

Multi-FSR Silicon Photonic Flex-LIONS Module for Bandwidth-Reconfigurable All-to-All Optical Interconnects

Original

Multi-FSR Silicon Photonic Flex-LIONS Module for Bandwidth-Reconfigurable All-to-All Optical Interconnects / Xiao, X., Proietti, R., Liu, G., Lu, H., Zhang, Y., Yoo, S.J.B.. - In: JOURNAL OF LIGHTWAVE TECHNOLOGY. - ISSN 0733-8724. - STAMPA. - 38:12(2020), pp. 3200-3208. [10.1109/JLT.2020.2982902]

Availability:

This version is available at: 11583/2972253 since: 2022-10-12T12:04:29Z

Publisher:

IEEE / Institute of Electrical and Electronics Engineers

Published

DOI:10.1109/JLT.2020.2982902

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2020 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Multi-FSR Silicon Photonic Flex-LIONS Module for Bandwidth-Reconfigurable All-to-All Optical Interconnects

Xian Xiao, Roberto Proietti, Gengchen Liu, Hongbo Lu, Yu Zhang, and S. J. Ben Yoo, *Fellow, IEEE*

(*Top-Scored Paper*)

Abstract—This paper proposes and experimentally demonstrates the first bandwidth-reconfigurable all-to-all optical interconnects using a multi-Free-Spectral-Ranges (FSR) integrated 8×8 SiPh Flex-LIONS module. The multi-FSR operation utilizes the first FSR (FSR₁) to steer the bandwidth between selected node pairs and the zeroth FSR (FSR₀) to guarantee a minimum diameter all-to-all topology among the interconnected nodes after reconfiguration. Successful Flex-LIONS design, fabrication, packaging, and system testing demonstrate error-free all-to-all interconnects for both FSR₀ and FSR₁ with a 5.3-dB power penalty induced by AWGR intra-band crosstalk under the worst-case polarization scenario. After reconfiguration in FSR₁, the bandwidth between the selected pair of nodes is increased from 50 Gb/s to 125 Gb/s while maintaining a 25 Gb/s/λ all-to-all interconnectivity in FSR₀.

Index Terms—Arrayed waveguide grating router, optical interconnections, optical switches, photonic integrated circuits, silicon photonics.

I. INTRODUCTION

Today's high performance computing (HPC) and datacenter systems are increasingly applying heterogeneous processor and memory nodes for better utilization of various resources (Fig. 1(a)) [1], [2]. The communication patterns in such systems tend to be spatially and temporally non-uniform, which means that the hotspots and coldspots simultaneously created at different locations of the network could cause heavy congestion in some data links within the datacenter [3], [4]. However, today's interconnection networks based on electronic switches have a fixed interconnection topology, which is incapable of dynamically adapting the bandwidth between certain node pairs to the workloads. On the other hand, the capability of all-to-all interconnects is necessary for many applications such as deep neural network (DNN) [5], [6], map-reduce [7], and parallel sorting applications [8]. Then it would be desirable to design an all-to-all interconnection network with bandwidth steering so

that the network topology can be dynamically reconfigured to match with communication patterns [9].

In the past few years, several integrated bandwidth-reconfigurable switching fabrics have been demonstrated by using wavelength-and-space selective optical switching [10]–[16]. Among these works, we proposed and experimentally demonstrated SiPh Flex-LIONS (silicon photonic flexible low-latency interconnect optical network switch), enabled by arrayed waveguide grating router (AWGR)-based all-to-all wavelength routing, microring resonator (MRR) add-drop filters, and multi-wavelength spatial switches [14]–[16]. The Flex-LIONS architecture has the lowest number of switching elements and insertion loss, enabling better scalability and energy efficiency when compared with other solutions. However, one limitation of all the state-of-the-art bandwidth-reconfigurable switching fabrics, including the Flex-LIONS works in [14]–[16], is that the reconfigured bandwidth is 'borrowed' from the other optical links, negatively affecting the connectivity between the other nodes in the network. This could lead to higher latency for the traffic between node pairs that are not part of the hotspot due to the additional number of hops required to reach the destination nodes.

Here, we propose to leverage multiple free spectral ranges (FSRs) in a Flex-LIONS architecture to address the above-mentioned issues. The multi-FSR operation of AWGR is firstly proposed and demonstrated in [17]. Due to the cyclic nature of AWGRs, with the same device, the connectivity between each pair of nodes can be easily increased by exploiting multiple FSRs. Some FSRs of the core AWGR (e.g. FSR₀) guarantees a minimum diameter all-to-all topology among the N interconnected nodes before and after reconfiguration as shown in Fig. 1(a), while the other FSRs (e.g. FSR₁) can be freely used to boost the bandwidth between specific node pairs as shown in Fig. 1(b). In this case, bandwidth-reconfigurable all-to-all interconnects can be achieved as shown in Fig. 1(c).

In this paper, we extend the work presented in [15] and report the first experimental demonstration of bandwidth

Manuscript received November 29, 2019. This work was supported in part by DoD contract H98230-16-C-0820, NSF grant 1611560, and DOE award DE-SC0019526.

X. Xiao, R. Proietti, G. Liu, H. Lu, Y. Zhang, and S. J. B. Yoo are with the Department of Electrical and Computer Engineering, University of California,

Davis, CA 95616 USA (e-mail: xxxiao@ucdavis.edu; rproietti@ucdavis.edu; genliu@ucdavis.edu; hlu@ucdavis.edu; dyuzhang@ucdavis.edu; sbyoo@ucdavis.edu).

