

# Monolithic Silicon Integration of Scaled Photonic Switch Fabrics, CMOS Logic, and Device Driver Circuits

Benjamin G. Lee, Alexander V. Rylyakov, William M. J. Green, Solomon Assefa, Christian W. Baks, Renato Rimolo-Donadio, Daniel M. Kuchta, Marwan H. Khater, Tymon Barwicz, Carol Reinholm, Edward Kiewra, Steven M. Shank, Clint L. Schow, and Yurii A. Vlasov

**Abstract**—We demonstrate  $4 \times 4$  and  $8 \times 8$  switch fabrics in multistage topologies based on  $2 \times 2$  Mach–Zehnder interferometer switching elements. These fabrics are integrated onto a single chip with digital CMOS logic, device drivers, thermo-optic phase tuners, and electro-optic phase modulators using IBM’s 90 nm silicon integrated nanophotonics technology. We show that the various switch-and-driver systems are capable of delivering nanosecond-scale reconfiguration times, low crosstalk, compact footprints, low power dissipations, and broad spectral bandwidths. Moreover, we validate the dynamic reconfigurability of the switch fabric changing the state of the fabric using time slots with sub-100-ns durations. We further verify the integrity of high-speed data transfers under such dynamic operation. This chip-scale switching system technology may provide a compelling solution to replace some routing functionality currently implemented as bandwidth- and power-limited electronic switch chips in high-performance computing systems.

**Index Terms**—CMOS integrated circuits, optical switches, photonic integrated circuits.

## I. INTRODUCTION

WIDESPREAD use of point-to-point optical interconnects in the current generation of high-performance computers has been a significant factor in the continued improvements and record-setting performances of many of today’s top machines [1]–[3]. It is envisioned that the next generation of systems will leverage an even greater number of optical components that have also been scaled in bandwidth, density, and power efficiency beyond those of today. As the bandwidth bottleneck at the electro–optical interface of future systems de-

scends into deeper levels of the package (necessitating board-level optics, carrier-level optics, and eventually chip-level optics), performing some amount of switching in the optical domain becomes more attractive.

Consequently, an electronically controlled optical switching technology such as micro-electro-mechanical systems (MEMS) provides the potential to bypass the power density issues and pin limitations that plague the future scalability of electronic switch chips. While MEMS provides an unparalleled port count scaling in comparison to other optical switching technologies, its millisecond-scale reconfiguration time limits the set of interesting application spaces to which it can offer performance enhancements [4], [5]. On the other hand, the silicon photonic platform has the potential to realize dense and low-power optical switches with nanosecond-scale reconfiguration times, while interconnecting a low-to-moderate number of ports. This capability more broadly targets a widespread set of demands within computer communications.

Previously, scaled switch fabrics (defined here as having four or more input and four or more output ports) have been implemented in the silicon material system using both thermo-optic and electro-optic phase modulators [6]–[9]. We also previously reported an electro-optic scaled switch fabric [10], which additionally was used to demonstrate a proof-of-concept photonic switching system that employed flip-chip integration to join the fabric with digital CMOS switch drivers [11], [12]. Since then, the development of IBM’s 90 nm silicon integrated nanophotonics technology [13] has made possible the full monolithic integration of a similar switching system with enhanced function. Monolithic integration enables increasingly sophisticated optoelectronic designs and potentially reduces system cost. We have previously made preliminary reports of monolithically integrated CMOS drivers and CMOS-photonic switch fabrics [14], [15]. Here, we expand upon these works by reporting additional details and improved performances.

## II. SWITCHING ELEMENTS

We first describe the performance of the CMOS-driven switching elements, based on Mach–Zehnder (MZ) interferometers, which become the building blocks in scaled switch fabrics. Fig. 1(a) shows a die photo of a test site that includes four digital CMOS drivers, one of which is wired to a  $2 \times 2$  MZ-based photonic switch [see Fig. 1(b)]. The MZ switch

Manuscript received June 20, 2013; revised August 5, 2013; accepted August 20, 2013. Date of publication August 30, 2013; date of current version January 10, 2014. This work was supported by the Defense Advanced Research Projects Agency (DARPA) under Contract HR0011-08-C-0102.

B. G. Lee, A. V. Rylyakov, W. M. J. Green, S. Assefa, C. W. Baks, R. Rimolo-Donadio, D. M. Kuchta, M. H. Khater, T. Barwicz, C. Reinholm, E. Kiewra, C. L. Schow, and Y. A. Vlasov are with the IBM T. J. Watson Research Center, Yorktown Heights, New York, NY 10598 USA (e-mail: bglee@us.ibm.com; sasha@us.ibm.com; wgreen@us.ibm.com; sasssefa@us.ibm.com; cbaks@us.ibm.com; renato.rimolo@us.ibm.com; kuchta@us.ibm.com; mkhater@us.ibm.com; tymon@us.ibm.com; reinholm@us.ibm.com; kiewre@us.ibm.com; cschow@us.ibm.com; yvlasov@us.ibm.com).

S. M. Shank is with the IBM Systems and Technology Group, Microelectronics Division, Essex Junction, VT 05452 USA (e-mail: smshank@us.ibm.com).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JLT.2013.2280400

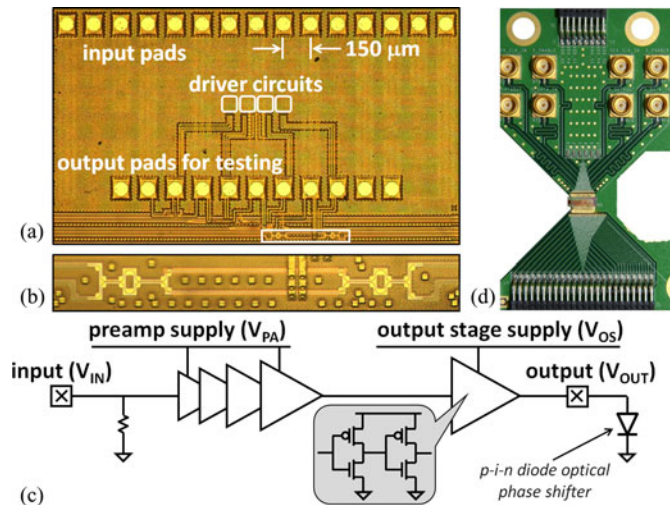


Fig. 1. (a) Die image showing a four-channel driver test site. (b) Magnified image of the WIMZ switch outlined in (a). (c) Schematic of the digital switch driver. (d) Custom test card.

