Architectural Design Exploration of Chip-Scale Photonic Interconnection Networks Using Physical-Layer Analysis

Johnnie Chan, Gilbert Hendry, Aleksandr Biberman, and Keren Bergman
Department of Electrical Engineering, Columbia University, 500 West 120th Street, New York, New York 10027
johnnie@ee.columbia.edu

Abstract: We conduct an architectural exploration of three chip-scale photonic interconnection networks in a novel simulation environment, exploring insertion loss, crosstalk, and energy. The impact of these metrics is evaluated in the context of network performance.

OCIS codes: (130.6750) Integrated Optics, Systems; (200.4650) Optics in Computing, Optical Interconnects

1. Introduction

Recent advancements in silicon nano-photonic technology have opened the possibility of integrating photonics for chip-scale interconnection networks. In comparison to electronics, photonics has the potential to offer higher-bandwidth connections by leveraging data parallelism offered by wavelength-division-multiplexing (WDM). Additionally, energy dissipated in photonic signaling is effectively distance independent, enabling greater energy efficiency for global chip- and board-scale communications. Although these advantages exist, optical signaling is incapable of practical in-flight processing or buffering without optical-electronic-optical conversion. Also, signal regeneration in optics cannot be easily accomplished on the CMOS-compatible silicon photonic platform. Photonic messages must therefore be able to propagate the length of the transmission path without accumulating significant optical loss. In light of these constraints, many novel photonic interconnect designs have been proposed for enabling optical data transmission in the chip-scale domain [1–3].

We report, to the best of our knowledge, the first detailed physical-layer analysis of chip-scale photonic interconnection networks. Although many photonic topologies have been proposed in an effort to improve computing performance, less emphasis has been placed on understanding whether such designs are feasible from a physical-layer standpoint. Since it is not currently practical to test full network topologies in a laboratory environment due to fabrication yield limitations, we implement physically-accurate simulation models for this analysis. We model the previously proposed *Torus* topology [3], and introduce two new topologies, *TorusNX* and *Square Root*. These three networks are analyzed in terms of three physical-layer metrics (that play a critical role in determining the overall scalability and performance of the network design: insertion loss, crosstalk, and energy.

2. Chip-Scale Photonic Networks

We investigate *space-routed* photonic networks, which are designed to use actively-controlled silicon ring-resonator-based broadband switches to route WDM messages, composed of a set of wavelength channels, from source to destination. The ring resonators are electro-optic devices that can be manipulated to be in an off-resonance through state allowing signals to pass by (Fig. 1a), or in an on-resonance drop state which shifts the signal onto another waveguide (Fig. 1b). An electronic control plane, mirroring the photonic network layout, is necessary to control each broadband switch through a circuit-switching protocol. When a photonic connection is being provisioned, a path-setup message on the control plane will trace out an optical path on the photonic plane by reserving and configuring the appropriate optical switches. This form of routing can fully utilize the optical spectrum by leveraging WDM to create extremely high-throughput links. This method contrasts with *wavelength-routed* networks, which leverage filters and wavelength selectivity to perform routing.

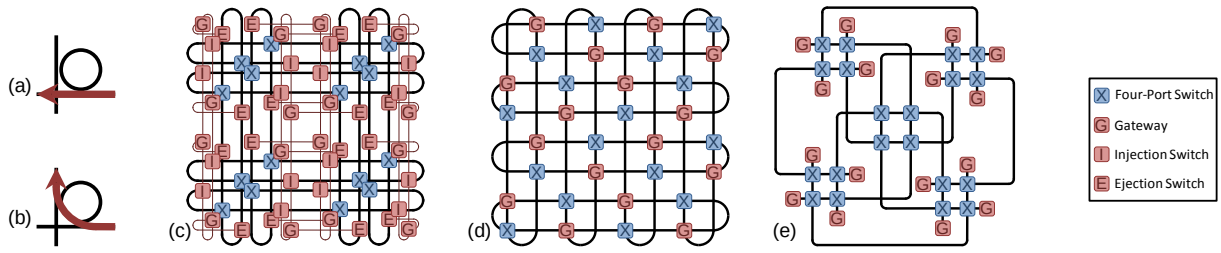

Fig. 1. The broadband ring switchs are essential parts of chip-scale photonic networks. Ring switches control the flow of optical signals by being in either an (a) off-resonance or (b) on-resonance state. Topologies that use this switch include (c) Torus, (d) TorusNX, and (e) Square Root.