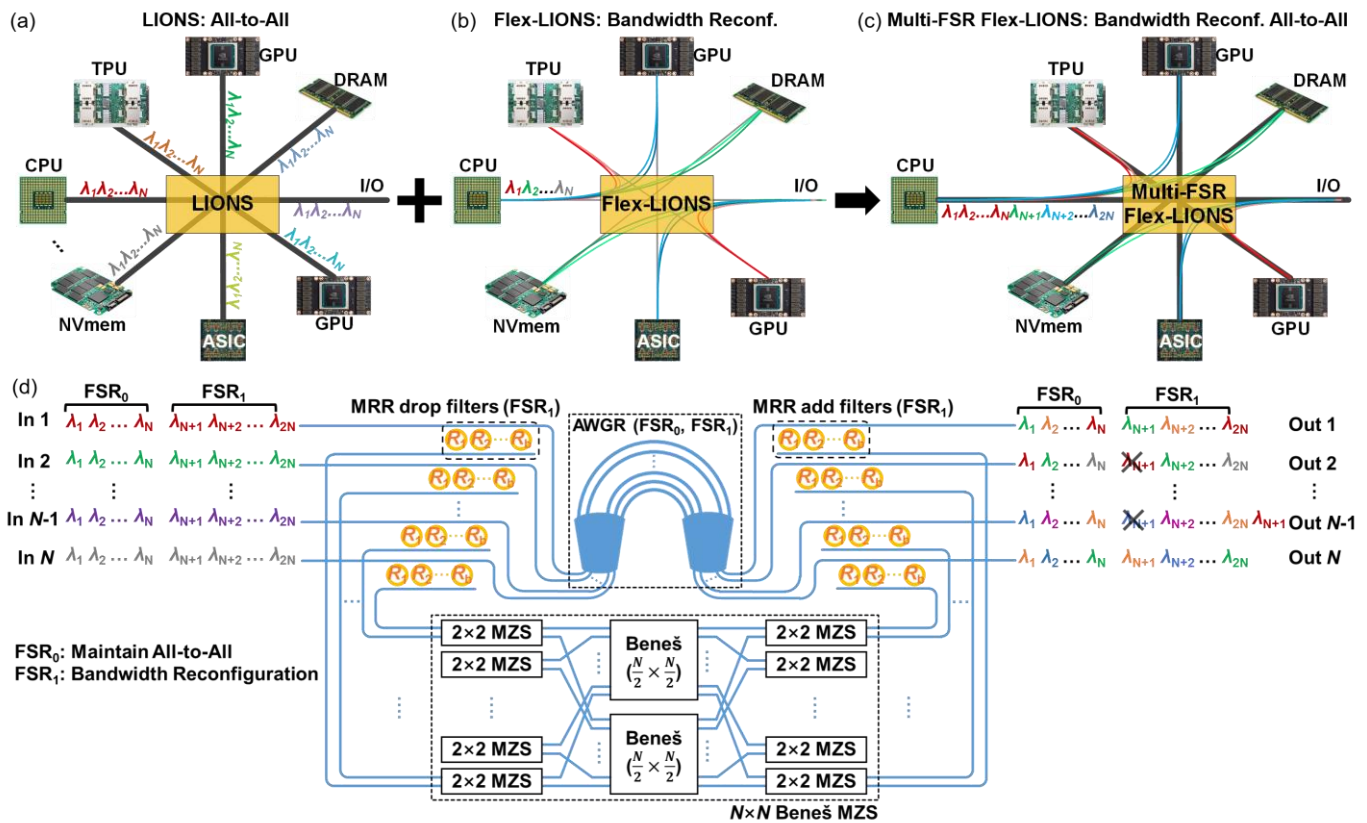


Fig. 1. Heterogeneous processor and memory nodes with: (a) LIONS (all-to-all interconnects); (b) Single-FSR Flex-LIONS (bandwidth reconfigurable interconnects); (c) Multi-FSR Flex-LIONS (bandwidth-reconfigurable all-to-all interconnects). (d) $N \times N$ multi-FSR Flex-LIONS architecture with $N \times N$ AWGR, b MRR add-drop filters at each input and output ports, and $N \times N$ Beneš MZS network. FSR₀ is used for maintaining all-to-all interconnectivity and FSR₁ is used for bandwidth reconfiguration.

reconfigurable all-to-all interconnects using a multi-FSR integrated 8×8 SiPh Flex-LIONS module. By using a broadband Beneš Mach-Zehnder switch (MZS) network as the spatial switch [18]–[20], this architecture exhibits lower complexity compared with the state-of-the-art architectures including our previous works in [14]–[16]. Before reconfiguration, experimental system testing shows error-free all-to-all interconnects for both FSR₀ and FSR₁ with a 5.3-dB power penalty induced by AWGR intra-band crosstalk under worst-case polarization scenario. After reconfiguration in FSR₁, the bandwidth between selected pair of nodes is increased from 50 Gb/s to 125 Gb/s while 25 Gb/s/ λ error-free all-to-all interconnects in FSR₀ is maintained.

The remainder of the paper is organized as follows. Section II introduces the architecture and principle of multi-FSR Flex-LIONS and compares Flex-LIONS with the state-of-the-art bandwidth-reconfigurable switching fabrics. Section III details the design, fabrication, and packaging of an 8×8 200-GHz-spacing SiPh Flex-LIONS with a wide-band Beneš MZS network as the spatial switch. Section IV reports the system testing of the integrated module showing bandwidth-reconfigurable all-to-all interconnects. Section V concludes the paper.

II. FLEX-LIONS ARCHITECTURE AND PRINCIPLE

Fig. 1(d) shows the architecture of the $N \times N$ multi-FSR Flex-LIONS which contains an $N \times N$ cyclic AWGR at the core, b

MRR add/drop filters at the input/output ports of the AWGR ($b < N$), and a broad-band $N \times N$ Beneš MZS network (rearrangeably non-blocking) at the bottom. $2N$ wavelength division multiplexing (WDM) signals within two adjacent FSRs of the core AWGR (FSR₀ and FSR₁) are loaded into each input port. For uniform-random traffic, both the FSR₀ ($\lambda_1, \lambda_2, \dots, \lambda_N$) and FSR₁ ($\lambda_{N+1}, \lambda_{N+2}, \dots, \lambda_{2N}$) are used for all-to-all interconnects based on the wavelength routing function of the AWGR so that the bandwidth between each pair of nodes is $2B$ (B is the bandwidth carried by single wavelength) for Flex-LIONS using two FSRs. For resolving hotspots, up to b of the N wavelengths in FSR₁ from each input port can be dropped by the MRR drop filters and spatially switched to a selected output port by the Beneš MZS network so that the bandwidth between a specific node pair can be increased to $(b+2)B$. Since FSR₀ is untouched, even if some pairs of nodes lose connectivity due to the reconfiguration in FSR₁, they can still maintain a minimum bandwidth of single- λ interconnections through FSR₀. For example, assuming λ_{N+1} from input port 1 (which is initially used for interconnecting with output port 2) is reconfigured to output port $N-1$, both input port 1 to output port 2 and input port 4 to output port $N-1$ will lose one wavelength (λ_{N+1}) as shown in Fig. 1(d). However, the connectivity between these two pairs of nodes is maintained by using λ_1 in FSR₀.

Compared with the state-of-the-art architectures with bandwidth reconfiguration capabilities, Flex-LIONS with Beneš MZS network (this work) has the lowest number of switching elements as shown in Fig. 2. At the radix of 128, our

architecture saves $62.4\times$ of the number of switching elements compared with SiPh echelle gratings + MEMS arrays [11] and SiPh multi-wavelength selective crossbar [12]. InGaAsP/InP AWGR plus SOA approach in [10] has a similar number of switching elements compared with this work, but the high on-chip insertion loss prevents it from scaling up to high radix. Compared to Flex-LIONS with multi-wavelength MRR crossbar [14]–[16], the number of cascaded MRRs on the path of the reconfigured channels in this work is reduced from three to two so that the bandwidth-narrowing effect is reduced. Besides, this work has lower architectural complexity as shown in Fig. 2 since the number of switching elements of the Beneš MZS network is $M\log_2 N - N/2$ while that of multi-wavelength MRR crossbar is N^2 .