is referred to as a wavelength-insensitive MZ (WIMZ) because it incorporates a broadband directional coupler design, as previously reported [16]. The switch contains in one arm a forward-biased horizontal-junction p-i-n diode driven directly from the output of the driver. The entire switch occupies  $0.02 \text{ mm}^2$ . The driver [see Fig. 1(c)] employs a five-stage digital buffer with a fanout of 2. Each stage is comprised of two inverters with the same transistor dimensions previously reported in a bulk CMOS (electronic only) implementation [12]. The measured output impedance of the driver is less than  $10 \Omega$  for both pMOS and nMOS transistors at supply voltages greater than  $0.9 \text{ V}$ . Low output impedance, particularly for the pMOS transistor, is necessary in order to efficiently source the required on-state current from the supply rail to the diode with a minimal voltage drop. The power supply is segmented such that the first four buffer stages receive power from the preamplifier supply rail, while the final driver stage—and thus the device—is powered by the output stage (OS) rail.

The chip, as in other test sites reported subsequently, is assembled on a custom printed circuit board (PCB) [see Fig. 1(d)]. A  $2 \times 8$  pin connector at the top of the PCB provides supply voltage and ground connections, while surface-mount SMP connectors provide the moderate-speed digitized control signals to the switch. The board contains cutouts near the center to allow fiber access for edge coupling. Tapered lensed fibers are used with light oriented to the transverse electric (TE) polarization. (The MZ-based switching elements are designed for single polarization; thus polarization-diversity or polarization-management schemes must be included in the high-level design.) A  $2 \times 24$  pin connector located at the bottom of the PCB delivers thermal tuning voltages used for controlling individual switch bias points within the  $4 \times 4$  and  $8 \times 8$  fabrics discussed later on, but not specifically used in the  $2 \times 2$  WIMZ switch reported in this section.

The spectral response of the CMOS-driven WIMZ switch, plotted in Fig. 2(a), is measured using a broadband infrared light source, a polarizer, and an optical spectrum analyzer (OSA).

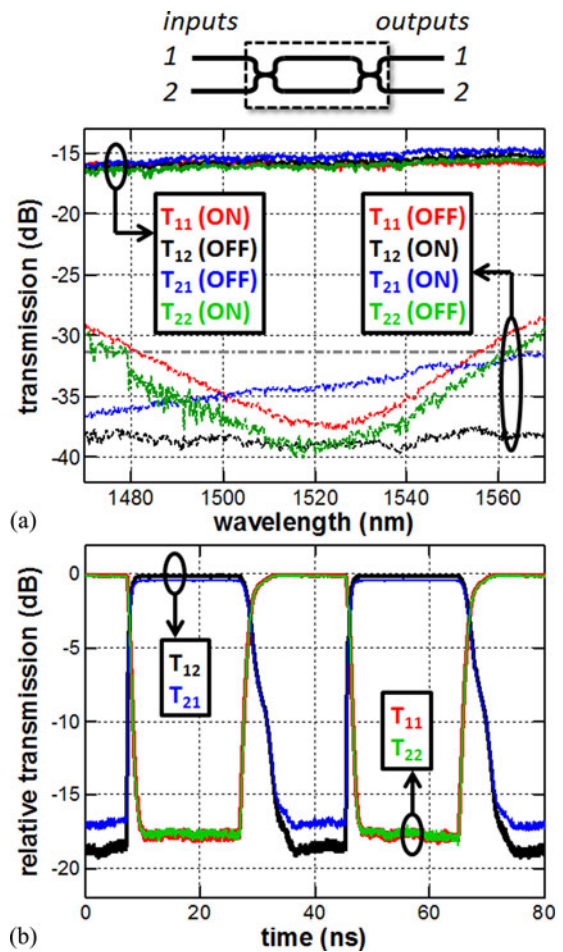


Fig. 2. (a) Spectral and (b) transient responses of the CMOS-driven WIMZ switching element, plotted for the four input/output port configurations ( $T_{11}$ ,  $T_{12}$ ,  $T_{21}$ , and  $T_{22}$ ) in both the on and off states. The gray dash-dot line in (a) denotes the  $-15$ -dB crosstalk threshold.

The OSA employs a  $1$ -nm resolution bandwidth and records the average of eight internal sweeps. The  $y$ -axis in Fig. 2(a) plots the output intensity relative to the intensity in the input fiber. The crosstalk for all states and configurations remains below  $-15 \text{ dB}$  over a  $75$ -nm spectral bandwidth centered near a wavelength of  $1520 \text{ nm}$ . The broadband directional coupler design is successful in making the switch well-suited to multiwavelength operation.

The transient response of the  $2 \times 2$  WIMZ switch is shown in Fig. 2(b). Here, a  $25$ -MHz square-wave with  $50\%$  duty cycle is applied to the input of the driver,  $V_{IN}$ , while a continuous-wave optical signal with wavelength of  $1530 \text{ nm}$  is injected into the two input ports of the switch alternately. The light egressing from the switch's two output ports is amplified with an erbium-doped fiber amplifier, received using a  $10$ -GHz photodetector, and viewed on a sampling oscilloscope. The turn-off and turn-on transients (measured between  $10$  and  $90\%$  of the maximum) are  $1.2$  and  $4.5 \text{ ns}$ , respectively. The short transition times can allow a scaled fabric comprised of these switches to be reconfigured on a packet-by-packet basis.

