

High-Speed 2×2 Switch for Multiwavelength Silicon-Photonic Networks—On-Chip

Benjamin G. Lee, *Student Member, IEEE*, Aleksandr Biberman, *Student Member, IEEE*,
Nicolás Sherwood-Droz, *Student Member, IEEE*, Carl B. Poitras, *Member, IEEE*,
Michal Lipson, *Senior Member, IEEE*, and Keren Bergman, *Fellow, IEEE*

Abstract—We report the fabrication and experimental verification of a multiwavelength high-speed 2×2 silicon photonic switch for ultrahigh-bandwidth message routing in optical on-chip networks. The structure employs only two microring resonators in order to implement the bar and cross states of the switch. These states are toggled using an optical pump at 1.5- μm wavelengths in-plane with the waveguide devices, though electronic, rather than optical, control schemes are envisioned for more complex systems built from these devices. Experiments characterize bit-error-rate performance in the bar and cross states during static and dynamic operation. The all-optical demonstration exhibits the ability of the switch to implement ultra-short transition times (< 2 ns), high extinction ratios (> 10 dB), and low power penalties (~ 1 dB) at a data rate of 10 Gb/s. Further performance improvements are expected by using electronic carrier injection via *p-i-n* diodes surrounding the ring waveguides. The 2×2 switching functionality facilitates the design of more complex routing structures, allowing the implementation of high-functionality integrated optical networks.

Index Terms—Electro-optic switches, multiprocessor interconnection, optical resonators, optical switches, photonic switching systems, wavelength-division multiplexing (WDM).

I. INTRODUCTION

THE field of silicon photonics in recent years has progressed rapidly, making the vision of applying long-haul optical networking paradigms to ultrashort-reach applications possible, such as inter and intrachip interconnects. These advances have opened the possibilities of exploiting the benefits of optical technologies (e.g., the immense bandwidth supported by wavelength-division multiplexing (WDM), low latencies corresponding to the optical time-of-flight, and low power consumption provided by transparent message routing) over a wider range of interconnect applications. It is envisioned that ultrashort-haul optical networks, implemented by using silicon photonics, will provide a practical solution to the existing power and bandwidth limitations of global electrical interconnects in future chip multiprocessor (CMP) computing

systems [1]. As CMPs scale the number of processing cores to obtain computational performance improvements, they will require an increasingly efficient high-bandwidth and low-power communications infrastructure in order to achieve the desired level of connectivity. Photonic networks-on-chip (NoCs) offer a solution while simultaneously supplying enormous offchip bandwidth as well.

The silicon-on-insulator (SOI) platform is an attractive material system for realizing integrated photonic networks. Its inherent high index contrast reduces the necessary footprint of optical circuits, and its compatibility with the complementary metal-oxide-semiconductor (CMOS) fabrication process enables low-cost high-quality production of optical devices in conjunction with sophisticated electrical circuitry [2], [3]. Furthermore, silicon microring resonators have become an essential building block for these systems and have been employed in circuits achieving complex tunable filters [4]–[6], fixed multibit delays [7], as well as electro-optic [8]–[10] and all-optical [11] modulators.

Complementing these elements, a high-speed, low-area switch designed for multiwavelength operation, which can be implemented by using a low-power switching scheme, will be required in order to leverage the full bandwidth advantage of the optical domain for on-chip message routing. Fast multiwavelength switches have been reported in III-V material systems [12]–[14], where the direct bandgap can provide more natural photonic functionalities. However, a lack of integrability with the CMOS process, and often device size, may prohibit their implementation within CMPs. Recently, fast multiwavelength switching devices have also been presented in silicon. Dong *et al.* proposed and demonstrated a comb-switching technique by using a single-ring resonator [15]. This 1×2 multiwavelength switch was shown to have low interchannel crosstalk, and simultaneous switching of 20 wavelengths with nanosecond-scale transition times has been demonstrated [16], [17]. Also, a similar 1×2 switch, based on a higher order coupled-resonator system, was reported [18].

In this paper, a first-order 2×2 switch utilizing the same comb-switching technique of [15]–[18] is demonstrated and characterized, extending on the work performed in [19]. The device is demonstrated with nanosecond-scale switching times and extinction ratios of up to 11.5 dB. Bit-error rates (BERs) are measured at 10-Gb/s data rates for static and dynamic operation in order to characterize the various sources of power penalty in the device. In addition to providing network designers with the imperative 2×2 switching functionality, this device, which is diagramed in Fig. 1(a), can also be an important building

Manuscript received August 29, 2008; revised February 18, 2009. Current version published July 09, 2009. This work was supported in part by the National Science Foundation (NSF) under Grant ECS-0725707, in part by the NSF CAREER Grant 0446571, and in part by the NSF under Grant ECS-0335765.

B. G. Lee, A. Biberman, and K. Bergman are with the Department of Electrical Engineering, Columbia University, New York, NY 10027 USA (e-mail: benlee@ee.columbia.edu; biberman@ee.columbia.edu; bergman@ee.columbia.edu).

N. Sherwood-Droz, C. B. Poitras, and M. Lipson are with the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14853 USA (e-mail: nrs35@cornell.edu; cbp8@cornell.edu; lipson@ece.cornell.edu).

