Two-dimensional photonic-switched high-speed interconnects for AI-driven data centre networks

Rui Ma[(1)], Lingzhi Luo[(1)], Tongyun Li[(1)], Dimitris Syrivelis[(2)], Jing Zhang[(3)], Wajahat Ali[(1)], Ziyao Zhang[(1)], Peng Bao[(1)], Günther Roelkens[(3)], Richard Penty[(1)], Qixiang Cheng[(1)] *

[(1)] Centre for Photonic Systems, Electrical Engineering Division, Department of Engineering, University of Cambridge, Cambridge CB3 0FA, qc223@cam.ac.uk
[(2)] NVIDIA, Ermou 56, Athens 10563, Greece
[(3)] Photonics Research Group, Ghent University-imec, 9052 Ghent, Belgium

Abstract *We present two-dimensional photonic-switched high-speed interconnects with a silicon 4×4×8λ space-and-wavelength selective photonic switch for data centre networks. The switch enables 100Gbps wavelength-space routing for InfiniBand and supports a 1Gbps real-time video transmission over Ethernet, demonstrating its potential for AI-driven, high-bandwidth, agile networks.*

2025 European Conference on Optical Communications (ECOC) | 979-8-3315-9531-9/25/$31.00 ©2025 IEEE | DOI: 10.1109/ECOC66593.2025.11262923

Introduction

The explosive growth of artificial intelligence (AI) is fundamentally reshaping the architectural and performance demands of modern data centres. These workloads require massive bandwidth, ultra-low latency, and highly scalable interconnect fabrics to support parallel processing and real-time data exchange across distributed accelerator clusters. Electronic switches—limited by bandwidth density, power consumption, and thermal constraints—are increasingly inadequate to meet these evolving demands [1].

Beyond compute scaling, emerging paradigms such as memory disaggregation also demand agile, high-performance interconnects. By centralizing memory into pooled resources, disaggregated architectures improve resource utilization but impose stringent requirements on network fabrics for high bandwidth and low latency. Optical circuit switches (OCS) have been proposed to address these challenges by enabling high-throughput, low-latency communication and dynamic reconfiguration. For instance, NVIDIA has integrated OCS to enhance resilience and reconfigurability in its DGX platform [2].

Silicon photonics presents a compelling solution to these challenges. By leveraging the inherent advantages of optical communication—including high speed, low latency, and energy efficiency—photonic switching architectures offer scalable and reconfigurable interconnects tailored to the demands of next-generation data centres. Among these, the space-and-wavelength selective switch (SWSS) provides enhanced flexibility by independently routing multiple wavelength channels across spatial ports. This capability enables fine-grained traffic control and dynamic bandwidth provisioning [3].

In this work, we demonstrate a silicon-based 4×4×8λ SWSS designed for scalable data centre interconnects, implemented using commercially available off-the-shelf (COTS) components. The switch core employs a wavelength-selective routing architecture capable of dynamically switching multiple wavelength channels among four input/output ports. As illustrated in Figure 1(a), the system interconnects multiple server nodes equipped with NVIDIA InfiniBand network interface cards (NICs) to a centralized FPGA-controlled photonic switch. Each NIC transmits data over one or more wavelength channels (e.g., λ_1, λ_2, λ_3), using standardized grids such as DWDM, CWDM, CW-WDM, or LWDM, thereby enabling flexible, reconfigurable communication paths between server racks through a shared optical switching fabric.

Switch design and characterization

Figure 1(b) shows the schematic of the proposed SWSS, which uses a modified dilated Banyan topology to cancel first-order in-band crosstalk. Signal fan-out and fan-in are implemented with 2×2×8λ switch elements (SEs) [4] and passive combiners. The design combines the high tuning efficiency and wavelength selectivity of micro-ring resonators (MRRs) with the broad bandwidth and stability of Mach-Zehnder interferometers (MZIs), making it suitable for high-performance, multi-dimensional switching. An analytical method to evaluate filter performance and supported wavelength channels is detailed in [5].

