

Design-Space Exploration for CMOS Photonic Processor Networks

Vladimir Stojanović^a, Ajay Joshi^b, Christopher Batten^a, Yong-Jin Kwon^c, Scott Beamer^c,
Sun Chen^a and Krste Asanović^c

^aMassachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, Massachusetts 02139

^bBoston University, Boston, Massachusetts

^cUniversity of California, Berkeley, California

*Email: vlada@mit.edu

Abstract: Monolithically integrated dense WDM photonic network topologies optimized for loss and power footprint of optical components can achieve up to 4x better energy-efficiency and throughput than electrical interconnects in core-to-core, and 10x in core-to-DRAM networks.

©2009 Optical Society of America

OCIS codes: (200.4650) Optical interconnects; (250.5300) Photonic integrated circuits;

1. Introduction

This paper presents a review of recent advances [1,2] in building high-throughput, energy-efficient photonic networks for core-to-core and core-to-DRAM communication in manycore processor systems. To sustain the performance scaling in these systems, the increase in core count has to be followed by the corresponding increase in energy-efficiency of the core, the interconnect, and bandwidth density [3,4]. Due to pin-density, wire-bandwidth and power dissipation limits, electrical DRAM interfaces are not expected to supply sufficient bandwidth with reasonable power consumption and packaging cost, and similar issues also limit energy-efficiency and bandwidth density of global on-chip wires. With potential for energy-efficient modulation/detection and dense wavelength division multiplexing (DWM), silicon-photonic interconnect technology is well suited to alleviate the bottleneck, however its application has to be carefully tailored to both the underlying process technology and the desired network topology.

2. Monolithic CMOS photonic network design

Recently developed infrastructure for photonic chip design and post-fabrication processing methodology [5,6] enabled for the first time monolithic integration of polysilicon and silicon-based photonic devices in a standard bulk CMOS and thin BOX SOI fabrication flows commonly used for processors. Based on this technology and the tight interaction between design of photonic interconnect components (waveguides, ring-resonators, modulators, photo-detectors, waveguide crossings) and network topology, in [1] we have proposed an efficient hybrid electro-optical core-to-DRAM shared memory network (local mesh global switch – LMGS) shown in Fig. 1, which provides a near ten-fold improvement in throughput compared to optimized electrical networks projected to 22 nm process node and a 256 core processor. To provide a balance between the bandwidth and latency/link utilization, the traffic from several tiles is aggregated via local electrical mesh into a point-to-point dense WDM interconnect with wavelength addressing to a part of a DRAM space. External buffer chip receives the optical signals and arbitrates requests from several core groups to the same DRAM module. This relatively simple interconnect results in significantly reduced number of optical components in the network compared to, for example, high-radix optical crossbar [7], minimizing the thermal tuning costs as well as losses along the optical path. To relax the loss specifications on integrated photonic devices within a required optical power envelope, Fig. 3a, the physical layout of the network follows a U-shape. Balancing the mesh bandwidth with degree of tile aggregation and optical bandwidth enables efficient utilization of raw energy-efficiency advantage of photonic over electrical interconnect (both across die and die-to-die), as shown in Fig. 4. Similar methodology was used to optimize the core-to-core network to provide a more uniform access for a variety of traffic patterns, by utilizing dense, energy-efficient photonic interconnects to realize otherwise expensive non-blocking Clos network, Fig. 2. Again, aggregation is used to decrease the radix of the network and balance the electrical and optical power of the network. The Clos photonic layout also follows the U-shape to relax the photonic device loss requirements, Fig. 3b. The Clos achieves significantly better latency and throughput uniformity compared to a concentrated mesh network, Fig. 5, across a variety of traffic patterns.

Acknowledgements

We would like to acknowledge 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. Kärtner, J. L. Hoyt, R. J. Ram, and H. I. Smith. This work was supported by DARPA awards W911NF-08-1-0134 and W911NF-08-1-0139.

