Scalability of Silicon Photonic Enabled Optically Connected Memory

Takashi Shiraishi[1,3], Yang Liu[2], Qi Li[1], Xiaoliang Zhu[1], Kishore Padmaraju[1], Ran Ding[2], Tom Baehr-Jones[2], Michael Hochberg[2,4,5], Keren Bergman[1]

1: Department of Electrical Engineering, Columbia University, 500W 120th St, New York, New York 10027
2: Department of Electrical and Computer Engineering, University of Delaware, 110 DuPont Hall, Newark, Delaware 19716
3: Fujitsu Laboratories Ltd, 10-1 Morinosato-Wakamiya, Atsugi 243-0197, Japan
*4: Institute of Microelectronics, A*STAR (Agency for Science, Technology and Research), 11 Science Park Road, Singapore Science Park II, Singapore 117685, Singapore*
5: Department of Electrical & Computer Engineering, National University of Singapore, Singapore 117576, Singapore
ts2821@columbia.edu

***Abstract*—We investigate the performance of a high-speed optical switch for optically connected memory and study free-carrier absorption losses as the switch scales. The bit-rate dependent optical power budget is also analyzed for optically connected memory.**

I. Introduction

Future CPU-memory communication requires larger capacity and higher bandwidth than is possible using current existing electrical interconnects. Leveraging wavelength division multiplexing (WDM), optically connected memory (OCM) has been demonstrated to enable high-bandwidth interconnectivity to large-capacity and physically-distant memory systems. Additionally, high-speed optical switches can provide flexible network connectivity for memory systems [1].

Integrated silicon photonic technologies are of interest for OCM due to their CMOS-compatible fabrication and low-power operation. To quantify OCM system scalability, it is essential to understand the link performance of silicon photonic networks.

In this work, we analyze the bit-rate dependent optical power budget of a silicon photonic WDM link by using device parameters optimized to maximize the power budget. We then discuss the scalability of the OCM system when incorporating an optical switch.

II. System & Scalable Optical Switch

The WDM link for a memory system is presented in Fig. 1. The optical link includes cascaded silicon microring modulators on the transmitting side and demultiplexing microring filters on the receiving side. We assumed the simple case in which one CPU communicates with many OCMs through an optical switch. The separated OCMs are connected with the CPU by optical fibers.

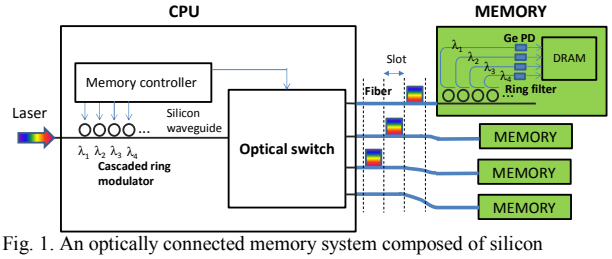

Fig. 1. An optically connected memory system composed of silicon photonic devices.

A high-speed Mach-Zehnder Interferometer (MZI) switch is a key component for the system. The methods for implementing control of the switch in a networked system have demonstrated in previous work using a FPGA testbed [2]. As seen in Fig. 2(a), cascading of several 2×2 MZI switches allows scalability of the one-to-many connection scheme. The basic 2x2 MZI silicon photonic switch evaluated for our scalability analysis was fabricated through OpSIS [3]. Both thermal heaters and p-i-n junctions were integrated on each MZI arm to control the optical phase as shown in Fig. 2(b).

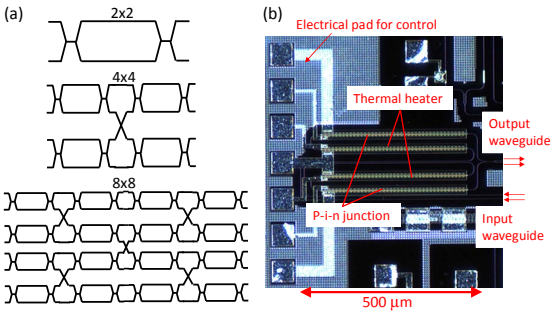

Fig. 2. (a) Structures of cascaded MZI switches. (b) A photograph of the 2×2 silicon photonic switch utilized in this measurement. A thermal heater and a built-in p-i-n junction are fabricated on each arm.

Fig. 3 shows switching characteristics using the p-i-n junction. A 0.64-Vpp PRBS signal with 0.94-V forward bias was applied to one arm of the MZI. Measured 10-90% rise and fall times were 2.2 and 3 ns, respectively.

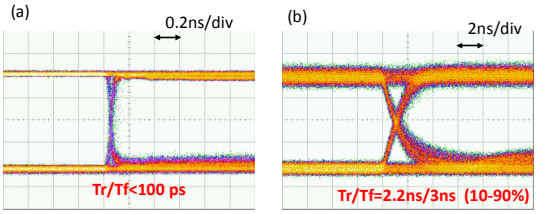

Fig. 3. (a) Electrical input voltage to the p-i-n junction of the switch. (b) Optical output of the switch.

Our measurement of the optical loss was focused on the free-carrier absorption (FCA) loss, a critical parameter for switch scalability. To reduce the FCA loss, we adopted a control method using both p-i-n junctions

on each MZI stage. With the initial phase position adjusted to a quadrature state using the thermal heater. Fig. 4(a) shows a reference control method using one p-i-n junction with an initially off state. Fig. 4(b) illustrates the switch condition using both p-i-n junctions. Fig. 4(c) shows the normalized output power of both control conditions. The measured FCA loss was reduced to be 0.5 dB in condition (b) in comparison to a FCA loss of 1.1 dB in condition (a).

