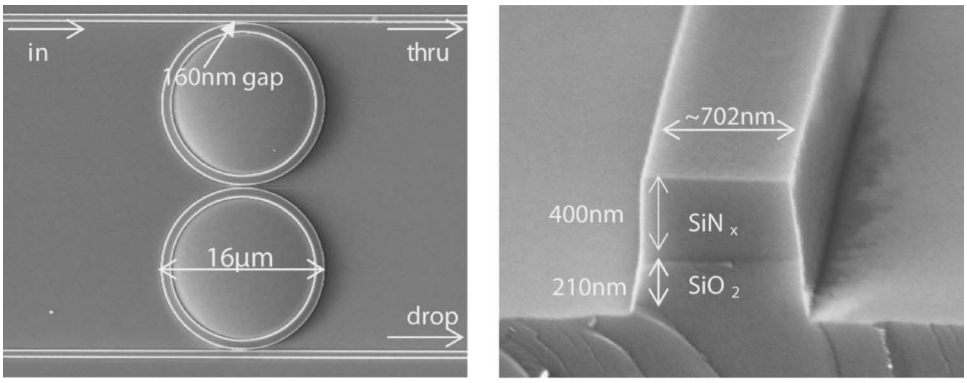

(a) (b)

Fig. 1. (a) Top view of a second-order microring resonator filter. (b) Cross-section view of the bus waveguide showing smooth vertical sidewalls.

charging from the electron beam. After developing the PMMA, Ni was evaporated onto the device, and a lift-off process was used to make a hard mask. The waveguides and filters were created by reactive ion etching (RIE) the sample with a mixture of CHF_3 and O_2 gases to etch through the 400 nm layer of SiN_x. The etch gas was then switched to CHF_3 only in order to etch 200 nm into the SiO_2 undercladding layer. Both of the RIE etch recipes were optimized to produce smooth, vertical sidewalls in each material as shown in Fig. 1(b). Due to the lack of an etch stop the actual fabricated device had a slight over-etch of 10 nm, which was within design tolerances.

Controlling the resonant frequency of HIC microring resonator filters requires an extremely high level of dimensional control. The resonant frequency of a microring resonator is determined by Eq. (1), where m is an integer, λ is the resonant wavelength, R is the ring radius, and n_{eff} is the effective index of refraction of the ring waveguide. It turns out that both R and n_{eff} can be controlled with lithography. R is simply controlled by changing the radius of the ring in the design layout whereas n_{eff} can be controlled by changing the width of the ring waveguide. For the design used the resonant frequency shifts 30 GHz for every nanometer increase in average ring waveguide width. A major problem that is encountered with changing these parameters in the design layout is that one is limited to discrete steps due to the pixel size used by the SEBL system. For the writing technique used, the

$$m\lambda = n_{\text{eff}} 2\pi R \tag{1}$$

discrete step sizes are limited to 6 and 12 nm for the ring radius and ring waveguide width, respectively. This results in discrete frequency shifts of 104 GHz steps when changing the radius and 360 GHz shifts when changing the ring waveguide width. This means that to achieve our objective of 150 GHz channel spacing a different technique needs to be used to control the resonant frequency that is not limited by the pixel size of the SEBL system. The technique used is dose modulation, an extension of the concept of dose compensation, which has previously been used to correct for proximity effects and coupling induced frequency shifts in microring resonators as explained in [16]. Dose modulation is the practice of slightly changing the electron dose that the rings are written with allowing for the average ring bus width to be changed on the tens of picometer scale. The frequency shift induced by dose modulation is dependent on many factors including the waveguide design, resist, base dose, and developer. Thus dose modulation needs to be very precisely calibrated for both the design and the fabrication process to be used. Calibration was done by fabricating a series of devices that were composed of a waveguide coupled to two microring resonator filters. One of the microring filters serves as a reference point while dose modulation is used on the test microring filter. This is done for different magnitudes of dose modulation, and by comparing the relative frequency shifts the frequency dependence on dose can be calculated. From the calibration it is found that the frequency shift is linearly dependent on dose modulation, shifting −18 GHz for each percent of the base dose increased, for increases less than 8% of the base dose. It was also found during calibration that, to a first approximation, frequency shifts due to dose modulation and radius changes can be assumed to add linearly. This technique allows us to overcome

the limitation of discrete frequency shifts permitting the fabrication of any desired channel spacing. This technique combined with changing the radius was used to make 8-channel second-order filter banks with channel spacing of 90, 104, 120, 140, 145, 150, 155, 160, and 180 GHz to demonstrate our ability for fabricating filter banks with precisely controlled channel spacings.

