Scalability Assessment of O-band SOA-based Broadcast and Select Switch with 100 Gb/s LWDM Commercial Transceivers

M.P.G. Rombouts,[1,*] and N. Calabretta,[1]

[1] *Eindhoven University of Technology, PO Box 513, 5600 MB Eindhoven, The Netherlands*

[*] *m.p.g.rombouts@tue.nl*

Abstract: Scalability and operation of O-band SOA-based Broadcast & Select switches are experimentally assessed using 100 Gb/s commercial transceivers. Results show error-free operation for a 32-port switch with < 1.8 dB power penalty and 10^{-8} for a 64-port switch.

1. Introduction

The integration of Machine Learning (ML) models in our daily lives has caused an increasing demand in compute power. With the yearly doubling of compute for large scale models, the model size and trained parameters increase proportionally [1]. Developments in accelerator hardware support this growth, but large-scale models such as Google's PaLM [2] with 540 billion parameters, require a large-scale distributed approach involving significant traffic streams through the network. Given that the state-of-the-art ML accelerators have a 3.6 Tb/s output capacity [3] and with optics getting closer to the core [4], there is an opportunity to put hundreds of optically interfaced ML accelerators in a cluster using data center technologies such as O-band wavelength grids. The network infrastructure needs to evolve to accommodate the higher bandwidths whilst keeping power consumption minimal. Limiting O/E/O conversions by providing a transparent path using fully optical switches could help to reduce latency, improve bandwidth, be transparent to protocols, and reduce power consumption.

High port count optical space switches are commonly built using technologies such as MEMS [6] or piezo-electric actuators (Polatis). They are fully transparent, generally have hundreds of ports and exhibit low optical losses, but are limited to optical circuit switching applications due to the low switching speeds. Faster optical switches are demonstrated using MZI-based 2×2 building blocks, integrated in a Benes, crossbar, PILOSS or other topology [7]. These devices are compromised on the port count by the limited optical power budget that is available due integrated component losses. Besides, those switches operate in the C-band and only a few O-band switches have been presented [8, 9]. Low-loss and broadband component design in the O-band is an important field of research to further develop such switches. SOA-based switches have shown switch speeds in the order of nanoseconds and have the added benefit of providing gain to the optical input signal, relieving the limitations in the optical power budget. However, SOAs introduce Amplified Spontaneous Emission (ASE) noise, limiting the Optical Signal-to-Noise Ratio (OSNR). Cascading multiple SOAs in a switched network might reduce the OSNR below acceptable limits. Therefore, proper design of the network and switch architectures can help to minimize the amount of cascaded SOAs and thus the OSNR degradation. Fig. 1a shows such network architecture as OPSquare

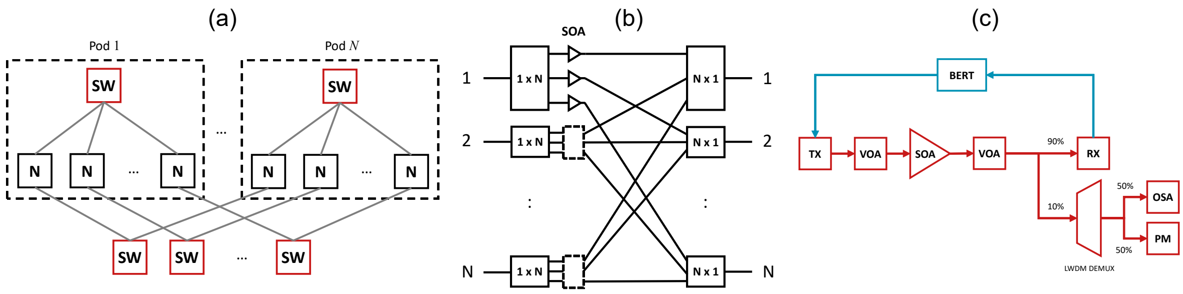

Fig. 1. (a) A flat optical network architecture such as OPSquare [5] where SOA-based switches are used. (b) Broadcast and Select topology for an $N \times N$ SOA-based switch. (c) Measurement setup to test the switch scale using discrete components.

that can interconnect N^2 nodes (which can be a rack or GPU, etc.) by using distributed (and not cascaded) $N \times N$ optical switches [5]. The $N \times N$ switch is based on a broadcast and select switch topology (see Fig. 1b) which minimizes the number of cascaded SOAs to switch any input to any output, limiting OSNR degradation. As the scalability of the switch determines the amount of interconnected nodes, it is essential to investigate the viability and scalability of O-band SOA-based space switch to switch high-capacity (> 100 Gb/s) multi-wavelength data signals.