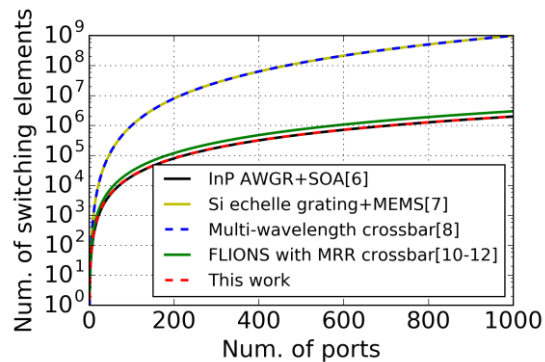


Fig. 2. The number of switching elements with varied number of ports among the state-of-art bandwidth-reconfigurable switching fabrics and Flex-LIONS.

III. SILICON PHOTONIC 8×8 FLEX-LIONS DESIGN, FABRICATION, AND PACKAGING

This section presents the design, fabrication, and packaging of the SiPh 8×8 Flex-LIONS ($N=8$, $b=3$) chip.

A. Design

The SiPh Flex-LIONS device is designed on a multi-layer platform on a silicon-on-insulator (SOI) wafer as shown in Fig. 3(a) [21]. The buried oxide of the SOI wafers is 3- μm thick. The bottom 220-nm Si layer contains MRR add-drop filters and Beneš MZS network. Ridge Si waveguides with 500-nm width are used for single transverse electric (TE) mode transmission and low propagation loss. Above the Si layer is the 200-nm silicon nitride (SiN) waveguide layer which contains the 200-GHz-spacing 8×8 low-crosstalk SiN AWGR. The SiN layer vertically interfaces with the Si layer through inverse-tapered evanescent couplers with a 600-nm gap [21], [22]. Ridge SiN waveguides with 2- μm width are used for low propagation loss and a relatively large bending radius. The silicon oxide cladding of the SiN layer is 3- μm thick. An oxide cladding window is etched to 1.2 μm above the Si layer for higher thermo-optical (TO) tuning efficiency of the switching elements. On top of the oxide cladding are the 400-nm-thick Ti heater layer and 800-nm-thick Au contact metal layer.

Fig. 3(b) shows the 8×8 SiPh Flex-LIONS chip layout. Edge coupler arrays with 127- μm -pitch are used for low coupling loss from the fiber array to the chip. The edge coupler contains a SiN inverse taper from 2 μm to 200 nm and an evanescent coupler from the SiN layer to the Si layer. Two loop-back

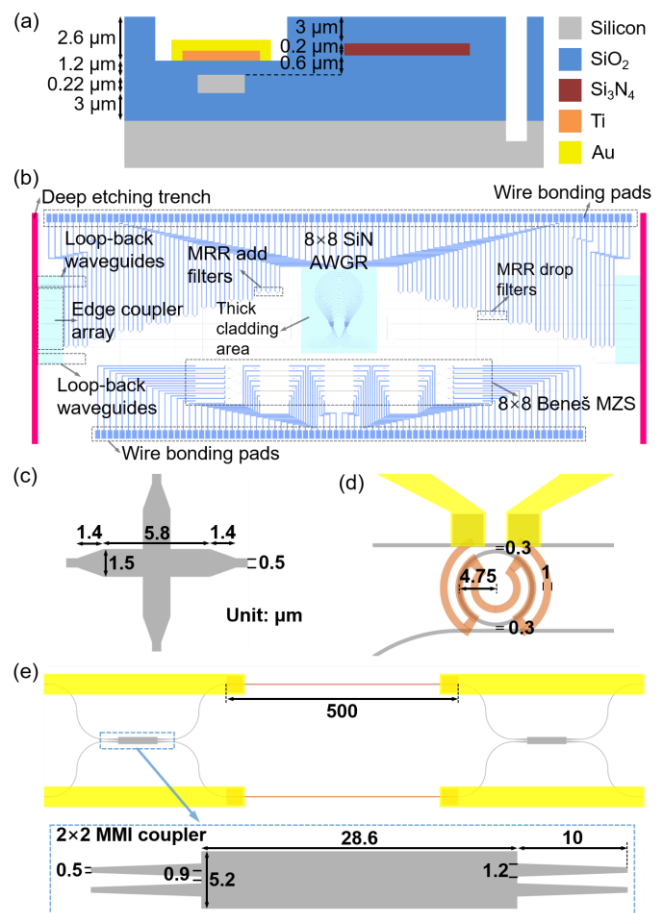


Fig. 3. (a) Cross section of the multi-layer platform. (b) 8×8 SiPh Flex-LIONS layout. (c) Design of MMI based waveguide crossing. (d) Layout of MRR add-drop filter. (e) Layout of 2×2 MZ switching element (arm length not to scale). waveguides are placed on both sides of the edge coupler array for fiber array alignment. SiPh multimode interference (MMI) waveguide crossings are designed to lower the overall insertion loss as shown in Fig. 3(c). The detailed design and simulation results of the MMI waveguide crossing can be found in [16].

The radius and the gap of the MRR add-drop filters are fabrication-calibrated to be 4.75 μm and 0.3 μm , respectively. Spiral resistive heaters along the MRR waveguide are designed to increase the TO tuning efficiency as shown in Fig. 3(d). The width of the heaters is 1 μm . Fig. 3(e) shows the layout of 2×2 MZS as the building block of the Beneš MZS network. The 2×2 MZS contains two 2×2 MMI couplers and two 500- μm -long arms. The 2×2 MMI couplers are designed and fabrication-calibrated for low insertion loss and high power balance [23]. The width and length of the 2×2 MMI couplers are 5.2 μm and 28.6 μm , respectively. The center-to-center distance between the two access waveguides is optimized to be 1.8 μm . The input and output waveguides are linearly tapered to 1.2 μm in a length of 10 μm . In order to achieve minimum TO tuning power, heaters are placed on both arms of the MZS. The width and length of the heaters are 1 μm and 500 μm , respectively.

