Insertion Loss Analysis in a Photonic Interconnection Network for On-Chip and Off-Chip Communications

Johnnie Chan, Aleksandr Biberman, Benjamin G. Lee, and Keren Bergman
Department of Electrical Engineering, Columbia University, New York, NY 10027
johnnie@ee.columbia.edu

***Abstract*—An on-chip photonic interconnection network is simulated to determine statistical insertion losses for different network sizes and non-blocking switch layouts. For an 8×8 folded-torus network, we obtain an optical link loss budget of 15.5 dB.**

INTRODUCTION

Networks-on-chip (NoC) have taken on an increasingly important role in the overall performance of computing systems. As the number of cores grows in chip multiprocessor designs, the supporting interconnection network is challenged by the increasing bandwidth it must support for inter-node communications. The significance of the NoC is exemplified by the fact that designers are now scaling up the numbers of cores to increase overall performance as uniprocessor design has reached the limits of clock frequency scaling and enhancements from instruction level parallelism.

Photonics has become a viable solution to a variety of problems facing modern computer chip design [1]. Foremost, photonics provide a high bandwidth, low latency medium for interconnection networks. The use of photonics on-chip also has the potential to be used as an interface to I/O and off-chip memory [2]. In addition to providing low latency and high bandwidths, optical interconnects also have the potential to deliver better power efficiency for the communications infrastructure. The power dissipated in an electronic bus is largely dependent on the data rate. The electronic bus consumes a significant portion of the power budget and is becoming a limiting factor in the drive towards increasing communication bandwidth. Optics has the potential to address this critical constraint and provide a scalable path toward power efficient high-bandwidth bit-rate transparent data communications.

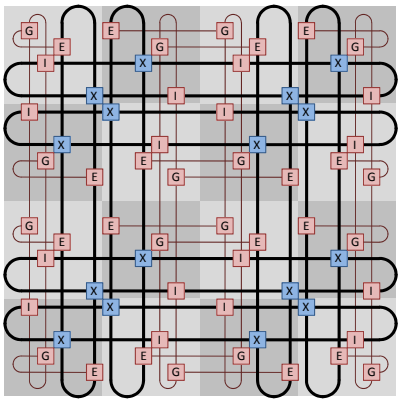

Fig. 1 Structure of a 4×4 node photonic layer NoC. The waveguides that make up the torus network are shown as thick lines, and the gateway access network for injecting packets to and ejecting packets from the network shown as thin lines. The blocks represent the following: gateway switch (G), injection switch (I), ejection switch (E), and a 4×4 non-blocking switch (X).

Previous work in simulations of photonic on-chip interconnection networks has been done using the OMNeT++ programming environment which focuses on high level network issues such as latency, routing and flow control [3]. Here, we extend this body of work to include physical characterizations of the components composing the NoC to determine network-level insertion loss. The drive is to design low-power photonic networks that can transmit data across the chip, and even off-chip, without requiring power hungry amplifiers. These parameters also determine how well the network can scale in size, as a photonic packet must travel through more waveguides and switches in larger networks. As a result, the power budget available for light propagation needs to be taken into account in the network design. An analysis of this network-level link budget is presented here.

SYSTEM CONFIGURATION

The photonic NoC described in [3] consists of an optical interconnection network with electronic routing control for inter-node communications. The photonic network provides a high bandwidth interconnect between nodes. The modulators and switches throughout the network (Fig. 1) are designed using ring resonator based electro-optic devices [4]. Additional broadband switches are placed in the network to allow packets to enter (injection), route, and exit (ejection) the interconnection network topology. Photonic packets cannot be easily switched without electrical control. Therefore, a separate electronic plane provides the necessary functions to arbitrate a complete circuit switched optical path from source node to destination node. Once the optical path is setup, the high bandwidth, low latency path is available to the node. The electronic plane also has the secondary functionality of providing a low bandwidth packet switched network.

THE SIMULATOR

The simulation consists of three planes: the photonic plane, the electronic plane, and the processing element plane. The photonic plane simulates the light path mechanisms. The electronic plane manages all of the low level network control and arbitration. The processing element plane represents the sources of data that communicate on the network.

The photonic plane is constructed from a set of optical building blocks which include a basic 2×2 and 1×2 photonic switching element, straight waveguide, 90° bend waveguide, and a waveguide crossing. All higher order functional components are derived from this set of elements. The assumed insertion loss parameters are listed in the table below. The values are obtained from reported devices and predictions for future scaling.

Component	**Insertion Loss Parameter**
Propagation Loss	1.5 dB/cm [5]
Waveguide Crossing	0.05 dB [6]
Waveguide Bend (radius = 2.5 μm)	0.005 dB [5]
Passing a Ring	0.005 dB [4]
Coupling with a Ring	0.5 dB [4]

978-1-4244-1932-6/08/$25.00 ©2008 IEEE

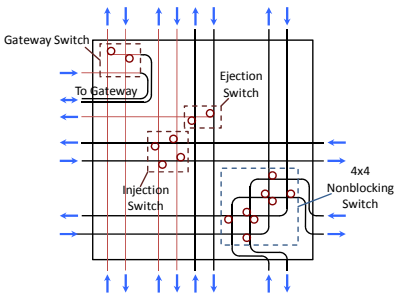

Fig. 2 Layout of a tile in the torus network.

An important issue not considered in previous work is the spatial layout of the optical components. The layout can significantly affect the available power budget of each optical signal. The simulator assumes a tile size of 2.0 mm × 1.5 mm, the size of a single core in Intel's 80-core chip [7]. Fig. 2 shows a typical layout of a tile in the photonic plane. Each tile consists of a gateway switch, injection switch, ejection switch, 4×4 non-blocking switch, and several optical paths to form the torus and gateway access network.