OThX4.pdf

Table 1. Simulation Insertion Loss Parameters

Parameter	Value
Propagation Loss in Silicon [6]	1.5 dB/cm
Crossing [7]	0.15 dB
Waveguide Bend [6]	0.005 dB
Ring Through Port [8]	0.005 dB
Ring Drop Port [8]	0.5 dB

Table 2. Photonic Energy Dissipation Parameters

Parameter	Value
Modulators (Dynamic)	85 fJ/bit
Modulators (Static)	30 µW
Photodetectors	50 fJ/bit
Broadband Ring Switch (Dynamic)	375 fJ/bit
Broadband Ring Switch (Static)	400 µW

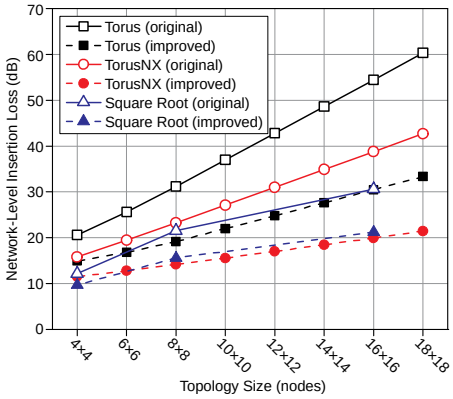

Fig. 2. Insertion loss performance of the photonic topologies.

Fig. 1c, 1d, and 1e illustrate the three networks in 4×4 configurations (*X*×*Y* expresses the number of nodes in the *X* and *Y* dimension). The blocks labeled 'G' denote gateways, locations on each node where a network user (e.g. processor core) can initiate or receive data transmissions. The thicker lines represent two waveguides used for bi-directional data transmission and the blocks marked 'X' represent four-port non-blocking photonic switches which are composed of ring-resonators and are used to effectively route data through the network [9]. Together, they form the main network through which data is routed. The Torus (Fig. 1c) requires an additional access network, represented by thinner lines (additional waveguides) and the blocks denoted by 'I' (injection) and 'E' (ejection) to facilitate entering and exiting the main network [3]. TorusNX (Fig. 1d) improves the Torus topology by introducing new a switch design that eliminates the need for the access network and directly integrates the gateway into the main topology. Square Root (Fig. 1e) is an alternative hierarchical topology optimized to reduce the required number of waveguide crossings and switching points. Due to the recursive nature of constructing the Square Root, the number of nodes along the *X* and *Y* dimension of the topology must be equal and a positive integer power of two (i.e. 2, 4, 8, 16, …). TorusNX and Square Root were both designed in response to preliminary physical-layer shortcomings of the Torus, since insertion losses due to waveguide crossings and the large number of switches have a dramatic impact on system performance.

3. Physical-Layer Simulation and Results

We simulate the networks using *PhoenixSim*, a physical-layer simulator developed in the OMNeT++ simulation environment [4], which incorporates detailed physical models of basic photonic building blocks such as waveguides, modulators, photodetectors, and switches. More complex photonic circuits and full topologies can be created by properly arranging these building blocks. These composite structures can then be analyzed within the simulator to determine the overall performance characteristics. Electronic energy performance is based on the ORION router model [5]. Relevant parameters for insertion loss and energy are listed in Tables 1 and 2.

The maximum possible network-level insertion loss that an optical signal will incur within each of the three networks is plotted in Fig. 2. This plot shows the loss performance based on the realistic parameters listed in Table 1 (labeled 'original') and also the performance of a system with an hypothetically improved crossing loss of 0.05 dB to represent a potential way to improve overall performance (labeled 'improved'). In both the 'original' and 'improved' cases, the two new topologies achieve drastically better insertion losses than Torus due to better designs. For the 16×16 topologies with the 'original' parameters, TorusNX exhibits 15.7 dB lower loss and Square Root exhibits 23.9 dB lower loss. For the same size topology, the improved crossing loss results in a 44% improvement in the Torus, 49% improvement in TorusNX, and 31% improvement in Square Root. This dramatic reduction in total loss through better crossings is indicative of the fact that crossing loss is a major contributor, and further development of this building block has a significant impact on the scalability of the network.

We relate the insertion loss to system scalability by considering the optical power budget, which measures the amount of loss that can be incurred before a signal can no longer be received properly. Since signals will incur higher losses as the photonic network scales to support more gateways, there will be a maximum size allowed by the budget. The optical power budget is calculated from the maximum allowable signal power in the network and the photodetector sensitivity. With a 30 dB budget in the 'original' case, the maximum size networks are a 6×6 Torus (36 nodes), 10×10 TorusNX (100 nodes), and 8×8 Square Root (64 nodes). With improved crossings, Torus can achieve a 14×14 topology (more than 4-fold increase in nodes). TorusNX and Square Root are each capable of operating at sizes of 18×18 and 16×16, respectively (these were the largest networks simulated in this work).

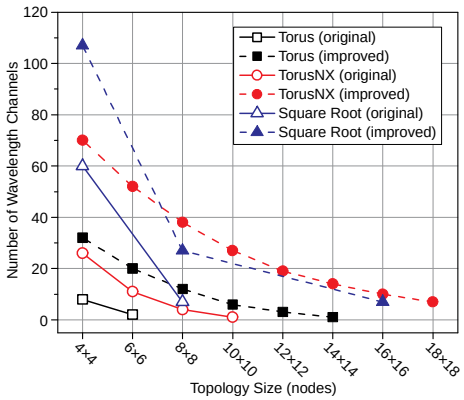

Fig. 3. Maximum possible number of wavelengths for each topology.

