

Integrated SiPh Flex-LIONS Module for All-to-All Optical Interconnects with Bandwidth Steering

Xian Xiao, Roberto Proietti, Gengchen Liu, Hongbo Lu, Yi-Chun Ling, Yu Zhang, and S. J. Ben Yoo

Department of Electrical and Computer Engineering, University of California, Davis, One Shields Ave., Davis, California 95616 USA
xxiao@ucdavis.edu, shyoo@ucdavis.edu

Abstract: We experimentally demonstrate the first all-to-all optical interconnects with bandwidth steering using an integrated 8×8 SiPh Flex-LIONS module. Experimental results show a 5-dB worst-case crosstalk penalty and 25 Gb/s to 100 Gb/s bandwidth steering. © 2020 The Author(s)

1. Introduction

Heterogeneous processor and memory nodes are applied in today's high performance computing (HPC) and datacenter system for better resource utilization (Fig. 1(a)). The communication patterns in such a system tend to be spatial and temporally non-uniform which means the hotspots and coldspots simultaneously created in different locations of the network can result in heavy congestions in some data links [1]. However, today's interconnection networks based on electronic switches and optical fibers are with fixed topology and fixed connectivity which is incapable of dynamically adapt the bandwidth between certain node pairs to the workloads. On the other hand, the capability of all-to-all interconnects is also necessary for many applications such as deep neural network (DNN), map-reduce, and parallel sorting applications. 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. Indeed, wavelength-and-space selective optical switching fabric that can reconfigure the bandwidth between selected pair of nodes has been demonstrated on various platforms including InGaAsP/InP arrayed waveguide grating routers (AWGR) + semiconductor optical amplifiers (SOAs) [2], silicon photonic (SiPh) echelle gratings + microelectromechanical systems (MEMS) arrays [3], and SiPh multi-wavelength selective microring resonator (MRR) crossbar [4]. In our previous work [5,6], bandwidth-reconfigurable one-to-all multicast has been experimentally demonstrated with SiPh Flex-LIONS with multi-wavelength MRR spatial switch. However, all-to-all interconnects with bandwidth steering are yet to be demonstrated with a fully integrated module and the power penalty induced by AWGR crosstalk need be verified.

In this paper, we experimentally demonstrate the first all-to-all optical interconnects with bandwidth steering using a fully integrated 200-GHz-spacing 8×8 SiPh Flex-LIONS module. By using wide-band Beneš MZS networks as the spatial switch, this architecture exhibits lower complexity and higher bandwidth of the reconfigured channel compared with our previous work [5,6]. System testing show error-free all-to-all interconnection with 5-dB power penalty induced by AWGR intra-band crosstalk under worst-case polarization scenario. Bandwidth steering from 25 Gb/s to 100 Gb/s has also been demonstrated between selected pairs of nodes.

2. Principle, Design, Fabrication, and Single Components Measurements

Figure 1(b) shows the architecture of the $N \times N$ 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 wide-band $N \times N$ Beneš Mach-Zehnder




Fig. 1. (a) Modern HPC systems with heterogeneous processor and memory nodes. (b) $N \times N$ 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 networks.