Optical transmission measurements were performed on the fabricated filter banks using a tunable laser in a fiber-to-fiber optical characterization setup. Figure 2 shows the measured transmission response for the filter bank fabricated with a target bandwidth of 50 GHz and channel spacing of 150 GHz. The actual filter bank response had a channel bandwidth of 46.5 GHz and an average channel spacing of 159 GHz with a standard deviation of 5 GHz. This offset in channel spacing is due to a 4.2% higher frequency dependence on radius changes than expected, since it is evident in all of the filter banks with different channel spacing. This error can be greatly reduced in future fabrication runs through careful calibration experiments similar to the ones performed for dose modulation. The standard deviation is due to stochastic variations that cannot be eliminated. It is likely that this number will be lower in future fabrications since the tip for the SEBL system had been recently replaced, and the electron beam current had not completely stabilized by the time of fabrication. Overcladding and heater designs are being investigated so that thermal trimming can be used to make slight adjustments to the resonant frequency of the microring resonator filter to correct for small errors due to calibration errors and stochastic variations. The drop loss of the filters, averaged over 40 filters to reduce errors from coupling variations, is calculated to be 1.5 ± 0.5 dB. Simulations indicate that this drop loss can be accounted for by a propagation loss of approximately 8 dB/cm. Loss measurements performed on ridge waveguides fabricated in SiN_x showed that it has an intrinsic absorption loss of 7.5 ± 2 dB/cm at 1540 nm. Therefore the scattering loss due to sidewall roughness in our fabricated devices is a maximum on the order of 1 dB/cm.

One important consideration when designing microring resonator filters for filter banks is the adjacent channel cross talk. For the proposed interconnect, as well as many other applications, it is important to minimize the adjacent channel cross talk. The fabricated filters give −30 dB cross talk for 150 GHz channel spacing as predicted from the design. For the designed microring resonator filters the cross talk will remain at −30 dB as the bandwidth and channel spacing are scaled down proportionally. A common way to decrease the adjacent channel cross talk is to use higher-order filters. The problem with higher-order filters is that they are much more sensitive to small defects. A better way to improve the adjacent channel cross talk is to use a multistage filter design, since it is much less sensitive to small defects than higher-order filters as described in [14]. In a two-stage filter layout both the through and drop signals go through two sets of second-order filters. Figure 3 shows the through and drop response measured for a two-stage second-order filter fabricated with the same basic filter design used for the filter banks. As seen in the figure, by multistaging the filters it is possible to greatly decrease the adjacent channel cross talk without changing the filter design. The cross talk is measured to be more than 40 dB down but cannot be

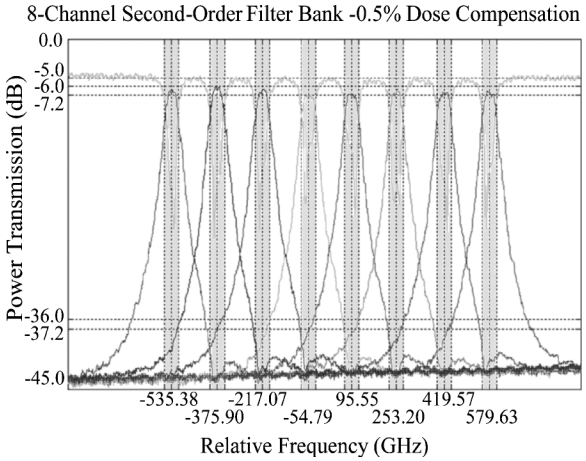

Fig. 2. Through- and drop-port characteristics of filter bank with second-order single stage filters and 159 GHz channel spacing. Frequencies are all relative to 1540 nm.

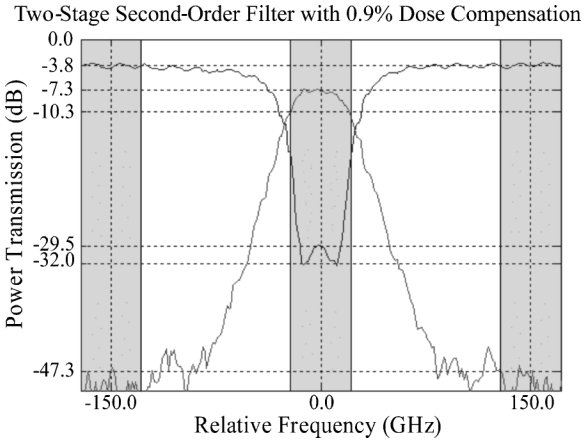

Fig. 3. Through- and drop-response of a two-stage second-order microring resonator filter with a bandwidth-to-channel ratio of 1:3.

accurately determined, since it is below the noise floor of the measurement system used. The only major disadvantage with using the two-stage design is that the overall drop loss is increased because the drop signal must now pass through two filters.