Figure 1(c) shows a microscope image of the fabricated switch, produced via an IMEC SOI MPW run. The 20 mm² switch includes 12 SEs in a 0.6 mm² footprint, with 200 µm spacing between MRRs to reduce thermal crosstalk. To demonstrate independent wavelength control, signals from a QSFP28 transceiver (O-band) are routed from port I_1 to O_1. As shown in Figure 2, all four wavelength channels achieve over 27 dB extinction ratio and 65 GHz optical bandwidth.

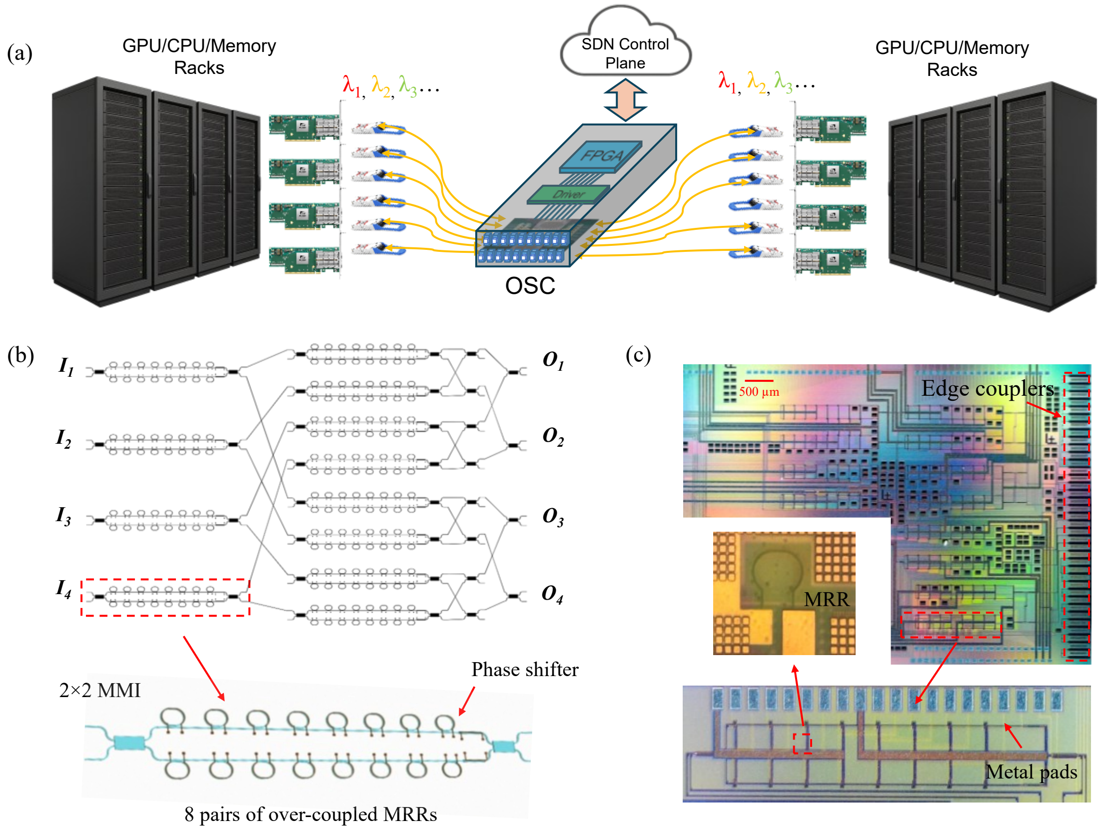

Fig. 1: (a) An SWSS for data centre interconnects. (b) Schematic of the 4×4×8λ switch. (c) Images of the 4×4×8λ switch.

InfiniBand Characterization

Figure 3(a) is a BER characterization using the NICs and transceivers. We connected an NVIDIA Mellanox NIC to a PC via the PCI-e interface, with a pair of 100G QSFP28 transceivers inserted into the NIC. These transceivers operate in the O-band and support up to 40 km reach using four optical wavelengths spaced 4.5 nm apart. The optical signal was transmitted through 500 meters of single-mode fibre (SMF), routed through our photonic switch, and then looped back. On the receiving end, the signal was captured by the corresponding QSFP28 transceiver captured the signal. We use The NVIDIA® Firmware Tools (MFT) to debug and monitor the link. The NIC provided full transceiver control and was configured to run the PRBS31 test mode at a data rate of 4 × 25.78 Gbps/ λ. Raw error counts were collected from the NIC and used to calculate the BER for different spatial paths and wavelengths after transmission through the switch.