Digital Object Identifier 10.1109/JLT.2009.2019256

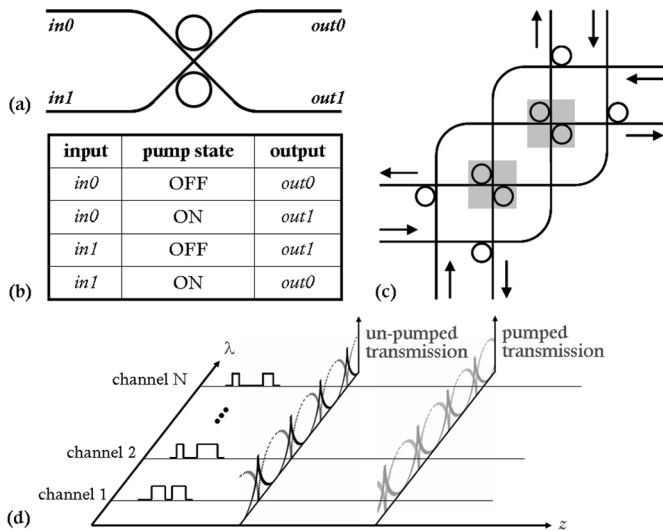


Fig. 1. (a) Schematic of the 2×2 switch layout. (b) Truth table for the switching states. (c) Four-port bidirectional multiwavelength routing structure proposed in [1] requiring two of the 2×2 switches reported here (shaded in gray). (d) Diagram of multiwavelength switching scheme used in the reported switch. The solid curves represent the bar state, while the dotted curves represent the cross state. The blue-shift observed between the curves from the unpumped state to the pumped state provides an exchange between high transmission and low transmission for each wavelength.

block in more complex routing structures. For example, the nonblocking four-port switch [Fig. 1(c)] proposed in [1] and further evaluated in [20] achieves much higher functionality than the 2×2 switch alone, and yet relies only on a combination of the previously reported 1×2 switch and the 2×2 switch described here.

The remainder of this paper is organized as follows. Section II discusses the device fabrication, geometry, and operation. The experimental procedures and techniques are outlined in Section III. Section IV describes static and dynamic measurement results, and Section V is the conclusion.

II. SWITCH DESCRIPTION

The device was fabricated at the Cornell Nanofabrication Facility on an SOITec SOI substrate with a $3\text{-}\mu\text{m}$ -thick buried oxide layer, using electron-beam lithography followed by reactive-ion plasma etching. The resulting structure is covered by a $3\text{-}\mu\text{m}$ -thick SiO_2 overladding layer deposited by plasma-enhanced chemical vapor deposition. The waveguides have cross sections of $250\text{ nm} \times 450\text{ nm}$ (height \times width), and inverse-taper-mode converters are used on the waveguides at each edge of the chip. The switch geometry [Fig. 1(a)] consists of a crossing between two straight waveguides with two ring waveguides coupled to vertical angles of the intersection. The waveguides undergo adiabatic tapers, expanding over a length of $30\text{ }\mu\text{m}$ to a width of $2\text{ }\mu\text{m}$, before the perpendicular waveguide intersection, in order to reduce reflection losses at the crossing. Presently, the tapers add no area to the device since the diameters of the rings are $100\text{ }\mu\text{m}$.

Ideally, the two rings have a series of resonator modes which coincide in wavelength. Without an applied pump, the input signal wavelength is set to overlap one of these modes. As the light enters from port $in0$ ($in1$), it is coupled into the upper

(lower) ring, through the resonator to the alternate waveguide, where it exits from port $out0$ ($out1$). Therefore, without an applied pump, the switch implements the bar state (i.e., $in0$ to $out0$ and $in1$ to $out1$). A control signal is used to change the state of the switch by shifting the wavelength of the resonator mode away from the signal's carrier wavelength. (The mechanism for this is explained in further detail in Section III-B.) Thus, when the control, or pump, is applied, the input lightwave from port $in0$ ($in1$), operating at the same wavelength as when no pump is applied, is no longer affected by either of the two ring resonators and, as a result, propagates through the waveguide crossing to port $out1$ ($out0$). Consequently, when the pump is enabled, the switch implements the cross state (i.e., $in0$ to $out1$ and $in1$ to $out0$). The switching state truth table is listed in Fig. 1(b).

By using rings with a relatively large diameter ($100\text{ }\mu\text{m}$), the free-spectral range (FSR) of the resonators is reduced to 1.6 nm . This allows encoding of the input signal over multiple wavelengths, which utilize many resonator modes of the ring by aligning a single wavelength channel to each mode, thus increasing the throughput bandwidth of the switch [Fig. 1(d)]. Furthermore, the switching power remains constant regardless of the number of wavelength channels employed in the signal spectrum. Therefore, the switching power per message bandwidth can be drastically decreased, allowing the largest advantages of the optical domain to be most efficiently leveraged. In this manner, all wavelength channels are routed through the switch cohesively, propagating from the source to destination as a single consolidated optical message. Consequently, the switches are intended to operate in networks which implement wavelength-parallel packet formatting in order to increase the message bandwidth, such as in [1], [20], and [21], among others, rather than networks which route messages on a single variable wavelength, such as those that employ the wavelength domain to alleviate packet contentions [22].

The resonator modes have 3-dB bandwidths of 0.1 nm (12.5 GHz) on average, designed to accommodate 10-Gb/s optical data. Future systems implementing the same FSR, but designed to handle 40- or 100-Gb/s data rates, may require higher order resonator structures, such as coupled-resonator optical waveguides (CROWs) [23], as implemented in [7] and [18]. The transfer functions of these systems provide sharper rolloff and flatter passbands than single-ring structures and, therefore, avoid degrading the extinction ratio of the switch when the resonator bandwidth is widened. Also, these structures have the additional advantage of being more robust to temperature fluctuations, such as might be expected in a CMP environment [18]. The present device may also be designed with a much smaller footprint by implementing the resonators in noncircular geometries, because the SOI waveguides permits very small bending radii [2], [3]. Moreover, the adiabatically tapered waveguide crossing may be replaced with more compact crossings, such as the one demonstrated in [24].