In this work, we experimentally assess the viability and scalability of the SOA-based broadcast and select space switch architecture with commercial 100 Gb/s LWDM transceivers to introduce non-idealities commonly found in low-cost optical interconnects, such as the lack of polarization control, channel power deviation, a fixed and limited optical output power of the transceivers and integrated WDM channel filters. By using discrete components, the scalability can be investigated, providing a performance upper bound for future photonic integrated designs. Error-free operation of a 32-port optical switch with < 1.8 dB power penalty and 64-port optical switch with error rate $< 10^{-8}$ for the worst channel have been experimentally verified.

2. Experimental setup

The $N \times N$ switch topology is shown in Fig. 1b. A 1 to N splitter splits the power of the incoming signal to N SOAs, which operate as gates as well as amplifiers. The SOA provides gain to passed signals and blocks signals when turned off with a contrast ratio higher than 40 dB. This is important as it minimizes the crosstalk of the leakage signals that combine via the N to 1 combiner at the same output port. The selected signals are combined using an N to 1 combiner. Scaling the switch port count leads to an increased attenuation of $3\log_2 N$ dB per splitter or combiner, requiring higher amplification and causing OSNR degradation. Since all optical paths face an equal set of components and hence losses, the experimental setup can be generalized to what is shown in Fig. 1c. The output signal of a 100G QSFP28 LWDM transceiver is attenuated by a Variable Optical Attenuator (VOA), modeling the 1 to N split losses. An O-band SOA (Philips CQF882) is used as a gate and amplifier, followed by another VOA to model the N to 1 combiner. The SOA has a polarization dependent gain of less than 0.5 dB, a gain up to 19 dB at 350 mA (Fig. 2a) and ASE as shown in Fig. 2b. The second VOA emulates the combiner attenuation and is set to the same value as the first VOA. Using a 90/10 splitter and a LWDM AWG with a 4 nm channel bandwidth as a filter, the OSNR and receiver input power are measured for each of the 4 channels using an optical spectrum analyzer and an optical power meter.

A BERT generates four decorrelated 25 Gb/s NRZ-OOK PRBS-13 streams simultaneously. The received data stream is analyzed per channel to measure the channel power independently. For each tested switch configuration, the SOA power is set such that all channels are observed to be error free.

3. Results

Using the setup shown in Fig. 1c, five switch configurations are tested, from 4×4 up to 64×64. For each N, the spectra, OSNR and Bit Error Rate (BER) are recorded. Fig. 2c shows the optical spectra at the output of the switch. With increasing attenuation of the broadcast and combine stages (up to 36 dB), the SOA current is increased with increased ASE as a consequence. Fig. 3b shows the BER curves for channel 2 ($\lambda = 1300.05$ nm). All configurations up to and including the 32×32 switch manage to get error free performance (BER $< 10^{-9}$). The 64×64 switch gets to $< 10^{-8}$ (well below the Forward Error Correction (FEC) limit of 10^{-3}). Using a higher gain SOA might result in error-free performance. The other channels show similar BER curves, although the ordering of switch configurations might be different due to the non-flat ASE spectrum and the OSNR. This effect is shown in the power penalty for error-free performance in Fig. 3b by means of the shaded area. The average penalty increases

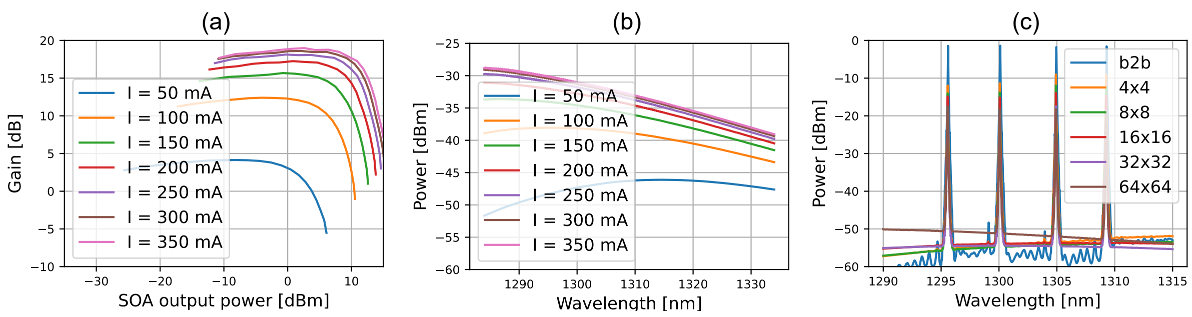

Fig. 2. (a) Gain curves for various drive currents for the gate SOA. (b) ASE spectra for various drive currents for the gate SOA. (c) Switch output spectrum for switch port counts up to 64 × 64.