The network uses a folded-torus topology and X-Y dimensional ordered routing. The simulation uses uniformly distributed generated transmission requests with exponentially distributed interpacket spacing. Although insertion loss is independent of network congestion, an arbitrary constant message length of 50 ns, equivalent to an 8 kb size packet at 160 Gb/s [4], is used. Layout differences and losses due to the on-chip routing of continuous-wave light and off-chip messages into each gateway are ignore for this simulation.

SIMULATION RESULTS

The analysis here investigates how insertion loss is affected by changes in topology size and different switch designs. Tori of size 4×4, 6×6, and 8×8 are considered. The different switch designs, described later, are labeled (A), (B), and (C) (Fig. 3).

Fig. 4 shows the distribution of insertion loss that a packet will experience when propagating from source to destination. Minimum losses for each switch layout remains constant for differing network sizes. For every additional two nodes in each dimension, 3.89 dB, 3.66 dB, and 3.36 dB of loss is added to the maximum loss for switch design A, B, and C, respectively. This is a result of the fact that the minimum length path from any two nodes remains the same while the maximum length will change with number of nodes. The three, five, and seven peaks that appear in the distribution for the 4×4, 6×6, and 8×8 node torus networks, respectively, equates to the maximum number of 4×4 non-blocking switches an optical packet must travel through, which rises as the node count scales up.

Next, we explore the performance of three different 4×4 non-blocking switch designs (Fig. 3): (A) is a design first introduced in [8]. (B) contains a reduced number of waveguide crossings while keeping the number of ring resonator structures at eight. (C) differs from the previous two designs by allowing packets that require a straight path to propagate through without requiring

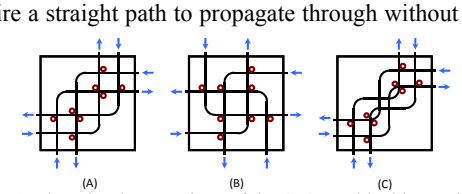

Fig. 3 Three implementations of the 4×4 non-blocking switch.

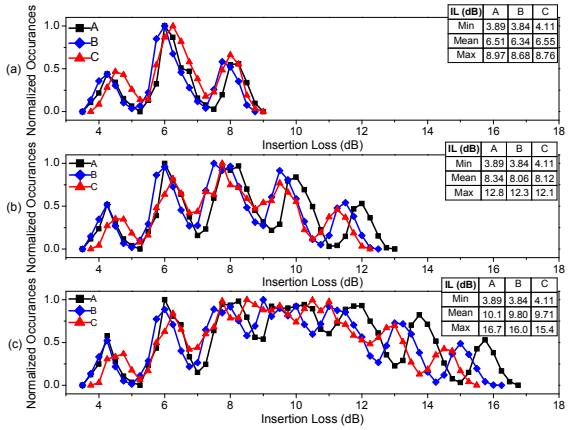

Fig. 4 Insertion loss distribution for folded torus topologies of size (a) 4×4, (b) 6×6, and (c) 8×8. Each graph contains plots of three differing switch designs. Inset within each graph is a table of minimum, mean and maximum insertion losses observed for each case.

a turn at a ring. Each design is non-blocking when no u-turns are allowed.

Each plot in Fig. 4 shows the general trend of the different switch designs. (B) has both a lower maximum and lower minimum loss, in comparison to (A), as expected. (C) consistently has a higher loss for the lower bound of the distribution. Although (C) exhibits higher maximum loss in the 4×4 node network than (B), it shows lower loss at sizes of 6×6 nodes and higher. This is attributed to the fact that even though the minimum insertion loss for this 4×4 switch design is higher than the others, the straight path (from north to south, east to west, or vice versa) has a lower loss because no rings are encountered. In contrast the paths in (A) and (B) that do not pass through any ring resonators implement a turn. The performance improvement noticed with switch (C) is a consequence of using dimensional ordered routing makes a single turn in any optical path, and mostly straight propagation through the switches.

CONCLUSIONS

The optical insertion losses through various photonic NoC configurations have been simulated. Additionally, the performance of three routing switch configurations was explored. Considerations such as these are crucial to the interconnection network design in order to optimize power efficiency.

The authors acknowledge partial support for this work from the Interconnect Focus Center, one of five research centers funded under the Focus Center Research Program, a Semiconductor Research Corporation and DARPA program.

REFERENCES

[1] A. Shacham, K. Bergman, L. P. Carloni, Networks-on-Chip Symposium (NOCS) 2007, paper 2.1 (May 2007).
[2] K. Bergman, L. P. Carloni, J. A. Kash, and Y. Vlasov, High Performance Embedded Computing (HPEC) 2007*, (Sept. 2007).*
[3] A. Shacham, K. Bergman, L. P. Carloni, *IEEE Trans. Computers*. (in press) **57** (9) (Sept. 2008).
[4] B. G. Lee, A. Biberman, P. Dong, M. Lipson, K. Bergman, *IEEE Photon. Technol. Lett.* **20** (10) 767–769 (May 15, 2008).
[5] F. Xia, L. Sekaric, Y. Vlasov, *Nature Photonics*, **1**, 65–71 (Jan. 2007).
[6] W. Bogaerts, P. Dumon, D. V. Thourhout, R. Baets, *Optics Letters*, **32** (19) 2801–2803 (Oct. 1, 2007).
[7] S. Vangal, *et al.*, International Solid-State Circuits Conference (ISSCC) 2007, 98–589, (Feb. 2007).
[8] A. Shacham, B. G. Lee, A. Biberman, K. Bergman, L. P. Carloni, Hot Interconnects (HOTI) 2007, 29–36 (Aug. 2007).