Overall the design and fabrication of microring resonator filter banks used for the proposed interconnect shows great promise. Through the combination of changing the ring radius and using dose modulation it has been shown that the resonant frequency of the filters can be precisely controlled to meet any desired target channel spacing. In future fabricated microring resonator filter banks, heaters will be used to thermally correct for small frequency spacing errors due to calibration errors and stochastic variations. For a cascade of a large number of filters to be feasible, through-port losses must be negligible. In the filters presented here, nonadjacent channels have losses of $<0.2\%$ by design and can be further improved. Experimentally, such losses were too low to characterize, but a definite upper bound is $\sim 0.5\%$ per filter stage. Scaling to narrower bandwidths and channel spacing will require low propagation losses and lower sensitivity of resonance frequencies to width to ensure aligned resonances in high-order filters.

3. Hitless switching and tuning of microphotonic filters

Tunable, switchable WDM photonic networks call for dynamically reconfigurable tunable add–drop filters. The filters then need to have a hitless tuning capability. Hitless tuning requires that other (express) wavelength channels remain undisturbed (on the granularity of single bits) at all times before, during, and after the reconfiguration of an add–drop filter from an initial wavelength to a target wavelength or to an off state [17,18]. A general approach to hitless tuning entails providing a means of substantially disabling both the amplitude and phase response of an add–drop filter in the through-port (and possibly also in the add–drop ports), tuning the center wavelength in the off state, and reactivating the filter at the target wavelength. Aside from interconnect networks addressed in this paper, hitless tuning is also an important requirement in filters for telecom applications.

One general approach to hitless switching is to provide a controllable bypass interferometer within which the filter is embedded [19,20]. Within this framework, one such implementation is shown in Fig. 4, comprising a pair of variable optical couplers (here, $\Delta\beta$ switches) in a Mach–Zehnder (MZ) configuration, and a ring–resonator filter in one MZ arm [20]. In one state of the variable couplers, the WDM spectrum entering the input port is routed through the filter, the through-port channels further passing through the second switch to the output port. In a second state, the switches are configured to pass the input spectrum fully through the top interferometer arm, effectively bypassing the filter. In the latter state, the filter may be reconfigured without effect on the express channels. Further, for hitless switching, the interferometer arms are coherently aligned by a π-radian differential phase shift, so that at all intermediate states representing the transition of signal from the filter arm to the bypass arm (or back), all through-port channels coherently recombine in the output port. Figure 5 shows spectra in the on state, off state, and transition states over a wavelength range

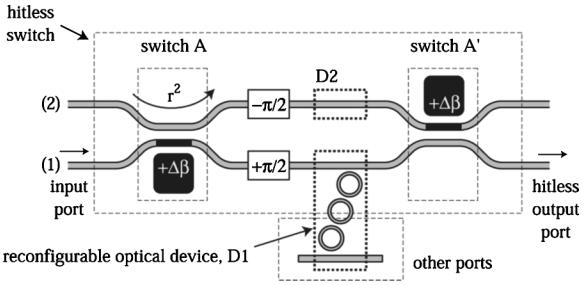

Fig. 4. Hitless tunable ring–resonator filter architecture based on a bypass interferometer.

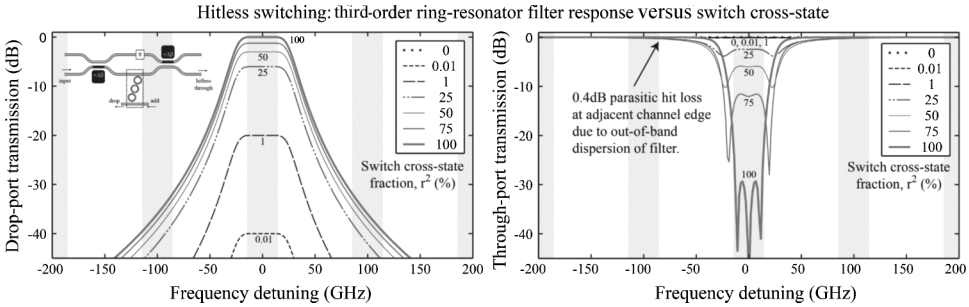

Fig. 5. Drop- and through-port response using a typical third-order microring–resonator filter with a 40 GHz passband and suitable for 100 GHz channel spacing, in the on state, the off state, and intermediate switch states, showing hitless switching operation.

near a typical filter passband. Owing to symmetries in the design, the geometry in [20] permits simple actuation, attainable operating bandwidths of hundreds of nanometers, and low sensitivity to dimensional variations, losses and even asymmetries between the two switches [21]. MZ interferometer switches used in [19] enable thermo-optic actuation, but one may also employ microelectro-mechanical systems (MEMS)-based switches [20,22], which permit a small footprint in HIC and low steady-state power consumption.