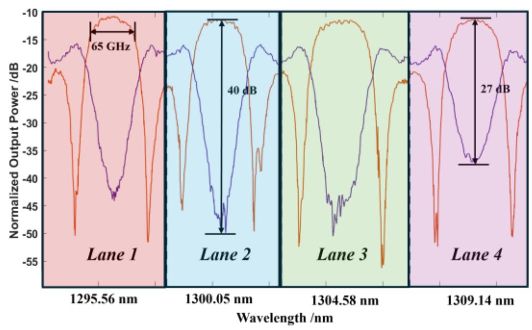

Fig. 2: Transmission spectra of four wavelength channels from I_1 to O_1.

Figure 3 (b) shows the measured BER across all combinations of 4 input ports (I_1–I_4), 4 output ports (O_1–O_4), and 4 wavelengths. Overall, BER values range from 10^{-9} to 10^{-6} across all switching paths. The observed performance indicates that our photonic switch supports data rates of up to 25.78 Gbps per channel. Insertion loss emerged as the primary factor influencing the BER, with each path experiencing approximately ~20 dB of total loss. The loss components include fiber coupling losses, switch insertion losses, and waveguide-related losses, which are inherent to the PIC design.

Ethernet Characterization

To further validate the practical performance of the SWSS for Ethernet traffic, we implemented a real-world Ethernet packet transmission test using 10Gbps SFP+ transceivers and commercial Ethernet switches, as shown in Figure 4(a). A pair of Ethernet switches with SFP+ ports was used to interface the photonic switch, where electrical-to-optical conversion was handled by the switches. The FPGA control board was connected to the switch driver to configure the

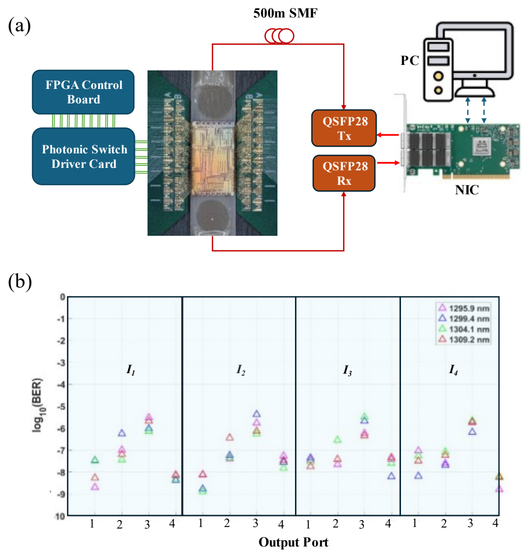

Fig. 3: (a) BER Measurement Using NIC and QSFP28 Modules. (b) BER performance for different input and output port combinations for four different wavelengths.

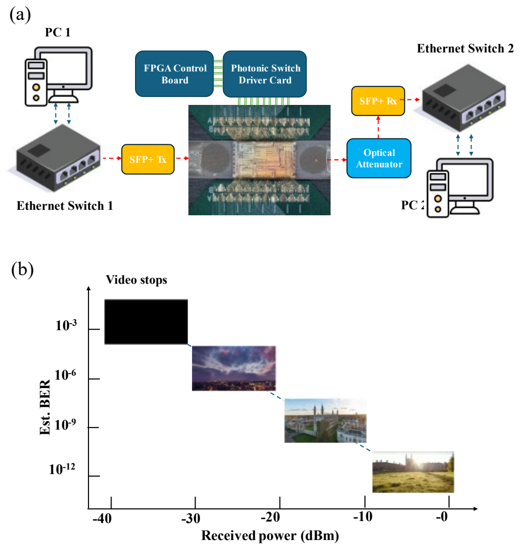

Fig. 4: (a) Ethernet Packet Transmission Using SFP+ and Ethernet Switches. (b) Relationship between the received optical power at the transceiver and the resulting BER.

switching states.

We conducted a series of 4K video transmission tests with a pair of 60 km, 1270 nm SFP+ transceivers. Since 4K video streaming has relatively low bandwidth requirements and the PCs were equipped with a 1G NIC, we used VLC media player software to combine multiple video streams, generating a total data rate of approximately 1 Gbps before sending the aggregated stream to the 10 G Ethernet switch with pluggable SFP+ transceivers.

