

Co-design of a Low-latency Centralized Controller for Silicon Photonic Multistage MZI-based Switches

Yule Xiong¹, Felipe Gohring de Magalhães^{2,3},
Gabriela Nicolescu², Fabiano Hessel³, and Odile Liboiron-Ladouceur¹

¹Dept. of Electrical and Computer Eng., McGill University, Montreal, Canada;
²Ecole Polytechnique de Montreal, Canada; ³PPGCC/PUCRS, Porto Alegre, Brazil
Email: yule.xiong@mail.mcgill.ca

Abstract: An FPGA-based centralized controller architecture for silicon photonics switches is experimentally demonstrated achieving scheduling decision in one clock cycle. The FPGA simultaneously operates as the controller, and the traffic payload generator with error detection.
OCIS codes: (200.4650) Optical interconnects; (250.5300) Photonic integrated circuits; (060.6718) Switching circuit;

1. Introduction

Optical switches and interconnects are promising approaches providing high throughput and potentially lower power consumption for high-performance computing and data center systems [1]. Particularly, silicon photonics (SiP) is an attractive platform to lower process costs and enable compact circuit footprint for greater bandwidth density. Through co-design with IC drivers and controlling ASICs, SiP further provides various functionalities to high-speed systems based on Optical Networks-on-chips (ONoCs) [2, 3]. While design improvement has been demonstrated at the physical-level devices, e.g., in [4], practical and successful deployment of SiP switches require further development of the controllers [5, 6]. Here, we experimentally demonstrate a co-design of an FPGA-based centralized controller with a 4×4 SiP Mach-Zehnder interferometer (MZI)-based switch. The controller based on a look-up table approach makes its routing decision within one single clock cycle, and resolve contention. The practical validation of the optical switch is further pushed by simultaneously using the high-speed features of the FPGA (Altera Stratix IV) as an 8 Gb/s packet generator and error detector. This significant step demonstrates the prospect of SiP towards small radix switch for modern computing platforms.

2. Co-design of the Controller

The co-design setup for the optical switch in Fig. 1 includes a high-speed FPGA (Altera Stratix IV), a 4×4 SiP switch, an optical modulator and photodetector (PD). The FPGA realizes the Look-Up Centralized Controller (LUCC) along with simultaneously being an 8 Gb/s payload generator and error detector. The SiP 4×4 switch employs a Spanke-Beneš topology with five integrated 2×2 balanced MZI switch elements directly controlled by LUCC. Carrier injection tuning method is employed to bias one arm of the 2×2 balanced MZI switch (inset of Fig. 2)) for efficient high-speed switching. The SiP chip measured $V_{\pi} \cdot L$ and switching time are approximately 0.18 V·mm and 6 ns, respectively.

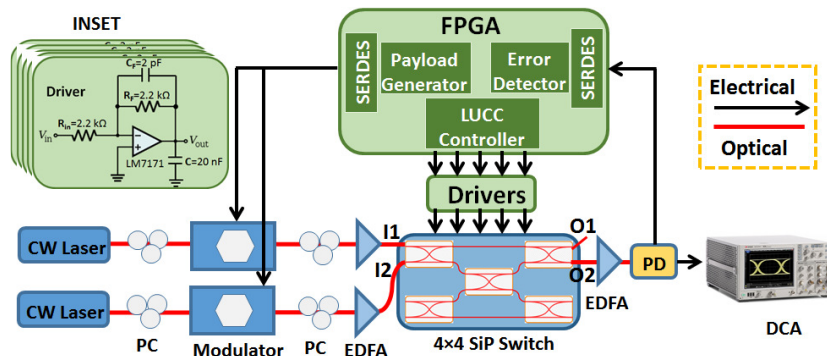


Figure 1: Schematic of the FPGA-based LUCC controller co-designed with the MZI-based multistage switch. CW: Continuous Wave; PC: Polarization Controller; EDFA: Erbium Doped Fiber Amplifier; PD: Photodetector; DCA: Digital Communication Analyzer; INSET: interfacing driving circuit between the FPGA and each MZI of the optical switch.