A generalization of this approach yields a fairly broad class of interferometers with similar properties [21]. It shows that the variable coupler switches in Fig. 4 may be replaced by arbitrary, substantially reflectionless splitter and combiner devices with two input and two output ports, which may include switches, resonators, and even nonreciprocal elements. Provided the splitter and combiner are oriented properly, the input spectrum can be split in an arbitrary time- or wavelength-dependent manner between the two interferometer arms, and then fully recombined in one output port by symmetry. This class of devices enables a general set of hitless switching configurations like those in [20], as well as other novel devices such as a combined lattice filter–ring filter design enabling dispersion-free free spectral range (FSR) (and tuning range) multiplication for the ring–resonator filter [21]. Experimental demonstration of these approaches for the hitless tuning of high-index contrast filters remains to be accomplished.

Other approaches to hitless tuning of resonators may make use of disabling a resonant response directly. First, loss switching of resonators has been investigated [7], e.g., by carrier injection, as a switching method, though it provides substantial residual losses in all but the highest finesse filters. Second, another approach to hitless tuning may make use of a MEMS-controlled waveguide-ring coupling, as shown in an all-pass filter in [23]. An advantage of such an approach is that a high extinction ratio may be obtained because the coupling changes exponentially as the gap is linearly varied, due to the evanescent field shape. The drawbacks include the fabrication complexity of such devices.

4. High-speed silicon modulator

One of the most important electro-optic devices for future electronic–photonic integrated circuits, and for the envisioned photonic interconnect in particular, is a silicon-

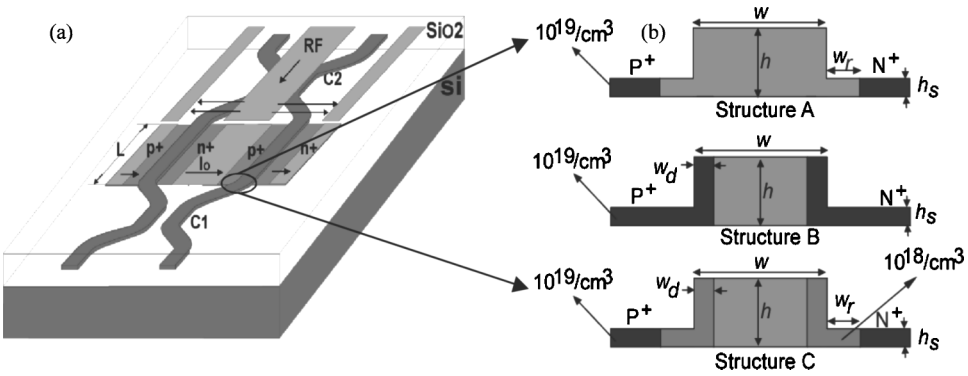

Fig. 6. MZ modulator shown in (a) top view and (b) cross section.

based high-speed electro-optic modulator that converts electronic signals into the optical domain. Plasma and thermal effects in silicon have long been known as an efficient means to achieve such modulation [24]. Early on, p^+-i-n^+ structures based on the plasma effect have been proposed and realized for silicon-based electro-optic phase and amplitude modulators with switching speeds up to 10 MHz [24,25]. The response speed of these devices is limited by carrier recombination. Very recently, a modulator based on a metal-oxide semiconductor (MOS) structure has been reported as operating at frequencies up to 10 GHz [26]. Here we investigate high-speed electro-refractive MZ modulators based on HIC silicon strip waveguides and forward-biased pin diodes [27].

A schematic of the modulator, which is a conventional MZ structure, is shown in Fig. 6(a). Figure 6(b) shows a schematic cross section through one of the electrically driven p-i-n–phase modulator sections. The light intensity is modulated by the phase shift induced in each arm of the interferometer due to carrier injection. The modulator was fabricated using optical lithography in a standard complementary MOS (CMOS) process. The fabricated waveguide width and height are 520 and 220 nm, respectively. Aluminum is used as electrical contact on top of the highly doped contact regions to minimize the contact resistance. The waveguide sidewalls are slightly doped to reduce resistance without incurring excessive optical loss. If the intrinsic region of the modulator is doped, it can also be used in reverse bias by modulating the width of the depletion region thereby sweeping the carriers in and out of the waveguide.