For the above measurements, the switch is operated with supply voltages near  $1.0 \text{ V}$ . Fig. 3 illustrates the measured sensitivity to supply settings. Due to the sinusoidal nature of the MZ's






Fig. 3. Tolerance to the (a) supply voltage and (b) supply current, measured at a wavelength of 1520 nm.




Fig. 4. Power consumption of the CMOS-driven WIMZ switch versus switching frequency.

amplitude versus phase relationship, deviation from the supply setting that provides an optimal phase shift rapidly degrades the crosstalk. (The same data are plotted in both Fig. 3(a) and (b) as a function of supply voltage and current, respectively.) Noise-tolerant photonic switch designs based on cascaded MZ structures have been realized previously, and could alleviate this problem [17].

The average power dissipation for the switch and driver is plotted as a function of square-wave frequency in Fig. 4. Both supplies are set to 1.0 V. Near zero frequency, the power consumption is the average of the on-state and off-state values. The power that is incurred by changing between states, which

TABLE I  
CMOS-DRIVEN  $2 \times 2$  WIMZ EXTRACTED POWER AND ENERGY PERFORMANCE

Supply Rail	Static Power <sup>1</sup> (mW)	Switching Energy <sup>2</sup> (pJ/cycle)
PA	0.4	9
OS	1.2	7
Total (PA+OS)	1.6	16

<sup>1</sup>Average power assuming a 50% state probability.

<sup>2</sup>Includes turn-on and turn-off dissipations.

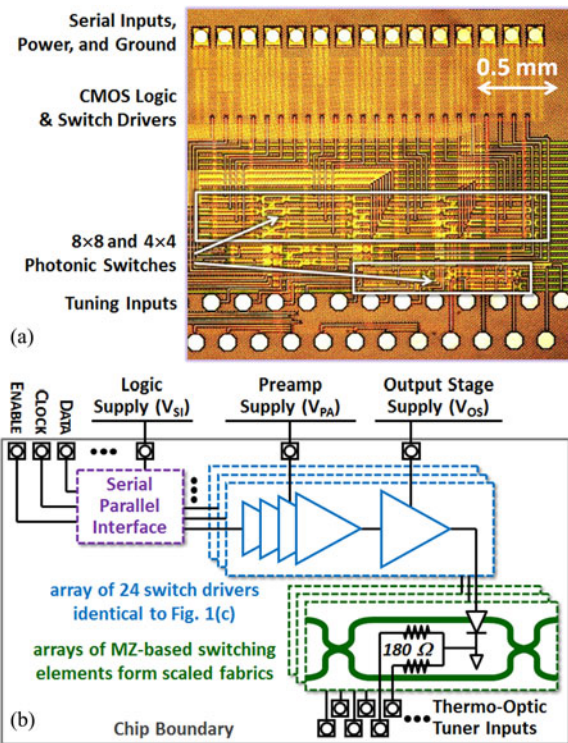


Fig. 5. (a) Die image of a scaled switching system-on-chip. (b) Schematic diagram of the CMOS logic and switch driver.

increases linearly as a function of frequency, then adds to the zero-frequency value. From these measurements, the per-cycle (turn-on plus turn-off) switching energy for the CMOS-driven WIMZ is inferred to be 16 pJ, and the average holding power is 1.6 mW. The split contributions coming from the two supply rails can be seen in Table I.

### III. SWITCH FABRICS

Next, we describe a test site implementing scaled switch fabrics laid out in multistage topologies employing MZ-based  $2 \times 2$  switching elements as building blocks. Fig. 5(a) shows a die image of the site, while Fig. 5(b) displays the schematic. Because of the large number of switch drivers used in this site, the inputs to each driver are addressed through a serial-to-parallel interface consisting of a 28-bit shift register. The chip receives the serial electrical inputs (DATA, CLOCK, and ENABLE) which control an array of parallel buffers, each accessing a digital inverter-based CMOS switch driver identical to those described in Section II. Each driver is connected to a




Fig. 6. (a) Topological arrangement and (b) micrograph of the  $4 \times 4$  fabric. (c) Topological arrangement and (d) micrograph of the  $8 \times 8$  fabric. The dashed lines in (a) and (c) represent  $2 \times 2$  MZ-based switching elements.

single forward-biased horizontal-junction p-i-n diode arranged in one arm of the MZ, which comprises one element of a larger multiport photonic switch. These MZ switches are also identical to the one described in Section II, except that these utilize a standard directional coupler, not specifically designed for wavelength insensitivity, in order to reduce complexity and risk in the scaled fabrics. Each MZ is also equipped with dual thermo-optic tuners for phase trimming.

Static measurements on an  $8 \times 8$  switch fabric and dynamic measurements on a  $4 \times 4$  switch fabric are described in this section. The  $4 \times 4$  fabric is arranged in a two-stage binary tree topology, and the  $8 \times 8$  in a custom four-stage topology for proof-of-concept (see Fig. 6). The area occupied by the serial/parallel interface logic, switch drivers,  $4 \times 4$  fabric, and  $8 \times 8$  fabric is  $0.007 \text{ mm}^2$ ,  $0.015 \text{ mm}^2$  per channel,  $0.165 \text{ mm}^2$ , and  $0.675 \text{ mm}^2$ , respectively.

#### A. $8 \times 8$ Mach-Zehnder-Based Switch Fabric

We demonstrate proper static operation of the  $8 \times 8$  switch fabric by reporting the spectral characteristics of all signal and crosstalk paths in two illustrative states: when all switching elements are off (see Fig. 7), and when all switching elements are on (see Fig. 8). The supply voltages for the  $8 \times 8$  fabric are held at 1.3 V. The total static power consumed by the fabric is 4.4 mW in the off state and 32.2 mW in the on state in addition to thermo-optic tuning power. The tuners are optimized for each of the 16 MZ elements such that for each MZ the cross state is implemented when the driver is in the off state (same as Section II). A custom power card provides a computer controlled interface that generates up to 32 analog tuning voltages and delivers them to the test card. The power required to tune the 16 MZ switches into the cross state averages 14 mW per MZ with maximum and minimum values of 29 and 1 mW, respectively, corresponding to a near uniform distribution of phase differences across the 16 pairs of interferometer arms. By optimizing the MZs (in the off state) to implement the least power option of either the cross or the bar state, average thermo-optic tuning power could be reduced by approximately a factor of 2 over that reported here.