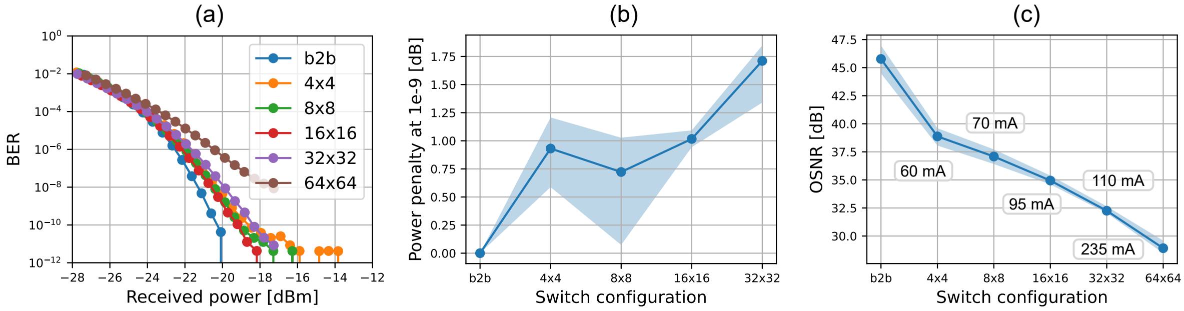

Fig. 3. (a) BER curves for channel 2 (λ = 1300.05 nm). (b) The average channel power penalty for various switch configurations where the BER is set to 10^{-9}. The minimum and maximums are shown by the shaded area. (c) The average channel OSNR for various switch configurations.

in small steps for increasing N and remains below 1.8 dB. The OSNR is shown in Fig. 3c, where it is clear that the OSNR reduces for each N. The differences in OSNR among channels is shown by the shaded area, the labels show the drive currents set to the gate SOA.

4. Conclusion

We have experimentally demonstrated the viability of large, SOA-based, polarization insensitive, optical space switches for commercial 100 Gb/s multichannel interconnects in the O-band. Error-free performance has been demonstrated for switch configurations up to 32×32 with a power penalty of less than 1.8 dB and $< 10^{-8}$ for a 64-port switch, which is well below the FEC limit. This indicates that there are still margins to scale to larger port counts before reaching the FEC limit. The results presented here provide an upper bound. Considering practical photonic integrated realization of the space switches, possible additional losses can occur such as waveguide crossings, (small) extra losses from multimode interference (MMI) couplers commonly used as splitters, and fiber to chip coupling losses. Therefore, practical photonic integrated switches with large port count require higher gain to overcome the optical losses. The FEC margin can be used to compensate additional noise due to the required higher amplification.

Acknowledgements

This work has been partially supported by the EU research and innovation program ALLEGRO grant agreement 101092766.

References

1. J. Sevilla, L. Heim, A. Ho, T. Besiroglu, M. Hobbhahn, and P. Villalobos, “Compute Trends Across Three Eras of Machine Learning,” (2022). ArXiv:2202.05924 [cs].
2. A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts *et al.*, “PaLM: Scaling Language Modeling with Pathways,” (2022). ArXiv:2204.02311 [cs].
3. “NVIDIA H100 Tensor Core GPU Architecture Overview,” https://resources.nvidia.com/en-us-tensor-core/gtc22-whitepaper-hopper (2022).
4. K. Hosseini, E. Kok, S. Y. Shumarayev, C.-P. Chiu, A. Sarkar, A. Toda *et al.*, “8 Tbps Co-Packaged FPGA and Silicon Photonics Optical IO,” p. 3 (2021).
5. F. Yan, W. Miao, O. Raz, and N. Calabretta, “Opsquare: A flat DCN architecture based on flow-controlled optical packet switches,” J. Opt. Commun. Netw. **9**, 291–303 (2017). Conference Name: Journal of Optical Communications and Networking.
6. R. Urata, H. Liu, K. Yasumura, E. Mao, J. Berger, X. Zhou, C. Lam *et al.*, “Mission Apollo: Landing Optical Circuit Switching at Datacenter Scale,” p. 13 (2022).
7. S. J. Ben Yoo, “Prospects and Challenges of Photonic Switching in Data Centers and Computing Systems,” J. Light. Technol. **40**, 2214–2243 (2022).
8. H. R. Mojaver, A. S. Dhillon, R. B. Priti, V. I. Tolstikhin, and O. Liboiron-Ladouceur, “Lossless Scalable Optical Switch Design in a SiP/InP Hybrid Platform,” in *2019 24th OECC and 2019 PSC,* (IEEE, Fukuoka, Japan, 2019), pp. 1–3.
9. R. B. Priti and O. Liboiron-Ladouceur, “A Broadband Rearrangable Non-blocking MZI-based Thermo-optic O-band Switch in Silicon-on-Insulator,” in *Advanced Photonics 2017 (IPR, NOMA, Sensors, Networks, SPPCom, PS),* (OSA, New Orleans, Louisiana, 2017), p. PM4D.2.