III. EXPERIMENTAL PROCEDURE

A. Thermal Tuning

The ideal assumption is that the resonator modes of the two rings in the switch overlap in wavelength may not always be re-




Fig. 2. OSA traces recorded from port *out1*, while the output from a broadband ASE source was injected into port *in0*. With thermal tuning applied to the incorrect ring, the resonator modes of the two rings are slightly detuned (solid curve). Applying the thermal tuning to the correct ring, the modes can be made to overlap (dashed curve). The extinction ratio depicted above is limited by the OSA resolution bandwidth of 0.06 nm.

alized due to fabrication tolerances. Improvements may be made with careful fabrication calibrating techniques [25], and by the addition of ohmic heaters placed over the ring waveguides to provide localized thermal tuning as in [5] and [6] at the expense of additional power consumption. For our experiments, it is sufficient to shine light from a diode-pumped solid-state (DPSS) laser over one of the two rings by using an optical fiber that is vertically incident to the chip. In this manner, the appropriate ring is heated so that its resonator modes are red-shifted to overlap those of the unheated ring.

The DPSS laser, with a wavelength of 532 nm, is chosen based on availability, and because it delivers photon energies above the bandgap of silicon. The light is guided to the desired location by a multimode fiber (MMF) mounted on a 3-D translation stage. The cleaved end of the fiber is placed just above the ring that corresponds to the shorter wavelength mode. By adjusting the output power of the DPSS laser and the placement of the fiber facet, the resonator modes can be well aligned, as observed in Fig. 2.

B. Pump Configuration

The pump signal leverages the free-carrier plasma dispersion effect [26] to modulate the effective index of the ring waveguide and, hence, shift the wavelengths on which the resonator modes occur. This effect may be carried out by injecting carriers into the waveguide optically [11], [15]–[18] or electronically [4], [8]–[10]. To simplify the fabrication of the device, optical carrier injection is chosen for the initial characterization of our device. Furthermore, carriers are optically injected most efficiently when the pump photons have energies above the bandgap of silicon. Nevertheless, the process may also be performed by inducing two-photon absorption (TPA) within the waveguides with the advantage of being able to utilize components operating within the telecommunications wavelength band near 1.5 μm in order to generate the pump. This is the method used in this work;

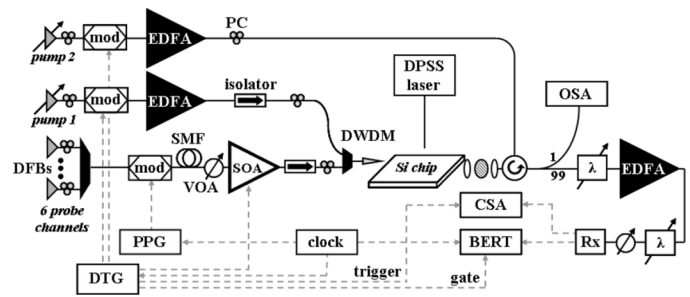


Fig. 3. Diagram of the experimental setup. Solid lines represent optical fiber, while dashed lines represent electrical cable.

however, in actual implementations aimed at the intended application, electronic carrier injection is required. Not only would the optical pumping scheme be impractical in structures, such as the one in Fig. 1(c), but electronic carrier injection is expected to improve the speed, insertion loss, and power requirements of the switch [8], while also providing a more stable switching scheme (see Section IV-B). The pump must be applied to both rings simultaneously for proper operation. Therefore, we employ one copropagating and one counter-propagating pump signal (with respect to the probe signal) in order to inject carriers into both rings simultaneously.

C. Experimental Setup

The experimental setup (Fig. 3) consists of the following four functions: 1) probe signal generation; 2) pump signal generation; 3) fiber-waveguide coupling; and 4) test and measurement. The probe signals are generated by distributed-feedback (DFB) lasers and occupy six ITU C-band channels: C22 (1559.8 nm), C24 (1558.2 nm), C33 (1550.9 nm), C35 (1549.3 nm), C37 (1547.7 nm), and C43 (1542.9 nm). The outputs are combined with a 32-channel dense wavelength-division multiplexer (DWDM) that has 100-GHz channel spacings, and are then simultaneously encoded with a 10-Gb/s nonreturn-to-zero (NRZ) signal that carries a $2^{31} - 1$ pseudorandom bit sequence (PRBS) using a LiNbO₃ modulator (mod in Fig. 3), which is driven by a signal from a pulse pattern generator (PPG). The wavelength channels are decorrelated by 25 km of single-mode fiber (SMF) achieving nearly seven bits of delay between the most closely spaced wavelength channels.

The pump signals are generated, modulated, and amplified by using tunable lasers, LiNbO₃ modulators, and erbium-doped fiber amplifiers (EDFAs), respectively. Pump 1, which copropagates with the probe channels, occupies channel C51 (1536.6 nm). It is combined with the probes by using a single-channel 100-GHz DWDM. Pump 2, located on C18 (1563.1 nm), is injected by using an optical circulator such that it counterpropagates with the probes. The pump waveforms are provided by a data-timing generator (DTG) which operates at 625 MHz. Measured before injection into the tapered fiber, pump 1 has an average power of 18 dBm; moreover, the DWDM device provides spectral filtering of the amplified spontaneous emission (ASE) from the EDFA. Similarly measured before injection into the fiber-collimating lens, pump 2 exhibits an average power of 17 dBm; however, no spectral filtering is

performed, so out of band ASE contributes somewhat to the average power of pump 2.

The probe signals, along with pump 1, are coupled to the inverse-tapered waveguides using a tapered fiber, and are extracted using a lens, polarizer (transmitting the quasi-TM polarization), and fiber-collimating lens. The polarizer extracts the undesired polarization, which is required due to the polarization dependencies that result from the rectangular waveguide cross sections. Polarization controllers (PCs) are used throughout the setup to enable independent alignment of all probe and pump polarizations. The chip is mounted on a thermoelectric cooler (TEC) for global tuning and stabilization of the resonator modes. A 532-nm DPSS laser is used to provide localized tuning of one ring with respect to the other.