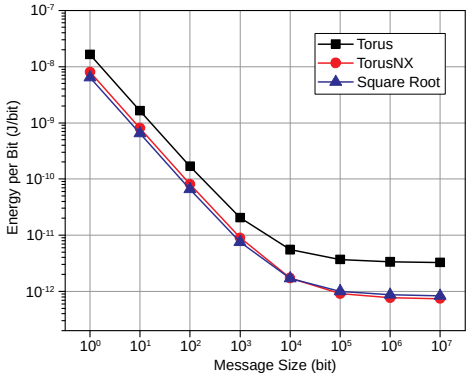

Fig. 4. Transmission efficiency of the photonic topologies.

Throughput is directly proportional to the number of wavelengths used in the WDM signal, and is also affected by the insertion loss and optical power budget. The expression $P_{budget} \geq IL_{max} + 10log_{10}n$ relates these three parameters where P_{budget} is the optical power budget, IL_{max} is the worst-case network insertion loss, and n is the number of wavelengths being used. P_{budget} and IL_{max} are expressed in decibel units. The $10log_{10}n$ accounts for each of the discrete signals in the WDM message which must transmit an equal fraction of the power limit that was defined for the optical budget. This relationship of scalability and link throughput is shown in Fig. 3.

Network energy performances of the 8×8 topologies are shown in Fig. 4, based on the photonic parameters listed in Table 2. These results assume networks configured to use the maximum number of wavelengths assuming a 30-dB budget and each operating at a line rate of 10 GHz, and operate in saturation using uniform random traffic. For the largest message size, Torus achieves an efficiency of about 3.26 pJ/bit, while TorusNX and Square Root achieve efficiency of 743 fJ/bit and 831 fJ/bit. This dramatic improvement is attributed to the lower-loss network designs which enable better bandwidth utilization and reductions in the number of required switches. While from an efficiency standpoint, larger message transmissions clearly perform better, crosstalk simulations indicates that the OSNR also decreases with increased message size. For message sizes larger than 10^5 bits, Square Root has the best crosstalk performance at about 16 dB, while Torus and TorusNX have OSNRs of about 11 dB and 12 dB, respectively. An ideal optimal binary receiver with orthogonal signaling can achieve a bit-error-rate of 10^{-12} with an OSNR of 16.9 dB, which is above all three networks. This indicates that in order to maintain the high energy efficiency that photonics provides, a scheme must be in place to either correct or mitigate these errors.

6. Conclusions

We have reported the architectural implications of physical-layer metrics for three chip-scale photonic networks in simulation and derived their impact on system performance. Insertion loss, crosstalk, and energy are shown to have a profound impact by affecting scalability, performance, and efficiency of the network.

This work was supported in part by the Interconnect Focus Center, one of five research centers funded under the Focus Center Research Program, a Semiconductor Research Corporation and DARPA program, and by DARPA MTO under a subcontract with IBM (prime contract HR0011-08-C-0102). The views, opinions, and/or findings contained in this article/presentation are those of the author/presenter and should not be interpreted as representing the official views or policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the Department of Defense.

8. References

[1] C. Batten, et al., “Building many-core processor-to-DRAM networks with monolithic CMOS silicon photonics,” IEEE Micro **29 (**4) 8-21 (2009).

[2] D. Vantrease, et al., “Corona: System implications of emerging nanophotonic technology,” in ISCA ’08: Proc. Of the 35th International Symposium on Computer Architecture, pp. 153-164 (2008).

[3] A. Shacham, K. Bergman, and L. P. Carloni, “Photonic networks-on-chip for future generations of chip multiprocessors,” IEEE Trans. on Computers **57** (9) 1246-1260 (September 2008).

[4] A. Varga, "OMNeT++ Discrete Event Simulation System," available at http://www.omnetpp.org.

[5] A. Kahng, B. Li, L. Peh and K. Samadi, "ORION 2.0: A Fast and Accurate NoC Power and Area Model for Early-Stage Design Space Exploration" in DATE: Proc. of Design Automation and Test in Europe (2009).

[6] F. Xia, L. Sekaric, Y. Vlasov, "Ultracompact optical buffers on a silicon chip," Nature Photonics **1**, 65-71 (2006).

[7] W. Bogaerts, P. Dumon, D. V. Thourhout, R. Baets, "Low-loss, low-cross-talk crossings for silicon-on-insulator nanophotonic waveguides," Opt. Lett. **32** (19) 2801-2803 (2007).

[8] B. G. Lee, A. Biberman, P. Dong, M. Lipson, K. Bergman, “All-optical comb switch for multiwavelength message routing in silicon photonic networks,” IEEE Photon. Technol. Lett. **20**, 767-769 (2008).

[9] H. Wang, M. Petracca, A. Biberman, B. G. Lee, L. P. Carloni, and K. Bergman, "Nanophotonic Optical Interconnection Network Architecture for On-Chip and Off-Chip Communications," in OFC ’08: Optical Fiber Communication Conference, JThA92 (2008).