B. Fabrication

The Flex-LIONS chip was fabricated on a 220-nm silicon on insulator (SOI) wafer with 3- μm -thick buried oxide using the micro and nanoscale fabrication facilities at the University of

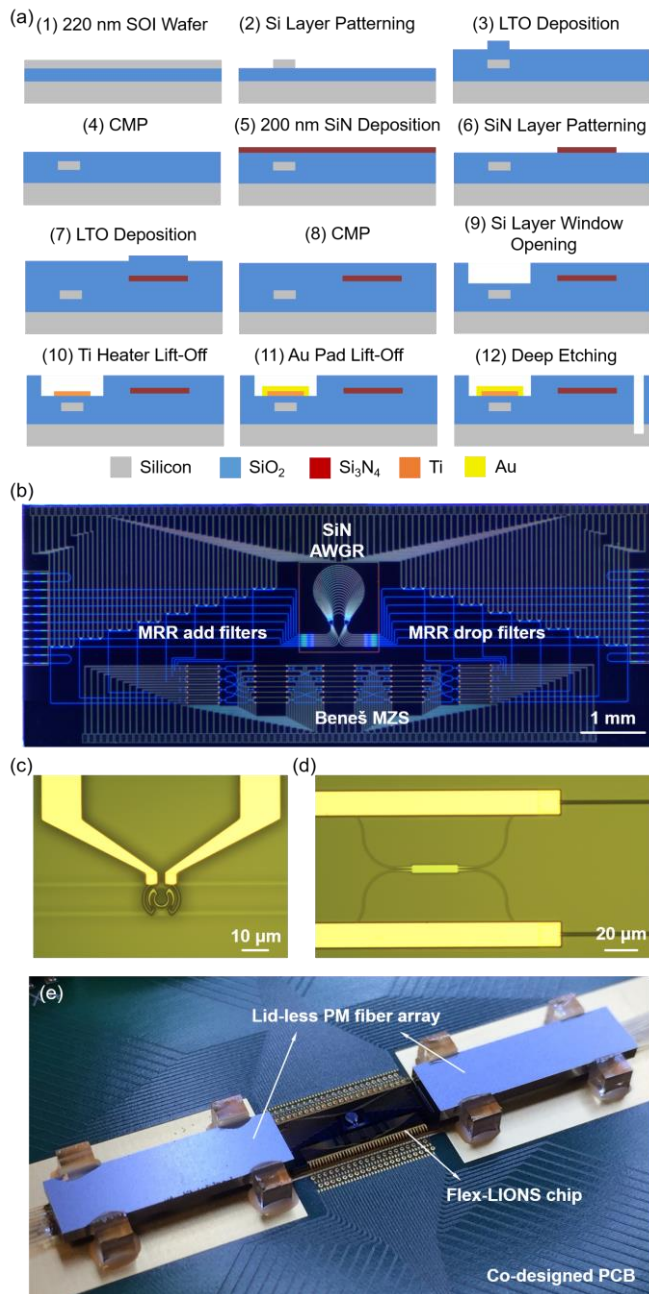


Fig. 4. (a) Fabrication flow charts for the 8×8 SiPh Flex-LIONS. (b) Microscope image of the fabricated 8×8 SiPh Flex-LIONS ($N=8$, $b=3$) chip. (c) Microscope image of MRR add-drop filter. (d) Microscope image of part of 2×2 MZS. (e) Photograph of the integrated Flex-LIONS module with lid-less PM fiber arrays on a co-designed PCB. (Courtesy of Optelligent, LLC).

California at Davis and Berkeley. Fig. 4(a) shows the fabrication flow charts. Firstly, the Si layer is defined by deep-UV projection lithography and inductive coupled plasma (ICP) etching. Then a 1000-nm low-temperature oxide (LTO) was deposited by low-pressure chemical vapor deposition (LPCVD) and then planarized to 800 nm by chemical mechanical planarization (CMP). Following the deposition of a 200-nm SiN layer by LPCVD, the AWGR was patterned by deep-UV lithography and ICP etching. Then a 3-μm LTO cladding was deposited and planarized. Subsequently, the oxide cladding window is opened by ICP etching. The 400-nm-thick Ti heater layer and 800-nm-Au contact metal layer were then fabricated




Fig. 5. (a) Transmission spectra of 8×8 SiN AWGR from input port 4. (b) Linear fitting of the normalized transmission of Si MMI waveguide crossing for insertion loss calculation.

by E-beam evaporation and lift-off. Finally, a 140-μm deep etching trench is fabricated using ICP etching. Fig. 4(b-d) show the microscope images of the fabricated chip, MRR add-drop filter, and the 2×2 MZS. The total chip size is 10 mm × 4 mm.

C. Packaging

The fabricated chip with 176 electrical pads on the edge was wire-bonded to a co-designed printed circuit board (PCB) for electrical fan-out. Two lid-less 16-channel 127-μm-pitch polarization-maintaining (PM) fiber arrays were attached to the input and output of the chip using index-matching UV epoxy. Flexible flat cable (FFC) connectors are surface-mounted on the PCB for a compact footprint. The coupling loss from the PM fiber array to the chip after packaging is 4.7-5.7 dB/facet. Fig. 4(e) shows the photograph of the integrated Flex-LIONS module.

IV. EXPERIMENTAL DEMONSTRATION OF BANDWIDTH-RECONFIGURABLE ALL-TO-ALL OPTICAL INTERCONNECTS

This section presents the detailed characterization of the single switching elements and an experimental demonstration of bandwidth-reconfigurable all-to-all optical interconnects using the fabricated Flex-LIONS module and two FSRs.

A. Single Elements Characterization

The transmission spectra of the 8×8 SiN AWGR within two FSRs are measured by an optical vector network analyzer (OVNA) system as shown in Fig. 5(a). The free spectral range (FSR), channel spacing, and full-width-at-half-maximum (FWHM) of the AWGR is 12.8 nm, 1.6 nm (200 GHz), and 1.07 nm respectively. The adjacent channel crosstalk is < -18 dB, the non-adjacent channel crosstalk is < -28 dB, and the insertion loss is < 3.5 dB. The eight wavelength channels in FSR₁ (λ₉, λ₁₀, ..., λ₁₆) are for bandwidth reconfiguration while the eight

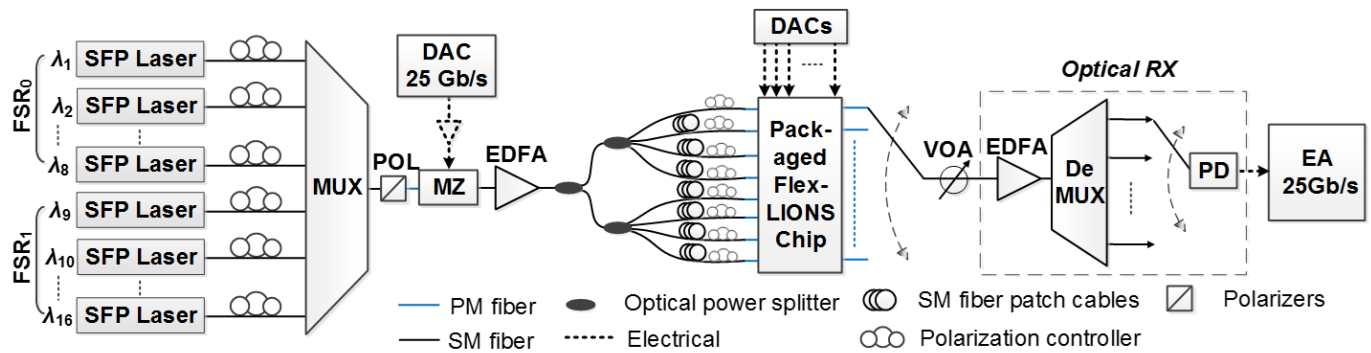


Fig. 7. Experimental setup. SFP: small form pluggable; MUX: multiplexer; MZ: Mach Zehnder; EDFA: erbium-doped fiber amplifier; DAC: digital to analog converter; VOA: variable optical attenuator; DeMUX: demultiplexer; PD: photodetector; EA: error analyzer.

wavelength channels in FSR₀ ($\lambda_1 = \lambda_9 - \text{FSR}$, $\lambda_2 = \lambda_{10} - \text{FSR}$, ..., $\lambda_8 = \lambda_{16} - \text{FSR}$) are for maintaining basic all-to-all connectivity. All the

wavelength channels match with the dense wavelength division multiplexing (DWDM) ITU grid. The insertion loss of the Si MMI waveguide crossing is measured as 0.08 dB through the linear fitting of the normalized transmission of four cascaded waveguide crossing structures as shown in Fig. 5(b).

