

Towards a Fast Centralized Controller for Integrated Silicon Photonic Multistage MZI-based Switches

Yule Xiong¹, Felipe Gohring de Magalhães^{2,3}, Bahaa Radi¹,
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@mcgill.ca

Abstract: An FPGA-based centralized controller architecture for silicon photonics MZI-based switches is experimentally demonstrated achieving scheduling decision in one clock cycle. To counter process variation affecting the switching bias voltages, a pulse-width modulation technique is presented.

OCIS codes: (200.4650) Optical interconnects; (250.5300) Photonic integrated circuits; (060.6718) Switching circuit

1. Introduction

Traditional copper-based electrical interconnects become a bottleneck due to power consumption constraints and throughput limitation. Optical interconnects are promising approaches providing larger bandwidth with potentially lower power consumption [1]. Specifically, silicon photonics (SiP) attracts broad attentions due to its existing process fabrication infrastructure leading to potentially lower cost process, and its high integration density. Furthermore, by co-packaging with IC drivers and controlling ASICs, SiP can provide versatile functionalities to high-speed systems based on Optical Networks-on-chips (ONoCs) [2]. In order to dynamically reconfigure the resources in an ONoC platform, however, large port count and low latency optical switching fabrics are required. While controllers for optical switching fabrics have been demonstrated [3,4], further development of the controllers is necessary for practical and successful deployment of integrated optical switching fabrics.

Here, we experimentally demonstrate a prototype of a centralized controller co-designed for a 4×4 SiP Mach-Zehnder interferometer (MZI)-based switch. The controller, named LUCC for Look-Up Table Centralized Controller, makes its decision in one single clock cycle, and resolve contention. LUCC is implemented in an FPGA and routes 10 Gb/s optical payload.

2. Co-design of the Controller

The schematic of the prototyping controller architecture for the optical switch is shown in Fig. 1(a). It includes an FPGA-based controller, the LUCC, a 4×4 optical switch, transmitter (TX) and receiver (RX) nodes. The integrated SiP switch used in the co-design is a 4×4 optical switch is a Spanke-Beneš topology with five integrated 2×2 MZIs directly controlled by LUCC. Carrier injection tuning method is employed to bias one arm of the MZI for high-speed and efficient switching. The SiP chip was fabricated by the IME foundry and the measured V_{π} -L and switching time are ~0.18 V-mm and ~6 ns, respectively.

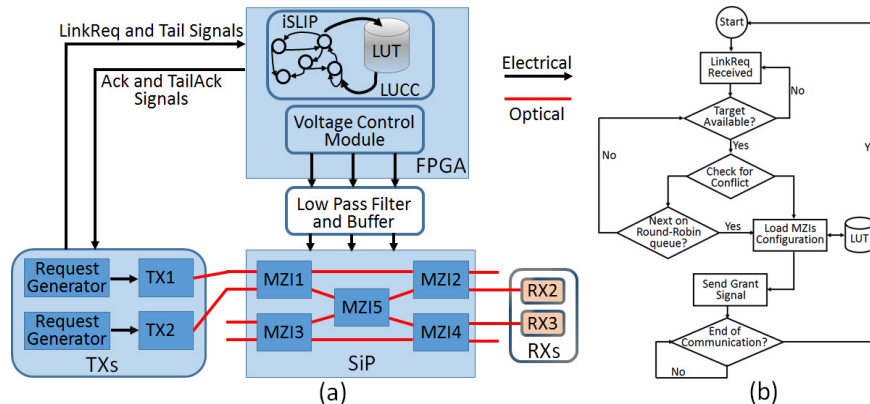


Figure 1 (a) Schematic of the FPGA-based Look-Up Table Centralized Controller (LUCC) co-designed with the MZI-based multistage switch; (b) Decision flow chart of the LUCC.

LUCC is specifically designed to achieve low latency scheduler decision and can be configured for different interconnect topology [5]. LUCC enables all MZIs connections to be established within the clock pulse. This

enables minimum delay when the switch configuration dynamically changes in packet-based applications. In order to achieve low-latencies, LUCC is designed relying on look-up tables (LUTs) with a modified iSLIP algorithm for decision, and for conflict resolution. The Dijkstra algorithm for shortest possible routes is used [5]. For contention, a Round-Robin (RR) algorithm is used to determine which TX is granted. LUCC receives requests, solves conflicts and grants access to the network in one clock cycle. The flow chart in Fig. 1(b) shows the decision making process of LUCC. Transmitters (TXs) firstly send their link requests (LinkReq) to the LUCC, which then checks the availability of the target optical paths. In case no path is available, the data packets at TXs wait in the Round-Robin queues, where the controller handles link requests from TXs in a circular manner [5]. In case of a granted request, LUCC configures the multiple MZIs to the desired states (bar or cross bar) through a LUT corresponding to the switch topology. As such, LUCC's architecture can be programmed through its LUT to account for different topology. Its scalability limitation in terms of the number of ports and stages that a topology may have is mainly limited by the memory available for the LUT. Once LUCC configures the switch, an acknowledgement signal (Ack) is fed back to the TXs enabling the optical packet generations.