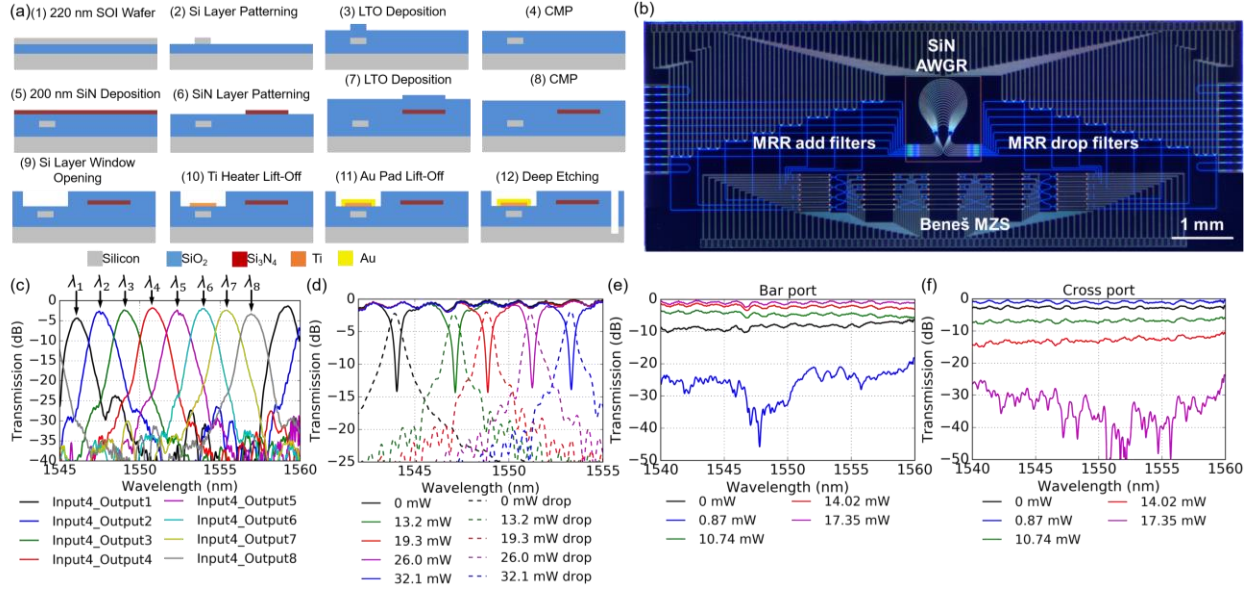


Fig. 2. (a) Fabrication flow charts for the 8×8 SiPh Flex-LIONS. (b) Microscope image of the fabricated 8×8 SiPh Flex-LIONS ($N=8$, $b=3$) chip. (c) Transmission spectra of 8×8 SiN AWGR from input port 4. (d) Transmission spectra of drop port of MRR add-drop filter with different TO tuning power. (e-f) Transmission spectra of 2×2 MZS at different TO tuning power for the cross port and the bar port.

switch (MZS) network (rearrangeable non-blocking) at the bottom. For uniform-random traffic, all the input and output ports can be all-to-all interconnected with N wavelength division multiplexing (WDM) signals based on the wavelength routing function of the AWGR. For resolving hotspots, up to b of the N wavelengths 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 them can be increased by a factor of b . Compared with our previous work [5,6], the number of cascaded MRRs on the path of the reconfigured channels is reduced from three to two so that the MRR filtering effect is reduced. Besides, the architectural complexity of this work is lower since the number of switching elements of the spatial switch is reduced from N^2 to $M\log_2 N - N/2$.

The SiPh Flex-LIONS device is designed on a multi-layer platform. The bottom layer is the 220-nm Si layer which contains MRR add-drop filters and Beneš MZS network. Above the Si layer is the 200-nm SiN layer which contains the 8×8 AWGR. An oxide cladding window is etched on top of the Si area for higher thermos-optical (TO) tuning efficiency. On top of the oxide cladding are the 400-nm-thick Ti heater layer and 800-nm-thick Au contact metal layer. The SiN layer vertically interfaces with the Si layer through inverse-tapered evanescent couplers with a 600-nm gap. The Flex-LIONS chip was fabricated on 220-nm silicon on insulator (SOI) wafer as shown in Fig. 2(a). 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, the AWGR was patterned by deep-UV lithography and ICP etching, followed by a 3- μ m LTO deposition and planarization. Then the oxide cladding window is opened by ICP etching. Then the heater and contact pad are fabricated by E-beam evaporation and lift-off. Finally, a 140- μ m deep etching trench is fabricated using ICP etching. Fig. 2(b) shows the microscope images of the fabricated chip after deep etching.

Figure 2 (c) shows the transmission spectra of the 8×8 cyclic SiN AWGR measured by an optical vector network analyzer (OVNA) system. 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. Figure 2(d) shows the transmission spectra of the through and drop ports of MRR add-drop filters with different TO tuning power values. The insertion loss for the drop port is 1.5 dB, the FWHM is 0.71 nm, and the TO tuning efficiency is 0.3 nm/mW (67 mW/FSR). Figure 2(e) and (f) show the transmission spectra of MZS at the bar and cross port with different TO tuning power. The insertion loss is 0.3 dB and the TO power to switch between cross and bar state is 16.5 mW.

3. Experimental Demonstration of Bandwidth-Reconfigurable All-to-All Optical Interconnects

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, as shown in Fig. 3(a). Two lid-less 16-channel 127- μ m-pitch polarization-maintaining (PM) fiber arrays were attached to the input and output of the chip using index-matching UV epoxy. Figure 3(b)

