Limits and Opportunities for Designing Manycore Processor-to-Memory Networks using Monolithic Silicon Photonics

Ajay Joshi[*], Christopher Batten[†], Yong-Jin Kwon[‡], Scott Beamer[‡], Imran Shamim[†]
Krste Asanović[‡], Vladimir Stojanović[†]

[*] *Department of ECE, Boston University, Boston, MA*
[†] *Department of EECS, Massachusetts Institute of Technology, Cambridge, MA*
[‡] *Department of EECS, University of California, Berkeley, CA*

1. Introduction

To sustain the historic performance improvement in VLSI systems, while remaining within the power envelope, the trend has moved towards designing multiple cores on a single die. However, if designed using current and/or projected electrical solutions, these systems would quickly get bandwidth-limited due to bandwith density limitations and power constraints. It is therefore necessary to explore alternate interconnect technologies like silicon photonics that could provide high bandwidth density and energy-efficient data transmission. Here, we summarize the results from our study [1] to determine the limits and opportunities for using silicon photonic technology for designing core-to-memory i.e. off-chip interconnect networks in manycore systems. We focus on the memory bandwidth aspect of the manycore design as this would be the primary bottleneck in extracting maximum performance out of a manycore system. There have been other approaches to designing both core-to-memory networks [2, 5] and core-to-core networks [3, 4, 6] using silicon photonics.

2. Electrical vs Photonic links

Figure 1 shows an example unified on-chip/off-chip photonic link that can be used for core-to-memory communication. Light waves from an external broadband laser source are coupled onto a chip using a vertical coupler, that guides the light waves into the waveguide. These light waves are then modualted using ring modulators that are driven by modulator drivers. Here, the data gets converted from electrical to optical medium. The modulated light waves travel to the other chip through a single-mode fiber, where the light waves are filtered using ring filters and are absorbed using a photodetector that converts light into current, which is then fed to the electrical receiver.

The two key advantages of photonic links is energy-efficient modulation and detection that is independent of the length of the wire, and high bandwidth density due to dense wavelength division multiplexing (DWDM). Assuming monolithic integration of photonic devices, for 22 nm technology, a unified on-chip/off-chip photonic link provides more than 20× advantage in terms of the data-dependent energy compared to the total energy spent in the on-chip and off-chip electrical links. However, the photonic links have a significant static energy cost in terms of energy in thermal tuning circuits and optical laser source, which can be much higher than the static energy cost of electrical links. Hence, networks designed using photonic links need to have high utilization to offset the large static energy overhead. In addition to potential energy savings, photonic links also provide 30× higher bandwidth density through DWDM.

3. Core-to-memory photonic network

For our case study we considered a 64-tile system with 16 groups that used a local meshes to global switches (LMGS) topology designed using a monolithically integrated unified on-chip/off-chip photonic link for core-to-memory communication. The static energy component of a photonic network is highly dependent on the physical layout of the waveguides and rings. Figure 2 and Figure 3 show the ring matrix and u-shaped layout of the waveguides and rings for implementing the LMGS network topology. Light waves traveling in one direction on the waveguides were used for transmitting data from the cores to memory, while light waves traveling in the opposite direction were used for transmitting data from memory to core. Along the optical path, the two main components of optical loss are waveguide loss and through ring loss. Figure 4 shows the waveguide loss and through loss limits for a desired optical power budget in the u-shaped and ring matrix layout. For a optical power budget of 2 W, a waveguide loss of less 0.5 dB/cm and 1.5 dB/cm is required for ring matrix and u-shaped layout, respectively, which is very challenging to achieve using current silicon photonic technology. Similarly, the through ring loss requirement of less than 0.01 dB/ring and 0.5 dB/ring is required for ring matrix and u-shaped layout respectively. In addition, the u-shaped layout has half the number of rings than that in the ring matrix, hence the thermal tuning energy required in u-shaped layout is half of that in the ring matrix. We therefore chose the u-shaped layout for our network.

A detailed cycle-level simulator was used to study the power and performance of the core-to-memory photonic network for the 64-tile system with 16 DRAM modules, running at 2.5 GHz and designed using 22 nm technology for a uniform random traffic pattern. We assumed a total power budget of 20 W for both on-chip and off-chip network. Figure 5 shows the plot of latency and power versus bandwidth for various electrical and photonic network configurations. We used on-chip network overprovisioning (increasing network link widths) and tile grouping to maximize the advantages of photonic links. From Figure 5(a), for pure electrical network, it can be observed that an increase in the number of groups increased the saturation throughput (≈3× for 16–group case). The latency also reduced due to the reduction in the on-chip hop count. The use of on-chip network

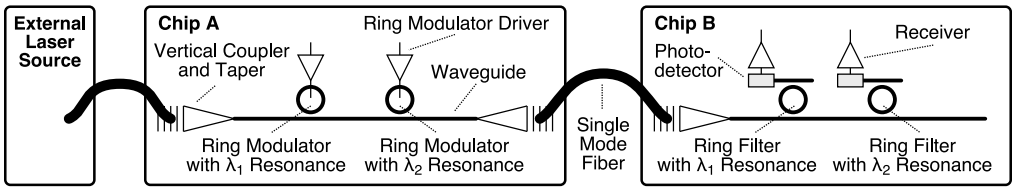

Figure 1: Unified on-chip/off-chip silicon photonic link.