Fig. 6(a) shows the transmission spectra of the through and drop ports of MRR add-drop filters with different TO tuning power. All the spectra are normalized to the reference waveguide. The insertion loss for the drop port, FWHM, and FSR are 1.4 dB, 0.71 nm, and 20.2 nm, respectively. Fig. 6(b) shows the linear fitting of the resonance wavelength shifting with TO tuning power. The measured TO tuning efficiency of the MRR add-drop filter is 0.3 nm/mW (67 mW/FSR). Fig. 6(c) shows the transmission spectra of MZS at the bar and cross port with different TO tuning power applied on the upper heater. An initial bias of 0.87 mW is required to achieve the cross state due to phase errors induced by fabrication imperfection. The insertion loss is 0.3 dB and the TO power to switch between cross and bar state is 16.5 mW. The crosstalk in the wavelength range of 20 nm is lower than -20 dB while the minimum crosstalk is lower than -40 dB.

B. Experimental Demonstration of Bandwidth Reconfiguration Using Two-FSR Flex-LIONS

Fig. 7 shows the experimental setup for demonstrating bandwidth-reconfigurable all-to-all optical interconnects using the integrated SiPh Flex-LIONS module. Here, two-FSR Flex-LIONS is demonstrated so that FSR₁ can be used for bandwidth steering while FSR₀ maintains basic all-to-all connectivity after reconfiguration.

Sixteen DWDM small form pluggable (SFP) lasers provide the sixteen 200-GHz-spacing WDM signals ($\lambda_1 = 1533.47$ nm, $\lambda_2 = 1535.04$ nm, ..., $\lambda_{16} = 1557.36$ nm). All the WDM signals are multiplexed and modulated by a MZ modulator at 25 Gb/s. The electrical driving signals are $2^{11}-1$ PRBS signals generated by a high-speed digital to analog converter (DAC). Sixteen polarization controllers (PCs) before the multiplexer (MUX) and a polarizer before the MZ modulator are used for polarization alignment. The modulated signal is boosted by an erbium-doped fiber amplifier (EDFA) and then split by a 1x8 splitter. The eight split signals are decorrelated by single-mode fiber catch cables with different lengths and aligned to the polarization of the PM fiber array by a PC before entering the Flex-LIONS module. The output signals from the chip are




Fig. 6. (a) Transmission spectra of through and drop ports of MRR add-drop filter with different TO tuning power. (b) TO tuning efficiency of MRR add-drop filter. (c) Transmission spectra of 2x2 MZS at different TO tuning power for the cross port and the bar port.

