Multi-Stage 8×8 Silicon Photonic Switch based on Dual-Microring Switching Elements

Article *in* Journal of Lightwave Technology · October 2019

DOI: 10.1109/JLT.2019.2945941

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JLT.2019.2945941, Journal of Lightwave Technology

Multi-Stage 8×8 Silicon Photonic Switch based on Dual-Microring Switching Elements

Yishen Huang, Qixiang Cheng, *Member, IEEE*, Yu-Han Hung, Hang Guan,
Xiang Meng, *Member, IEEE, Member, OSA*, Ari Novack, Matthew Streshinsky,
Michael Hochberg, and Keren Bergman, *Fellow, IEEE, Fellow, OSA*

(*Invited paper*)

***Abstract*—We demonstrate the first multi-stage 8×8 silicon photonic switch with switching elements based on dual add-drop microrings with a compact footprint of 4 mm^2. This device leverages co-design of the switch architecture and the switching elements with a well-balanced set of performance metrics. The switching elements are designed to have a 3-dB optical passband of 165 GHz, exhibiting off- and on-resonance losses of 0.67 dB and 2 dB, respectively. Full characterization of all switch paths shows an end-to-end on-chip loss between 4.4 and 9.6 dB, with worst-case crosstalk leakage averaged at -16 dB. Owing to the efficient doped waveguide thermo-optic phase shifters, the device features a tuning efficiency of 48.85 GHz/mW. The reconfiguration time of the switch fabric is measured to be 1.2 µs and 0.5 µs at the rise and fall edge, respectively. The dual-microring switching element together with the multi-stage architecture preserves an end-to-end passband over 55 GHz. We validate the switch performance with optical paths of varying numbers of on- and off-resonance switching elements – less than 2 dB power penalties are obtained for all data routings at 32 Gbps.**

***Index Terms*— Optical switches, photonic integrated circuits, silicon photonics, microring resonators**

I. Introduction

The exponential increase of datacenter traffic has motivated dynamic optical connectivity in datacenter interconnects. As the capacity scaling of current electronic switch ASICs faces stringent challenges for distributed data applications, optical switching can potentially deliver high-bandwidth and modulation indifferent routing in datacenter networks [1]. To enable emergent datacenter designs such as disaggregated hardware and application-dependent bandwidth allocation, dynamically reconfigurable networks need to be realized with the help of efficient and economic optical switching mechanisms [2]. Free-space optics-based optical switch modules have been commercialized based on micro-electromechanical systems (MEMS)-actuated mirrors and beam-steerers; however, the necessary and meticulous alignment and stabilization systems contribute to the high cost per port of these implementations, which could stymie their wide adoption in data center applications [1]. Integrated photonic switch fabrics in silicon can offer dramatic miniaturization of device footprint and reduction in cost per port via mass production [1,3]. In addition, silicon photonics can utilize its compatibility with existing CMOS platforms to achieve monolithic integration with driving electronics [4,5], which could further drive down device packaging cost and improves device performance.

The highly confined guided mode in silicon waveguides, owing to the large core-cladding index contrast, as well as its temperature- and carrier-dependent index, provide significant flexibility in designing building blocks of a switch fabric. Integrated photonic switches on silicon are typically built by a fabric of 2×2 switching elements (SE). The SEs can be realized using tunable Mach-Zehnder interferometers (MZI) [6,7], MEMS couplers [8,9], microring resonators (MRR) [10–15], and MRR-assisted MZIs [16–18]. By arranging and interconnecting a number of the 2×2 SEs in certain architectures, larger networks can be formed to connect more inputs and outputs, and the network architecture dictates the type of connectivity and the routing control of the switch device. MZI- and MEMS-based SEs rely on couplers and phase shifters with lengths of tens to hundreds of microns [6,8,9]. While they are used to build a few impressive demonstrations of large-scale integrated photonic switches, these high-port count devices typically have a die area of 50 – 150 mm^2 [6–8, 18], which translates directly to increased cost per port. By using MRRs directly for through- and drop-state switching, the SEs can leverage traveling wave cavity dynamics to reduce individual SE footprint by 50-100 times, and thus enabling much higher integration density with greatly shrunk chip area. Table I compares the performance of recent notable demonstrations of switch fabrics incorporating MRRs.

[submission date] This work was supported by the ARPA-E ENLITENED program under Grant DEAR00000843. (*Corresponding author: Qixiang Cheng*)

Y. Huang, Q. Cheng, Y-H. Hung, Xiang Meng, and K. Bergman are with the Department of Electrical Engineering, Columbia University, New York, NY 10027, USA (e-mail: yh2785@columbia.edu; qc2228@columbia.edu; yh3128@columbia.edu; xm2137@columbia.edu; bergman@ee.columbia.edu).

H. Guan, M. Streshinsky, and M. Hochberg are with Elenion Technologies LLC, 171 Madison Ave. STE 1100, New York, NY 10016, USA (e-mail: hang.guan@elenion.com; matthew.streshinsky@elenion.com; michael.hochberg@elenion.com).

A. Novack is with both the Department of Electrical Engineering, Columbia University and Elenion Technologies LLC (e-mail: ari.novack@elenion.com).

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

Digital Object Identifier []

TABLE I
NOTABLE DEMONSTRATIONS OF MRR-BASED SILICON PHOTONIC SWITCHES

Port count	Architecture	Switching element	On-chip loss [dB]	Crosstalk [dB]	Extinction Ratio [dB]	Switching Speed [μs]	Bandwidth [GHz]	Source
4×4	Switch-and-select	1st-order MRR	1.8 – 20.4	-50.7 to -31.6	>50	1.2 to 14.3	24	[11]
8×4	Crossbar	2nd-order MRR	6 – 14	-21	32	2.5 – 15.9	100	[15]
8×7	Crossbar	5th-order MRR	2 – 10	N/A	19.5 – 23.4	4 – 17	>75	[12]
16×16	Beneš	MRR-assisted MZI	16.2	-20.5	30	1.65e-3	40.7	[18]
8×8	Omega	Dual MRR	4.4 – 9.6	-16.75	14.7 – 18.8	0.4 – 1.2	55	This work

