Automated Routing of a Spatial-and-Wavelength Selective Switching Fabric

Brett George[1,*], Liang Yuan Dai[1], Kaylx Jang[1], Zhenguo Wu[1], Priyanka Dilip[1], Michael Cullen[1], and Keren Bergman[1]

[1] *Department of Electrical Engineering, Columbia University, 500 W 120th St., New York, New York, USA. 10027*

[*] *brett.george@columbia.edu*

Abstract: We developed a fully automated routing algorithm and demonstrate end-to-end control of a spatial-and-wavelength selective switching fabric. We show full reconfigurability for all channels with error-free signal recovery for modulated signals up to 16 Gbps.

1. Introduction

The rapid expansion of artificial intelligence (AI) applications has highlighted the need for more efficient networking solutions. Optical switches offer a compelling solution due to their high bandwidth, low power consumption, and capability to dynamically adapt network topology to traffic demands. Specifically, spatial-and-wavelength selective switches (SWSS) have been proposed to efficiently route massively parallel wavelengths with fine granularity in dense wavelength division multiplexed (DWDM) communication systems [1, 2]. Prior studies on SWSS-based network architectures [3, 4] have demonstrated improved job completion times for AI training and high-performance computing workloads using network simulations and small-scale testbed experiments. Scaling up these architectures requires efficient control of these SWSS's to realize low end-to-end overhead.

In this work, we develop and demonstrate a fully automated wavelength routing algorithm for a SWSS fabric. Our end-to-end control operations include 1) constructing an input demand matrix, 2) formulating a graph-based vertex coloring problem for wavelength channel assignment, and 3) automatic thermal tuning of each individual component within the SWSS. We evaluate the feasibility of our approach by performing the above three steps in an experimental testbed where 16 wavelengths are automatically routed in a $4 \times 4 \times 4\lambda$ switch, demonstrating error-free transmission in all cases.

2. Switch Control Algorithm and Results

Prior works [5] have demonstrated the feasibility of Mach-Zehnder Interferometer (MZI) based switching fabrics and its potential to enhance AI training performance. This work demonstrates how a microring resonator (MMR) based switch improves upon prior work with higher scalability [1] while reducing the spatial footprint of the switching elements [2]. In addition, we develop a fully automated control algorithm for switch reconfiguration.

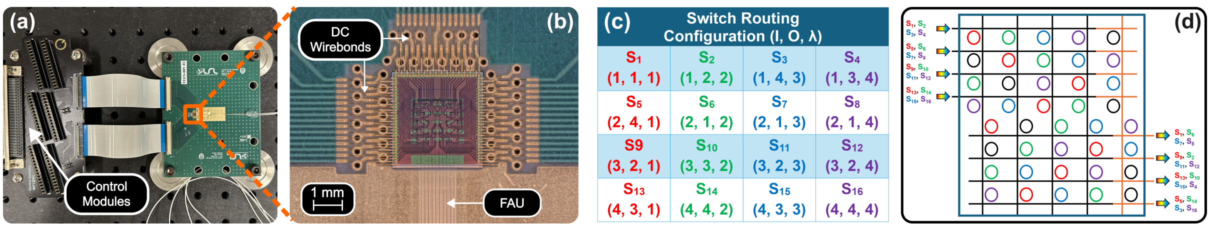

Fig. 1. (a) $4 \times 4 \times 4\lambda$ switch test setup. (b) Wirebonded $4 \times 4 \times 4\lambda$ switch fabricated by AIM Photonics. (c) Example input demand matrix. (d) Block-diagram of a fully configured switch.

Our switch control algorithm first receives an input demand matrix (Fig. 1c) where each tuple entry (I, O, λ) represents a signal with input I, output O, and wavelength λ. The algorithm then converts this configuration into a graph $G(V,E)$, where each node $(v \in V)$ represents a signal tuple, and each edge $(e \in E)$ denotes the *association* between two signals. An edge exists between two nodes if they share the same I, O, or λ. We apply the Vertex-Coloring algorithm on G to configure the switch channels where the color assignments are mapped to each column bus in Fig. 1d. This assignment maps each signal and its corresponding wavelength to specific MMR channels on the switch from input to output, while ensuring that no two identical wavelengths share the same bus. We then apply a pre-calibrated voltage tuning function that calculates and sets the thermal tuning voltage to the MRR switch channels, aligning their resonance to the desired wavelength λ based on the wavelength mapping.

