

# Photonic On-Chip Networks for Performance-Energy Optimized Off-Chip Memory Access

Gilbert Hendry<sup>†</sup> (grh2111@columbia.edu),  
Daniel Brunina<sup>†</sup> (dsb2135@columbia.edu), Johnnie Chan<sup>†</sup> (johnnie@ee.columbia.edu),  
Luca Carloni<sup>§</sup> (luca@cs.columbia.edu), Keren Bergman<sup>†</sup> (bergman@ee.columbia.edu)  
<sup>†</sup>Lightwave Research Lab, Columbia University, New York, NY  
<sup>§</sup>Computer Science Department, Columbia University, New York, NY

## Motivation

As chip multiprocessors scale to increasing numbers of cores and commensurate on-chip computational power, the gap between the available off-chip bandwidth that is required to appropriately feed the processors continues to widen under current memory access architectures. It is debatable as to whether or not electronic off-chip signaling will be able to sustain the invariably increasing performance trends and power requirements. With recent advancements in 3D Integration (3DI) technology [1], we envision a dedicated plane available for integration of silicon nano-photonic devices, enabling the design of an electronically controlled hybrid photonic network [2]. Silicon nano-photonics offers significant benefits to problems related to off-chip signaling in three key areas all related to the end-to-end nature of optical transmission:

- **Distance.** Optical transmission of data can be made agnostic to distance at the chip-scale; photonics effectively does not pay in energy per millimeter.
- **Datarate.** Photonic waveguides crossing the chip boundary are bitrate transparent, providing a natural bandwidth match between compute cores and the memory subsystem.
- **Density.** In addition to waveguide pitches being significantly smaller than electronic wires, using Wavelength Division Multiplexing (WDM) provides dramatically higher bandwidth density off-chip.

In this work, we investigate how to provide access to memory from an existing hybrid photonic network-on-chip framework [2]. We propose how to alter the peripheral network gateway switches to setup photonic circuits off chip, enabling all-to-all communication between cores and DRAM banks. We test this memory access architecture using a well-known parallel FFT algorithm, with various on-chip network architectures. The different on-chip network designs play a key role in how data is transferred from external memory to the cores. We find that characteristics of the on-chip network can drastically affect the overall execution time and energy consumption in the application performance, clearly pointing to the need for designing on-chip optical interconnection networks with memory access as a primary concern, and not as an afterthought.

## Memory Access Architecture

Under the hybrid photonic NoC framework that we consider [2], which uses an electronic control plane to circuit switch a photonic ring-resonator based network,

alternate circuit pathways must be available for transferring data in and out of the chip. We assume off-chip photonic signaling achieved through lateral coupling [3] [4], where the optically encoded data is brought in and out of the photonic links through inverse-taper optical mode converters which expand the on-chip optical cross section to match the cross section of the external guiding medium. To accommodate lateral coupling for optical waveguides in the simplest way, we envision memory access points at each network node on the periphery, as seen in Figure 1. Each peripheral gateway connects both electronic wires and optical waveguides to their respective external media. External distributed memory controllers provide access to multiple DRAM modules by arbitrating and translating requests, as well as refreshing the DRAM cells.

Under this memory access architecture, the network provides photonic circuit paths between DRAM and any core in the grid. The design of this network is therefore critical.

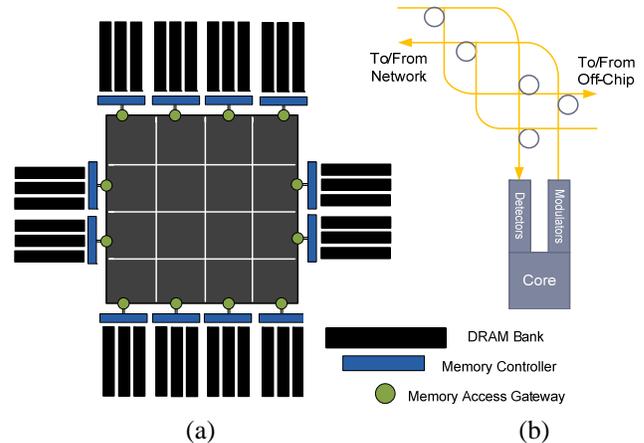


Figure 1 – (a) Memory access architecture, (b) peripheral Memory Access Gateways (green) that connect DRAM to the on-chip network

## On-Chip Networks

Performance and power will be dictated by how many access points are available, topology and routing, and transmission technology. One of the main contributions of this work is the comprehensive investigation of how these characteristics of the on-chip network impact the overall execution time and energy consumption of off-chip communication.

We assume a CMP with 64 processors arranged in a 2D planar fashion. Although we do not simulate the processors themselves, we assume simple in-order 2mm × 2mm cores

with local store memories. We simulate various networks, including an Electronic Mesh, Electronic Flattened Butterfly, Photonic Torus, Concentrated Photonic Torus, and a Concentrated Photonic Flattened Fat Tree. In our hybrid photonic networks, the number of wavelengths available for transmission depends on the worst case optical loss experienced through the network. For this reason, we introduce a new topology shown in Figure 2 known as the Photonic Square Root, which attempts to keep one of the largest contributors of optical loss, waveguide crossings, to a minimum.