The multichannel output is monitored by an optical spectrum analyzer (OSA), while one probe channel propagates through a tunable grating filter (λ in Fig. 3), an EDFA, a second tunable filter, and a variable optical attenuator (VOA), before being received by a *p-i-n* photodiode integrated with a transimpedance and limiting amplifier pair (Rx in Fig. 3). The output is observed on a communications signal analyzer (CSA) and analyzed by using a BER tester (BERT), which is directly synchronized to the PPG. No clock recovery is performed. For the dynamic BER measurements only, a semiconductor optical amplifier (SOA) is used to gate the optical signal into 9.6-ns packets, recurring every 12.8 ns. When the SOA is used, it is driven by a signal generated from the DTG, and the DTG sends an additional signal to gate the BERT over the arrival of the packetized data. A VOA precedes the SOA to prevent saturation and minimize crosstalk between the WDM channels.

IV. RESULTS

Initially, the device is tested in the steady state. That is, the switching state is toggled by manually enabling and disabling the pump, rather than providing modulated pump signals. Dynamic measurements are discussed in Section IV-B. The significance of the static measurements should not be undervalued, as many interconnection networks which stand to benefit from this type of switch operate under a circuit-switched (or fast circuit-switched) control scheme, where the switching states remain constant over the duration of several optical messages. However, the promising results of the high-speed measurements indicate that the switch may also be valuable in more dynamically routed interconnection networks.

A. Static Results

The power penalty is measured for the six-channel WDM signal, which is injected into input port *in0* and exits from either *out0* [Fig. 4(a)] or *out1* [Fig. 4(b)]. When the signal exits from port *out0*, the pump signal is not activated; therefore, we expect the power penalty in the static case to be mostly a result of the insertion loss, spectral filtering, and nonlinear phase induced by the ring resonator on each wavelength channel. When the signal exits from port *out1*, the pump signal is activated. Since the probe signals do not couple to the blue-shifted resonator modes, only a simple waveguide crossing is encountered in the optical

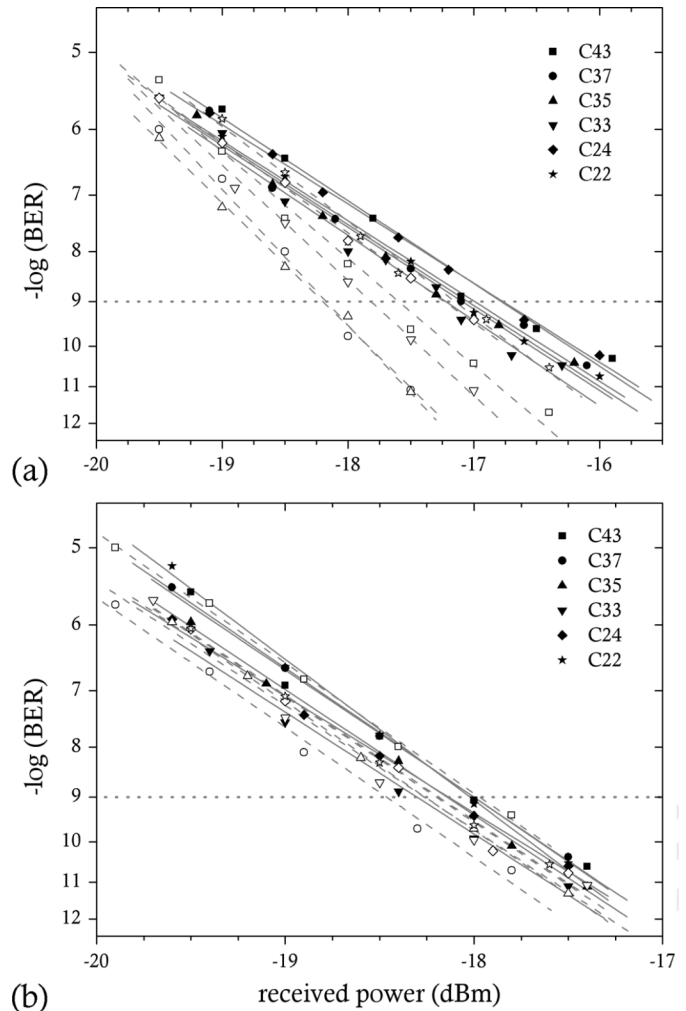


Fig. 4. BER curves depicting the static power penalty of a six-channel WDM signal in (a) the bar state and (b) the cross state. Each character corresponds to measurements on an ITU channel as shown in the legend. The back-to-back measurements are shown by using dashed lines and open characters, while the switched measurements are shown using solid lines and filled characters.

pathway. We therefore expect the power penalty to result mostly from the nonlinear effects induced by the pump signals (e.g., TPA) in addition to degradation from the ASE noise.

In order to measure the *in0-out0* path, the input signal wavelengths are aligned along the resonator modes so that they exit port *out0* when no pump is applied, and six BER curves are recorded, one for each channel. The back-to-back curves are obtained by cooling the TEC on which the chip is mounted until the resonator modes are shifted away from the signal wavelengths so that the signals no longer pass through the ring. The output is then gathered from port *out1*, and six back-to-back BER curves are recorded. An attenuator is inserted prior to the EDFA for the back-to-back curves to compensate the ring insertion losses ensuring consistent power levels entering the EDFA. This is done because errors resulting from optical signal-to-noise-ratio (OSNR) degradation, which can occur in the EDFA due to low input signal powers, are independent from errors incurred by degradation from the device and, thus, should not factor into the device power penalty.