Nevertheless, the scalability of MRR-based fabric can be limited on three fronts: insertion loss of resonator SEs in both on- and off-resonance states, successive passband narrowing of cascaded MRRs, and control complexity of the large number of resonators. To the best of our knowledge, the demonstrations of MRR-based optical switches to date achieve a record of 8×8 connectivity [12,13], and both adopt the cross-bar architecture that requires only one SE to be controlled to connect an input to an output. However, the number of SEs in a cross-bar architecture scales poorly as N^2 for an N×N device, which poses a tremendous challenge in on-chip wiring and packaging complexity. In addition, a lightpath in a cross-bar switch can traverse between 1 to $2N-1$ SEs, which means both the worst-case insertion loss and the variations of path-dependent power penalties would grow quickly with increasing port count [12]. In contrast, multi-stage architectures, such as the Beneš and Omega topologies, can provide a balance in the trade-offs between the total number of SEs and the number of cascaded switching stages [1,3]. Designing MRR-based multi-stage switch requires co-optimizing the SE performance with the selection of architecture [14]. Individual MRR SE needs to possess a balanced set of performance metrics that meet the targets in both insertion loss and switching bandwidth, while the selection of architecture allows ease of control and limits the number of cascaded stages to preserve the end-to-end switch passband. Hence, we opt for the Omega architecture in this initial demonstration for a modest-scale multi-stage MRR-based switch fabric. This design trades off the non-blocking connectivity for a much-reduced number of switch stages and simplified routing control.

In this work, we present the design and characterization of the first multi-stage silicon switch with 8×8 connectivity implementing dual add-drop MRRs, leveraging a commercial silicon photonics process and design flow at Elenion Technologies [19]. We highlight the combination of low on-chip loss, wide passband, and high tuning efficiency of the switch device as the key enablers to optically-switched datacenter network designs [1]. We extend our work [20] by discussing the operation of the dual-MRR SEs and presenting comprehensive switching performance and usability analysis. The paper is organized as follows: Section II describes the switch fabric device and architecture; Section III discusses the design and performance of the MRRs and dual-MRR SEs. Section IV reports the end-to-end performance of the full switch fabric. The paper concludes in Section V.

II. SWITCH ARCHITECTURE AND DESIGN

The Omega architecture, as a Banyan-type network originally proposed for high-performance computer networks, is also attractive for high-speed electronic and optical switching applications. A Banyan-type network is defined as a class of multistage networks that have exactly one path from any input port to any output port. Generally, a Banyan-type switching fabric with N ports is constructed from $\frac{N}{d}(log_d N)$ d×d switching elements arranged in $log_d N$ stages, which is also referred to as d-nary switch [21]. In the optical domain, more attention has been focused on binary switching fabrics (d=2). In particular, the Omega architecture is defined by its perfect shuffle connection of SEs between adjacent stages, which interleaves each half of the previous stage's output ports.

To achieve connectivity between 8 inputs and 8 outputs, the switch arranges 12 SEs into an Omega network illustrated in Fig. 1A. Each of the 2×2 SEs can be independently controlled as Bar state or Cross state, as shown in Fig. 1B. The dual-MRR configuration differs from a single-MRR SE [10,11] by operating two parallel-coupled resonators and achieving a broadened passband as discussed in Section III. For a binary N×N Omega network, with N being a power of 2, the total number of 2×2 SEs is $\frac{N}{2}log_2 N$, and the number of cascaded switching stages is $log_2 N$. Fig. 2 compares the worst-case insertion loss among three representative optical switch architectures – cross-bar, Omega, and Beneš [22]. It is evident that the increment of insertion loss in multi-stage architectures is much slower comparing to the cross-bar architecture as port counts increase. This is due to the number of bypass MRRs, i.e. off-resonance rings, increases linearly with the switch port count in Crossbar network and thus the accumulated through-MRR loss dominates over the drop-MRR loss. Omega network has a lower increment in loss than Beneš because it contains about half of the total stage counts. Reduced number of cascaded stages is also critical to preserving the lightpath passband [14].

Since Omega networks have the property of exactly one connecting path from any input to any output, they can take advantage of self-routing. Unlike Beneš, whose routing configurations need to be iteratively computed [23] or pre-determined as a look-up table (LUT) [24,25], the Omega network can be controlled solely with destination-based routing, by configuring each SE along the lightpath directly from the output label. Fig. 3 illustrates an example of such routing procedure. The reduced routing control complexity of

Omega network eschews the need for complicated routing logic and the associated control and computation overhead.

Figure. 4A shows the micrograph of the switch device. The entire switch chip, consisting of 24 thermo-optic (TO) MRRs, 8 co-integrated monitor photodiodes (PD), and 46 electrical bonding pads, has a footprint of 4 mm^2. Shown in Fig. 4B, the chip is die- and wire-bonded to a chip carrier placed on a custom printed circuit board (PCB) to allow electrical access of the thermos-optic phase shifters and PDs. An array of 18 fibers at 127-micron pitch are grating coupled to the chip to provide optical access to the switch inputs and outputs.

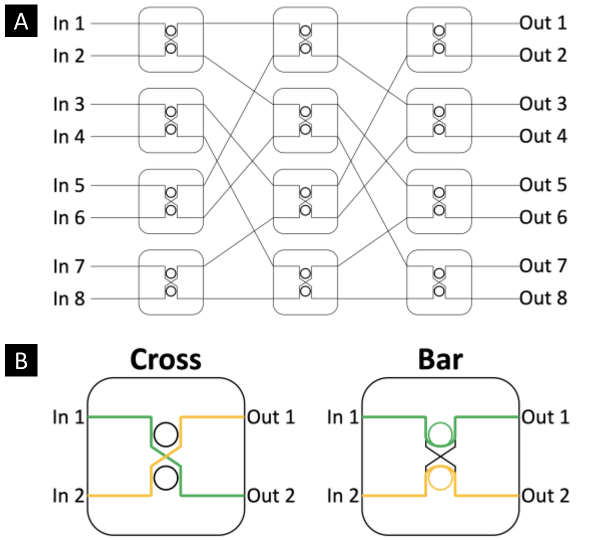

Fig. 1. (A) Schematic of the 8×8 Omega network implemented with 12 2×2 SEs. (B) Illustration of the SE's Cross and Bar states and their corresponding dual-MRR configurations.

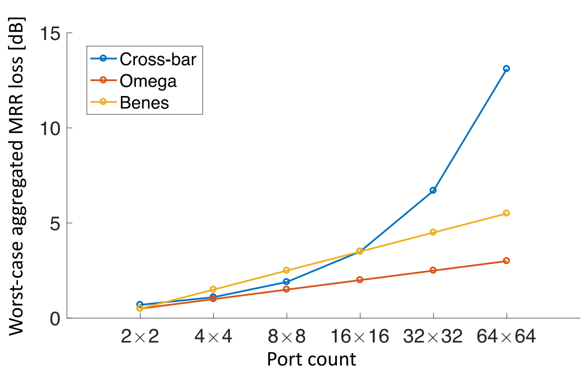

Fig. 2. Comparison of the worst path insertion loss on-chip for Cross-bar, Omega, and Beneš architectures with MRR SEs, based on the record-low on- and off-resonance MRR loss values from [11].