The electrical-to-optical frequency response $|S_{21}|^2$ of a MZ modulator with 0.25 mm long p^+-i-n^+ sections in forward-bias operation has been measured [see Fig. 7(a)]. For the forward-bias regime, the electrical bandwidth f_{3dB} is only about ~200 MHz. In this mode of operation, the bandwidth is carrier lifetime limited. Due to surface recombination, the lifetime of this device is expected to be $\tau \approx 1$ ns, $f_{3dB} \approx 1/2\pi\tau = 160$ MHz, which is consistent with the measurement. The optimum bias voltage is ~1 V at a current of ~10 mA. The electrical power consumption for a 0.5 mm long device at 25% modulation depth is shown in Fig. 7(b). At frequencies below 10 MHz, only 1 μW of rf drive power is necessary to achieve 25% modulation depth. This drive power increases for high frequencies due to the reduction in diffusion capacitance and contact resistances to about 200 mW for operation at 10 GHz.

Since the necessary drive power in the low frequency range is so low, one can obtain a modulator with 10 GHz bandwidth from this device using a high-pass filter in the drive circuit to equalize the frequency response up to 10 GHz. If such a modulator was used in an optical interconnect operating at 10 Gbits/s, a drive energy per bit of about 20 pJ/bit would be necessary. This is prohibitively high for interchip and intrachip communications. However, the MZ modulator demonstrated so far is broadband, which is not necessary for modulation of a narrowband signal. Thus resonant enhancement of the carrier injection effect can be employed to reduce the drive power, as was recently demonstrated [28]. Using a resonance with a free spectral range of 6 THz and linewidth of 20 GHz, i.e., a finesse of 300, a resonant enhancement factor of 100 can be expected. This enhancement factor is large enough to bring the drive energy per bit down to 200 fJ/bit, which is comparable to the currently most advanced electronic drive circuits used for interchip communications.

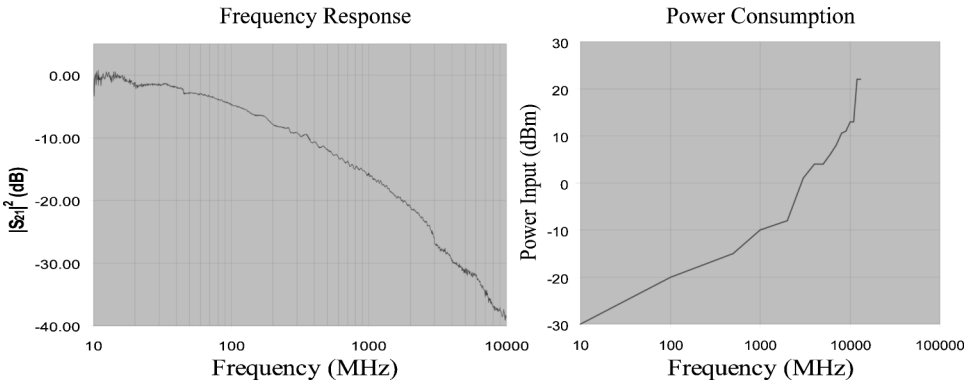

Fig. 7. (a) Frequency response of a 0.25 mm long MZ modulator operated in forward bias mode. (b) Power consumption of a 0.5 mm long MZ modulator under forward bias with 25% modulation depth.

5. Germanium photodiode

Strained germanium films have the required responsivity and speed to serve as photodiodes up to the 1.55 μm wavelength range. Thus integrating germanium films onto silicon substrates (Ge/Si) into a CMOS-compatible process is an attractive goal for making arrays of on-chip detectors. In this paper, Ge films are deposited by LPCVD in an applied materials epitaxial reactor, and the optical and electrical properties of Ge photodiodes are examined.

It has been demonstrated in ultra-high-vacuum chemical vapor deposition (UHVCVD) systems that depositing a low-temperature Ge layer (seed layer), followed by the deposition of a high-temperature layer (cap layer) with subsequent annealing, can create a smooth, planar Ge film on a (100) silicon substrate with threading dislocation density on the order of 10^7 cm^{-2} [29]. This two-step deposition process has been successfully adapted to a LPCVD system. The seed layer is grown at 335°C, 30 Torr, and the cap layer is grown at 700°C, 30 Torr. After annealing at 900°C for 30 min the blanket Ge films have threading dislocation densities of $\sim 2 \times 10^7$ cm^{-2}. These Ge films experience a slight tensile strain after removal from the deposition chamber due to the mismatch between the thermal expansion and contraction coefficients for the germanium film and underlying silicon substrate. This tensile strain lowers the direct bandgap for the Ge film from 0.8 to 0.78 eV, giving these Ge films high absorption coefficients for wavelengths up to 1.6 μm [30,31].