Spectral measurements were obtained in the same manner as described in Section II. With all switching elements in the off state, the crosstalk at band center (1460 nm) is  $-20$  dB below the signal level for each path. Across all 64 traces displayed in Fig. 7, there is a 20-nm window (1460 nm  $\pm$  10 nm) over which crosstalk remains  $-15$  dB below the signal. The measured crosstalk significantly increases when all the switching elements are turned on. The crosstalk measured on a few paths shown in Fig. 8 (for input ports 1, 2, 7, and 8) is as large as  $-5$  to  $-8$  dB within the same bandwidth. In particular, coherence phenomena are clearly identifiable on these traces, which represent switch configurations where two crosstalk paths combine to beat against each other causing additional wavelength dependences. Variation across the diodes' voltage-current relationships, the diodes' current-phase relationships, and/or the drivers' output impedances causes each of the 16 MZ switching elements to be optimized at slightly different OS supply voltage settings. Since there is only one OS supply voltage that powers all of the 16 MZ switching elements simultaneously (depending only on the driver's input state), many of the switching elements must be operated far from their ideal supply setting, deteriorating crosstalk performance for the fabric. A more tolerant photonic switch design or a more flexible electronic driver design can mitigate this limitation in future implementations. The figure does demonstrate, however, the correct static routing operation of the switch fabric with most paths adhering to crosstalk below  $-15$  dB. In addition to optical crosstalk, crosstalk arising from the electrical or thermal sources may also deteriorate switch performance. However, these sources of crosstalk are expected to be secondary compared to optical crosstalk, and therefore will be reserved for a future investigation.

#### B. Dynamic Routing in a $4 \times 4$ Switch Fabric

We demonstrate dynamic routing functionality using the simpler  $4 \times 4$  switch fabric. First extensive static spectral measurements were recorded for the  $4 \times 4$  fabric in similar fashion to those recorded for the  $8 \times 8$  fabric. However, only one  $2 \times 2$  element was enabled at a time so that the OS supply voltage could be optimized for each element within the  $4 \times 4$  fabric. In this manner, all relevant input/output traces were recorded for five states corresponding to all the elements being disabled and each of the four elements being enabled one at a time. The results are only summarized here for brevity. Optimal supply voltage settings ranged from 1.5 to 1.6 V for the four MZ switching elements. The phase tuners were optimized once and held constant for all measurements, consuming an average of 8.6 mW per MZ. The switch fabric displays a center wavelength near 1470 nm, and demonstrates a bandwidth of 30 nm in which crosstalk remains below  $-14$  dB. The total insertion loss including fiber coupling is recorded for two switch states: when all MZs are off (the signal passes through two MZs in the off state) and when all MZs are on (the signal passes through two MZs in the on state). The measured off-state insertion loss for all four signal paths ( $T_{14}$ ,  $T_{22}$ ,  $T_{33}$ , and  $T_{41}$ ) is  $14.4 \pm 0.3$  dB, and the measured on-state insertion loss for all four signal paths ( $T_{11}$ ,  $T_{23}$ ,  $T_{32}$ , and  $T_{44}$ ) is  $16.0 \pm 0.3$  dB. Note that the total losses are similar



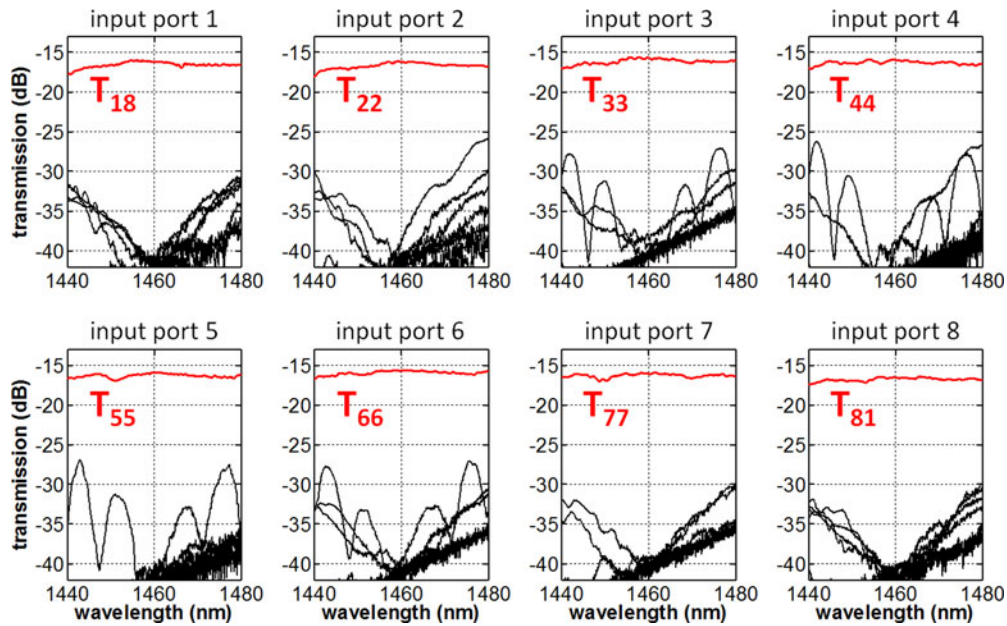


Fig. 7. Spectral characteristics of all signal (red) and crosstalk (black) paths for the  $8 \times 8$  fabric when all MZs are in the off state. For each plot, light is injected into one input port while all eight output ports are recorded in succession.