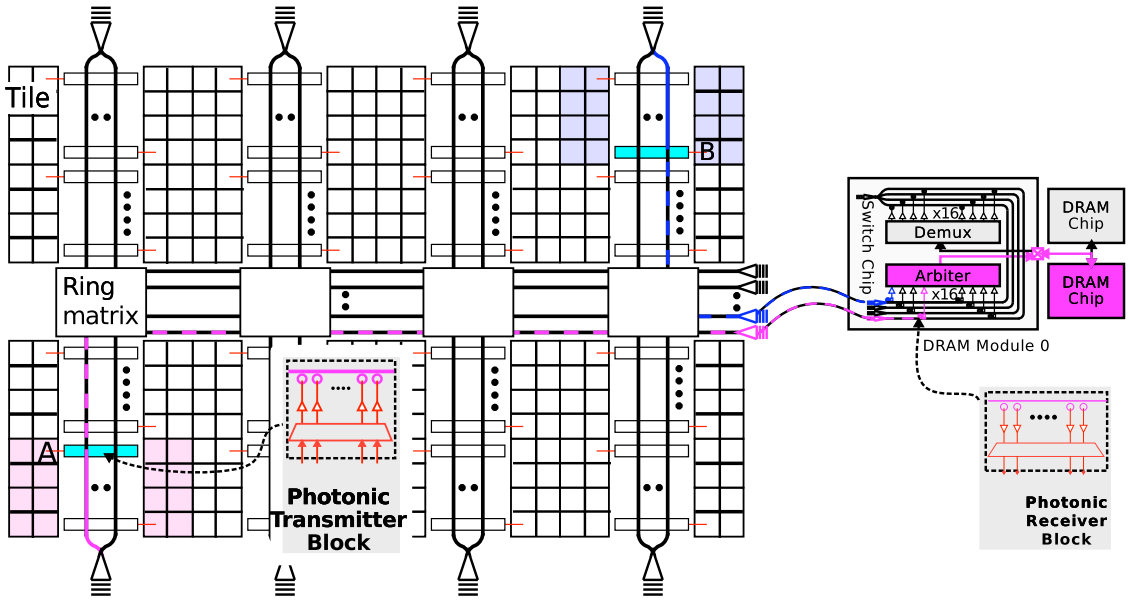

Figure 2: LMGS network design for a 64-tile system with 16 groups implemented using a ring matrix physical network. Two different requests to the same DRAM are highlighted.

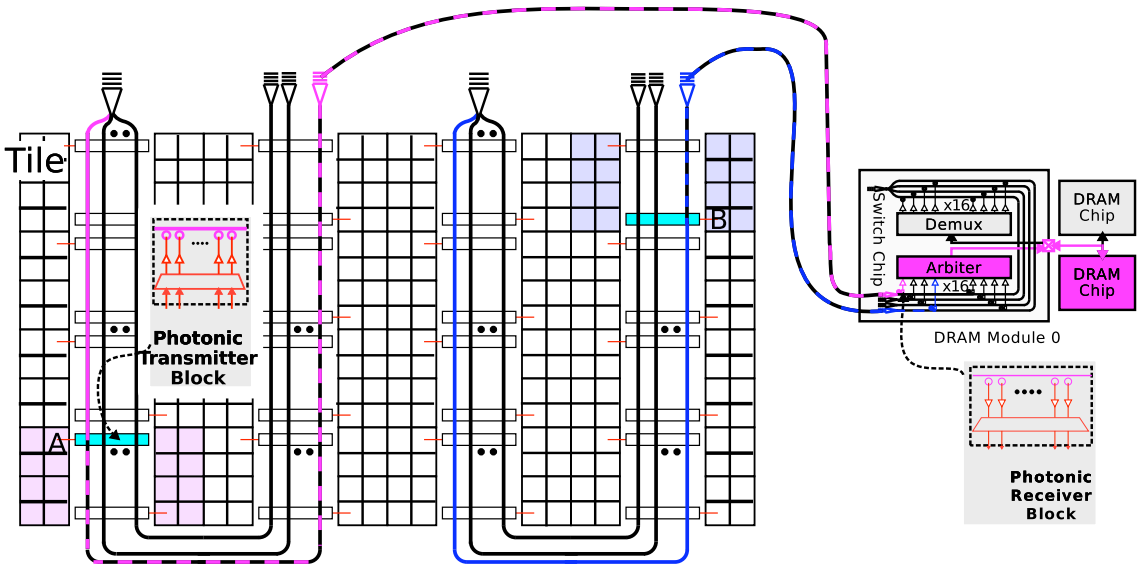

Figure 3: LMGS network design for a 64-tile system with 16 groups implemented using a u-shaped physical network. Two different requests to the same DRAM are highlighted.