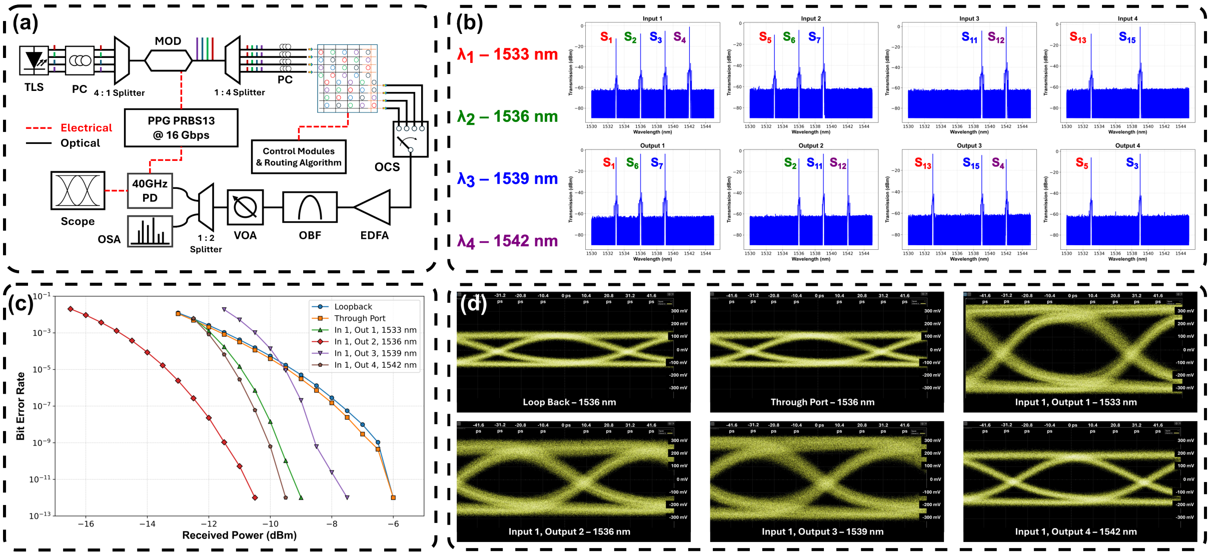

Fig. 2. (a) Schematic of the experimental setup. (b) Input and output transmission spectrum across spatial ports. (c) Bit Error Rate (BER) performance of configured switch channels. (d) Modulated signal eye diagrams for all spatial ports and wavelength channels.

We built our test setup shown in Fig. 2a to transmit and then recover a modulated pseudo-random bit sequence $10^{13} - 1$ bit (PRBS13) signals at 16 Gbps per wavelength. We used four WDM channels starting from 1533 nm to 1542 nm spaced at 3 nm intervals shown in the top row of plots in Fig. 2b. By using a subset of the signals from Fig. 1c, we observe the distribution of signals after the algorithm has completed routing in the second row of plots in Fig. 2b. Our results show low BER and open eye diagrams in Fig. 2(c, d) respectively which leads us to conclude error-free signal recovery of PRBS13 signals at 16 Gbps after switch routing and reconfiguration.

This demonstration of end-to-end control of SWSS interconnects offers an efficient solution to current bandwidth and power consumptions bottlenecks in network architectures. Our implementation enables the co-integration of SWSS interconnects with electronic network infrastructure that is highly scalable and is fully automated to meet changing network traffic demands.

3. Conclusion

We developed an automated routing algorithm that demonstrates full reconfigurability across all spatial and wavelength channels. Using the $4 \times 4 \times 4\lambda$ architecture, we showed that we can achieve error-free signal recovery for modulated signals up to 16 Gbps after being rerouted through several switch channels.

References

1. L. Y. Dai, V. Gopal, and K. Bergman, "Ultra-scalable microring-based architecture for spatial-and-wavelength selective switching," in *2023 IEEE Silicon Photonics Conference (SiPhotonics)*, 2023, pp. 1–2.
2. X. Chen, J. Lin, and K. Wang, "A review of silicon-based integrated optical switches," *Laser & Photonics Reviews*, vol. 17, no. 4, p. 2200571, 2023. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/lpor.202200571
3. G. Liu, R. Proietti, M. Fariborz, P. Fotouhi, X. Xiao, and S. B. Yoo, "Architecture and performance studies of 3d-hyper-flex-lion for reconfigurable all-to-all hpc networks," in *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 2020, pp. 1–16.
4. Z. Wu, L. Y. Dai, Y. Wang, S. Wang, and K. Bergman, "Flexible silicon photonic architecture for accelerating distributed deep learning," *J. Opt. Commun. Netw.*, vol. 16, no. 2, pp. A157–A168, Feb 2024. [Online]. Available: https://opg.optica.org/jocn/abstract.cfm?URI=jocn-16-2-A157
5. A. V. Kumar, A. Devraj, D. Bunandar, and R. Singh, "A case for server-scale photonic connectivity," in *Proceedings of the 23rd ACM Workshop on Hot Topics in Networks*, ser. HotNets '24. New York, NY, USA: Association for Computing Machinery, 2024, p. 290–299. [Online]. Available: https://doi.org/10.1145/3696348.3696856

Acknowledgments: This work was sponsored in part by the U.S. Advanced Research Projects Agency–Energy under ENLITENED Grant DE-AR000843, and in part by the Center for Ubiquitous Connectivity (CUbiC), sponsored by the Semiconductor Research Corporation (SRC) and DARPA under the JUMP 2.0 program.