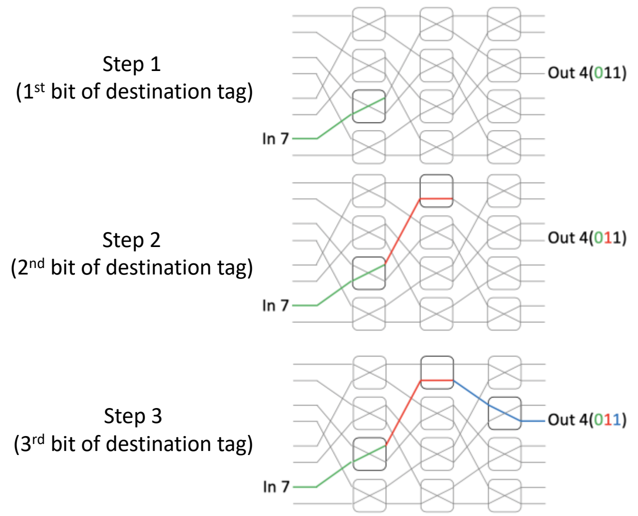

Fig. 3. Example of destination-based routing of Omega network. The output's binary label indicates the configuration states of the three SEs along the lightpath – in the sequence of the SEs traversed by the optical signal, the bit 0 selects the upper output of the corresponding SE, and the bit 1 selects the lower output. Hence, a label 011 indicates setting the first SE to Cross, second SE to Bar, and third SE to Cross, to connect a link from Input 7 to Output 4.

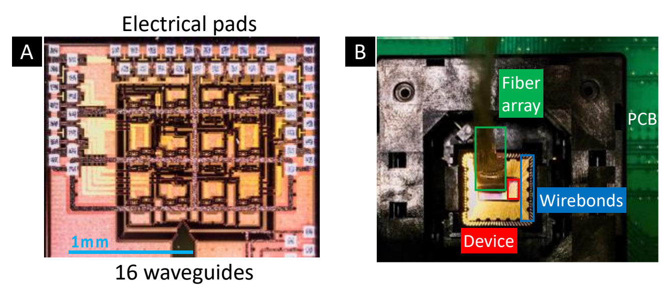

Fig. 4. (A) Micrograph of the switch device showing a device footprint of 4 mm^2 including electrical pads. (B) Photo of the switch device package.

III. Switching Element Design and Characterization

Each of the 12 SEs integrated in the switch device contains 2 racetrack MRRs coupled to 2 parallel waveguides on either side. Fig. 5A illustrates the arrangement of MRRs and the waveguide crossing in an SE. After each SE, 1% power is tapped on both SE output waveguides into an on-chip photodiode (PD), which can be used to infer if optical power arrives at the corresponding SE. We characterize the through-MRR transmission spectra for the 8 MRRs in the last stage of the switch and demonstrate highly consistent resonance profile as shown in Fig. 5B. Designed for operating with 120 GHz of passband, each single MRR shows an extinction ratio of about 9.5 dB. By operating in dual-MRR switching mode, as discussed later in this section, the extinction is extended to about 14.7 dB for Bar state and 18.8 dB in Cross state. To study the free spectral range (FSR) of an MRR, we measure the transmission spectrum of the lightpath connecting Input 1 and Output 8, which traverses three SEs containing six MRRs. A bias voltage of 2.6 V is applied to the phase shifter of the last MRR while leaving the other five MRRs unbiased. It is evident from Fig. 5C that the single biased MRR shows an FSR of about 1.831 THz (shallow troughs), and unbiased MRR shows good alignment of resonances (deep troughs).

The resonance of each MRR can be adjusted by applying a DC voltage across the N-doped portions of the resonator, which behave as a resistive heater and thermo-optically change the roundtrip phase of the MRR. We show the resonance shift as a function of bias voltage between 2.6 V – 2.82 V for a single ring in Fig. 6A. The thermo-optic tuning efficiency can thus be extracted, as shown in Fig. 6B, to be about 0.39 nm/mW or 48.85 GHz/mW, which corresponds to a P_π of about 18.7 mW. We also characterize the switch reconfiguration times by measuring the optical time-domain response of the thermo-optic switch. With a 150 KHz electrical square-wave signal at 50% duty cycle applied to the path from Input 7 to Output 1, we observe a switching rise time of 1.2 μs and fall time of 0.5 μs (0% to 100 %), as shown in Fig. 7. For doped waveguide heaters, the heating process is typically longer than the cooling process. The longer rise time is due to the slower temperature increase of the phase-shifter relative to decreasing its temperature via dissipation.

The dual-MRR configuration of the SE resembles two MRRs simultaneously coupled to two parallel waveguides, as illustrated in Fig. 8. The effective filter behavior is therefore a combined effect between both MRRs (Fig. 8A) and the larger cavity formed by the MRRs and the waveguides (Fig. 8B). When both MRRs in an SE are aligned in resonance, the dual-MRR switching widens the passband of the transmission. Fig. 9 compares the transmission spectra of single- and dual-MRR switching mechanisms, as well as the crosstalk spectrum from the large cavity resonance when both MRRs are off-resonance. The single-ring switching shows a passband of 120 GHz. By aligning the resonances of both MRRs with the resonance of the large cavity, the SE provides a boost in switching bandwidth to 165 GHz, as well as a modest improvement of 0.5 dB in peak switching power. The close agreement between bandwidths of the dual-MRR and the large cavity cases means that the cavity mode in dual-MRR switching largely exists in the large cavity. The large cavity, however, can still allow light to circulate even when both MRRs are at off-resonance, and therefore sets an extinction floor of about 15 dB. Future iterations of the device will address this issue by careful design of the waveguide section which detunes the peak of the larger cavity passband when the MRRs are off-resonance, or by inserting a phase shifter in the waveguide section for additional tunability between switching states. In the context of current 25 – 50 GHz per-channel baud rates for datacenter applications, the wide MRR passband can potentially eliminate the need for stabilization due to sufficient margin for thermal drifts in comparison to the signal bandwidth. In the following analysis, we evaluate the switch device's performance under dual-MRR switching mode because of both its extended passband and lower loss metrics.