Fig. 8. Spectral characteristics of all signal (red) and crosstalk (black) paths for the  $8 \times 8$  fabric when all MZs are in the on state. For each plot, light is injected into one input port while all eight output ports are recorded in succession.

to those reported for both the  $2 \times 2$  MZ switching element [see Fig. 2(a)] and the  $8 \times 8$  switch fabric (see Fig. 7), indicating that the bulk of these losses result from input/output fiber coupling. Unfortunately, test sites for accurately determining the switch loss per stage were not included in the design. The additional 1.6 dB of loss in the on state is attributed to free-carrier losses and is consistent with the previous studies [16].

To demonstrate dynamic reconfiguration of the switching system, the MZ states shown in Fig. 9(a) were chosen and assigned to four successive time slots. The experimental setup in Fig. 9(b) was employed to characterize the switching perfor-

mance under this dynamic operation with supply voltages fixed at 1.5 V. A continuous stream of 16 Gb/s data was injected into any one of the fabric's four optical input ports, while cycling through these four illustrative states in succession at the rate of 90 ns per time slot. The egressing data packets for each input/output port configuration across the four time slots are displayed in Fig. 10(a). It can be verified that the switch fabric does indeed route packets to their appropriate destination. Furthermore, even at a fixed supply voltage, the scope traces show no observable crosstalk within the interior of any data packet window. Two traces show crosstalk present between the time slots




Fig. 9. (a) States assigned to the four switching elements for each of the four time slots. (b) Experimental setup used to demonstrate dynamic routing. Solid and dashed lines represent optical and electrical connections, respectively. An inset eye diagram shows the performance at 16 Gb/s of the reference transmitter (bypassing the switch in the setup).




Fig. 10. (a) Scope traces for each input/output configuration demonstrating proper routing. (b) 16-Gb/s eye diagrams recorded at the center of each packet in (a) using time scales of 20 ps/div and a constant amplitude scale (0.7 mW full scale).

(i.e., during the reconfiguration transient) where interpacket guard bands would exist in a real system. This results from inserting a continuous stream of data rather than discrete packets. Eye diagrams are shown in Fig. 10(b) for the data at the center of each packet present in Fig. 10(a) providing visual verification of the fidelity of the transmitted data signal. Fig. 11 provides an example of the  $4 \times 4$  switch fabric's transition times, recorded by injecting continuous-wave light into the fabric while dynamically reconfiguring. The scaled fabric demonstrates transition times slightly longer than those shown for the  $2 \times 2$  switch, but verifies the fundamentally nanosecond-scale response time of the switching mechanism employed within fabrics such as this one.

The measured power and energy dissipation of the switching system is obtained by applying to all drivers a square wave with frequency swept from 1 to 4 MHz. A linear fit to the




Fig. 11. Example switching transients of the  $4 \times 4$  fabric measured through the path  $T_{44}$ . The traces are obtained by bypassing the modulator in Fig. 9(b) so that continuous-wave light is injected into the switch fabric. Parts (a) and (b) denote the rising and falling edges, respectively, of the corresponding packet in Fig. 10(a). The 10/90 transition times are labeled. The scope uses a time scale of 5 ns/div.

TABLE II  
CMOS-DRIVEN  $4 \times 4$  MZ FABRIC EXTRACTED POWER AND ENERGY PERFORMANCE

Supply Rail	Static Power <sup>1</sup> (mW)	Switching Energy <sup>2</sup> (pJ/cycle)
Logic	< 2.1	---
PA	2.4	82
OS	7.6	24
Thermo-Optics	34.4	---
Total	46.5	106

<sup>1</sup>Average power assuming a 50% state probability.

<sup>2</sup>Includes turn-on and turn-off dissipations.

measured power versus frequency data provides information about the average static power and dynamic switching energy as in Section II. Table II summarizes the extracted power and energy parameters. For the  $4 \times 4$  switch, the total dynamic energy of 106 pJ/cycle consumes less than 1 mW of power when changing states at a rate below 9 MHz, and thus adds negligibly to the 46.5 mW of average static power consumption. This static power, which is dominated by the thermo-optic contribution, could again be reduced by applying the scheme described in Section III-A, where a MZ switching element's off-state configuration (cross or bar) is assigned to be the one which consumes less thermo-optic tuning power. Conservatively, assuming that the switch facilitates four 25-Gb/s wavelength channels, the electrical contribution (including thermal) to the throughput-normalized energy can be estimated to be approximately 120 fJ/bit. Dong *et al.* demonstrate an order-of-magnitude reduction in thermo-optic tuning power by using



undercut structures [18]. Nevertheless, even without this improvement, contributions that arise from added laser power or optical amplification needed to overcome switch insertion losses are expected to dominate over the electrical contribution.

#### IV. CONCLUSION

We have realized  $4 \times 4$  and  $8 \times 8$  switch fabrics consisting of  $2 \times 2$  MZ interferometer switching elements integrated together with digital CMOS logic, device drivers, thermo-optic phase tuners, and electro-optic phase modulators. The photonic switching system is fabricated in IBM's 90 nm silicon integrated nanophotonics technology. We have shown that the building block switch elements are capable of delivering fast reconfigurability ( $<5$  ns), relatively low crosstalk ( $<-15$  dB over bandwidth), compact footprint ( $0.02 \text{ mm}^2$ ), low power dissipation ( $<2$  mW), and broad spectral bandwidth ( $>75$  nm). Furthermore, the combination of these building blocks into scaled fabrics in proof-of-concept topologies does not reveal any significant performance limitation that cannot be straightforwardly addressed in future designs. For the current implementation, improvements to optical losses (including fiber coupling losses) and further reductions in optical crosstalk may be the most necessary advancements required for the technology to have meaningful system impact.

#### ACKNOWLEDGMENT

The authors thank Dr. J. Shah of the Defense Advanced Research Projects Agency, Microsystems Technology Office, for inspiration and support. The views, opinions, and/or findings contained in this article are those of the authors and should not be interpreted as representing the official views or policies, either expressed or implied, of DARPA or the Department of Defense. Approved for public release, distribution unlimited.