OWI1.pdf

References

- [1] C. Batten, *et al*, "Building manycore processor to DRAM networks with monolithic CMOS silicon photonics," *IEEE Micro*, vol. 29, no. 4, pp. 8-21, 2009.
- [2] A. Joshi, *et al*, "Silicon-Photonic Clos Networks for Global On-Chip Communication," 3rd ACM/IEEE International Symposium on Networks-on-Chip, San Diego, CA, pp. 124-133, May 2008.
- [3] M. Tremblay, S. Chaudhry, "A Third-Generation 65nm 16-Core 32-Thread Plus 32-Scout-Thread CMT SPARC[®] Processor," *IEEE ISSCC*, pp. 82-83, 2008.
- [4] S. Bell *et al*, "TILE64 Processor: A 64-Core SoC with Mesh Interconnect," *IEEE ISSCC*, pp. 88-598, 2008.
- [5] J. S. Orcutt, *et al*, "Demonstration of an Electronic Photonic Integrated Circuit in a Commercial Scaled Bulk CMOS Process," *Optical Society of America - CLEO/QELS Conference*, San Jose, CA, 2 pages, May 2008.
- [6] C. W. Holzwarth, "Localized Substrate Removal Technique Enabling Strong-Confinement Microphotronics in Bulk Si CMOS Processes," *Optical Society of America - CLEO/QELS Conference*, San Jose, CA, 2 pages, May 2008.
- [7] D. Vantrease *et al*, "Corona: System Implications of Emerging Nanophotonic Technology," *ISCA*, June 2008.

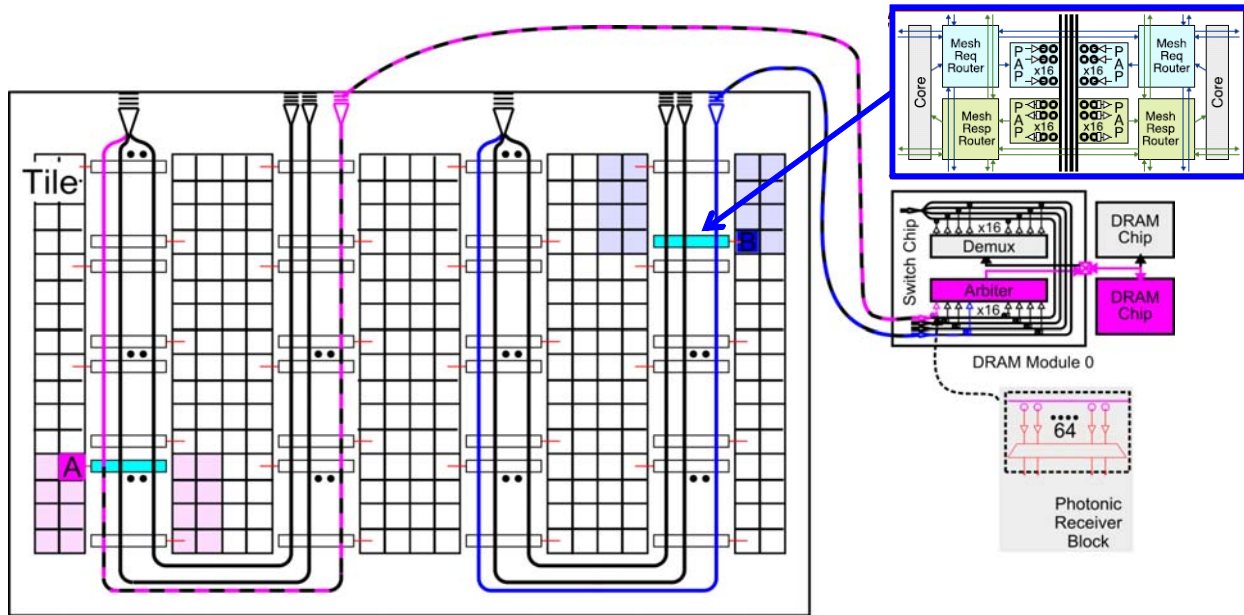


Fig. 1. Photonic network template for core-to-DRAM communication: LMGS network, 64 tiles (4 cores per tile), 64 waveguides (for tile throughput = 128 b/cyc), 256 modulators per group, 256 ring filters per group, total rings > 16K with 0.32W on thermal tuning

