Manycore Processor Networks with Monolithic Integrated CMOS Photonics

Vladimir Stojanović[a], Ajay Joshi[a], Christopher Batten[a], Yong-jin Kwon[b] and Krste Asanović[b]

[a]*Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, Massachusetts 02139*
[b]*University of California, Berkeley, California 94720*
[*]*Email: vlada@mit.edu*

Abstract: This paper presents an overview of advances in highly-integrated photonic networks for emerging manycore processors. It explores the tight interaction among logical and physical implementations of all-to-all core-to-core and core-to-DRAM networks, and underlying photonic devices.
©2008 Optical Society of America
OCIS codes: (200.4650) Optical interconnects; (250.5300) Photonic integrated circuits;

1. Introduction

Power-constrained process scaling is driving processor design towards increased levels of parallelism, with modern processors already at tens of cores on a single die [1,2] and projections of core count scaling into hundreds over the next decade. To keep scaling the performance, this increase in core count has to be followed by the corresponding increase in core-to-core and core-to-memory bandwidth. In addition to improving the core energy-efficiency, the big emerging problem is that electrical interconnect solutions (both on-chip and off-chip) cannot meet the increased bandwidth demand from growing number of cores. Due to pin-density, wire-bandwidth and power dissipation limits, the projected future enhancements of existing electrical DRAM interfaces, [3], are not expected to supply sufficient bandwidth with reasonable power consumption and packaging cost. Similar issues also limit energy-efficiency and bandwidth density of global on-chip wires [4]. To overcome this performance wall, an interconnect technology with both high energy-efficiency and bandwidth density is needed. With its potential for energy-efficient modulation and detection, and dense wavelength division multiplexing (DWDM), silicon-photonic interconnect technology is well suited to alleviate this bottleneck, however its application has to be carefully tailored to both the underlying process technology and the desired network topology.

2. Monolithic integration of CMOS photonics

Most initial efforts in silicon photonic interconnects have relied upon specialized processes, such as silicon-on-insulator (SOI) with several-micron thick buried oxide (BOX) [5], which are not compatible with processes used in processor fabrication. Apart from preserving the massive investment in standard fabrication technology, monolithic integration also reduces the area and energy costs of interfacing electrical and optical components and provides higher integrated bandwidth density. Recently developed infrastructure for photonic chip design and post-fabrication processing methodology [6,7] enabled for the first time a monolithic integration of polysilicon and silicon-based photonic devices in a standard bulk CMOS and thin BOX SOI fabrication flows commonly used for processors.
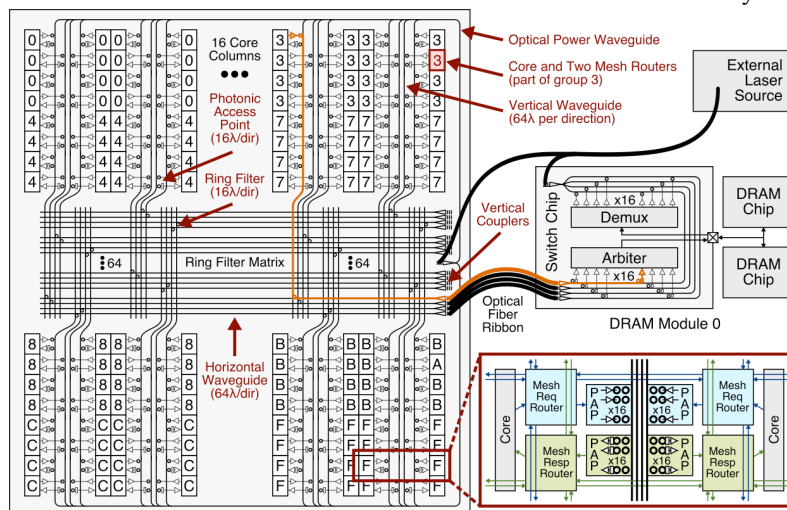

Fig. 1. A 256 core processor with a monolithic electro-optical core-to-DRAM shared memory network [8].

Based on this technology and tight interaction between design of photonic interconnect components (waveguides, ring-resonators, modulators, photo-detectors, waveguide crossings), choice of network topology and implementation, in [8] we have proposed an efficient hybrid electro-optical core-to-DRAM shared memory network 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.

