Nanophotonic Optical Interconnection Network Architecture for On-Chip and Off-Chip Communications

Howard Wang[1], Michele Petracca[2], Aleksandr Biberman[1], Benjamin G. Lee[1], Luca P. Carloni[2], and Keren Bergman[1]

1: Department of Electrical Engineering 2: Department of Computer Science
Columbia University, 500 W. 120th St., New York, New York 10027
howard@ee.columbia.edu

Abstract: An architecture for an integrated low-power, high-bandwidth optical interconnection network based on microring resonator technology is presented. The layout of the non-blocking network is described and a simulation-based performance evaluation is conducted.

OCIS Codes: (060.4250) Networks; (200.4650) Optical interconnects

1. Introduction

The recent trend toward chip multiprocessors (CMPs) in modern high-performance computing and communication systems has highlighted the need for a low-latency, high-bandwidth networking infrastructure. Switched on-chip global networks have been proposed as an attractive solution to mitigate the emerging communications bottleneck apparent in current generation electronic cross-chip point-to-point interconnects [1]. However, an unmanageable amount of power will necessarily be consumed by electronic interconnect solutions designed to meet the throughput demands of future CMP based systems [2]. Therefore, as power dissipation continues to be the ultimate limiting factor in contemporary designs, *performance-per-watt* becomes the defining figure of merit for the next generation of high-performance architectures.

The need for an interconnect solution providing low latencies and high bandwidths while maintaining minimal power dissipation can be met by leveraging the unique advantages provided by optical interconnection networks. Due to the inherently high bandwidth-distance product of photonic communication links and the bit rate transparency of optical switching elements, power scaling can essentially be decoupled from link distance and interconnect bandwidth, respectively. The possibility of a nanophotonic Network-on-Chip (NoC) would dramatically disrupt the current trend of power scaling in multiprocessor architectures by offering significant power savings in both global on-chip and off-chip communications over comparable electronic networks [2].

Recent advancements in nanoscale silicon photonics bring the integration of complete, functional photonic systems integrated using standard VLSI technologies closer to fruition [3]. Even so, the implementation of an all-optical NoC is intractable due to the immaturity of processing and buffering in the optical domain. Hybrid architectures have been introduced which circumvent these disadvantages via the utilization of a low-bandwidth topologically equivalent electronic network in tandem with a high-bandwidth photonic network [4].

The electronic network is used to carry small-size control (and data) packets while the photonic network transfers large-size data messages between pairs of cores as well as switches messages off-chip to external nodes and memory. The photonic messages are transmitted without buffering once the path has been acquired. This approach can be seen as a form of optical circuit switching: the established paths are, in essence, optical circuits (or transparent lightpaths) between cores enabling low-power, low-latency, and high-bandwidth communications. In this work, we introduce a novel 4×4 non-blocking switch element design based on ring-resonator building blocks and leverage the increased functionality of the switching node to develop a strictly non-blocking network architecture. The envisioned layout of this network is developed with a discussion of its physical layer considerations. Finally, we present a performance evaluation of the proposed network architecture under various traffic conditions.

2. Nanophotonic Switch Architecture

The physical limitations imposed by integration necessitate a compact, low-loss nanophotonic switch design, while maintaining a maximum level of functionality and flexibility for application in photonic on-chip interconnection network architectures. Microring resonator devices based on the silicon-on-insulator (SOI) platform have been demonstrated as high-speed optical modulators achieving data rates of up to 12.5 Gb/s [6]. More recently, the capability of similar ring-resonator structures to switch broadband multi-wavelength signals with sub-ns speeds have been reported [7]. The small footprint of these devices and their capacity for processing high-bandwidth wavelength division multiplexed (WDM) data in a highly power efficient manner make them ideal candidates as the fundamental switching element utilized in a photonic NoC.

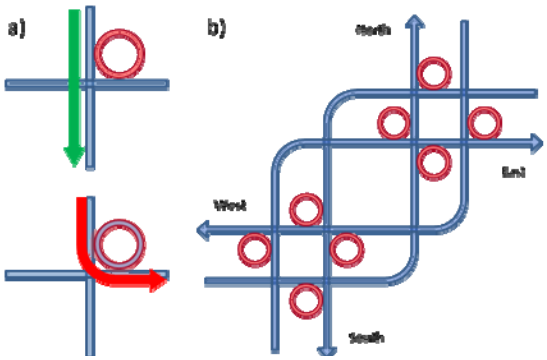

Fig. 1: (a) Atomic switching element depicting message propagation in two possible states. (b) 4×4 non-blocking nanophotonic switching node.