Approximately 2 μm thick intrinsic Ge films were deposited on p$^+$ (100) Si substrates, and capped with a 0.2 μm N+ polysilicon layer to create a vertical p-i-n photodiode (see Fig. 8 inset). The I–V characteristics of the photodiodes fabricated with

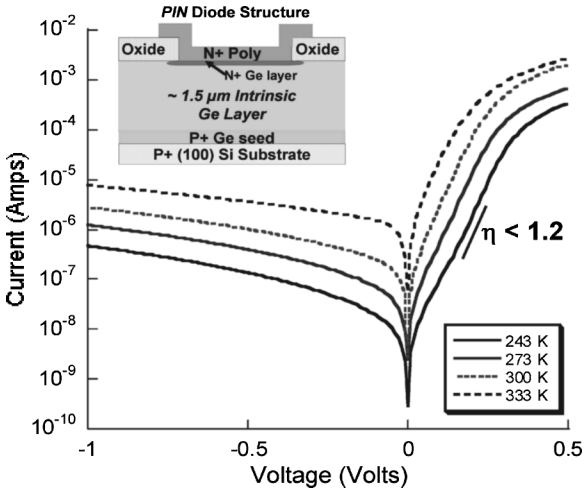

Fig. 8. I–V characteristics of a 100 μm $\times$ 100 μm diode as a function of temperature. The diode has an ideality factor less than 1.2 at 300 K with a perimeter dominated reverse leakage current of $\sim$55 mA/cm^2 at –1 volt bias. Inset shows a schematic of a diode cross section.

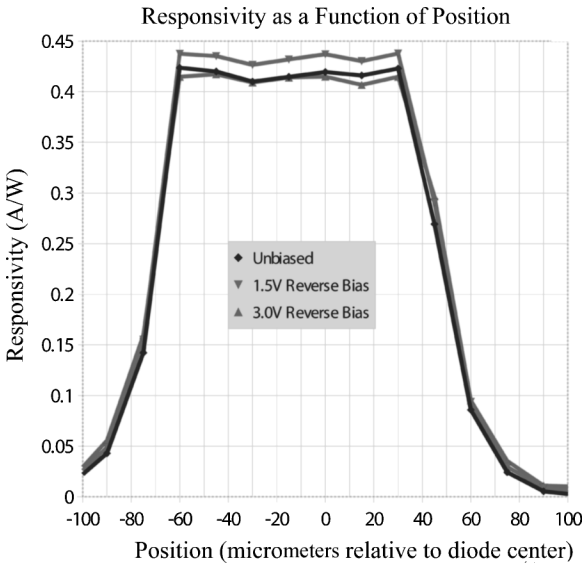

Fig. 9. Responsivity versus position for a 100 μm $\times$ 100 μm diode with no antireflective coating. A lensed fiber probe is translated across the surface of the detector to spatially resolve the responsivity.

the blanket Ge films are shown in Fig. 8. The reverse leakage current at 300 K and a −1 V bias is ∼40 mA/cm^2, with the reverse leakage current of the diodes being perimeter dominated for diodes less than 1 mm^2 in surface area. The perimeter dominated reverse leakage current is believed to be due to surface states introduced during the Ge passivation process. This hypothesis is supported by the spatially resolved responsivity (see Fig. 9), which shows a dramatic decrease in the diode responsivity along the perimeter of the diode.

At a 1.55 μm wavelength, the Ge film has a responsivity of 0.45 A/W. Below 1.2 μm, the p$^+$ silicon substrate begins to strongly absorb the incoming light causing a drop in responsivity. At a −5 V bias, the diode is measured to have a 3 dB frequency of 1.4 GHz. Impulse response measurements show a 100 ps rise time and a seven times slower fall time of 700 ps. The lack of symmetry in the rise and fall may be associated with interfacial trap states at the N$^+$-polysilicon, dielectric, and Ge boundaries. This long fall time is believed to be a limiting factor in the frequency response of the photodiodes.

In summary, photodiodes with ∼40 mA/cm^2 reverse leakage currents, greater than 0.4 A/W responsivities and GHz-range 3 dB frequencies at 300 K and −1 V bias have been demonstrated. It is believed that improved passivation of the Ge films will further improve the bandwidths and the reverse leakage currents of future devices. Initial interconnect architecture designs indicate that photodetectors with 7.5 GHz bandwidth, 0.5 A/W responsivity, and 10 fF capacitance will be required. Of the order of 3000 such photodetectors must be fabricated on a chip. The capacitance specification originates from the energy efficiency requirement. Initial results on vertically illuminated Ge-on-Si p-i-n photodiodes presented above indicate that the required specifications should be within reach, so long as the active area of the device is scaled to reduce the capacitance, and that leakage currents can be scaled through effective control of perimeter defect states.