Due to the impact of process variation in SiP fabrication and its impact on the phase delay difference of the MZI switches, the required switching voltages are actually different for each MZI within the same die. While heaters can be included to compensate for these inherent variations, it leads to non-trivial power consumption and increased complexity as the number of switch ports scales. To mitigate further control requirements at the interface to the switch, the centralized controller employs a simple electrical pulse-width modulation (PWM) method to convert the controlling signal to the desired bias voltage for each individual MZI switch [6]. This approach mitigates the need of thermo-optic phase trimmers leading to a more energy-efficient solution. As illustrated in the FPGA implementation in Fig. 1(a), the repetitive voltage pulse trains of the PWM are generated by the voltage control module and applied to a low-pass filter and an operational amplifier buffer. The accurate switching voltages are extracted based on the duty cycle of the pulse trains. Hence, the desired bias voltages can be accurately obtained by changing the duty cycle of the voltage pulse train, taking into account the impact of process variations and potentially changes in the MZI bias voltages due to temperature changes.

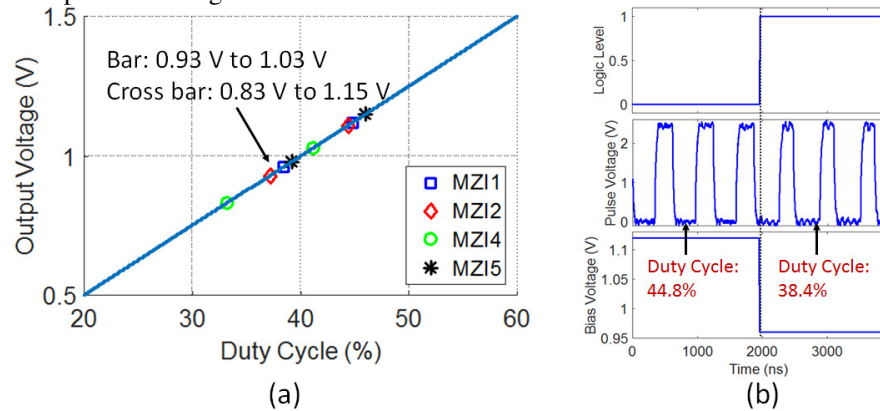


Figure 2 Bias voltage adjustment technique using the PWM method: (a) the bias voltage is linearly proportional to the duty cycle. (b) Output high and low voltages from the controller to one MZI with corresponding duty cycle.

Fig. 2 illustrates the method used to accurately bias the MZIs for optimum extinction ratio between the bar and cross bar configuration. Essentially, the bias voltage to the MZI switch is linearly proportional to the duty cycle of the voltage pulse train from the voltage control module. In this current implementation, the adjustable resolution is 80 mV in the range between 0 (0% duty cycle) and 2.5 V (100% duty cycle). The resolution can be increased by adding bits and four bits are used here. As illustrated in Fig. 2(b) to enable the switching of the first MZI switch (MZI1), the applied bias voltage is 0.96 V (Bar state) and 1.12 V (Cross bar state). This was enabled by adjusting the duty-cycle of the voltage pulse trains to 44.8% for the cross bar state which corresponds to a logical '0' in the decision made by LUCC. The bar state voltage (0.96 V) is obtained by adjusting the duty-cycle to 38.4%. With such approach, the technique can account for changes in the switching voltage due to process variations and/or temperature changes. For this specific SiP switch, the variations seen in the required bias voltages of the other MZI switches are +/- 50 mV for the bar state, and +/-160 mV for cross bar state [Fig. 2(a)]. One of the MZI switch (MZI3) is not shown as it was not functioning due to a short in the wire bonding of the SiP switch. In a deployed solution, the controller would be calibrated with all the required switch voltages while a feedback control can be implemented within the control to take into account changes due to temperature.

3. Results and Discussion

The proof of concept for the prototype was shown in Fig. 1(a). Two transmitter nodes (TX1 and TX2), and two destination nodes (RX2 and RX3) are considered. Experimentally, the 10 Gb/s optical payload is injected into TX1 while TX2 is not injected due to component limitations. At the destination node RX, the output signal is monitored to verify the switching performances. The time diagrams are detailed in Fig. 3(a) where the MZI's states (0) and (1) denote bar and cross bar states of the MZI, respectively. The request state (-1) of the TXs denotes the idle state. After one request (LinkReq), LUCC takes only one clock cycle to send an Ack signal triggering the start of the optical communication between the source and destination ports. The switches configuration is set within that same clock cycle. The order of the four bits of the controller digital signals LinkReq, Ack, Tail and TailAck in Fig. 3(a) corresponds to TX4, TX3, TX2 and TX1, respectively, where (0) denotes the “off” state and (1) denotes the “on” state.