3. Interaction of network topology, implementation and critical device properties

Both the choice of logical network and its physical implementation have a large impact on network performance and cost. In addition to the performance impact at the application level, this choice involves balancing various loss components, such as ring-resonator losses, waveguide crossings and waveguide loss, to minimize the optical power, along with minimizing the electrical power through optimized temperature control, modulation and detection energy-efficiency. In Fig. 2, we illustrate several physical topologies for an all-to-all butterfly network among 16-core clusters on a 256-core die. These topologies, however, can have very different optical power requirements, that even scale differently with network capacity (mapped into number of waveguides, through rings per waveguide, waveguide crossings, etc). For example, the network in Fig. 2(a) will require between 0.4 and 33 W of optical power for the capacity range of 40 - 300 Tb/s (~ 8 - 60 Bytes/cycle/core), while those in Fig. 2(b) and (c) would require 1.4 - 11 W and 0.5 - 4 W, for same capacity range, respectively. The example losses are calculated based on optimized waveguide crossing designs with approx 0.05 dB/crossing [10], ring-resonator through loss of 0.0001 dB [9] and waveguide loss of 1 dB/cm, all very challenging photonic device specifications. These large variations in required optical power mandate careful co-design of optical network topology, implementation and underlying photonic devices. They also indicate that power requirements scale differently depending both on the network topology, network size and the required throughput.
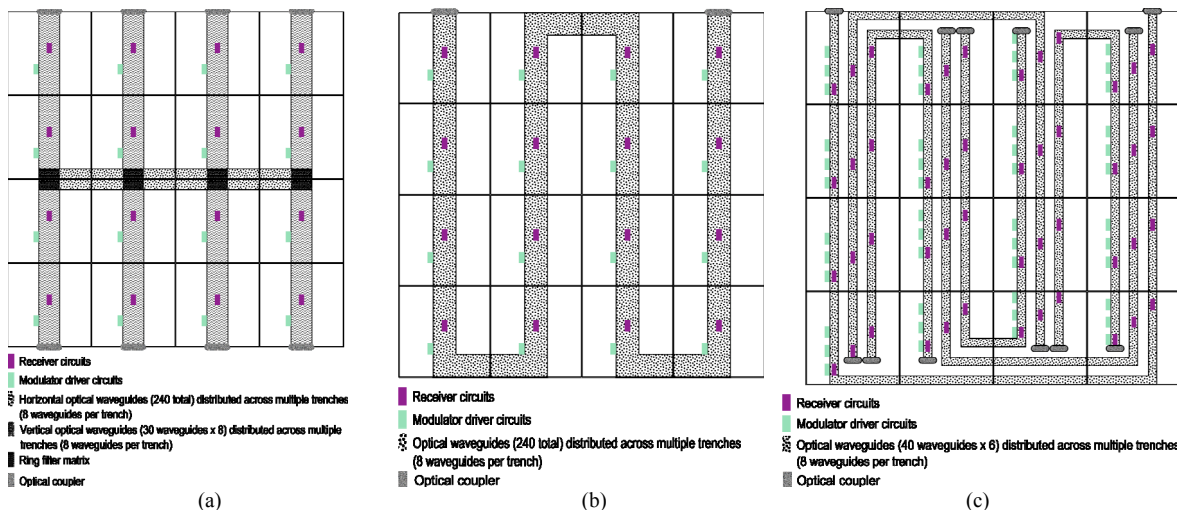

Receiver circuits
Modulator driver circuits
Horizontal optical waveguides (240 total) distributed across multiple trenches (8 waveguides per trench)
Vertical optical waveguides (30 waveguides x 8) distributed across multiple trenches (8 waveguides per trench)
Ring filter matrix
Optical coupler

Receiver circuits
Modulator driver circuits
Optical waveguides (240 total) distributed across multiple trenches (8 waveguides per trench)
Optical coupler

Receiver circuits
Modulator driver circuits
Optical waveguides (40 waveguides x 6) distributed across multiple trenches (8 waveguides per trench)
Optical coupler

(a)   (b)   (c)

Fig. 2. Photonic network layouts for a butterfly network (a) ring-filter matrix [8], (b) serpentine, (c) loss-optimized

References

[1] 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.

[2] S. Bell *et al*, "TILE64 Processor: A 64-Core SoC with Mesh Interconnect," *IEEE ISSCC*, pp. 88-598, 2008.

[3] J. Kennedy *et al*, "A 3.6-Gb/s point-to-point heterogeneous-voltage-capable DRAM interface for capacity-scalable memory subsystems," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 1, pp. 233-244, Jan. 2005.

[4] R. Ho, K.W. Mai and M.A. Horowitz "The future of wires," *Proceedings of the IEEE* vol. 89, no. 4, pp. 490-504, 2001.

[5] C. Gunn, "Fully integrated VLSI CMOS and photonics 'CMOS photonics'," in *IEEE Symposium on VLSI Technology*, pp. 6-9, 2007.

[6] 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, May 2008.

[7] C. W. Holzwarth *et al*, "Localized Substrate Removal Technique Enabling Strong-Confinement Microphotonics in Bulk Si CMOS Processes," *Optical Society of America - CLEO/QELS Conference*, San Jose, CA, May 2008.

[8] C. Batten *et al*, "Building manycore processor to DRAM networks with monolithic silicon photonics," *IEEE Symposium on High-Performance Interconnects,* Stanford, CA, pp. 21-30, August 2008.

[9] D. Vantrease *et al*, "Corona: System Implications of Emerging Nanophotonic Technology," *ISCA*, June 2008.

[10] M.A. Popović *et al*, "Low-loss Bloch waves in open structures and highly compact, efficient Si waveguide-crossing arrays," in Proc. 20th An. Meet. IEEE Lasers and Electro-Optics Soc. (LEOS), Oct 2007.