Figure 4(b) illustrates the relationship between the received optical power at the transceiver and the resulting BER, which directly affects the quality of transmitted video. As the received optical power decreases, the BER increases, leading to a degradation in video quality. To evaluate the system's tolerance to additional loss, we inserted an optical attenuator and gradually increased the attenuation. The system was able to maintain stable 4K video transmission even with an extra 13 dB of attenuation, indicating a total dynamic range of at least 13 dB. This confirms that the photonic switch can operate reliably under significant optical power variations, making it suitable for deployment in practical optical networks.

It is worth noting that although the current setup is limited to a maximum data rate of 1 Gbps, this limitation arises from the Ethernet switch interfaces and not the photonic switch itself. Given the measured 65 GHz optical bandwidth of the switch wavelength channels, it is theoretically capable of supporting over 60Gbaud/s, enabling data rate of 200Gbps/λ, or even beyond, assuming future integration with higher-speed IMDD or coherent transceivers and compatible network hardware.

Conclusions

We have experimentally demonstrated a silicon photonic 4×4×8λ SWSS aimed at scalable, low-latency interconnects for modern data centres. BER measurements across InfiniBand links confirmed reliable operation at up to 100 Gbps, with BER values ranging from 10^{-9} to 10^{-6} across all port and wavelength combinations.

To further validate real-world applicability, we implemented an Ethernet-based video transmission test using 10G SFP+ transceivers over 60 km fibre links. Despite the switch's inherent ~20 dB insertion loss, the system successfully sustained high-quality 4K video streaming even under an additional 13 dB of optical attenuation. This demonstrates a total dynamic range of at least 13 dB, confirming robust performance under variable link budgets.

With a measured 3 dB bandwidth exceeding 65 GHz, the SWSS can support future high-speed modulation formats (e.g., 200 Gbps/λ and beyond), contingent on advances in transceiver technology. These results establish the SWSS as a promising platform for high-throughput, reconfigurable optical networks in AI-centric, disaggregated data centre architectures.

Acknowledgements

This work was funded by the UKRI-EPSRC, project QUDOS (EP/T028475/1); and the European Union's Horizon Europe Research and Innovation Program, project PUNCH (101070560) and project INSPIRE (101017088).

References

[1] A. Singh, J. Ong, A. Agarwal, G. Anderson, A. Armistead, R. Bannon, S. Boving, G. Desai, B. Felderman, P. Germano and A. Kanagala, "Jupiter rising: A decade of clos topologies and centralized control in google's datacenter network," ACM SIGCOMM computer communication review, vol. 45, no. 4, pp.183-197,2015, DOI: 10.1145/2829988.2787508.

[2] G. Patronas, N. Terzenidis, P. Kashinkunti, E. Zahavi, D. Syrivelis, L. Capps, Z.A. Wertheimer, N. Argyris, A. Fevgas, C. Thompson and A. Ganor, "Optical switching for data centers and advanced computing systems," Journal of Optical Communications and Networking," vol. 17, no. 1, pp. A87-A95, 2025, DOI: 10.1364/JOCN.534317.

[3] Z. Zhang, M. Chen, R. Ma, B. Sun, A. Wonfor, R. Penty, and Q. Cheng, "Dilated space-and-wavelength selective crosspoint optical switch," Photonics Research, vol. 13. no. 4, pp.924-934,2025, DOI: 10.1364/PRJ.543402.

[4] R. Ma, L. Luo, M. Chen, C. Yao, R. Penty, and Q. Cheng, "A reconfigurable multi-channel optical filter based on a microring-assisted MZI," In 2024 IEEE Photonics Conference (IPC), pp. 1-2, 2024, DOI: 10.1109/IPC60965.2024.10799678.

[5] L. Luo, R. Ma, R. Penty and Q. Cheng, "Unlocking Electro-optic Resonant Phase Shifting for Multi-dimensional, Ultra-dynamic Photonic Switches," Journal of Lightwave Technology, vol. 43, no. 5, pp. 2238-2247, 2025, DOI: 10.1109/JLT.2024.3486100.