6. Conclusions

Integration of HIC optical devices on the Si technology platform together with electronics may lead to advanced interconnect architectures in the near future. In this paper, we discussed the key components of such a system such as filter banks, modulators, and detectors. While the components in a photonic link can satisfy 10 Gbits/s data rate relatively easily, the biggest challenge is the energy efficiency of the link, which—as stated in Section 1—should be lower than 100 fJ/bit for purely on-chip photonic link. The most critical component from the energy-efficiency standpoint is the modulator (current energy cost per bit is ∼5–40 pJ/bit [27] for modulation rates of up

to 10 Gbits/s). In addition to energy-efficiency, creating a fast tunable (<10 ns) add–drop filter is a very challenging task, which, if achieved would enable flexible network topologies and efficient arbitration. Our approach to meeting the above challenges has been to focus on high-index contrast waveguide structures. This paradigm leverages the nanometer-scale precision of CMOS lithography to realize the necessary performance for next generation interconnects.

Acknowledgments

The authors acknowledge support from the National Science Foundation, Analog Devices, Inc., the Semiconductor Research Corporation, and a generous donation from Applied Materials, Inc. John Yasaitis (Analog Devices, Inc.) has been especially instrumental in guiding and encouraging the work on Ge photodetectors. The assistance of Gary Riggott and the staff and facilities of the Microsystems Technology Laboratory at MIT are also acknowledged. The DARPA EPIC Program under contract W911NF-04-1-0431 made this work possible.

References and Links

1. D. C. Pham, T. Aipperspach, D. Boerstler, M. Bolliger, R. Chaudhry, D. Cox, P. Harvey, P. M. Harvey, H. P. Hofstee, C. Johns, J. Kahle, A. Kameyama, J. Keaty, Y. Masubuchi, M. Pham, J. Pille, S. Posluszny, M. Riley, and D. L. Stasiak, "Overview of the architecture, circuit design, and physical implementation of a first-generation cell processor," IEEE J. Solid-State Circuits **41**, 179–196 (2006).
2. P. Kongetira, K. Aingaran, and K. Olukotun, "Niagara: A 32-Way Multithreaded SPARC Processor," IEEE Micro **25**, 21–29 (2005).
3. R. Ho, K. W. Mai, and M. A. Horowitz, "The future of wires," Proc. IEEE **89**, 490–504 (2001).
4. R. Ho, K. Mai, and M. Horowitz, "Efficient on-chip global interconnects," *Symposium on VLSI Circuits, Digest of Technical Papers* (IEEE, 2003), pp. 271–274.
5. D. Schinkel, E. Mensink, E. A. Klumperink, E. van Tuijl, and B. Nauta, "A 3-Gbits/s/ch transceiver for 10-mm uninterrupted RC-limited global on-chip interconnects," IEEE J. Solid-State Circuits **41**, 297–306 (2006).
6. B. Mukherjee, "WDM-based local lightwave networks. I. Single-hop systems," IEEE Network **6**, 12–27 (1992).
7. B. E. Little, H. A. Haus, J. S. Foresi, L. C. Kimerling, E. P. Ippen, and D. J. Ripin, "Wavelength switching and routing using absorption and resonance," IEEE Photon. Technol. Lett. **10**, 816–818 (1998).
8. S. J. Emelett and R. Soref, "Design and simulation of silicon microring optical routing switches," J. Lightwave Technol. **23**, 1800–1807 (2005).
9. E. A. J. Marcatili, "Bends in optical dielectric guides," Bell Syst. Tech. J. **48**, 2103–2132 (1969).
10. B. E. Little, S. T. Chu, H. A. Haus, J. Foresi, and J.-P. Laine, "Microring resonator channel dropping filters," J. Lightwave Technol. **15**, 998–1005 (1997).
11. B. E. Little, J. S. Foresi, G. Steinmeyer, E. R. Thoen, S. T. Chu, H. A. Haus, E. P. Ippen, L. C. Kimerling, and W. Greene, "Ultra-compact Si-SiO_2 microring resonator optical channel dropping filters," IEEE Photon. Technol. Lett. **10**, 549–551 (1998).
12. H. A. Haus, "Microwaves and Photonics," in *OSA Trends in Optics and Photonics (TOPS), Vol. 23 Symposium on Electro-Optics: Present and Future* (Optical Society of America, 1998), pp. 2–8.
13. B. E. Little, "Advances in microring resonators," in *Proceedings of the Integrated Photonics Research Conference*, (Optical Society of America, 2003), pp. 165–167.
14. M. A. Popovic, T. Barwicz, M. R. Watts, P. T. Rakich, L. Socci, E. P. Ippen, F. X. Kärtner, and H. I. Smith, "Multistage high-order microring-resonator filters with relaxed tolerances for high through-port extinction," in *Proceedings of Conference on Lasers and Electro-Optics* (IEEE, 2005), pp. 266–268.
15. T. Barwicz, M. A. Popović, P. T. Rakich, M. R. Watts, H. A. Haus, E. P. Ippen, and H. I. Smith, "Fabrication and analysis of add/drop filters based on microring resonators in SiN," *Proceedings of the Optical Fiber Conference* (IEEE 2004), paper TuL4.
16. T. Barwicz, M. A. Popović, P. T. Rakich, M. R. Watts, F. X. Kaertner, E. P. Ippen, and H. I. Smith, "Fabrication control of the resonance frequencies of high-index-contrast microphotonic cavities," presented at Integrated Photonics Research and Applications, April 24-26, 2006, Uncasville, Connecticut, USA (Optical Society of America), paper JWA3.
17. T. Shiragaki, N. Henmi, M. Fujiwara, M. Misono, T. Shiozawa, and S. Suzuki, "Optical cross-connect system incorporated with newly developed operation and management system," IEEE J. Sel. Areas Commun. **16**, 1179–1189 (1998).
18. C. R. Doerr, L. W. Stulz, D. S. Levy, R. Pafchek, M. Cappuzzo, L. Gomez, A. Wong-Foy, E. Chen, E. Laskowski, G. Bogert, and G. Richards, "Wavelength add-drop node using silica waveguide integration," J. Lightwave Technol. **22**, 2755–2762 (2004).
19. R. E. Scotti, C. K. Madsen, C. H. Henry, G. Lenz, Y. P. Li, H. Presby, and A. White, "A hitless reconfigurable add-drop multiplexer for WDM networks utilizing planar waveguides,