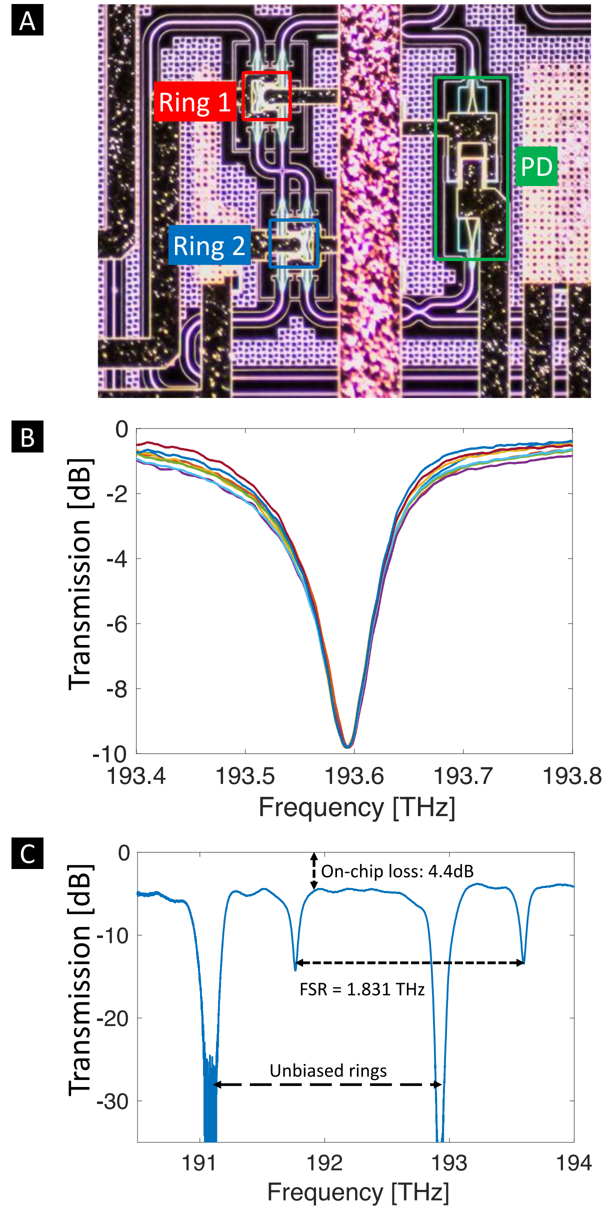

Fig. 5. (A) Micrograph of a single SE showing configuration of both MRRs, waveguide crossing, and on-chip PD. (B) Bus waveguide transmission of 8 MRRs on the last switching stage, showing consistent filter profiles across MRRs. (C) Lightpath transmission spectrum through six MRRs, with one MRR biased at 2.6 V and the rest unbiased, showing an MRR FSR of 1.831 THz.

To characterize individual SE performance, we examine 16 paths of the switch which differ pairwise by the state of one SE each. The paths studied connect Input 1 – Outputs 7/8, Input 2 – Outputs 3/4, Input 3 – Outputs 7/8, Input 4 – Outputs 3/4, Input 5 – Outputs 5/6, Input 6 – Outputs 1/2, Input 7 – Outputs 5/6, Input 8 – Outputs 1/2. For the same input, the SE in the last stage is toggled between Bar and Cross, allowing the routed signal and leakage power levels to be measured for both states, and their extinction ratio and crosstalk levels to be determined. Given a single input into a 2x2 SE in a particular switching state, the signal is the power level measured at the designated output, and the leakage is the power level measured at the undesignated output; their difference is the crosstalk level of that particular state. We define a state's extinction ratio as the difference between its signal power and the other state's

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JLT.2019.2945941, Journal of Lightwave Technology

leakage arriving at the same output. Fig. 10 shows the signal and leakage power levels for the 16 paths; an average extinction ratio of 18.8 dB for the Cross state and 14.7 dB for the Bar state are observed. Crosstalk levels are similar between Cross and Bar states, averaging to -16.75 dB.

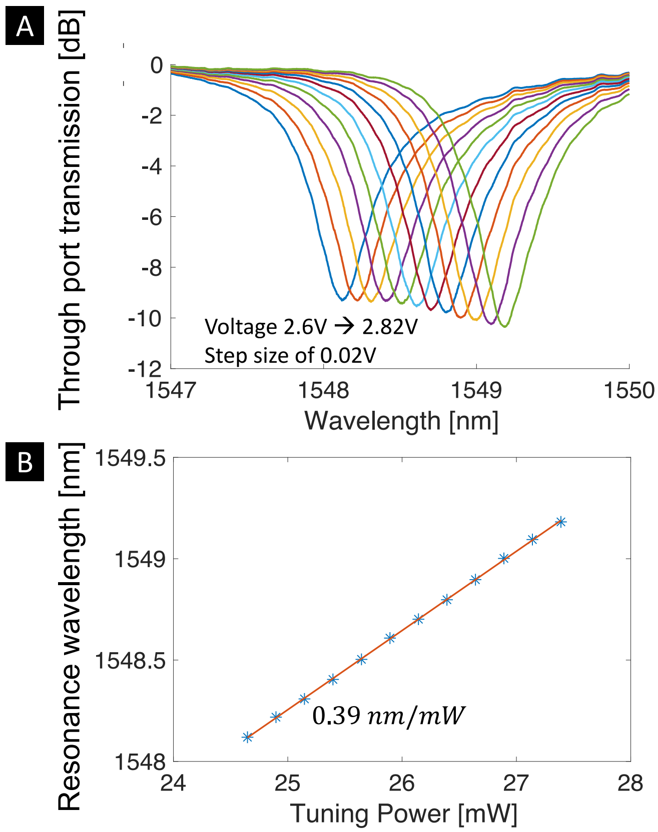

Fig. 6. (A) Change in resonance wavelength of the MRR as heater bias is swept from 2.6 V to 2.82 V in 0.02 V step size. (B) Extracted resonance tuning power efficiency showing a linear trend around the wavelength range shown in Fig. 6A.