TABLE I
STATIC POWER PENALTIES

ITU Channel	Output Port / Switch State	Power Penalty (dB) at a BER of 10^{-9}
C22	<i>out0 / bar</i>	0.1
	<i>out1 / cross</i>	0.2
C24	<i>out0 / bar</i>	0.4
	<i>out1 / cross</i>	0.1
C33	<i>out0 / bar</i>	0.6
	<i>out1 / cross</i>	0.0
C35	<i>out0 / bar</i>	1.1
	<i>out1 / cross</i>	0.1
C37	<i>out0 / bar</i>	1.2
	<i>out1 / cross</i>	0.5
C43	<i>out0 / bar</i>	0.8
	<i>out1 / cross</i>	0.0

Inputs are injected into port *in0* with approximately -6 dBm of average power per channel (measured before the tapered fiber).

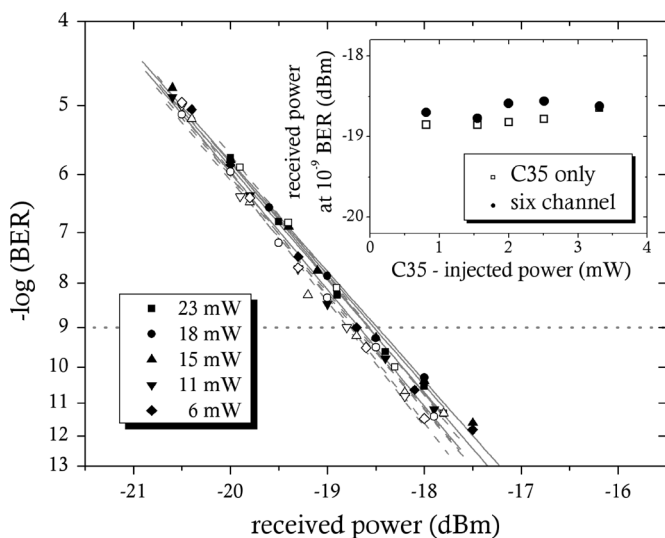


Fig. 5. BER curves depicting the wavelength crosstalk for varying multi-channel injection powers (shown as different shapes in the legend). Curves are shown for a six-channel (solid lines, filled symbols) and single-channel (dashed lines, open symbols) signal propagating from *in0* to *out0*. Measurements are taken on channel C35 in both cases. The inset plots the points at which the trend-lines intersect a BER of 10^{-9} versus the injection power of C35.

The *in0-out1* path is measured using the same channels at the same wavelengths. The pump is activated, shifting the resonator mode wavelengths, and the signals are collected from port *out1*. Six BER curves are recorded. The back-to-back curves are obtained in the same manner as before, by disabling the pumps, inserting an attenuator, and detuning the TEC controller so that the signals remain on port *out1* after the pumps are off. The resulting power penalties are tabulated for the six WDM channels for both output ports (Table I).

Next, the interchannel crosstalk within the ring is characterized as a function of injected power. Since nonlinear effects can be produced at low powers in ring resonators, we expect that the paths *in0-out0* and *in1-out1* will have the strongest opportunity for interchannel wavelength crosstalk. Therefore, BER curves are taken for channel C35, both in the presence of the other five wavelength channels and alone, as they propagate through the path *in0-out0* (Fig. 5).




Fig. 6. CSA traces showing the switched CW probe through the path *in0-out0* (left) and *in0-out1* (right). The CSA uses a time and amplitude scale of 10 ns/div and 100 μ W/div, respectively, with a 16-point average. The zero-power level is depicted by overlaying a trace that was taken with the input disconnected.

For this measurement, an EDFA and a second VOA replace the SOA in the experimental setup (Fig. 3). The first VOA is set to provide appropriate input-power levels to the EDFA, while the second is varied over five settings. A BER curve is recorded for channel C35 for each of the five VOA settings. Then, all channels are disabled except C35, and a BER curve is recorded for five new VOA settings which match the five previous injection powers of channel C35. (The five VOA settings differ slightly between the one-channel and six-channel experiments, because the EDFA provides more gain to C35 when it is alone.) The power values listed in Fig. 5 represent the average power over a 50% duty cycle. The independent axis of the inset represents the power of channel C35 before injection into the tapered fiber, while the values listed in the legend show the sum of the powers over the six WDM channels, again before injection into the tapered fiber. The inset shows that the receiver sensitivity curve shifts inconsequentially (about 0.1 dB on average) when the five additional channels are disabled, indicating negligible interchannel crosstalk for injection powers up to 23 mW. The measurement was limited in maximum injection power by the saturation output power of the available multichannel EDFA.

B. Dynamic Results

In order to take measurements under dynamic operation of the switch, the two pump signals are modulated with 12.8-ns square pulses with a periodicity of 102.4-ns (12.5% duty cycle), and are synchronized to coincide in time at the switch. By inserting a single-channel continuous-wave (CW) probe signal into the device, we may observe the envelope profile of the switched signal by using the CSA. The paths *in0-out0* and *in0-out1* are used for these measurements. (The other two paths may be ignored due to symmetry.) From the traces (Fig. 6), it can be seen that the rise and fall times are less than 2 ns each for both switch states. The ringing in the signal, seen just after the pumps are disabled, is attributed to undershoot in the waveform of the drive voltage supplied to the modulator. The ringing observed over the duration of the state in which the pumps are enabled is caused by instability in the optical pumping scheme. As the high-power pump changes the ring waveguide's index of refraction, it provides switching of the probe signal, but also simultaneously decouples itself from the resonator, because the pump is of a fixed wavelength. As a result, the ring index returns to its previous state when the injected carriers decay, and the pump is again re-coupled, etc.