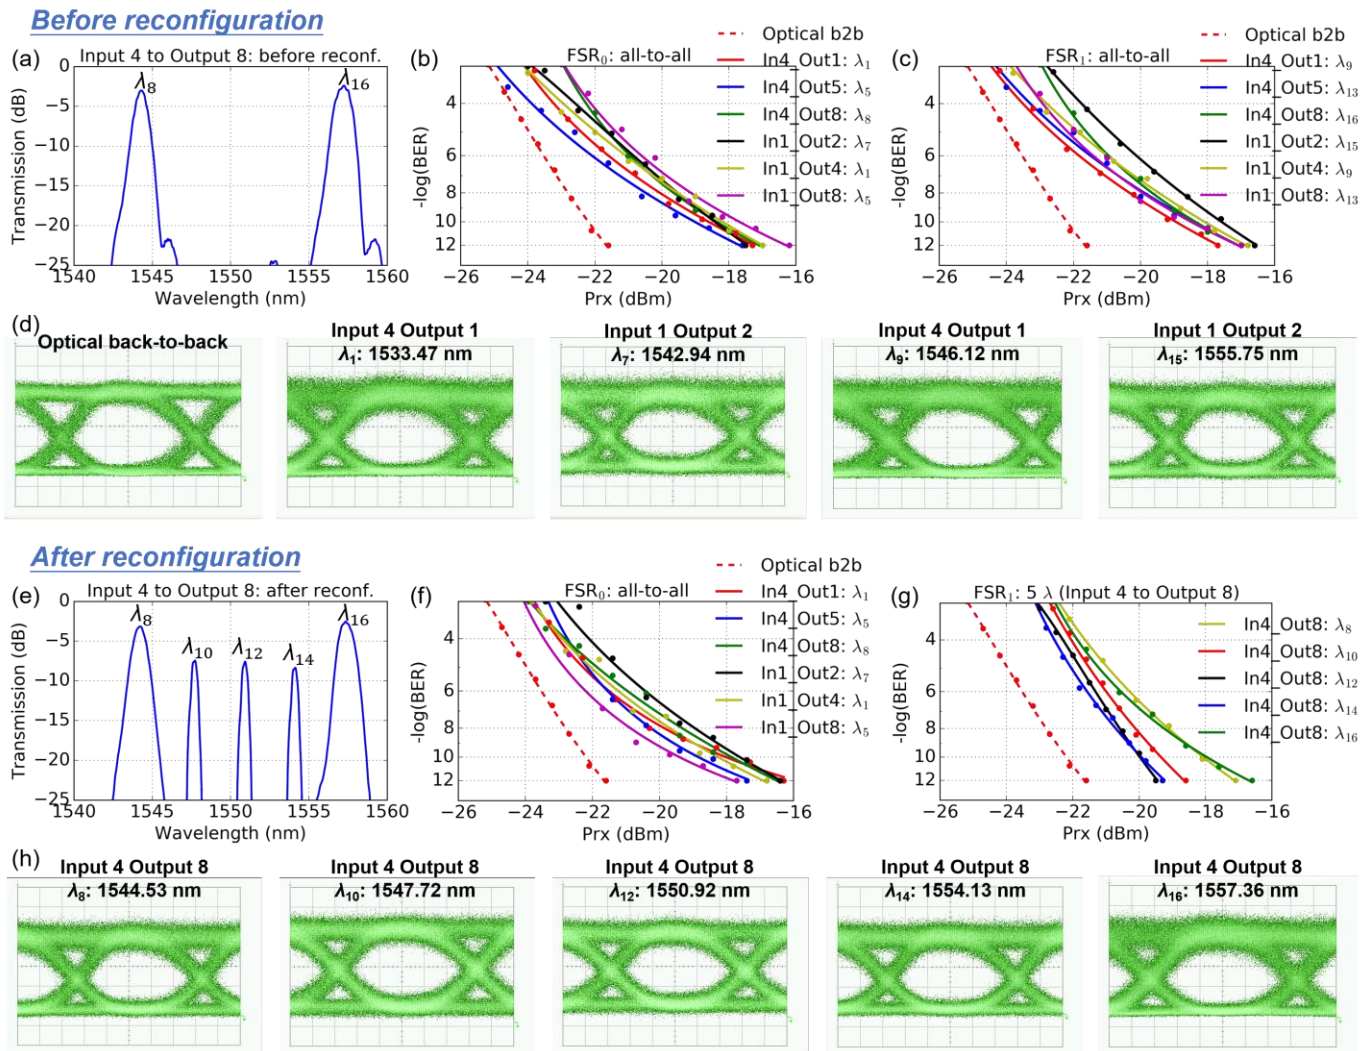


Fig. 8. (a) Transmission spectrum from input port 4 to output port 8 before reconfiguration. (b) BER curves of all-to-all interconnects through FSR_0 before reconfiguration. (c) BER curves of all-to-all interconnects through FSR_1 before reconfiguration. (d) 25 Gb/s eye diagrams for back-to-back and selected input and output ports. (e) Transmission spectrum from input port 4 to output port 8 after reconfiguration. (f) BER curves of all-to-all interconnects through FSR_0 after reconfiguration. (g) BER curves of input port 4 to output port 8 after reconfiguration (λ_8 in FSR_0 , λ_{10} , λ_{12} , λ_{14} , λ_{16} in FSR_1). (h) Eye diagrams of input port 4 to output port 8 using λ_8 , λ_{10} , λ_{12} , λ_{14} , and λ_{16} after reconfiguration.

received by an optically pre-amplified receiver (RX). A real-time error analyzer (EA) performs BER measurements as a function of the RX input power, which is measured by the optical power monitor of the variable optical attenuator (VOA).

Before bandwidth reconfiguration, both FSRs implement all-to-all optical interconnects based on AWGR's wavelength routing property so that the bandwidth between each pair of input and output ports is $2\lambda \times 25 \text{ Gb/s}/\lambda = 50 \text{ Gb/s}$. The total system capacity is $25 \text{ Gb/s}/\lambda \times 16 \lambda \times 8 = 3.2 \text{ Tb/s}$. Fig. 8(a) shows the transmission spectrum from input port 4 to output port 8 with AWGR channel λ_8 and λ_{16} . Fig. 8 (b) and (c) show the BER curves from center and side input ports through FSR_0 and FSR_1 which both demonstrates error-free all-to-all optical interconnects. Comparing with the back-to-back curve (no crosstalk signals added), the measured power penalty under the worst-case crosstalk scenario (aligned polarization for all the input signals) is in the range of 3.9 to 5.3 dB at $BER=10^{-12}$. Such power penalty is mainly induced by the intra-band crosstalk of the AWGR since the crosstalk from cascaded MRR add-drop filters is a second-order crosstalk. The measured power penalty

is slightly lower than the theoretically calculated value [24] due to the polarization of the input signals not being perfectly aligned. Lower crosstalk penalty can be achieved by optimized AWGR design and fabrication [25]. Fig. 8(d) shows the eye diagrams for the back-to-back and selected input and output ports.

After bandwidth reconfiguration, three wavelengths in FSR_1 from input port 4 (λ_{10} , λ_{12} , and λ_{14}) are dropped by the MRR add-drop filter and then routed to output port 8 by the Beneš MZS network. Together with two wavelength channels from the AWGR (λ_8 and λ_{16}), the total number of wavelengths channels from input port 4 to output port 8 is increased to 5 as shown in Fig. 8(e). Note that, the dropping of any wavelength in FSR_1 will not cause any unwanted wavelength drop in FSR_0 since the FSR of the MRR add-drop filter is 12.6 times the channel spacing of the AWRG. The FWHM of AWGR channels (λ_8 and λ_{16}) are 1.05 nm and the FWHM of reconfigured channels (λ_{10} , λ_{12} , and λ_{14}) are narrower (0.42 nm) due to the filtering effect of two cascaded MRR add-drop filters. The insertion loss of the reconfigured channels is $< 8.4 \text{ dB}$

which consists of: 2.8 dB (2×1.4 dB) from the drop loss of the MRR add-drop filters, 0.32 dB (4×0.08 dB) from the insertion loss of the MMI waveguide crossings, 4.1 dB from the Beneš MZS network, and 1.2 dB from the propagation loss of routing waveguides. Error-free operations of all the five wavelength channels show that the bandwidth between input port 4 and output port 8 is increased by $2.5\times$ (50 Gb/s to 125 Gb/s) as shown in Fig. 8(f). Fig. 8(h) shows the eye diagram of these five channels. Note that λ_{10} , λ_{12} , and λ_{14} from input port 4 are initially used for interconnecting with output port 2, 4, and 6 before reconfiguration, respectively. Although these three wavelengths are routed to output port 8 after reconfiguration, all-to-all interconnects through FSR₀ are maintained (as shown in Fig. 8(g)) so that input port 4 can still interconnect with output port 2, 4, and 6 at 25 Gb/s through λ_2 , λ_4 , and λ_6 , respectively.

C. Switching Speed Characterization

The switching speed of the Flex-LIONS chip is characterized by measuring the temporal response of the switching elements. Fig. 9(a) shows the 5-kHz square-wave electrical driving signals that are applied to the MRR add-drop filters and the upper heater of the 2×2 MZS. The peak-to-peak drive voltage is 2 V. Fig. 9(b) and (c) show the measured optical waveform for the MRR add-drop filters and the 2×2 MZS, respectively. The dashed lines mark the 10% and 90% power levels. The measured rise/fall time of the MRR add-drop filters and 2×2 MZS are 7.6/13.6 μ s and 13.2/11.2 μ s, respectively. Faster

switching speed can be obtained by using electro-optical (EO) tuning in the future [19], [26].

D. Power consumption

Without tuning, the resonance of the MRR add-drop filters is designed to be located between λ_8 in FSR₀ and λ_9 in FSR₁ so that the required TO tuning power for reconfiguration is minimum. The average power consumption to correct the fabrication variation for each MRR add-drop filter is 4.23 mW. For the case shown in Section IV.B, the total power consumption is 141.81 mW, which includes 137.46 mW for tuning six MRR add-drop filters and 4.35 mW for switching five MZSs to the cross state. In the worst case, the total power consumption to reconfigure three wavelength channels between a pair of input and output ports is 320.81 mW, assuming the six MRR add-drop filters are tuned to drop the longest wavelength channels in FSR₁ and the five MZSs on the path are switched to bar state.