When a conflict occurs where the RX destination node is the same for two or more transmitters (TX1 and TX2 in Fig. 3), the data packet from TX2 is delayed in the Round-Robin (RR) queue until TX1 finishes its communication. Fig. 3(b) illustrates the contention resolution of LUCC where a CW light is injected into TX1 and TX2. The normalized light power is monitored while the injected traffic is the same as shown in Fig. 3(a), where TX1 has higher priority in the case of a contention. Then 10 Gb/s payload from TX1 is injected into the switch fabric via grating couplers, and converted into electrical signals at the RX2. The difference in optical power is due to non-uniform insertion loss of the Spanke- Beneš topology. Fig. 3(c) shows the large extinction ratio of the received electrical signals when the target optical path is either the “ON” or “OFF” state.

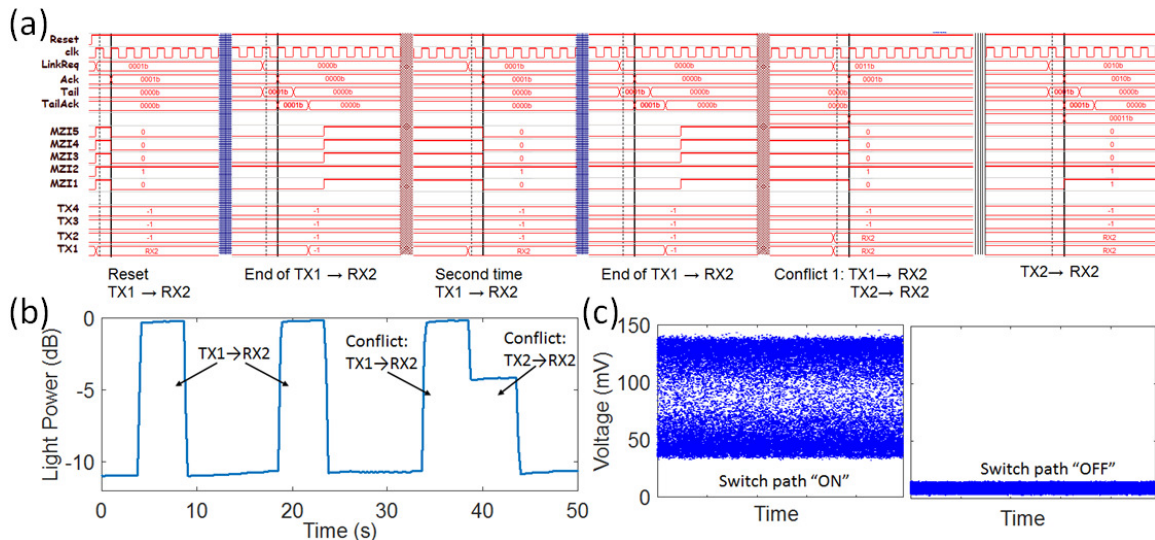


Figure 3 Time diagrams: (a) Main digital signals of LUCC; (b) Optical light injected at the TX1 and received at RX2; (c) Generated 10 Gb/s payload at the TX1 and received at RX2.

4. Conclusion

We experimentally presented an FPGA-based low latency controller architecture for a 4×4 silicon photonic multistage MZI-based switch. A modified iSLIP algorithm, LUCC, based on look-up tables is employed to achieve scheduling with contention resolutions in one clock cycle. To mitigate the impact of fabrication variations on the MZI switching voltage, a pulse-width modulation technique is used.

- [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] Fei Lou, *et al.*, "Towards a centralized controller for silicon photonic MZI-based interconnects," in *IEEE OIC*, paper WD4, May 2015.
- [4] N. Dupuis, *et al.*, "Modeling and Characterization of a Nonblocking 4 × 4 Mach-Zehnder Silicon Photonic Switch Fabric," *Journal of Lightwave Technology*, vol. 33, no. 20, pp. 4329-4337, Oct. 2015
- [5] F. Gohring de Magalhães *et al.*, "A Low-Latency Centralized Controller for MZI-based Optical Integrated Networks," *Photonics in Switching*, Sep. 2015.
- [6] M. Barr, "Pulse Width Modulation," *Embedded Systems Programming*, vol. 14, pp. 103-104, Sep. 2001.