#### REFERENCES

- [1] TOP500 Supercomputer Site. (2013). [Online]. Available: <http://top500.org/>
- [2] D. Chen, N. A. Easley, P. Heidelberger, R. M. Senger, Y. Sugawara, S. Kumar, V. Salapura, D. Satterfield, B. Steinmacher-Burow, and J. Parker, "The IBM blue gene/Q interconnection fabric," *IEEE Micro*, vol. 32, no. 1, pp. 32–43, Jan./Feb. 2012.
- [3] A. Benner, D. M. Kuchta, P. K. Pepeljugin, R. A. Budd, G. Hougham, B. V. Fasano, K. Marston, H. Bagheri, E. J. Seminaro, H. Xu, D. Meadowcroft, M. H. Fields, L. McColloch, M. Robinson, F. W. Miller, R. Kaneshiro, R. Granger, D. Childers, and E. Childers, "Optics for high-performance servers and supercomputers," in *Proc. Opt. Fiber Commun. Conf.*, Mar. 2010, pp. 1–3, Paper OTuH1.
- [4] L. Schares, X. J. Zhang, R. Wagle, D. Rajan, P. Selo, S. P. Chang, J. Giles, K. Hildrum, D. Kuchta, J. Wolf, and E. Schenfeld, "A reconfigurable interconnect fabric with optical circuit switch and software optimizer for stream computing systems," in *Proc. Opt. Fiber Commun. Conf.*, Mar. 2009, pp. 1–3, Paper OTuA1.
- [5] S. Kamil, L. Oliker, A. Pinar, and J. Shalf, "Communication requirements and interconnect optimization for high-end scientific applications," *IEEE Trans. Parallel Distrib. Syst.*, vol. 21, no. 2, pp. 188–202, Feb. 2010.
- [6] R. Ji, J. Xu, and L. Yang, "Five-port optical router based on microring switches for photonic networks-on-chip," *IEEE Photon. Technol. Lett.*, vol. 25, no. 5, pp. 492–495, Mar. 1, 2013.
- [7] L. Chen and Y. Chen, "Compact, low-loss and low-power  $8 \times 8$  broadband silicon optical switch," *Opt. Exp.*, vol. 20, no. 17, pp. 18977–18985, Aug. 2, 2012.
- [8] N. Sherwood-Droz, H. Wang, L. Chen, B. G. Lee, A. Biberman, K. Bergman, and M. Lipson, "Optical  $4 \times 4$  hitless silicon router for optical networks-on-chip (NoC)," *Opt. Exp.*, vol. 16, no. 20, pp. 15915–15922, Sep. 29, 2008.
- [9] P. DasMahapatra, A. Rohit, R. Stabile, and K. A. Williams, "Broadband  $4 \times 4$  switch matrix using fifth-order resonators," in *Proc. Opt. Fiber Commun. Conf.*, Mar. 2013, pp. 1–3, Paper OW3H.2.
- [10] M. Yang, W. M. J. Green, S. Assefa, J. Van Campenhout, B. G. Lee, C. V. Jahnes, F. E. Doany, C. L. Schow, J. A. Kash, and Y. A. Vlasov, "Non-blocking  $4 \times 4$  electro-optic silicon switch for on-chip photonic networks," *Opt. Exp.*, vol. 19, no. 1, pp. 47–54, Dec. 20, 2011.
- [11] A. V. Rylyakov, C. L. Schow, B. G. Lee, W. M. J. Green, S. Assefa, F. E. Doany, M. Yang, J. Van Campenhout, C. V. Jahnes, J. A. Kash, and Y. A. Vlasov, "Silicon photonic switches hybrid-integrated with CMOS drivers," *IEEE J. Solid-State Circuits*, vol. 47, no. 1, pp. 345–354, Jan. 2012.
- [12] B. G. Lee, C. L. Schow, A. V. Rylyakov, J. Van Campenhout, W. M. J. Green, S. Assefa, F. E. Doany, M. Yang, R. A. John, C. V. Jahnes, J. A. Kash, and Y. A. Vlasov, "Demonstration of a digital CMOS driver codesigned and integrated with a broadband silicon photonic switch," *J. Lightw. Technol.*, vol. 29, no. 8, pp. 1136–1142, Apr. 15, 2011.
- [13] S. Assefa, S. Shank, W. Green, M. Khater, E. Kiewra, C. Reinholm, S. Kamalpurkar, A. Rylyakov, C. Schow, F. Horst, H. Pan, T. Topuria, P. Rice, D. M. Gill, J. Rosenberg, T. Barwicz, M. Yang, J. Proesel, J. Hofrichter, B. Offrein, X. Gu, W. Haensch, J. Ellis-Monaghan, and Y. Vlasov, "A 90 nm CMOS integrated nano-photonics technology for 25 Gbps WDM optical communications applications," in *Proc. IEEE Int. Electron Devices Meet.*, Dec. 2012, pp. 33.8.1–33.8.3.
- [14] B. G. Lee, W. M. J. Green, A. V. Rylyakov, S. Assefa, M. Khater, T. Barwicz, C. Reinholm, E. Kiewra, S. M. Shank, C. L. Schow, and Y. A. Vlasov, "Monolithically integrated photonic switches driven by digital CMOS," presented at Proc. Conf. Lasers Electro-Optics, San Jose, CA, USA, Jun. 2013, Paper CTu1L.1.
- [15] B. G. Lee, A. V. Rylyakov, W. M. J. Green, S. Assefa, C. W. Baks, R. Rimolo-Donadio, D. M. Kuchta, M. Khater, T. Barwicz, C. Reinholm, E. Kiewra, S. M. Shank, C. L. Schow, and Y. A. Vlasov, "Four- and eight-port photonic switches monolithically integrated with digital CMOS logic and driver circuits," in *Proc. Opt. Fiber Commun. Conf.*, Mar. 2013, pp. 1–3, Paper PDP5C.3.
- [16] J. Van Campenhout, W. M. Green, S. Assefa, and Y. A. Vlasov, "Low-power,  $2 \times 2$  silicon electro-optic switch with 110-nm bandwidth for broadband reconfigurable optical networks," *Opt. Exp.*, vol. 17, no. 26, pp. 24020–24029, Dec. 16, 2009.
- [17] J. Van Campenhout, W. M. Green, S. Assefa, and Y. A. Vlasov, "Drive-noise-tolerant broadband silicon electro-optic switch," *Opt. Exp.*, vol. 19, no. 12, pp. 11568–11577, Jun. 06, 2011.
- [18] P. Dong, W. Qian, H. Liang, R. Shafiqi, D. Feng, G. Li, J. E. Cunningham, A. V. Krishnamoorthy, and M. Asghari, "Thermally tunable silicon race-track resonators with ultralow tuning power," *Opt. Exp.*, vol. 18, no. 19, pp. 20298–20304, Sep. 13, 2010.