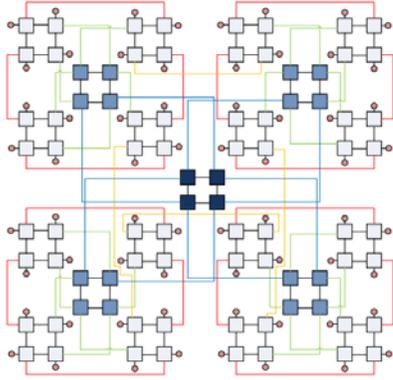


Figure 2 – Square Root topology

### Simulation Methodology

Computing the Fast Fourier Transform (FFT) of a set of data points is an essential signal processing algorithm which underlies many commercial and scientific processes. A typical way the FFT is computed in parallel is the Cooley-Tukey method [5], which is logically represented in Figure 3. We adopt a similar execution model assumed in previous work [6] and employ the results presented by Chow et al. [7] which characterizes the performance of a large single-precision FFT running on the Cell Broadband Engine using Bailey's FFT algorithm [8]. We extend previous work's investigation of the FFT execution on different networks by adding read and write memory stages at the beginning and end of the computation, respectively. Each block of data is 4MB. We use our comprehensive simulation environment [9], which uniquely integrates a physical-layer library of electronic and photonic components to model on chip networks.

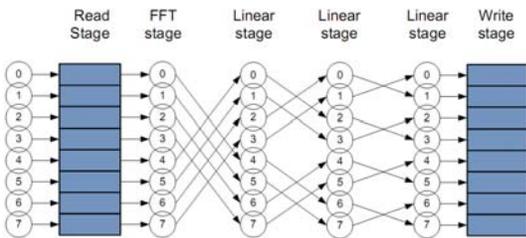
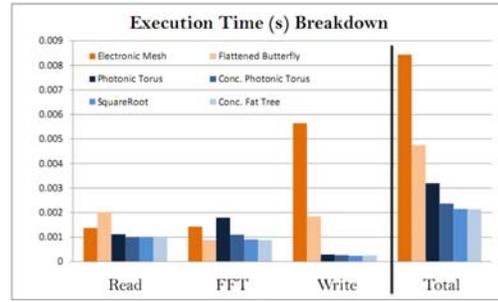


Figure 3 – Cooley-Tukey FFT algorithm with memory read and write stages

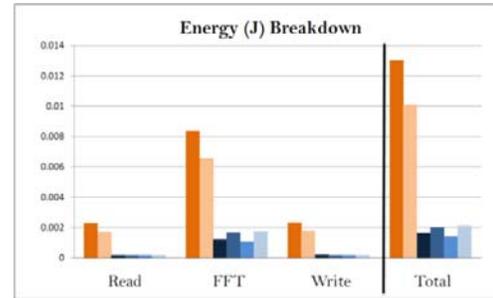
### Results

We report execution time and energy consumption for the networks studied running the FFT algorithm in Figure 4. The Concentrated Photonic Fat Tree and Photonic Square Root networks show a 2× gain in performance and a 10×

improvement in energy consumption over the Flattened Butterfly, due largely to off-chip electronic limitations.



(a)



(b)

Figure 4 – Breakdown for read, FFT, and write stages for execution time (a) and energy consumption (b).

### Acknowledgements

The authors gratefully acknowledge support for this work under DARPA MTO, ARL Grant W911NF-08-1-0127

### References

- [1] K. Bernstein *et al.*, "Interconnects in the Third Dimension: Design Challenges for 3D ICs." 2007. Design Automation Conference.
- [2] A. Shacham, K. Bergman, and L. P. Carloni, "Photonic Networks-on-Chip for Future Generations of Chip Multiprocessors." IEEE Transactions on Computers, s.l. : IEEE Computer Society, 2008, Issue 9, Vol. 57, pp. 1246-1260.
- [3] V. R. Almeida, R. R. Panepucci, and M. Lipson, "Nanotaper for Compact Mode Conversion." Optics Letters, Aug 2003, Issue 15, Vol. 28, pp. 1302-1304.
- [4] S. J. McNab, N. Moll, and Y. A. Vlasov, "Ultra-low Loss Photonic Integrated Circuit with Membrane-type Photonic Crystal Waveguides." Optics Express, November 2003, Issue 22, Vol. 11, pp. 2927-2938.
- [5] J. W. Cooley, J. W. Tukey, "An Algorithm for the Machine Calculation of Complex Fourier Series." Mathematics of Computation, 1965, Vol. 19, pp. 297-301.
- [6] M. Petracca, B. G. Lee, K. Bergman, L. P. Carloni, "Design Exploration of Optical Interconnection Networks for Chip Multiprocessors." 2008. 16th IEEE Symposium on High Performance Interconnects.
- [7] A. C. Chow, G. C. Fossum, D. A. Brokenshire, *A Programming Example: Large FFT on the Cell Broadband Engine*. IBM. 2005.
- [8] D. H. Bailey, "A High-Performance FFT Algorithm for Vector Supercomputers." International Journal of Supercomputer Applications, 1988, Vol. 2, pp. 82-87.
- [9] G. Hendry *et al.*, "Analysis of Photonic Networks for a Chip-Multiprocessor Using Scientific Applications." 2009. The 3rd International Symposium on Networks-on-Chip.