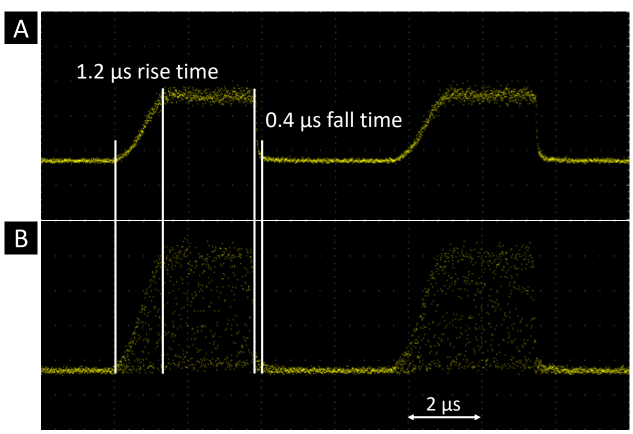

Fig. 7. (A) Rise and fall time of a continuous wave (CW) optical signal at 1548.5 nm through a single MRR. (B) Rise and fall time of an optical signal at 1548.5 nm modulated at 32 Gbps. Both are measured from 0% to 100%.

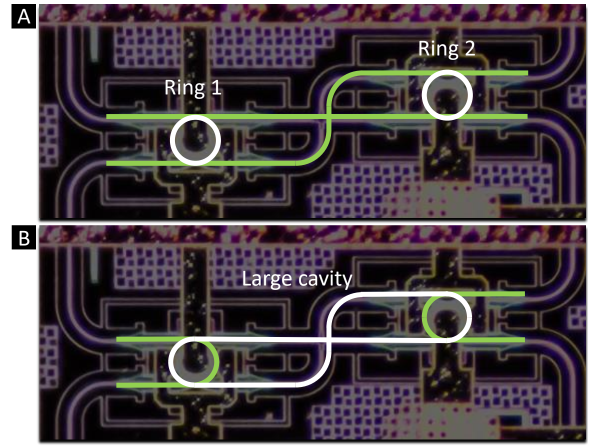

Fig. 8. Schematic of the parallel coupled resonators showing the two small cavities formed by the two rings (A), as well as the large cavity formed by halves of the two rings and the waveguides connecting them (B).

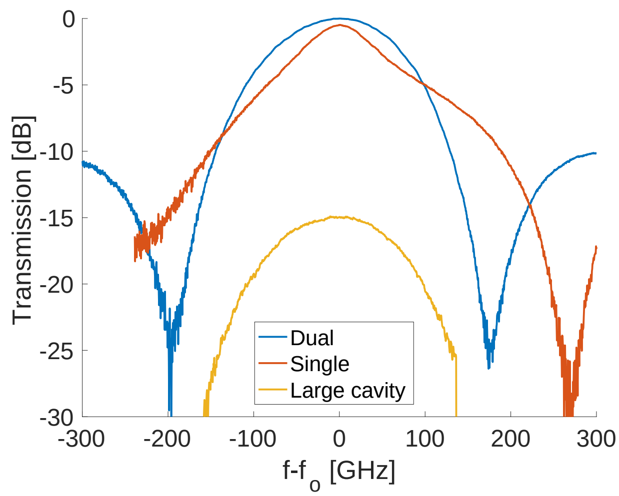

Fig. 9. Comparison of drop passbands and peak transmission for a dual-MRR SE under single- and dual-MRR switching. The large cavity transmission is taken with both MRRs far off-resonance. Transmission is normalized based on the dual-MRR case.

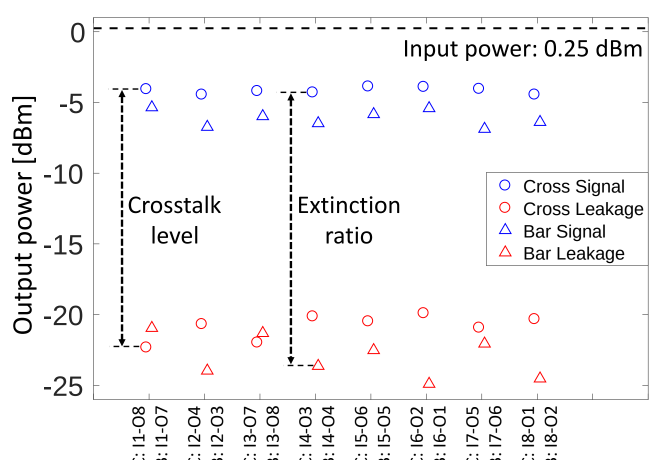

Fig. 10. Signal and leakage power levels for paths connecting 16 pairs of input-output as indicated, showing crosstalk and extinction ratio along each path. Each pair of the paths differ only in the state of one SEs, allowing per-SE extinction ratio and crosstalk levels to be extracted as indicated.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JLT.2019.2945941, Journal of Lightwave Technology

IV. Switch Fabric Performance

We first examine the switch's performance in two representative cases of operations – SEs configured in all-Cross and all-Bar states at 1548.5 nm, the peak transmission wavelength of the on-chip grating couplers. The switch fabric defaults to all-Cross with minimal loss along each lightpath with all MRRs unbiased and at off-resonance. In the all-Bar case, all MRRs are biased to resonate with the input wavelength, inducing the maximal attenuation on each lightpath due to loss associated with traversing in and out of the MRRs and attenuation in the doped waveguides. Shown in Fig. 11, the average on-chip loss among all paths is 4.4 dB in all-Cross and 8.4 dB in all-Bar. The off-resonance and on-resonance losses are estimated at 0.67 dB and 2 dB per SE. A breakdown of the component loss contributions of the device is shown in Table II. The performance of all switch paths from 8 inputs to 8 outputs is summarized in Fig. 12, in which we show a mean path insertion loss of 6.7 dB. The worst-case crosstalk for each input-output connection, which is taken at the non-signal port with the highest leakage power, is -16 dB on average. While a majority of the worst-case crosstalk levels are within -13 dB and -23 dB, we attribute a few cases (Input 6 to Outputs 3 and 7, and Input 8 to Output 7), where high worst-case crosstalk levels are observed, to fabrication variations of the MRR elements. We further characterize the change in passband as the number of on-resonance SEs increases along a lightpath. In this case, a broadband signal is injected through Input 7 of the switch, and routed to Outputs 5, 1, 3, 4 via 0 – 3 on-resonance SEs respectively, as illustrated in Figs. 13A-13D. We show the spectra of the signal and three 1st order crosstalks resulted from each of the 3 SEs traversed by the lightpath in each routing. The 2nd order crosstalks are suppressed below -35 dB and therefore omitted for clarity. It is evident from Figs. 13E-13H that, while cascaded MRR SEs increasingly narrows the switched passband, a lightpath traversing through 1 – 3 on-resonance SEs still maintains 147 GHz, 96 GHz, and 55 GHz of bandwidth, respectively. The power consumption per MRR is on average 0 mW and 25.6 mW for off- and on-resonance states, respectively.