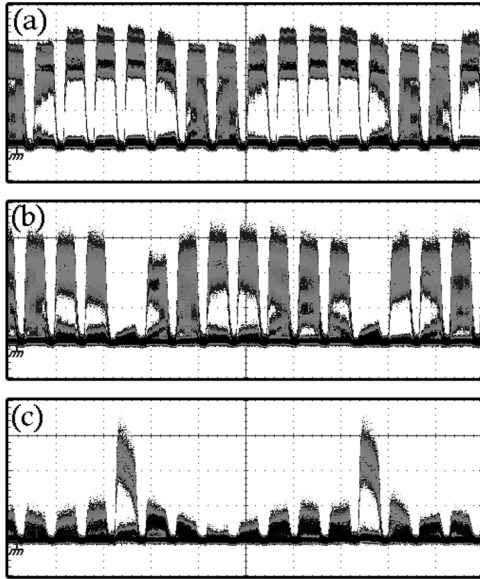


Fig. 7. CSA traces of packetized data encoded on C35 (a) before injection into port *in0*, (b) at the output of port *out0*, and (c) at the output of port *out1*. The time scale is 20 ns/div with infinite persistence enabled; no averaging is used. The apparent eye diagram is a sampling phenomenon of the CSA, and should be ignored as the data rate of 10 Gb/s is much faster than the displayed packet rate.

The extinction ratios and relative insertion losses of the switch are determined from the CSA traces. Extinction ratios of 11.5 dB (*in0-out0*) and 7.8 dB (*in0-out1*) are observed. The insertion loss of the cross state is about 4 dB higher than that of the bar state. This occurs mostly because of the TPA-induced free carrier absorption (FCA) generated by the high-power pumps, which induce higher losses for the probe section that coincides with the pump in time. Much of this loss can be avoided in electronic injection schemes, where the pump and probes will not propagate through the waveguides together. A certain amount of FCA will always exist where the carriers are injected, but it may be possible to match the insertion loss of the ring (about 1 to 2 dB [16]) to the combination of losses resulting from FCA and the waveguide crossing, thus equalizing the insertion loss for all possible pathways.

In addition to affecting the loss uniformity, the nonlinear loss also degrades the port-to-port crosstalk. Here, we define the port-to-port crosstalk for path *in0-out0* (*in0-out1*) as the ratio of the ON-state power through path *in0-out0* (*in0-out1*) to the OFF-state power through the path *in0-out1* (*in0-out0*). This metric is important for characterizing the effect of finite extinction ratios on other messages in the switch. It describes the amount of leakage power that may adversely affect another message in the switch. The crosstalk experienced for path *in0-out1*, which is the path that is degraded by the nonlinear loss, is 7.0 dB; in the case of path *in0-out0*, the crosstalk is 12.3 dB.

Finally, the dynamic power penalty of the switch is measured by using the same six-channel probe source incident on port *in0*. Here, the data are gated into 9.6-ns packets recurring in 12.8-ns slots by using an SOA. The same pump signal is used to switch one of every eight packets to port *out1*, while the rest travel to port *out0* (Fig. 7). The back-to-back curves for the power




Fig. 8. BER curves depicting the dynamic power penalty of a six-channel WDM signal propagating from *in0* to *out0* (●) and to *out1* (■). The back-to-back measurements, taken with the pumps disabled and the resonator modes realigned thermally, are shown by using dashed lines and open symbols, while the switched measurements are shown by using solid lines and filled symbols.

penalty measurements (Fig. 8) are taken with the pumps disabled and the resonances realigned to the wavelength channels by using the TEC controller. The BERT is gated over the arrival of the switched packets, separately for each measurement (*in0-out0* and *in0-out1*). The results show approximately 1.1 dB of difference between the two back-to-back curves, attributed mostly to narrowband filtering and consistent with the static result listed in Table I for C35 on port *out0*. An additional power penalty of 1.9 dB for path *in0-out0* is incurred under dynamic operation. We believe this is caused by slight power fluctuations imposed on the switched data packets, which exist under dynamic operation and result from small thermal variations. Since the thermal time constant is longer than the pump periodicity, these variations occur slowly across many pumping cycles, but affect parameters, such as the optimal receiver threshold setting over the course of measuring one BER curve.

For path *in0-out1*, an overall power penalty of 3.5 dB is measured. The causes of this penalty are divided between the nonlinear losses, the pump ASE noise, and dynamic operation. Note that the nonlinear losses and ASE accumulation were also present during the static BER measurements, which resulted in power penalties of less than or equal to 0.5 dB. However, the power fluctuations on the switched packet due to dynamic operation, which are manifest by the declining power observed across the packet duration [Fig. 7(c)], are now evident over the duration of a single packet. Therefore, the fluctuations must be attributed to carrier dynamics, as they occur much too fast for thermal effects. The pump signal is not able to maintain an adequate level of carrier injection over the entire length of the 9.6-ns packet. This is validated by the opposite shape of the suppressed packet in Fig. 7(b), and is clearly an example of the performance problems which can be mitigated by using electronic means of carrier injection. As in [8], voltage

waveforms may be more easily shaped to provide a constant level of charge over a given length of time. Furthermore, when electronic injection is used, the instabilities of the inplane optical pumping scheme are not encountered.

As a final note, measurements were not taken while simultaneously injecting two data signals on separate input ports due to the spatial proximity of the input waveguides and the size limitations of the coupling equipment in the current testbed. These measurements will be important for verifying consistent BER performance given the finite port-to-port crosstalk in future work. Furthermore, the cascability of the switches, in terms of their insertion loss and power penalty, is an important consideration in the design of large-scale NoCs, where signals need to traverse multiple switches as they propagate from source to destination. This work is currently being pursued in order to determine the most optimal switch configuration and network topology, based on the performance of the individual silicon photonic components [27].