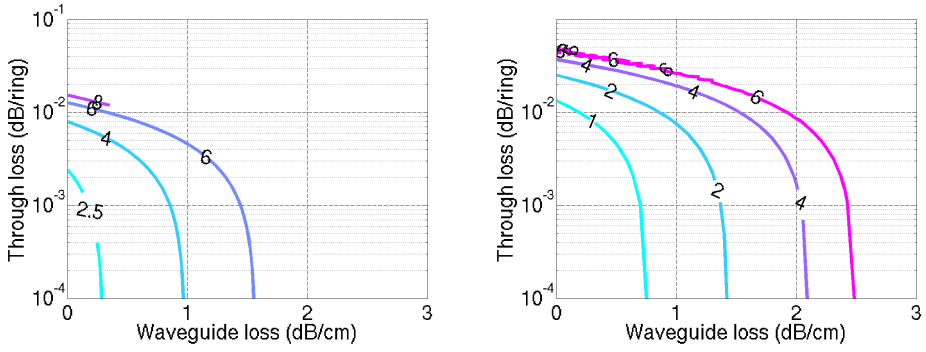

(a) Ring Matrix (b) U-shape

Figure 4: Optical power (in W) contour plot for the off-chip photonic LMGS network in a 64-tile system.

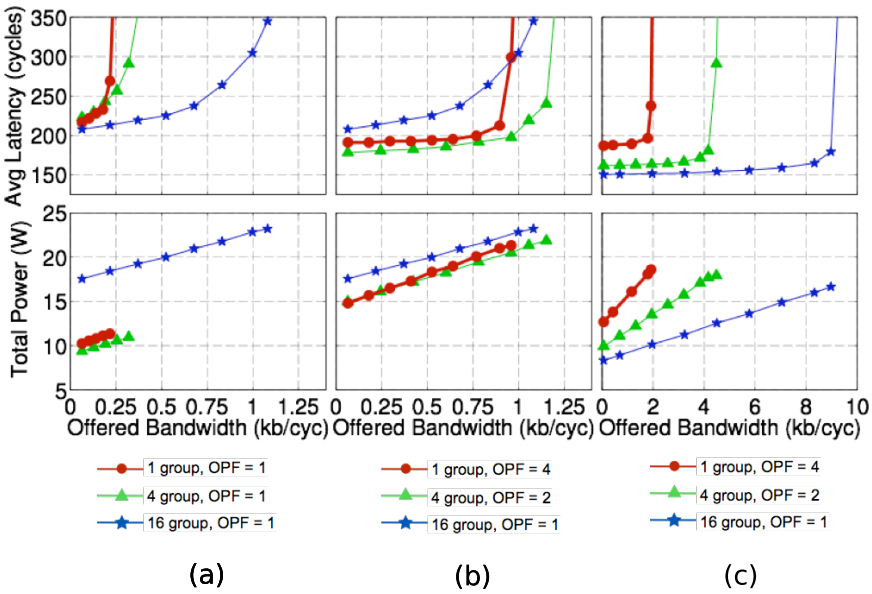

(a) (b) (c)

Figure 5: Simulated power and performance for various LMGS network configurations assuming (a) electrical link (on-chip and off-chip) with grouping, (b) electrical link (on-chip and off-chip) with grouping and overprovisioning and (c) electrical link (on-chip) photonic link (on-chip and off-chip) with grouping and overprovisioning. Link widths chosen such that network does not exceed power budget of 20 W. Thermal tuning power (1 μW/ring/K) included, laser power not included.

overprovisioning (Figure 5(b)) provided ≈3× improvement in the saturation throughput for 1-group and 4-group case, but with an increase in power dissipation. The replacement of power hungry off-chip and part of on-chip electrical links with the unified energy-efficient photonic links enabled the use of wider links for both on-chip and off-chip network, which improved the saturation throughput and reduced latency. Comparing the 16-group photonic case with overprovisioning factor of 1 and best case electrical (4-group with overprovisioning factor of 2) showed an order of magnitude higher saturation throughput and lower latency for the photonic case at comparable power dissipation.

Acknowledgments

We would like to thank all the members of the MIT photonic team, including J. Orcutt, A. Khilo, M. Popović, C. Holzwarth, B. Moss, H. Li, M. Georgas, J. Leu, F. X. Kartner, J. L. Hoyt, R. J. Ram, and H. I. Smith.

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] G. Hendry et al. Photonic on-chip networks for performance-energy optimized off-chip memory access. In *High Performance Embedded Computing*, 2009.

[3] A. Joshi et al. Silicon-photonic Clos networks for global on-chip communication. *Networks-on-Chip, International Symposium on*, 0:124–133, 2009.

[4] Y. Pan et al. Firefly: illuminating future network-on-chip with nanophotonics. In *ISCA*, pages 429–440, 2009.

[5] C. Schow et al. A ¡5mw/gb/s/link, 16x10gb/s bidirectional single-chip CMOS optical transceiver for board level optical interconnects. In *ISSCC*, 2008.

[6] A. Shacham et al. On the design of a photonic network-on-chip. In *NOCS*, pages 53–64, 7–9 May 2007.