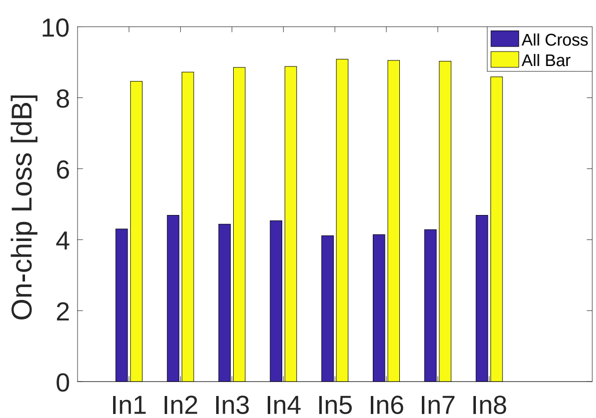

Fig. 11. On-chip loss of each input signal in all-Cross and all-Bar configurations.

TABLE II
Key Component loss

Item	Loss
Waveguide propagation loss	2 dB/cm
SE in Cross	0.67 dB
SE in Bar	2 dB
Grating coupler	3.6 dB/facet

Data transmission was performed using an Anritsu MP1900A Signal Quality Analyzer and a Thorlabs MX35E Reference Transmitter at 32 Gbps non-return-to-zero (NRZ) on-off keying (OOK) using 2^{31}-1 pseudo-random bit sequence (PRBS31). On-chip switch paths with varying number of on-resonance SEs all exhibit error free operations. The schematic of the data test is shown in Fig. 14A, and the same four paths described in Fig. 13 are tested. The optical carrier is launched from a tunable laser diode (TLD) at 1548.5 nm and modulated by a Mach-Zehnder modulator (MZM) with 0 dBm output power. The intensity modulated optical signal is then guided to the silicon photonic MRR-based switch via a polarization controller (PC). A power adjuster (PA) consisting of a variable attenuator (VOA) and an Erbium-doped fiber amplifier (EDFA) is used before the switch to compensate for coupling and propagation losses through the chip, ensuring -10 dBm of optical power exits the chip. This PA is also used to replicate the device insertion loss in the back-to-back (B2B) reference case. A second set of EDFA and VOA are used to adjust the receiver optical power for the BER measurement. An optical filter (OF) with 180 GHz passband is used to reject out-of-band amplified spontaneous emission (ASE) noise. The receiver consists of a Finisar XPRV2022A PD-transimpedance-amplifier (TIA) assembly, which performs the optical-to-electrical conversion and allows the data signal to be analyzed by the Anritsu error checker. All lightpaths examined are within 2 dB power penalty comparing to the B2B reference at 10^{-9} BER, shown in Fig. 14B. Despite going through 0 – 3 on-resonance SEs with successively narrowed optical passband, all four switch paths' show clear eye-openings shown in Fig. 14C and are with in 1 dB penalty variations at 32 Gbps, which indicates negligible path dependence of power penalties due to consecutive filtering of the MRR switch device.

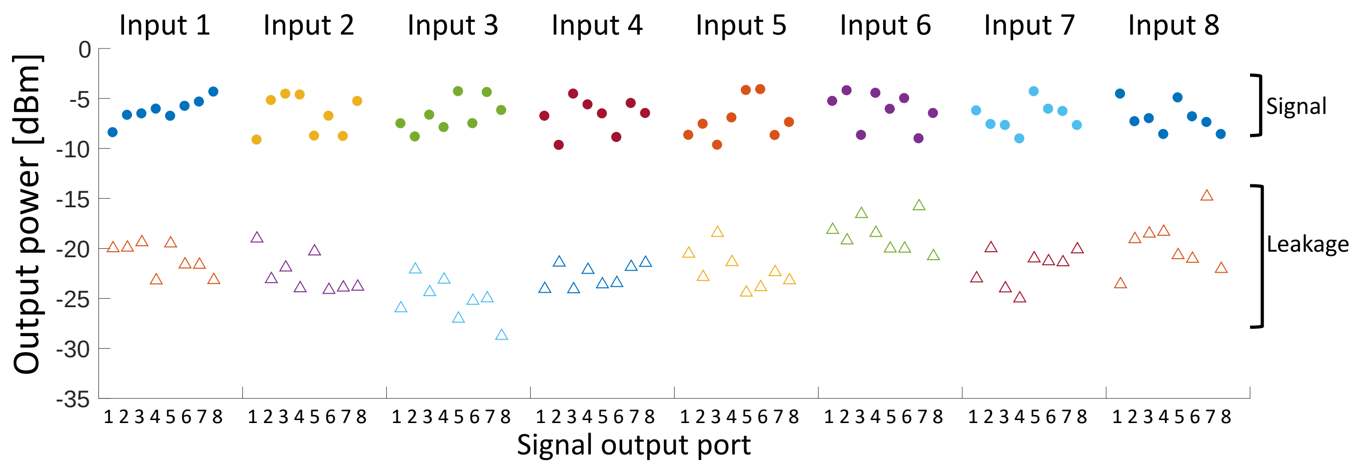

Fig. 12. On-chip signal power and worst-case leakage power levels for all switch paths connecting every input to every output. Grating coupler loss is compensated at injection. Worst-case crosstalk is taken as the difference between signal and leakage power for each connection.