The controller LUCC is specifically designed to achieve the required low-latency scheduler decision to better exploit the benefit of high-throughput and low-latency optical switches. The developed LUCC can be configured for various interconnect topologies and enables all MZIs connections to be established within the clock cycle [7]. This enables minimum delay when the switch routing path configuration dynamically changes in packet-based

applications. To achieve low-latency, LUCC is designed relying on look-up tables (LUTs) with a modified iSLIP algorithm for routing decision and conflict resolution. The Dijkstra algorithm for shortest possible routes is used [7]. For contention, a Round-Robin (RR) algorithm is used to determine which transmitter (e.g., I1 or I2 in Fig. 4) is granted. LUCC receives requests, solves conflicts and grants access to the network in one clock cycle. In the case of a packet conflict where no routing path is available, the conflicted data packets wait in the Round-Robin queues at the transmitter where the controller handles link requests from the transmitting nodes in a circular manner. For a granted request, LUCC configures the multiple MZIs to the desired states (bar or cross) through a LUT corresponding to the switch topology. As such, LUCC's architecture can be programmed through its LUT to account for the specific interconnection network topology. Topology scalability limitation in terms of the number of ports and stages is limited by the memory available for the LUT. This limitation is estimated to 64 ports for a Spanke-Beneš topology [7]. Once LUCC configures the switch, the transmitters generate the optical packets for routing through the SiP switch to its destination.

Besides the realization of the LUCC, the FPGA also implements a pattern generator as well as a detector for transmission error in the bits of the payload. The high-speed transceiver units of the FPGA serialize and de-serialize the 8 Gb/s payload while the error detector unit allocates temporary memory to store the generated and received payload for bit-level error detection purpose. As a result, this work demonstrates one step forward on the advances of silicon photonic switch fabrics on system-level design and integration. The impact of fabrication process variation in SiP is non-trivial particularly as it relates to the phase delay differences between the MZI element within the switch. This inherent situation leads to a requirement for different switching voltages for each MZI within the same die requiring a one-time calibration to provide the appropriate signal level to each MZI. Further, the current-limited digital control signal of the FPGA cannot drive the analog MZI switch directly requiring up to 11 mA. As such, circuit drivers using an operational amplifier interface the FPGA-based controller and each of the five MZI element, as shown in Fig. 1(INSET). By adjusting the power supply voltage of the operational amplifiers, each driver can provide different switching voltage for each MZI element in the 4x4 SiP switch.

3. Results and Discussion

Figure 2 shows the dynamic switching response for a 50 MHz gating signal for the bar and cross states of one 2x2 MZI switch element, respectively. For this measurement, both the 10 Gb/s payload (PRBS31) and 50 MHz gate signals were generated by pulse pattern generators. The measured rise and fall times are approximately 3 ns, while open eye diagrams are observed for both bar (from IN1 to OUT1) and cross (from IN1 to OUT2) states.




Figure 2: Switching response of the 2x2 MZI switch at 50 MHz. The top row shows the signal for the bar state and the bottom row for the cross state. The inset picture shows the schematic of the 2x2 balanced MZI switch. The bar and cross states are IN1-OUT1 and IN1-OUT2, respectively.

We then characterize the performance of the FPGA as an 8 Gb/s payload generator and error detector. A PRBS31 pattern generated by the FPGA is routed through the 4x4 MZI switch. The FPGA controller is disabled providing static voltage signals to configure the switch without any dynamic gating. A clear 8 Gb/s eye diagram for the PRBS31 pattern payload from the FPGA is shown in the inset of Fig. 3(c) demonstrating successful static routing through the optical switch. In Fig. 3(a), the bit error rate (BER) performance is reported for two static routing paths (I2-O1 and I2-O2). By comparing the received payload data against the generated data, the error counting is completed using the SignalTap II Logic Analyzer tool of the FPGA (Fig. 3(b)). The BER is the number of bit errors (Error Count in Fig. 3(b)) divided by the total number of transferred bits (sum of Good and Error Counts in Fig. 3(b)). The BER as a function of the average received power in the photodetector is illustrated in Fig. 3(c). As a reference for the back-to-back (B2B) BER measurement, the switch is replaced by an attenuator with the same loss corresponding to the insertion loss associated with the respective channels. A negligible power penalty is observed for the tested routes (I2-O1 and I2-O2) and the B2B route.

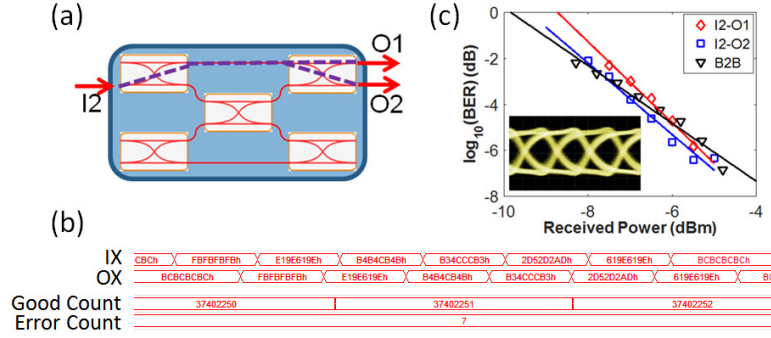


Figure 3: (a) Tested routing paths of the 4x4 MZI switch for BER measurement. (b) 8 Gb/s payload pattern, error and good counts recorded by the SignalTap II Logic Analyzer tool. (c) BER as a function of the average received optical power measured for routing path I2-O1, I2-O2, and the back-to-back configuration with respect to routing path I2-O1.