V. CONCLUSION

We have characterized a multiwavelength 2×2 switch that can be implemented in silicon-photonic NoCs in a manner allowing low power consumption and small footprint. The all-optical pumping scheme was used to demonstrate the ability of the device to achieve fast switching (< 2 ns) with high extinction ratios (> 10 dB). Very low power penalties were measured at 10 Gb/s under static operation, and the wavelength crosstalk was shown to be minimal even as the injected signal power scaled to values larger than 20 mW. Moderate power penalties were obtained during ultrafast dynamic switching. The implementation of electronic carrier injection in future devices, which will be required for realistic implementations, is expected to minimize the dynamic power penalties by providing more constant charge levels in the ring waveguides over the duration of the pump's enabled state. Furthermore, electronic injection is expected to minimize FCA losses in the cross state, giving more uniformity to the insertion losses, extinction ratios, and port-to-port crosstalk values between the bar and cross states of the switch. Finally, future devices may achieve larger channel bandwidths, higher extinction ratios, and more thermal stability by using higher order resonators, at the expense of additional complexity.

REFERENCES

- [1] A. Shacham, B. G. Lee, A. Biberman, K. Bergman, and L. P. Carloni, "Photonic NOC for DMA communications in chip multiprocessors," in *Proc. 15th IEEE Symp. High-Performance Interconnects*, Aug. 2007, pp. 29–36.
- [2] M. Lipson, "Guiding, modulating, and emitting light on silicon—Challenges and opportunities," *J. Lightw. Technol.*, vol. 23, no. 12, pp. 4222–4238, Dec. 2005.
- [3] B. Jalali, "Silicon photonics," *J. Lightw. Technol.*, vol. 24, no. 12, pp. 4600–4615, Dec. 2006.
- [4] T. Barwicz *et al.*, "Silicon photonics for compact, energy-efficient interconnects," *J. Opt. Netw.*, vol. 6, no. 1, pp. 63–73, Jan. 2007.
- [5] M. A. Popović, T. Barwicz, M. S. Dahlem, F. Gan, C. W. Holzwarth, P. T. Rakich, H. I. Smith, E. P. Ippen, and F. X. Kärtner, "Tunable, fourth-order silicon microring-resonator add-drop filters," presented at the Eur. Conf. Optical Communication, Berlin, Germany, Sep. 2007, paper 1.2.3.
- [6] L. Chen, N. Sherwood-Droz, and M. Lipson, "Compact bandwidth-tunable microring resonators," *Opt. Lett.*, vol. 32, no. 22, pp. 3361–3363, Nov. 15, 2007.
- [7] F. Xia, L. Sekaric, and Y. Vlasov, "Ultracompact optical buffers on a silicon chip," *Nature Photon.*, vol. 1, pp. 65–71, Jan. 2007.
- [8] Q. Xu, S. Manipatruni, B. Schmidt, J. Shakya, and M. Lipson, "12.5 Gbit/s carrier-injection-based silicon micro-ring silicon modulators," *Opt. Express*, vol. 15, no. 2, pp. 430–436, Jan. 22, 2007.
- [9] B. G. Lee, B. A. Small, Q. Xu, M. Lipson, and K. Bergman, "Characterization of a 4×4 Gb/s parallel electronic bus to WDM optical link silicon photonic translator," *IEEE Photon. Technol. Lett.*, vol. 19, no. 7, pp. 456–458, Apr. 1, 2007.
- [10] C. Li, L. Zhou, and A. W. Poon, "Silicon microring carrier-injection-based modulators/switches with tunable extinction ratios and OR-logic switching by using waveguide cross-coupling," *Opt. Express*, vol. 15, no. 8, pp. 5069–5076, Apr. 16, 2007.
- [11] S. F. Preble, Q. Xu, B. S. Schmidt, and M. Lipson, "Ultrafast all-optical modulation on a silicon chip," *Opt. Lett.*, vol. 30, no. 21, pp. 2891–2893, Nov. 1, 2005.
- [12] T. C. Huang, G. J. Simonis, and L. A. Coldren, "Directional coupler optical switch constructed from field-induced waveguides," *Electron. Lett.*, vol. 28, no. 25, pp. 2288–2289, Dec. 1992.
- [13] K. A. Williams, G. F. Roberts, T. Lin, R. V. Penty, I. H. White, M. Glick, and D. McAuley, "Integrated optical 2×2 switch for wavelength multiplexed interconnects," *IEEE J. Sel. Topics Quantum Electron.*, vol. 11, no. 1, pp. 78–85, Jan./Feb. 2005.
- [14] O. Fidaner, H. V. Demir, V. A. Sabnis, J.-F. Zheng, J. S. Harris Jr., and D. A. B. Miller, "Integrated photonic switches for nanosecond packet-switched optical wavelength conversion," *Opt. Express*, vol. 14, no. 1, pp. 361–368, Jan. 2006.
- [15] P. Dong, S. F. Preble, and M. Lipson, "All-optical compact silicon comb switch," *Opt. Express*, vol. 15, no. 15, pp. 9600–9605, Jul. 23, 2007.
- [16] B. G. Lee, A. Biberman, P. Dong, M. Lipson, and K. Bergman, "All-optical comb switch for multi-wavelength message routing in silicon photonic networks," *IEEE Photon. Technol. Lett.*, vol. 20, no. 10, pp. 767–769, May 15, 2008.
- [17] A. Biberman, B. G. Lee, P. Dong, M. Lipson, and K. Bergman, "250 Gb/s multi-wavelength operation of microring resonator-based broadband comb switch for silicon photonic networks-on-chip," presented at the Eur. Conf. Optical Communication, Brussels, Belgium, Sep. 2008, Paper P.2.23.
- [18] Y. A. Vlasov, W. M. J. Green, and F. Xia, "High-throughput silicon nanophotonic wavelength-insensitive switch for on-chip optical networks," *Nature Photon.*, vol. 2, pp. 242–246, Apr. 2008.
- [19] B. G. Lee, A. Biberman, N. Sherwood-Droz, C. B. Poitras, M. Lipson, and K. Bergman, "High-speed 2×2 switch for multi-wavelength message routing in on-chip silicon photonic networks," presented at the Eur. Conf. Optical Communication, Brussels, Belgium, Sep. 2008, Paper Tu.3.C.3.
- [20] H. Wang, M. Petracca, A. Biberman, B. G. Lee, L. P. Carloni, and K. Bergman, "Nanophotonic optical interconnection network architecture for on-chip and off-chip communications," presented at the Optical Fiber Communications Conf, San Diego, CA, Feb. 2008, Paper JThA92.
- [21] A. Shacham and K. Bergman, "Building ultralow-latency interconnection networks using photonic integration," *IEEE Micro*, vol. 27, no. 4, pp. 6–20, Jul./Aug. 2007.
- [22] S. J. B. Yoo, "Optical packet and burst switching technologies for the future photonic internet," *J. Lightw. Technol.*, vol. 24, no. 12, pp. 4468–4492, Dec. 2006.
- [23] A. Yariv, Y. Xu, R. K. Lee, and A. Scherer, "Coupled-resonator optical waveguide: A proposal and analysis," *Opt. Lett.*, vol. 24, no. 11, pp. 711–713, Jun. 1999.
- [24] W. Bogaerts, P. Dumon, D. V. Thourhout, and R. Baets, "Low-loss, low-cross-talk crossings for silicon-on-insulator nanophotonic waveguides," *Opt. Lett.*, vol. 32, no. 19, pp. 2801–2803, Oct. 1, 2007.
- [25] T. Barwicz, M. A. Popović, M. R. Watts, P. T. Rakich, E. P. Ippen, and H. I. Smith, "Fabrication of add-drop filters based on frequency-matched microring resonators," *J. Lightw. Technol.*, vol. 24, no. 5, pp. 2207–2218, May 2006.
- [26] R. Soref and B. R. Bennett, "Electrooptical effects in silicon," *IEEE J. Quantum Electron.*, vol. QE-23, no. 1, pp. 123–129, Jan. 1987.
- [27] J. Chan, A. Biberman, B. G. Lee, and K. Bergman, "Insertion loss analysis in a photonic interconnection network for on-chip and off-chip communications," presented at the Annu. Mtg. Lasers Electro-Optics Soc., Newport Beach, CA, Nov. 2008, Paper TuT3.