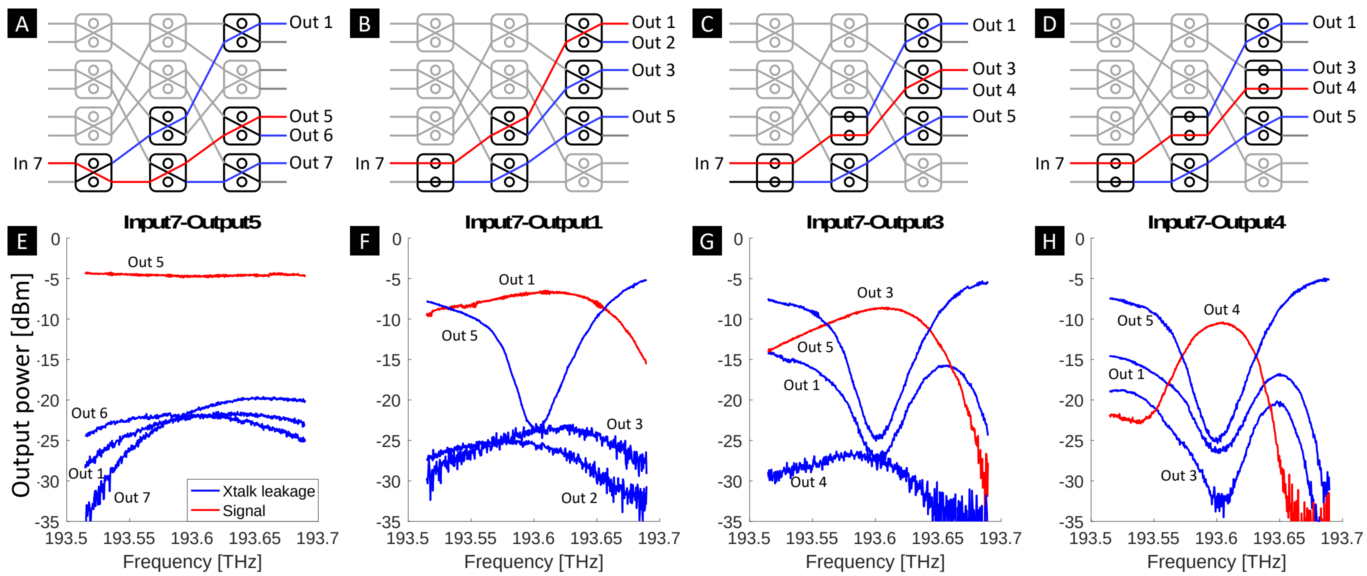

Fig. 13. (A-D) Illustration of the switch paths through 0 – 3 on-resonance SEs, respectively. (E-H) Signal and 1st order crosstalk spectra for the routing schemes shown in A-D.

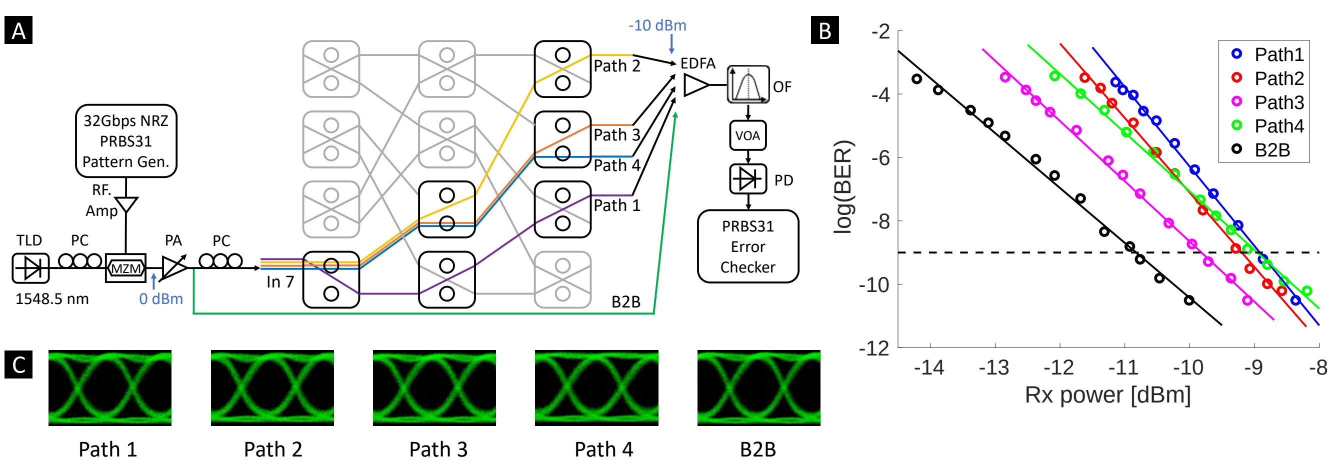

Fig. 14. (A) Data transmission setup showing the switch paths and the B2B link examined. (B) BER-Rx power relationship for the switch paths and B2B reference path at 32G NRZ-OOK PRBS31. (B) Open eye diagrams of the switch paths and B2B reference path at 32G NRZ-OOK PRBS31.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JLT.2019.2945941, Journal of Lightwave Technology

V. Conclusion

We present the first multi-stage 8×8 silicon photonic switch implementing dual add-drop MRR SEs with a compact footprint of 4 mm^2. This device is taped out using a commercial silicon photonics design flow at Elenion Technologies. The switch demonstrates a well-balanced set of performance metrics, showing off- and on-resonance SEs losses to be 0.67 dB and 2 dB, respectively, and an end-to-end on-chip loss ranging between 4.4 dB and 9.6 dB. The worst-case first-order switching crosstalk levels have an average of -16 dB, mostly ranging between -13 dB and -23 dB. Component characterizations of the switch show a 120 GHz passband per MRR and a 165 GHz passband for dual-MRR switching. A minimum passband of 55 GHz is observed after three-stage SE filtering. The switching speed of the thermally driven SEs is measured as 1.2 µs rise time and 0.5 µs fall time, with a thermal tuning power efficiency of 48.85 GHz/mW. We perform data transmission tests with 32 Gbps NRZ-OOK, showing less than 2 dB power penalty incurred by the switch routing. The collective of appealing characteristics makes the device suitable for agile functionalities such as bandwidth steering and network reconfiguration for 200G and 400G datacenter applications and pave way for future designs of optically-switched datacenter networks.