thermo-optic switches and UV-induced gratings," in *Proceedings of the Optical Fiber Conference* (IEEE, 1998), pp. 142–143.
20. H. A. Haus, M. A. Popović, and M. R. Watts, "Broadband hitless bypass switch for integrated photonic circuits," IEEE Photon. Technol. Lett. **18**, 1137–1139 (2006).
21. M. A. Popović, E. P. Ippen, and F. X. Kärtner, "Universally balanced photonic interferometers," Opt. Lett. **31**, 2713–2715 (2006).
22. M. M. Lee and M. C. Wu, "MEMS-actuated microdisk resonators with variable power coupling ratios," IEEE Photon. Technol. Lett. **17**, 1034–1036 (2005).
23. G. A. Magel, "Integrated optic devices using micromachined metal membranes," in *Integrated Optics and Microstructures III*, Proc. SPIE **2686**, 54–63 (1996).
24. R. A. Soref, "Silicon-based optoelectronics," Proc. IEEE **81**, 1687–1706 (1993).
25. P. Dainesi, A. Kung, M. Chabloz, A. Lagos, P. Fluckiger, A. Ionescu, P. Fazan, M. Declerq, P. Renaud, and P. Robert, "CMOS compatible fully integrated Mach-Zehnder interferometer in SOI technology," IEEE Photon. Technol. Lett. **12**, 660–662 (2000).
26. A. Liu, D. Samara-Rubio, L. Liao, and M. Paniccia, "Scaling the modulation bandwidth and phase efficiency of a silicon optical modulator," IEEE J. Sel. Top. Quantum Electron. **11**, 367–372 (2005).
27. F. Gan and F. X. Kärtner, "High-speed silicon electro-optic modulator design," IEEE Photon. Technol. Lett. **17**, 1007–1009 (2005).
28. Q. Xu, Br. Schmidt, S. Pradhan, and M. Lipson, "Micrometre-scale silicon electro-optic modulator," Nature **435**, 325–327 (2005).
29. H.-C. Luan, D. R. Lim, K. K. Lee, K. M. Chen, J. G. Sandland, K. Wada, and L. C. Kimerling, "High-quality Ge epilayers on Si with low threading-dislocation densities," Appl. Phys. Lett. **75**, 2909–2911 (1999).
30. D. D. Cannon, "Strain-engineered CMOS-compatible Ge photodetectors," Ph.D. dissertation (Massachusetts Institute of Technology, 2003).
31. D. D. Cannon, "Tensile strained epitaxial Ge films on Si(100) substrates with potential application in *L*-band telecommunications," Appl. Phys. Lett. **84**, 906–908 (2004).