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

In 2009, he became a Postdoctoral Researcher at IBM Thomas J. Watson Research Center, Yorktown Heights, New York, where he is currently a Research Staff Member. He is also an Assistant Adjunct Professor of Electrical Engineering at Columbia University, New York. His research interests include silicon photonic devices, integrated optical switches, and networks for high-performance computing systems and datacenters, and highly parallel multimode transceivers.

Dr. Lee is a member of the Optical Society and the IEEE Photonics Society, where he serves as an Associate Vice President of Membership. He has served on the technical program committees for the International Symposium on Networks-on-Chip and the Photonics in Switching Conference.

**Alexander V. Rylyakov** received the M.S. degree in physics from the Moscow Institute of Physics and Technology in 1989 and the Ph.D. degree in physics from the State University of New York (SUNY), Stony Brook, NY, USA, in 1997. From 1994 to 1999, he worked in the Department of Physics at SUNY Stony Brook on the design and testing of integrated circuits based on Josephson junctions. He joined IBM Thomas J. Watson Research Center, Yorktown Heights, New York, NY, USA, as a Research Staff Member in 1999. His current research interests include digital phase-locked loops and integrated circuits for wireline and optical communication.

**William M. J. Green** (M'03–SM'13) received the B.Sc. degree in engineering physics from the University of Alberta, Edmonton, Alberta, Canada, and the Ph.D. degree in electrical engineering from the California Institute of Technology, Pasadena, CA, USA, in 2005. He is currently a Research Staff Member at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, NY, USA. His research activities include the design of optical devices and integrated systems for terabit-per-second-class silicon nanophotonic interconnects. In addition, his work has extended the silicon photonic integrated circuit platform to the generation and processing of mid-infrared optical signals, for various applications in molecular spectroscopy, environmental sensing, and medical diagnostics. He received the 2012 IBM Corporate Award and the 2012 IEEE Photonics Society Young Investigator Award. He has served the technical organizing committees for numerous OSA and IEEE conferences. He is a member of the OSA.

**Solomon Assefa** (SM'13) received the B.S. degree in physics and EECS in 2001, the M.S. degree in EECS in 2001, and the Ph.D. degree in 2004, all from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA. He is currently a Research Scientist at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, NY, USA. He has worked on nanophotonics technology for optical interconnects, with responsibilities spanning research, development, and technology transfer to commercial foundry. His research interests include high-speed optical detectors, nanostructured platforms for biosensing, and quantum information processing. He has coauthored more than 40 scientific articles, 40 patents, and has appeared as a Guest Speaker at numerous conferences worldwide.

He received several awards including the Technical Accomplishment Award, the Corporate Recognition Award, and several Invention Achievement Awards from IBM. He was named one of the World's Top Young Innovators under 35 and received the TR35 awarded by MIT's Technology Review in 2011. He was honored by the World Economic Forum as a Young Global Leader in 2013. He has served as a Committee Chair and Workshop Organizer for various conferences including the Optical Fibers Conference (OFC) and Conference on Lasers and Electro-Optics (CLEO). He is a member of the Optical Society of America (OSA) and the American Physical Society (APS).

**Christian W. Baks** received the B.S. degree in applied physics from the Fontys College of Technology, Eindhoven, The Netherlands, in 2000, and the M.S. degree in physics from the State University of New York, Albany, NY, USA, in 2001. He joined the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, NY, USA, as an Engineer in 2001. His research interests include high-speed optoelectronic package and backplane interconnect design specializing in signal integrity issues.

**Renato Rimolo-Donadio** (S'08–M'11) received the B.S. and Lic. degrees in electrical engineering from the Technical University of Costa Rica (ITCR), Cartago, Costa Rica, in 1999 and 2004, respectively, and the M.S. degree in microelectronics and microsystems (with distinction), and the Ph.D. degree in electrical engineering (*summa cum laude*), in 2006 and 2010, respectively, both from the Technical University of Hamburg-Harburg (TUHH), Hamburg, Germany.

In 2012, he joined the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, NY, USA, as a Postdoctoral Researcher. From 2006 to 2011, he was with the Institute of Electromagnetic Theory (TET), the TUHH. His current research interests include system-level modeling and optimization of interconnects, and analysis of signal and power integrity problems at IC, package, and PCB levels.

**Daniel M. Kuchta** (SM'97) received the B.S., M.S., and Ph.D. degrees in electrical engineering and computer science from the University of California, Berkeley, CA, USA, in 1986, 1988, and 1992, respectively. He is currently a Research Staff Member in the Department of Communication Technology, IBM Thomas J. Watson Research Center, New York, NY, USA. He subsequently joined IBM at the Thomas J. Watson Research Center. His research interests include high-speed VCSEL characterization, multimode fiber links, and parallel fiber optic link research. He is the author/coauthor of more than 100 technical papers and inventor/coinventor of more than 15 patents.