The TO tuning efficiency of the MRR add-drop filters and MZSs can be further improved by reducing the heater-waveguide distance [27], using silicon doped heater [28], or removing the waveguide substrate and adding air trenches [29]. In addition to the power required for tuning the resonant wavelength, MRR add-drop filters also consume power for wavelength stabilization. A recent work in [30] reported a 65 nm CMOS circuit for MRR resonance auto-alignment and tracking that consumed 5.17 mW. Further reduction in power consumption can be achieved by replacing thermo-optical tuning elements with electro-optical tuning elements [19].

V. CONCLUSION

We propose and experimentally demonstrated the first bandwidth-reconfigurable all-to-all interconnects using a multi-FSR integrated 8×8 SiPh Flex-LIONS module. Device design, fabrication, packaging, and system testing results demonstrate error-free bandwidth reconfiguration from 50 Gb/s to 125 Gb/s between selected node pairs. After reconfiguration in FSR₁, error-free all-to-all interconnects are maintained through FSR₀ with a worst-case crosstalk penalty of 5.3 dB. For scaling up to higher radix, Thin-CLOS Flex-LIONS architecture can be used to overcome the limitation caused by AWGR crosstalk [16].

ACKNOWLEDGMENT

Fabrication of the devices utilized the facilities at the Marvell Nanofabrication Laboratory (Berkeley, CA) and the Center for Nano-Micro Manufacturing (Davis, CA). The authors acknowledge Yi-Chun Ling from the Next Generation Network System (NGNS) Laboratory at UC Davis, Paul Gaudette and David C. Scott from Optelligent, LLC for device packaging.

REFERENCES

- [1] M. J. Schulte *et al.*, "Achieving Exascale Capabilities through Heterogeneous Computing," *IEEE Micro*, vol. 35, no. 4, pp. 26–36, 2015.
- [2] S. Mittal and J. S. Vetter, "A Survey of CPU-GPU Heterogeneous Computing Techniques," *ACM Comput. Surv.*, vol. 47, no. 4, pp. 69:1–69:35, Jul. 2015.




Fig. 9. Time-domain optical response of the switch element. (a) Applied square-wave electrical drive signal. (b) Measured optical waveform for MRR add-drop filter. (c) Measured optical waveform for 2×2 MZS.

- [3] A. Roy, H. Zeng, J. Bagga, G. Porter, and A. C. Snoeren, "Inside the Social Network's (Datacenter) Network," *SIGCOMM Comput. Commun. Rev.*, vol. 45, no. 4, pp. 123–137, Aug. 2015.
- [4] Q. Zhang, V. Liu, H. Zeng, and A. Krishnamurthy, "High-resolution Measurement of Data Center Microbursts," in *Proceedings of the 2017 Internet Measurement Conference*, 2017, pp. 78–85.
- [5] H. Kwon, A. Samajdar, and T. Krishna, "A Communication-Centric Approach for Designing Flexible DNN Accelerators," *IEEE Micro*, vol. 38, no. 6, pp. 25–35, 2018.
- [6] K.-C. (Jimmy) Chen, M. Ebrahimi, T.-Y. Wang, and Y.-C. Yang, "NoC-Based DNN Accelerator: A Future Design Paradigm," in *Proceedings of the 13th IEEE/ACM International Symposium on Networks-on-Chip*, 2019.
- [7] Y. H. Ezzeldin, M. Karmoose, and C. Fragouli, "Communication vs distributed computation: An alternative trade-off curve," in *2017 IEEE Information Theory Workshop (ITW)*, 2017, pp. 279–283.
- [8] M. Ghaffari, "Distributed MIS via All-to-All Communication," in *Proceedings of the ACM Symposium on Principles of Distributed Computing*, 2017, pp. 141–149.
- [9] Y. Lu and H. Gu, "Flexible and Scalable Optical Interconnects for Data Centers: Trends and Challenges," *IEEE Commun. Mag.*, vol. 57, no. 10, pp. 27–33, 2019.
- [10] R. Stabile, A. Rohit, and K. A. Williams, "Monolithically Integrated 8×8 Space and Wavelength Selective Cross-Connect," *J. Light. Technol.*, vol. 32, no. 2, pp. 201–207, 2014.
- [11] T. J. Seok *et al.*, "MEMS-Actuated 8×8 Silicon Photonic Wavelength-Selective Switches with 8 Wavelength Channels," in *2018 Conference on Lasers and Electro-Optics (CLEO)*, 2018, pp. 1–2.
- [12] A. S. P. Khope *et al.*, "Multi-wavelength selective crossbar switch," *Opt. Express*, vol. 27, no. 4, pp. 5203–5216, Feb. 2019.
- [13] R. Proietti, G. Liu, X. Xiao, S. Werner, P. Fotouhi, and S. J. B. Yoo, "FlexLION: A Reconfigurable All-to-All Optical Interconnect Fabric with Bandwidth Steering," in *Conference on Lasers and Electro-Optics*, 2019, p. SM3G.2.
- [14] X. Xiao, R. Proietti, S. Werner, P. Fotouhi, and S. J. B. Yoo, "Flex-LIONS: A Scalable Silicon Photonic Bandwidth-Reconfigurable Optical Switch Fabric," in *2019 24th OptoElectronics and Communications Conference (OECC) and 2019 International Conference on Photonics in Switching and Computing (PSC)*, 2019, pp. 1–3.
- [15] X. Xiao, R. Proietti, G. Liu, H. Lu, Y. Zhang, and S. J. Ben Yoo, "Experimental Demonstration of SiPh Flex-LIONS for Bandwidth-Reconfigurable Optical Interconnects," in *2019 European Conference on Optical Communication (ECOC)*, 2019, pp. 1–4.
- [16] X. Xiao *et al.*, "Silicon Photonic Flex-LIONS for Bandwidth-Reconfigurable Optical Interconnects," *IEEE J. Sel. Top. Quantum Electron.*, vol. 26, no. 2, pp. 1–10, 2020.
- [17] P. Grani, G. Liu, R. Proietti, and S. J. Ben Yoo, "Bit-parallel all-to-all and flexible AWGR-based optical interconnects," in *2017 Optical Fiber Communications Conference and Exhibition (OFC)*, 2017, pp. 1–3.
- [18] S. Zhao, L. Lu, L. Zhou, D. Li, Z. Guo, and J. Chen, "16 \times 16 silicon Mach-Zehnder interferometer switch actuated with waveguide microheaters," *Photon. Res.*, vol. 4, no. 5, pp. 202–207, Oct. 2016.
- [19] L. Lu *et al.*, "16 \times 16 non-blocking silicon optical switch based on electro-optic Mach-Zehnder interferometers," *Opt. Express*, vol. 24, no. 9, p. 9295, 2016.
- [20] N. Dupuis and B. G. Lee, "Impact of Topology on the Scalability of Mach-Zehnder-Based Multistage Silicon Photonic Switch Networks," *J. Light. Technol.*, vol. 36, no. 3, pp. 763–772, 2018.
- [21] K. Shang, S. Pathak, B. Guan, G. Liu, and S. J. B. Yoo, "Low-loss compact multilayer silicon nitride platform for 3D photonic integrated circuits," *Opt. Express*, vol. 23, no. 16, pp. 21334–21342, Aug. 2015.
- [22] W. D. Sacher, Y. Huang, G. Lo, and J. K. S. Poon, "Multilayer Silicon Nitride-on-Silicon Integrated Photonic Platforms and Devices," *J. Light. Technol.*, vol. 33, no. 4, pp. 901–910, 2015.
- [23] L. B. Soldano and E. C. M. Pennings, "Optical multi-mode interference devices based on self-imaging: principles and applications," *J. Light. Technol.*, vol. 13, no. 4, pp. 615–627, 1995.
- [24] H. Takahashi, K. Oda, and H. Toba, "Impact of crosstalk in an arrayed-waveguide multiplexer on N/spl times/N optical interconnection," *J. Light. Technol.*, vol. 14, no. 6, pp. 1097–1105, 1996.
- [25] M. Piels, J. F. Bauters, M. L. Davenport, M. J. R. Heck, and J. E. Bowers, "Low-Loss Silicon Nitride AWG Demultiplexer Heterogeneously Integrated With Hybrid III–V/Silicon Photodetectors," *J. Light. Technol.*, vol. 32, no. 4, pp. 817–823, 2014.
- [26] C. Wang *et al.*, "Integrated lithium niobate electro-optic modulators operating at CMOS-compatible voltages," *Nature*, vol. 562, no. 7725, pp. 101–104, 2018.
- [27] P. Pintus *et al.*, "PWM-Driven Thermally Tunable Silicon Microring Resonators: Design, Fabrication, and Characterization," *Laser Photon. Rev.*, vol. 13, no. 9, p. 1800275, Sep. 2019.
- [28] M. R. Watts, W. A. Zortman, D. C. Trotter, G. N. Nielson, D. L. Luck, and R. W. Young, "Adiabatic Resonant Microrings (ARMs) with directly integrated thermal microphotronics," in *2009 Conference on Lasers and Electro-Optics and 2009 Conference on Quantum Electronics and Laser Science Conference*, 2009, pp. 1–2.
- [29] P. Dong *et al.*, "Thermally tunable silicon racetrack resonators with ultralow tuning power," *Opt. Express*, vol. 18, no. 19, pp. 20298–20304, Sep. 2010.
- [30] H. Li *et al.*, "A 25 Gb/s, 4.4 V-Swing, AC-Coupled Ring Modulator-Based WDM Transmitter with Wavelength Stabilization in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 50, no. 12, pp. 3145–3159, 2015.