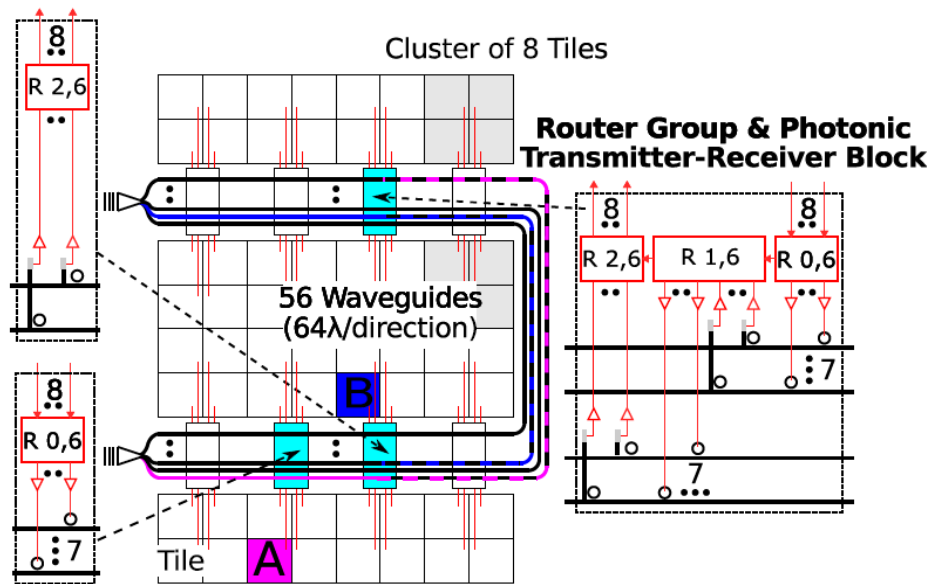


Fig. 2. Photonic network template for core-to-core communication: Clos Network, 64 tiles, 56 waveguides (for tile throughput = 128 b/cyc), 128 modulators per cluster, 128 ring filters per cluster, total rings \approx 28K with 0.56W on thermal tuning

OWI1.pdf

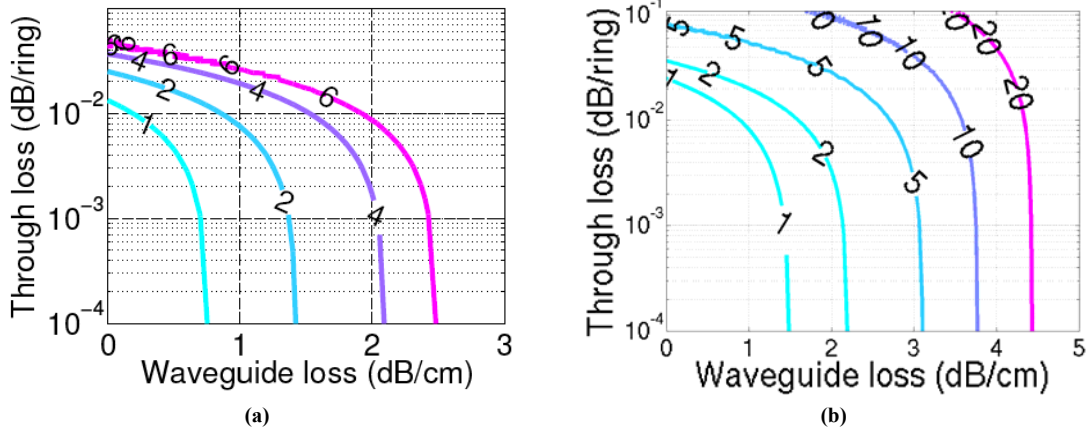


Fig. 3. Optical power requirement contours for: (a) core-to-DRAM LMGs network in Fig. 1, (b) core-to-core Clos network in Fig. 2,