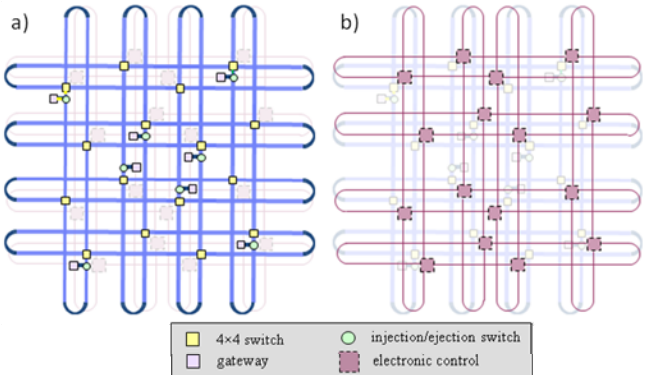

Fig. 2: (a) 8-port photonic NoC organization. (b) The associated electronic network

The atomic switching element utilized in the design of the nanophotonic switch and its two modes of operation are shown schematically in Fig 1a. The structure consists of a microring resonator positioned adjacent to a waveguide intersection. Switching is achieved through resonance modulation via carrier injection into the ring. Using the fundamental switching element introduced above, a 4×4 non-blocking nanophotonic switching node can be constructed (Fig. 1b) [5]. The switch design employs the minimum number of ring resonator devices (8) needed to provide the desired switching functionality while avoiding internal message blocking via the provisioning of exactly one ring per necessary switched path and the existence of non-overlapping internal paths from each input to every other output port of the node. This switch guarantees that a photonic message traversing it will never be blocked, assuming that there are no output contentions (ie. no two packets contending for the same output) and that the source and destination ports do not coincide. The switch maintains a compact and symmetrical structure yielding the ability to tile multiple nodes across a processor die.

4. Network Architecture and Layout

By maximizing the functionality of the atomic switching node, simplifications in the physical layout, network design, and routing algorithms are attainable. Utilizing the non-blocking 4×4 switch, we can now design a strictly non-blocking mesh-based network architecture for the photonic NoC. Figure 2 depicts the organization of the switching nodes and gateways for an 8-port system. Once again, a mesh topology is chosen due to the planar distribution of the network nodes, which aligns well with the tiled layout of current CMP architectures. Nodes are organized such that, at most, two nodes will occupy a given column or row in the network. Message routing though the network conforms to the simple XY dimension order minimum-distance algorithm. A message is injected into the network through gateway switches positioned adjacent to the switching nodes of the network and routed through the associated row dimension until it reaches the correct output column. From there, the message traverses the column switches until the desired gateway switch is reached, from which the packet is ejected from the network. The distribution of the gateway nodes ensures that, at most, two messages are transmitted along a dimension at any instant. Given the bidirectionality of the paths in the network, this guarantees that, other than output contentions, messages traversing the network will not block the paths of any other message.

The integrability of the proposed non-blocking photonic mesh network in an eight node implementation is

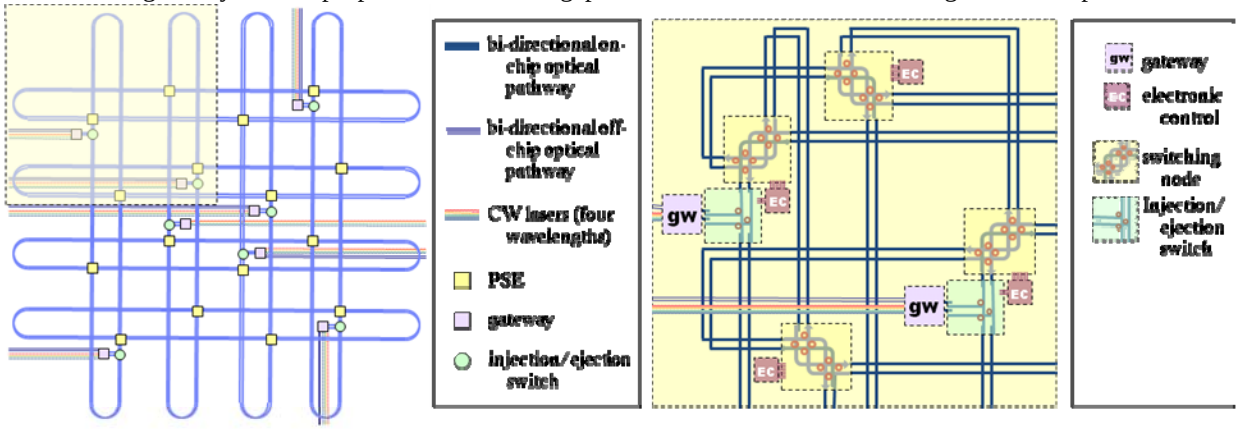

Fig. 2: Detailed layout of highlighted slice in the 8-port non-blocking photonic NoC

detailed in Figure 2. A detailed layout diagram of a slice of the proposed network, contained in the highlighted square, is depicted. The gateways contain the hardware necessary to interface with the photonic network, including the optical modulators and photodetectors. Externally generated continuous wave light is delivered to the optical modulators contained within each gateway. Each gateway also provides a high-bandwidth optical interface to off-chip devices via broadband waveguides extending to the periphery of the chip. The injection/ejection switches (highlighted in green) allow for reduced functionality and are thus less complex than the network switches.