**Marwan H. Khater** (M'00–SM'06) received the B.S. degree in physics from Middle East Technical University and the M.S. degree in physics from Texas State University, San Marcos, TX, USA, with emphasis on applied solid-state physics, and the Ph.D. degree in electrical engineering from the University of Texas at Dallas, Richardson, TX. His research interests include low temperature plasma sources and characterization for microelectronics processing. He joined IBM Corporation, Yorktown Heights, New York, NY, USA, in 2000, where he has been involved in various advanced research and technology projects. He has been a major contributor to the advancement and development of IBM's silicon-germanium heterojunction bipolar transistor (HBT) and BiCMOS technologies for RF and high-speed communication systems. His research also includes various projects on advanced CMOS devices for 22 nm node technology and beyond. These include FinFET and thin SOI fully depleted devices, as well as Schottky source/drain CMOS devices, with metal gates and high- $\kappa$  dielectrics. Most recently, he has been involved in the research and process integration development of photonic devices (waveguides, modulators, and germanium photodetectors) with CMOS to enable Silicon Nanophotonics technology for high-speed data transfer. He has published more than 53 technical publications and holds 49 patents in the field of microelectronic device design and fabrication. He received the George E. Smith Award by the IEEE Electron Devices Society in 2003 and 2007. He is a member of the American Vacuum Society.

**Tymon Barwicz** received the B.Eng. degree in engineering physics from the Ecole Polytechnique de Montreal, Montreal, QC, Canada, in 2000 and the joint M.Sc./Ph.D. degree in materials science and engineering at the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2005. His doctoral dissertation focused on accurate nanofabrication techniques for strong-confinement microphotonic devices. From 2005 to 2006, he was a Postdoctoral Associate at MIT, where he led nanofabrication of silicon photonic reconfigurable add-drop multiplexers. He joined the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, NY, USA, in 2006. From 2006 to 2010, he worked on shape and dimensional control of sub-10-nm silicon nanowires as well as their applications to FETs, NEMS, and photovoltaics. He currently leads silicon nanophotonic packaging research at IBM with a focus on design, fabrication, and assembly of low-cost nanophotonic waveguide to fiber interfaces and their high-volume integration with microelectronic packaging.



**Carol Reinholm**, biography not available at the time of publication.

**Edward Kiewra** received the B.S. degree in chemistry and chemical engineering from the Florida Institute of Technology, Melbourne, FL, USA, in 1981 and 1982, respectively and the M.S. and Ph.D. degrees from Rensselaer Polytechnic Institute, New York, NY, USA, in 1987 and 1988, respectively, both in chemical engineering.

In 1988, he joined IBM Systems and Technology Division, East Fishkill, New York, where he developed manufacturable integration processes for IBM's advanced DRAM and Logic products. He has also fabricated high-power 980-nm AlGaAs pump diodes at the IBM Zurich Research Lab, Zurich, Switzerland, and III-V channel MOSFET's at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York.

He is currently a Senior Engineering Manager on assignment at IBM Burlington, Essex Junction, VT, USA, where he is supporting integrated silicon photonics, phase change memory, and MRAM technology transfer for foundry manufacturing.

**Steven M. Shank** received the B.S. degree in physics from the State University of New York, Binghamton, NY, USA, in 1986 and the Ph.D. degree in optical engineering from the University of Rochester, Rochester, NY, in 1993. He is currently a Senior Engineer in the Microelectronics Division of the Systems and Technology Group, IBM, Essex Junction, VT, USA, and has been with IBM since 1999. His research interests include various aspects of semiconductor manufacturing including process design, development, and integration, device and yield characterization, and technology transfer. He has made significant contributions to the manufacturability of gate dielectrics, embedded DRAM, SRAM, CMOS, RF-CMOS, ASICs, and silicon photonics. He holds 31 U.S. patents.

**Clint L. Schow** (SM'10) received the Ph.D. degree in electrical engineering from the University of Texas at Austin, Austin, TX, USA, in 1999. He joined IBM, Rochester, MN, USA, assuming responsibility for the optical receivers used in IBM's optical transceiver business. From 2001 to 2004, he was with Agility Communications, Santa Barbara, CA, USA, developing high-speed optoelectronic modulators and tunable laser sources. In 2004, he joined the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, NY, USA, as a Research Staff Member and currently manages the Optical Link and System Design group responsible for optics in future generations of servers and supercomputers. He has directed multiple DARPA-sponsored programs investigating chip-to-chip optical links, nanophotonic switches, and future systems utilizing photonic switching fabrics. He has published more than 125 journals and conference articles, and has ten issued and more than 20 pending patents. He is a senior member of the OSA.

**Yurii A. Vlasov** received the M.S. degree from the University of St. Petersburg, St. Petersburg, Russia, in 1988, and the Ph.D. degree from the Ioffe Institute, St. Petersburg, in 1994, both in physics. He is currently a Manager of a Silicon Nanophotonics Project at the IBM Research, Yorktown Heights, New York, NY, USA. For the last 12 years at IBM, he has been focused on the development of silicon nanophotonics starting from its initial exploratory fundamental research followed by advanced device engineering and, most recently, manufacturing and product development. Currently, he leads the company-wide effort on transitioning the IBM Silicon Nanophotonics technology to commercial manufacturing for products relevant to IBM business. For a couple of years, he also served as an Adjunct Professor at Columbia University's Department of Electrical Engineering.

Prior to IBM, he developed semiconductor nanophotonics at the NEC Research Institute, Princeton and at the Strasbourg IPCMS Institute, France. For over a decade, he was also a Research Scientist with the Ioffe Institute of Physics and Technology, St. Petersburg, Russia, working on optics of nanostructured semiconductors.

Dr. Vlasov has published more than 300 journals and conference papers, filed more than 50 U.S. patents, and delivered more than 100 invited, plenary, and tutorial talks in the area of nanophotonics. He served on numerous organizing committees of conferences on nanophotonics under OSA, IEEE, APS, MRS, etc. He also contributed to the development of IEEE standards on 100G Ethernet optical links. He was elected as a Fellow of both the OSA and the APS. He was awarded several IBM Outstanding Technical Achievement Awards, the "Best of IBM" Award, and the IBM 2011 Corporate Award, as well as was named "Scientist of the Year" by the Scientific American journal.