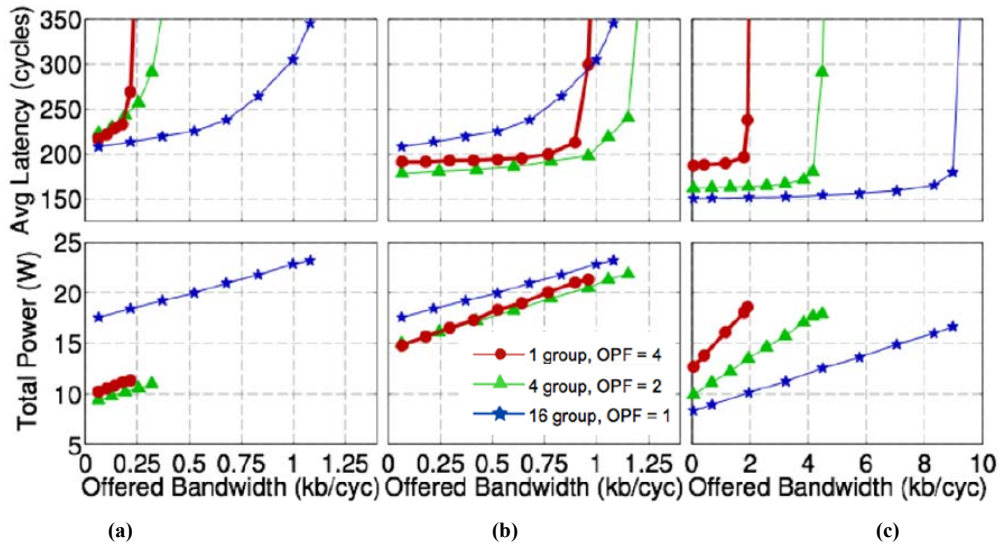


Fig. 4. Power and performance for various LMGs network configurations assuming (a) electrical interconnect with grouping, (b) electrical interconnect with grouping and overprovisioning and (c) photonic interconnect with grouping and overprovisioning. Link widths chosen such that network does not exceed power budget of 20 W.

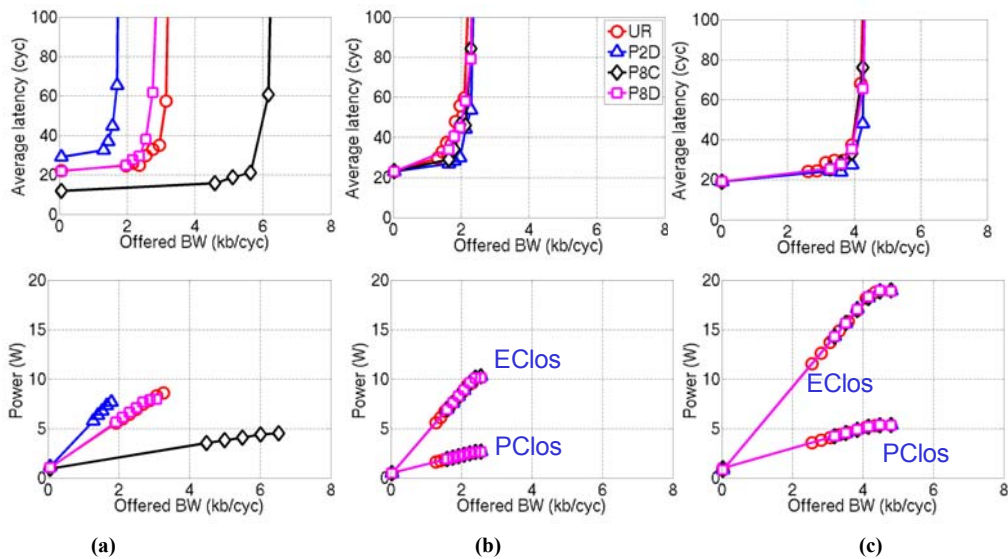


Fig. 5. Power and performance comparison of (a) electrical cmeshX2 (128b channel width) (b) electrical Clos and photonic Clos with 64b channel width and (c) electrical and photonic Clos with channel width of 128b