Benjamin G. Lee (S'04) received the B.S. degree in electrical engineering from Oklahoma State University, Stillwater, in 2004 and the M.S. and Ph.D. degrees in electrical engineering from Columbia University, New York, in 2006 and 2009, respectively.

Currently, he is a Postdoctoral Researcher with the IBM T. J. Watson Research Center, Yorktown Heights, NY. His research interests include silicon-photonic devices, integrated optical switches and networks for high-performance computing systems, and all-optical processing systems.

Dr. Lee is a member of the IEEE Photonics Society and the Optical Society of America.



Aleksandr Biberman (S'05) received the B.S. degree in electrical and computer and systems engineering from Rensselaer Polytechnic Institute, Troy, NY, in 2006, and the M.S. degree in electrical engineering from Columbia University, New York, in 2008, where he is currently pursuing the Ph.D. degree in electrical engineering.

His research interests include silicon nanophotonic devices for networks-on-chip and interchip communication, photonic interconnection networks for multicore chip architectures, optical networking

in high-performance computing systems, as well as silicon photonic devices for parametric optical processes and systems.



Nicolás Sherwood-Droz (S'01) received the B.S. (Hons.) and M.S. degrees in electrical and computer engineering from Worcester Polytechnic Institute, Worcester, MA, in 2002 and 2005, respectively, and is currently pursuing the Ph.D. degree in electrical and computer engineering from Cornell University, Ithaca, NY.



Carl B. Poitras (M'09) received the B.Sc. degree in physics from the University of Québec, Montréal, QC, Canada, in 1996, the M.Sc. degree from the National Institute of Telecommunications, Montréal, in 2001, and the Ph.D. degree in electrical engineering from Cornell University, Ithaca, NY, in 2006.

Currently, he is a Research Associate in the School of Electrical and Computer Engineering, Cornell University (in Prof. Michal Lipson's Nanophotonics group). His research is in photoluminescence, integrated silicon photonics, and rare earth-doped GaN.

Michal Lipson (SM'07) received the B.S., M.S., and Ph.D. degrees in physics from the Technion—Israel Institute of Technology, Haifa, Israel, in 1992, 1994, and 1998, respectively.

In 1998, she joined the Department of Material Science and Engineering at the Massachusetts Institute of Technology (MIT), Cambridge, as a Postdoctoral Associate. In 2001, she joined the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY, as an Assistant Professor. Her current research interests include novel on-chip nanophotonics devices. She is currently a Topical Editor of *Optics Letters*.

Dr. Lipson was the recipient of a National Science Foundation CAREER Award in 2004 and the IBM Faculty Award in 2006.



Keren Bergman (S'87–M'93–SM'07–F'08) received the B.S. degree from Bucknell University, Lewisburg, PA, in 1988, and the M.S. and Ph.D. degrees from the Massachusetts Institute of Technology, Cambridge, in 1991 and 1994, respectively, all in electrical engineering.

Currently, she is a Professor of Electrical Engineering in the Department of Electrical Engineering, Columbia University, New York, where she also directs the Lightwave Research Laboratory. Her research programs involve optical interconnection

networks for advanced computing systems, photonic packet switching, and nanophotonic networks-on-chip. She is the co-Editor-in-Chief of the recently launched IEEE/OSA JOURNAL OF OPTICAL COMMUNICATIONS AND NETWORKING.

Dr. Bergman is a Fellow of the Optical Society of America.