Xian Xiao received the B.S. and M.S. degrees from Tsinghua University, Beijing, China, in 2012 and 2015. He has been working toward the Ph.D. degree in electrical and computer engineering at the University of California, Davis, CA, USA, since 2015.

He was a research intern with Nokia Bell Labs in the summer of 2016 and 2017, with Lawrence Berkeley National Laboratory from 2017 to 2018, and with Hewlett-Packard Labs in the summer of 2018.

His current research interest includes silicon photonics, optical interconnects, 2.5D/3D photonic integration, neuromorphic computing.

Roberto Proietti received the M.S. degree in telecommunications engineering from the University of Pisa, Pisa, Italy, in 2004, and the Ph.D. degree in electrical engineering from Scuola Superiore Sant'Anna, Pisa, in 2009. He is a Project Scientist with the Next Generation Networking Systems Laboratory, University of California, Davis. His research interests include optical switching technologies and architectures for supercomputing and data center applications, high-spectral-efficiency coherent transmission systems, and elastic optical networking.

Gengchen Liu was born in Beijing, China, in 1993. He received the B.S. degree in Electrical Engineering from Central Michigan University, Mount Pleasant, MI, USA, in 2015. He is currently pursuing the Ph.D. degree in Electrical and Computer Engineering at University of California, Davis, CA, USA.

His research interest focuses on high-speed optical transceiver design, digital signal processing, and machine learning.

Hongbo Lu was born in Shanghai, China, in 1989. He received the B.S. degree in Communication Engineering from Fudan University, Shanghai, China, in 2011, and the M.S. degree in Electrical Engineering from Tokyo University, Tokyo, Japan, in 2014. He is currently pursuing the Ph.D. degree in Electrical

and Computer Engineering at University of California, Davis, CA, USA.

His research interest focuses on coherent optical communication and optical-wireless integrations.

Yu Zhang received his B. S. degree of Optoelectronics from Huazhong University of Science and Technology, Wuhan, China, in 2010 and the Ph.D. degree in Electronic and Computer Engineering from The Hong Kong University of Science and Technology, Hong Kong, in 2016. He is now an assistant project scientist in electrical and computer engineering at the University of California, Davis, CA.

His current research interest focuses on 3D integrated silicon photonic optical phased array, hybrid silicon laser, amplifier and phase modulators and silicon photonic integrated circuit enabled all-to-all interconnection for scalable high-performance computing systems.

S. J. Ben Yoo [S'82, M'84, SM'97, F'07] currently serves as a distinguished professor of electrical engineering at UC Davis. His research at UC Davis includes 2D/3D photonic integration for future computing, communication, imaging, and navigation systems, micro/nano systems integration, and the future Internet. Prior to joining UC Davis in 1999, he was a Senior Research Scientist at Bellcore, leading technical efforts in integrated photonics, optical networking, and systems integration. His research activities at Bellcore included the next-generation Internet, reconfigurable multiwavelength optical networks (MONET), wavelength interchanging cross connects, wavelength converters, vertical-cavity lasers, and high-speed modulators. He led the MONET testbed experimentation efforts, and participated in ATD/MONET systems integration and a number of standardization activities. Prior to joining Bellcore in 1991, he conducted research on nonlinear optical processes in quantum wells, a four-wave-mixing study of relaxation mechanisms in dye molecules, and ultrafast diffusion-driven photodetectors at Stanford University. Prof. Yoo received his B.S. degree in electrical engineering with distinction, his M.S. degree in electrical engineering, and his Ph.D. degree in electrical engineering with a minor in physics, all from Stanford University, California, in 1984, 1986, and 1991, respectively. He is Fellow of IEEE, OSA, NIAC and a recipient of the DARPA Award for Sustained Excellence (1997), the Bellcore CEO Award (1998), the Mid-Career Research Faculty Award (2004 UC Davis), and the Senior Research Faculty Award (2011 UC Davis).