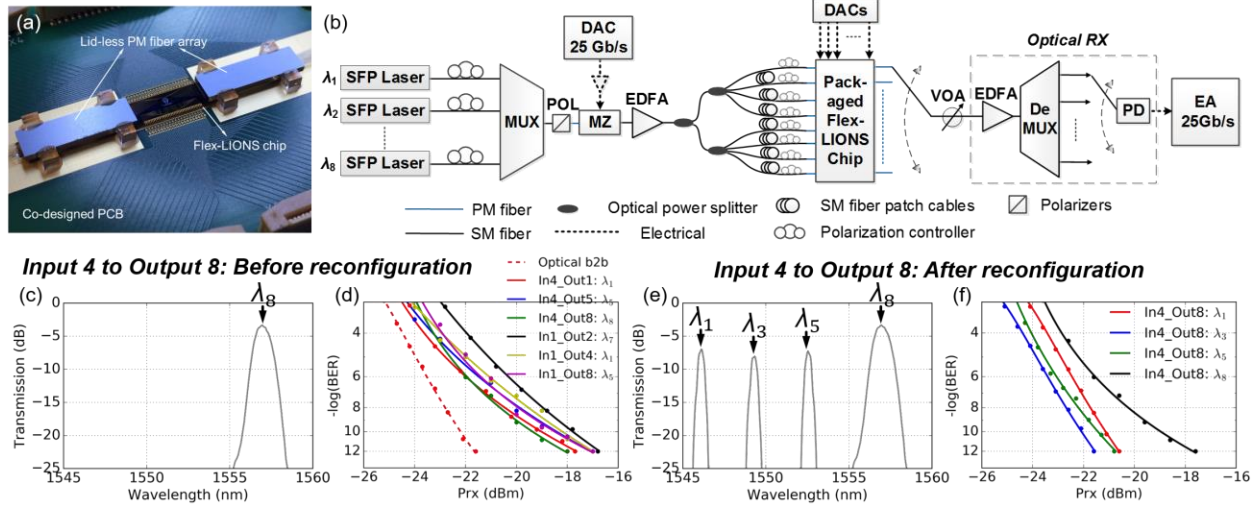


Fig. 3. (a) Photograph of integrated Flex-LIONS module (Courtesy of Optelligent, LLC). (b) Experimental setup. (c) Transmission spectrum from input port 4 to output port 8 before reconfiguration. (d) BER curves of all-to-all interconnects ($1\times$ bandwidth for all interconnected nodes). (e) Transmission spectrum from input port 4 to output port 8 after reconfiguration. (f) BER curves of input port 4 to output port 8.

shows the experimental setup. Eight small form pluggable (SFP) lasers provide the light source of the 200-GHz-spacing WDM signals which match with the AWGR channels. All the WDM signals are multiplexed and modulated by an MZ modulator at 25 Gb/s. The driven signals are $2^{11}-1$ PRBS signals generated by a high-speed digital to analog converter (DAC). Eight 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 split by a 1×8 splitter. Then the eight signals are decorrelated by single-mode fiber catch cables with different lengths and aligned to the polarization of the PM fiber by a PC before input into the packaged Flex-LIONS chip. The output signal from the chip is then 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 reconfiguration, the device provides all-to-all connectivity. Figure 3(c) shows the transmission spectrum from input port 4 to output port 8 with AWGR channel λ_8 . The power penalty from center and side input ports is measured under the worst-case crosstalk scenario (all the input signals aligned in polarization). Figure 3(d) shows the BER curves for selected input and output port combinations which show error-free all-to-all interconnects. The measured power penalty at $\text{BER}=10^{-12}$ is in the range of 3.9 dB to 5 dB compared with back-to-back (no crosstalk signal added). Figure 3(e) shows the transmission spectrum from input port 4 to output port 8 after reconfiguration. λ_8 channel is from the passband of AWGR while the other three channels (λ_1 , λ_3 , λ_5) are from the path through cascaded MRR add-drop filters and Beneš MZS network. Error-free operation of all the four channels demonstrates $4\times$ bandwidth steering (25 Gb/s to 100 Gb/s) between input port 4 and output port 8 as shown in Fig. 3(f).

4. Conclusion

This paper demonstrates the first all-to-all optical interconnects with the ability of bandwidth steering using a fully integrated Flex-LIONS module. The worst-case crosstalk power penalty is measured as below 5 dB. Error-free operation of bandwidth steering shows $4\times$ bandwidth enhancement between selected node pairs.

5. References

- [1] Q. Zhang et al., "High-resolution measurement of data center microbursts," Proceedings of the 2017 Internet Measurement Conference. ACM, 2017.
- [2] R. Stabile et al., "Monolithically Integrated 8×8 Space and Wavelength Selective Cross-Connect," J. Light. Technol., vol. 32, no. 2, pp. 201–207, 2014.
- [3] 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.
- [4] A. S. P. Khope et al., "Multi-wavelength selective crossbar switch," Opt. Express, vol. 27, no. 4, pp. 5203–5216, Feb. 2019.
- [5] X. Xiao et al., "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). IEEE, 2019.
- [6] X. Xiao et al., "Experimental Demonstration of SiPh Flex-LIONS for Bandwidth-Reconfigurable Optical Interconnects," in 2019 European Conference on Optical Communication (ECOC). IEEE, 2019.

This work was supported in part by DoD contract H98230-16-C-0820 and NSF grant 1611560. Device fabrication utilized the facilities at the Marvell Nanofabrication Laboratory (Berkeley, CA) and the Center for Nano-Micro Manufacturing (Davis, CA). The authors acknowledge Paul Gaudette and David C. Scott from Optelligent, LLC for device packaging.