We complete the demonstration of the co-design of the LUCC with optical switch by considering two scenarios: 1) only I1 sends a link request for the destination O2; 2) both I1 and I2 send link requests for the destination O2, simultaneously. For the first scenario, Fig. 4(a) illustrates that the 8 Gb/s payload is routed from I1 to O2. The received data packet is monitored by the DCA oscilloscope. As there is no conflict and I1 is granted access to its destination (O2), the controller LUCC completes its decision in one clock cycle (1 MHz for Altera Stratix IV). For the second scenario, a conflict occurs where the destination node O2 is the same for both requests by the two transmitters (I1 and I2 in Fig. 4(b)). Fig. 4(b) illustrates the successful contention resolution of the controller LUCC where the payload injected into I1 is given higher priority. The observable difference in optical power is due to the non-uniform optical insertion loss of the Spanke-Beneš topology for different routing paths. For this contention resolution, LUCC successfully takes only one clock cycle to resolve the contention and to grant the access to I1. The switch configuration is further established within that same clock cycle. The data packet from I2 is delayed in the Round-Robin (RR) queue and final transmitted after I1 completes its transmission to output port O2.

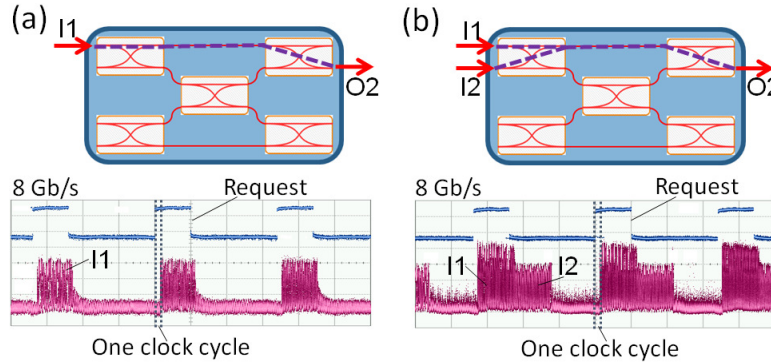


Figure 4: Validation of the LUCC controller co-design with the optical 4x4 MZI switch: (a) only input I1 sends a link request for destination O2; (b) both input I1 and I2 simultaneously send link requests to destination O2, leading to a contention resolved by the controller within one clock cycle.

4. Conclusion

We experimentally demonstrate the co-design of an FPGA-based low-latency controller for a low radix silicon photonic multistage MZI-based switch. A modified iSLIP algorithm is used by the controller LUCC which is based on look-up tables to achieve scheduling with contention resolutions in one clock cycle. Furthermore, the FPGA-based controller executes both the payload pattern and error detection for the initial physical layer validation. This work is an important step towards paving a way for advances of silicon photonic technology on system-level design.

[1] D. A. B. Miller, "Rationale and challenges for optical interconnects to electronic chips", *Proc. IEEE*, vol. 88, pp. 728-749, Jun. 2000.
[2] A. V. Krishnamoorthy, "Computer systems based on silicon photonic interconnects", *Proc. IEEE*, vol. 97, no. 7, pp. 1337-1361, Jul. 2009.
[3] C. Sun, *et al.*, "Single-chip microprocessor that communicates directly using light," *Nature*, vol. 528, no. 7583, pp. 534-538, Dec. 2015.
[4] N. Dupuis, *et al.*, "Modeling and characterization of a nonblocking 4 x 4 mach-zehnder silicon photonic switch fabric," *Journal of Lightwave Technology*, vol. 33, no. 20, pp. 4329-4337, Oct. 2015
[5] F. Lou, *et al.*, "Towards a centralized controller for silicon photonic MZI-based interconnects," in *IEEE OIC*, paper WD4, May 2015.
[6] Y. Xiong, *et al.*, "Towards a fast centralized controller for integrated silicon photonic multistage MZI-based switches," in *Optical Fiber Communication Conference*, paper W1J.2, March 2016.
[7] F.G. de Magalhães, *et al.*, "Design and modelling of a low-latency centralized controller for optical integrated networks," *IEEE Communications Letters*, vol. 20, no. 3, pp. 462-465, March 2016.