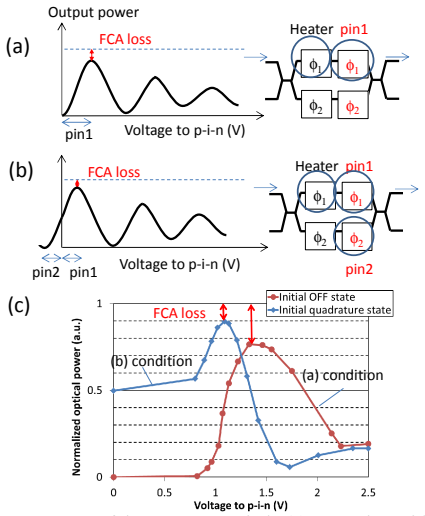

Fig. 4. Output powers of the 2x2 MZI switch. (a) Operating with a single p-i-n junction. Initial phase is adjusted to OFF state. (b) Operating with two p-i-n junctions. Initial phase is adjusted to quadrature state. (c) Measured FCA loss with different operating conditions.

III. LINK ANALYSIS & SCALABILITY

We analyzed the bit-rate dependent optical power budget of the silicon photonic link illustrated in Fig.1. Each device in the design was optimized to maximize the optical power budget at each given bit-rate. The design includes a depletion-type microring modulator with varying doping density and operating wavelength, and a microring demultiplexer with an optimal *Q* for the given bit-rate. Receiver sensitivity was calculated using noise parameters extracted from the leading reported technology on germanium photodiodes and CMOS receiver circuits. Fig. 5 shows the calculated result of bit-rate dependent power budget for a single wavelength. The power budget decreases rapidly beyond 30 Gb/s due to increased crosstalk and receiver noise.

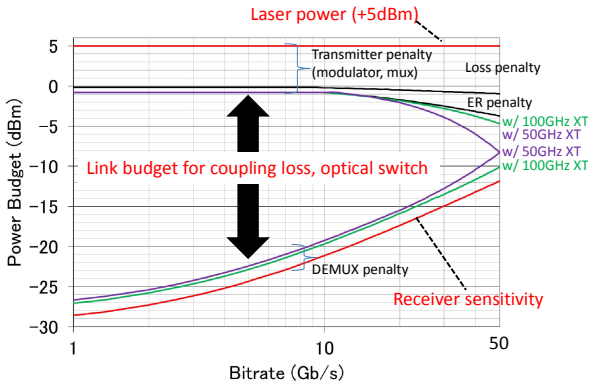

Fig. 5. Silicon photonic link power budget calculated from model of cascaded ring modulator, cascaded ring demultiplexer and receiver with germanium photodiode.

We then calculated scalability of an OCM system using this link budget. We assumed time-slotted messages to calculate the aggregate bandwidth (as described in [4]). Fig. 6 shows the aggregate bandwidth versus the number of connected memory. Fixed-size 256-Byte data and 5-ns guard time packet characteristics were used in the analysis. 16- and 32-channel WDM signals were used in 100-GHz and 50-GHz channel grid, respectively. The bit-rates of the wavelengths were adjusted to maximize the bandwidth. Higher bit-rate signals were able to be used when there were fewer connected memories because of the larger optical power budget.

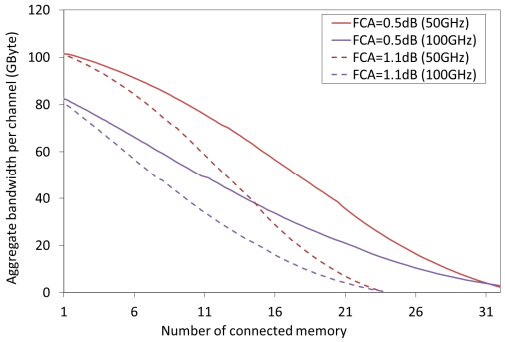

Fig. 6. Single channel bandwidth vs channel capacity. 16 channel WDM for 100-GHz grid 32 channel for 50 GHz.

The optical loss of the switch was calculated by using two FCA losses. As the number of memories increase, the FCA loss strongly influences the aggregate bandwidth. In the case of 16 connected memories, the aggregate bandwidth was improved to nearly double when assuming a 0.5-dB FCA loss.

The figure also reveals the greater scalability of the silicon photonic link compared to current electrically connected memory. For example, conventional 12.8-GByte DDR3 architecture enables having only 3 DIMMs per channel [5].

REFERENCES

[1] M. Madarbux et al., "Low latency scheduling Algorithum for shared memory communications over optical networks", 21st Hot interconnect, pp.83-86, (2013).

[2] T. Shiraishi et al., "A Reconfigurable and Redundant Optically-Connected Memory System using a Silicon Photonic Switch ", OFC2014, Th2A.10, (2014).

[3] http://opsisfoundry.org/

[4] A. Shacham et al., "Building ultralow-latency interconnection networks using photonic integration," IEEE Micro, 27(4), (2007).

[5] A DELL Technical white paper, "Dell™ PowerEdge™ Servers 2009Memory",http://www.dell.com/downloads/global/products/pedge/en/server-poweredge-11g-white-paper.pdf.