5. Performance Analysis

A brief performance analysis of the non-blocking mesh architecture in comparison to an integrated crossbar network and the previously proposed conventional mesh architecture of [5] is conducted to demonstrate the advantages of the presented network. Network performance is analyzed via simulation by evaluation of the *latency* and *throughput* vs. *offered load*. These metrics are functions of the path-setup overhead, which was determined to be the largest contributor to performance degradation in previous analysis [5]. We vary the *offered load* to the network, defined as the ratio of time a core is ready to transmit a message over the total simulation time. In this simulation, we assume 36 computing cores transmitting packets 16 kB in size over a photonic network achieving a line rate of 960 Gbps, corresponding to a message duration of 134 ns. The results are shown in Figure 4.

We observe that the non-blocking mesh achieves superior performance over that of the originally proposed blocking mesh architecture of [5]. The non-blocking behavior of the network guarantees connections between unoccupied nodes, eliminating contention-based latency degradation and thus reducing the overall path setup time. Although the overall network diameter of a non-blocking mesh is larger than that of a comparable blocking mesh network, the disadvantages are mitigated by the reduction of the number of gateway devices needed per compute node and the associated reduction in routing complexity.

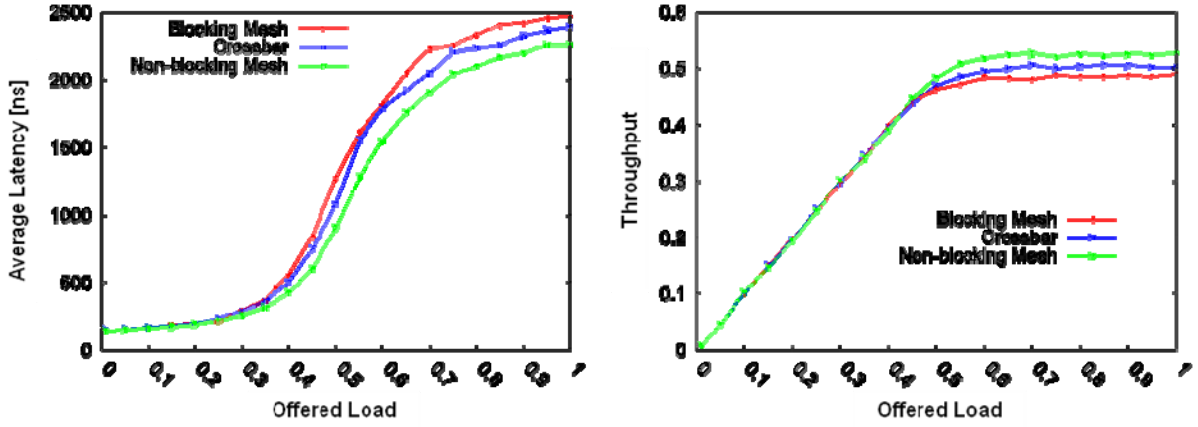

Fig. 4: Simulated latency and throughput performance evaluation for various 36-port network architectures.

6. Conclusions

We present a novel 4×4 non-blocking switch element based on microring resonator devices and introduce a strictly non-blocking network architecture based around the switching node design. A layout for the network is developed and a performance evaluation of the proposed network architecture under various traffic conditions is conducted.

We acknowledge support for this work from the NSF under grant CCF-0523771.

7. References

[1] W.J. Dally and B. Towles. “Route Packets, Not Wires: On-Chip Interconnection Networks,” *Proc. of the Design Automation Conf.*, pages 684-689, Jun 2001.

[2] A. Shacham, K. Bergman, L. Carloni. “The Case for Low-Power Photonic Networks-on-Chip,” *Proc. of the Design Automation Conf.*, pages 132-135, Jun 2007.

[3] C. Gunn. “CMOS Photonics for High-Speed Interconnects,” *IEEE Micro*, 26(2):58-66, Mar/Apr 2006.

[4] A. Shacham, K. Bergman, L. Carloni. “On the Design of a Photonic Network-on-Chip,” *Proc. of the First Intl. Symp. On Networks-on-Chips (NOCS)*, May 2007.

[5] A. Shacham, B.G. Lee, A. Biberman, K. Bergman, L.P. Carloni. “Photonic NoC for DMA Communications in Chip Multiprocessors,” *IEEE Symp. On High-Performance Interconnects*, Aug 2007.

[6] Q. Xu, S. Manipatruni, B. Schmidt, J. Shakya, M. Lipson. “12.5 Gbit/s Carrier-Injection-Based Silion Microring Silicon Modulators. *Optics Express*, 15(2):430–436, 22 Jan 2007.

[7] A. Biberman, P. Dong, B.G. Lee, J.D. Foster, M. Lipson, K. Bergman, "Silicon Microring Resonator-Based Broadband Comb Switch for Wavelength-Parallel Message Routing," *LEOS 2007* WG3 (Oct 2007).