References

[1] Q. Cheng *et al.*, “Photonic switching in high performance datacenters [Invited],” *Opt. Express*, vol. 26, no. 12, pp. 16022-16043, 2018.
[2] Y. Shen *et al.*, “Silicon photonics for extreme scale systems,” *J. Light. Technol.*, vol. 37, no. 2, pp. 245-259, 2019.
[3] Q. Cheng *et al.*, “Recent advances in optical technologies for data centers: a review,” *Optica*, vol. 5, no. 11, pp. 1354-1370, 2018.
[4] B. G. Lee *et al.*, “Monolithic silicon integration of scaled photonic switch fabrics, CMOS logic, and device driver circuits,” *J. Light. Technol.*, vol. 32, no. 4, pp. 743-751, 2014.
[5] C. Sun *et al.*, “Single-chip microprocessor that communicates directly using light,” *Nature*, vol. 528, pp. 534–538, 2015.
[6] L. Qiao *et al.*, “32 × 32 silicon electro-optic switch with built-in monitors and balanced-status units,” *Scientific Reports*, vol. 7, 2017.
[7] D. Celo *et al.*, “32×32 silicon photonic switch,” in *Proc. OECC*, Niigata, Japan, 2016.
[8] S. Han *et al.*, “Large-scale silicon photonic switches with movable directional couplers,” *Optica*, vol. 2, no. 4, pp. 370-375, 2015.
[9] T. J. Seok *et al.*, “240×240 wafer-scale silicon photonic switches,” in *Proc. OFC*, paper Th1E.5, San Diego, USA, 2019.
[10] N. Sherwood-Droz, *et al.*, “Optical 4x4 hitless silicon router for optical networks-on-chip (NoC),” *Opt. Express*, vol. 16, no. 20, pp. 15915–15922, 2008.
[11] Q. Cheng *et al.*, “Ultralow-crosstalk, strictly non-blocking microring-based optical switch,” *Photon. Res.*, vol. 7, no. 2, pp. 155–161, 2019.
[12] P. DasMahapatra, *et al.*, “Optical crosspoint matrix using broadband resonant switches,” *IEEE J. Sel. Top. Quantum Electron.*, vol. 20, no. 4, 2014.
[13] A. S. P. Khope *et al.*, “Elastic WDM optoelectronic crossbar switch with on-chip wavelength control,” in *Proc. Adv. Photon.*, Paper PTh1D.3, 2017.
[14] Q. Cheng, *et al.*, “Scalable microring-based silicon Clos switch fabric with switch-and-select stages,” *IEEE J. Sel. Top. Quantum Electron.*, vol. 25, no. 5, 2019.
[15] A. S. P. Khope *et al.*, “Multi-wavelength selective crossbar switch,” *Opt. Express*, vol. 27, no. 4, pp. 5203-5216, 2019.
[16] L. Lu, *et al.*, “Low-power 2 × 2 silicon electro-optic switches based on double-ring assisted Mach–Zehnder interferometers,” *Opt. Lett.*, vol. 39, no. 6, 2014.
[17] L. Lu, *et al.*, “4 × 4 silicon optical switches based on double-ring-assisted Mach–Zehnder interferometers,” *IEEE Photon. Technol. Lett.*, vol. 27, no. 23, 2015.
[18] Z. Guo, *et al.*, “16 × 16 Silicon Optical Switch Based on Dual-Ring-Assisted Mach–Zehnder Interferometers,” *J. Light. Technol.*, vol. 36, no. 2, 2018.
[19] M. Bohn *et al.*, “Next-generation silicon photonic interconnect solutions,” in *Proc. OFC*, paper M3J.3, San Diego, USA, 2019.
[20] Y. Huang *et al.*, “Dual-microring resonator based 8×8 silicon photonic switch,” in *Proc. OFC*, paper W1E.6, San Diego, USA, 2019.
[21] W. Kabacinski, *Nonblocking Electronic and Photonic Switching Fabrics.* Springer, 2005.
[22] V. E. Beneš, *Mathematical Theory of Connecting Networks and Telephone Traffic.* Academic Press, 1965.
[23] D. C. Opferman and N. T. Tsao-Wu, “On a class of rearrangeable switching networks, Part I: control algorithm,” *Bell Syst. Tech. J.*, vol. 50, no. 5, pp. 1579-1600, 1971.
[24] Q. Cheng *et al.*, “Fabric-wide, penalty-optimized path routing algorithms for integrated optical switches,” in *Proc. OFC*, paper Th3A.4, San Diego, USA, 2019.
[25] Q. Cheng *et al.*, “Advanced routing strategy with highly-efficient fabric-wide characterization for optical integrated switches,” in *Proc. ICTON*, Bucharest, Romania, 2018.

Yishen Huang received the B.Sc. degree in Engineering Physics from Queen’s University, Kingston, ON, Canada, and the M.S. degree in Electrical Engineering from Columbia University, New York, NY, USA, where he is currently working towards a Ph.D. degree in the Lightwave Research Laboratory. His research interests include large-scale silicon photonic switch fabrics and high-performance silicon photonic modulators, with specific focus on novel structures and control schemes to enable energy efficient optical link designs.

Qixiang Cheng (M’18) received the B.S degree from the Huazhong University of Science and Technology, Wuhan, China, in 2010, and the Ph.D. degree from the University of Cambridge, Cambridge, U.K., in 2014. In March 2015, he joined Shannon Laboratory, Huawei, China, as a research engineer studying future optical computing systems. He is currently a Research Scientist with Lightwave Research Laboratory, Columbia University, New York, NY, USA. His research interests include design, simulation, and characterization of large-scale optical integrated devices for data center and optical computing applications.

Yu-Han Hung received the M.S. and Ph.D. degrees from National Cheng Kung University, Tainan, Taiwan, in 2012 and 2016, respectively, both in nonlinear dynamics of semiconductor lasers and their applications. In June 2018, he joined Lightwave Research Laboratory as a Postdoctoral Researcher. He is currently a Postdoctoral Research Scientist in Columbia University, New York, NY, USA. His research interests include design/testing of photonic circuit integrations for data centers and test-bed design for high-performance computing architecture.

Hang Guan received the M.S. and Ph.D. degrees from Columbia University, New York, NY, USA in 2017 and 2018, respectively, all in Electrical Engineering. He is currently with Elenion Technologies, LLC. His research interests include high-speed silicon photonic transceiver design/testing, and silicon photonic devices design/testing.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/JLT.2019.2945941, Journal of Lightwave Technology

Xiang Meng (M'19) received the B.Sc. from the University of Saskatchewan, Saskatoon, SK, Canada, in 2011, and the Ph.D. degree in electrical engineering from Columbia University, New York, NY, USA, in 2017. His research interests include scientific parallel computing, design and analysis on emerging nanophotonic devices, ranging from nano-lasers, nano-sensors, to high-speed optical transceivers and energy efficient photonic interconnects.

Ari Novack received the B.Sc. degree in Engineering from Swarthmore College, Swarthmore, PA, USA and the M.S. degree in Electrical Engineering from Columbia University, New York, NY, USA. He is currently working towards his Ph.D. at Columbia University and employed at Elenion Technologies. His research interest is in advanced silicon photonic systems.

Matthew Streshinsky biography not available at the time of submission.

Michael Hochberg biography not available at the time of submission.

Keren Bergman (S'87–M'93–SM'07–F'09) received the B.S. degree from Bucknell University, Lewisburg, PA, USA, in 1988, and the M.S. and Ph.D. degrees from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 1991 and 1994, respectively, all in electrical engineering. Bergman is the Charles Batchelor Professor of Electrical Engineering at Columbia University, New York, NY, USA, where she also directs the Lightwave Research Laboratory. She leads multiple research programs on optical interconnection networks for advanced computing systems, data centers, optical packet switched routers, and chip multiprocessor nanophotonic networks-on-chip. Bergman is a Fellow of OSA and IEEE.