Circuit-Switched Memory Access in Photonic Interconnection Networks for High-Performance Embedded Computing

Gilbert Hendry*, Eric Robinson†, Vitaliy Gleyzer†, Johnnie Chan*,
Luca P. Carloni‡, Nadya Bliss† and Keren Bergman*

* Lightwave Research Lab, Department of Electrical Engineering, Columbia University, New York, NY
† Lincoln Laboratories, MIT, Lexington, MA
‡ Computer Science Department, Columbia University, New York, NY

Introduction

Many important application classes for embedded systems are characterized by a combination of real-time high-performance requirements and stringent energy constraints [1], [2]. These embedded systems include personal mobile devices, avionics, and image processing. As chip multiprocessors (CMPs) used in embedded computing scale to increasing numbers of cores and greater on-chip computational power, on-chip and IO communication becomes an increasingly important factor in determining power and realized performance.

Recent advances in silicon nanophotonic devices and integration have made it possible to consider optical transmission on the chip- and board-scale [3]. Microprocessor I/O signaling can directly benefit from photonics by using wavelength division multiplexing (WDM) for high bandwidth-density channels and distance-independent energy consumption.

In this work, we study the problem of designing a single circuit-switched hybrid photonic network-on-chip (NoC) architecture that supports both on-chip communication and off-chip memory access in a power-efficient way, both of which are necessary for programming models such as PGAS. We simulate our memory access architecture on a 64-node network using both random traffic and detailed traces of a high-performance embedded application, specifically the Matrix Projective Transform.

Circuit-Switched Memory Access

In a circuit-switched network, a control network provides a mechanism for setting up and tearing down energy-efficient high-bandwidth end-to-end circuit paths. This method effectively relaxes the relationship between electronic router buffer size, a large contributor to NoC power, and performance because router buffers do not become directly congested as communication demands grow. Because the higher bandwidth data plane is circuit switched end-to-end, it suffers from higher latency due to the circuit-path setup overhead, which must be amortized through a combination of larger messages and well-scheduled or time-division multiplexed communication patterns.

We propose to allow a circuit switched on-chip network to directly access memory modules, giving a single core exclusive access to a memory module for the duration of the transaction it requested. Access overhead is amortized using increased burst lengths as shown. The memory controller complexity can be greatly reduced because a memory module must sustain only one transaction at a time. The key difference is that each transaction is an entire message using long burst lengths, as opposed to small packets that must be properly scheduled.

To facilitate switching on-chip circuit paths off chip to memory modules, we place memory access points (MAPs) around the periphery of the chip connected to the network. These MAPs contain a memory controller that can service memory transactions and use the NoC to allow end-to-end communication between cores and DRAM modules.

A. Silicon Nanophotonic Technology

Circuit-switching photonic networks can be achieved using active broad-band ring-resonators whose diameter is manufactured such that its resonant modes directly align with all of the wavelengths injected into the nearby waveguide. The ring resonator can be configured to be used as a photonic switching element (PSE), as shown in Figure 1. By electrically injecting carriers into the ring, the entire resonant profile is shifted, effectively creating a spatial switch between the ports of the device [4].

We assume off-chip photonic signaling is achieved through lateral coupling [5], which is employed due to its slightly lower insertion loss compared to vertical coupling. Waveguide pitch at the chip edge can easily be on the order of 60 µm interfacing to off-chip arrayed waveguides or optical fiber. This photonic I/O pitch remains well below that of current electrical I/O pitch, illustrating the potential for vastly higher bandwidth density that is offered by using photonic waveguides when using WDM.

B. Circuit-Accessed Memory Module

Our proposed circuit-switched memory access architecture requires slightly different usage of DRAM modules.

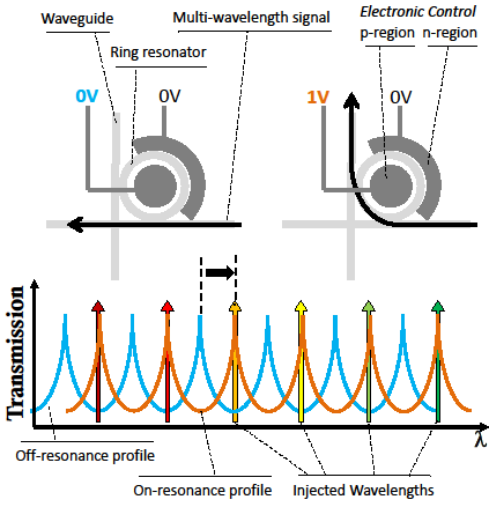

Figure 1: Broadband PSE operation

Individual conventional DRAM chips are connected via a local electronic bus to a central optical controller/ transceiver. The controller is responsible for demultiplexing the single optical channel into the address and data bus much in the same way as Rambus RDRAM memory technology.

Simulation Results

We investigate how silicon photonic technology and circuit-switching affect power efficiency in transporting data to and from off-chip DRAM by modeling different network configurations using PhoenixSim, a simulation environment which models both photonic and electronic network components [6]. We simulate a 2-D photonic mesh using two different switch designs, designated PS-1 and PS-2, which have different insertion loss and blocking characteristics. We compare them to a baseline packet-switched electronic mesh (Emesh) and a circuit-switched electronic mesh (EmeshCS), normalizing all solutions to topology, pin constraints, total memory, and CMOS technology.

The first memory access pattern we consider is uniform random, where a read or write (with equal probability) is generated at each core as a Poisson process with a uniformly random memory access point as its destination and a random address within the space of that access point. Read requests block further processing of the core until the data is returned.

In Figure 2, we see that there is little difference in total DRAM bandwidth between the different networks until of around 2 kB messages. At this point, the Emesh bandwidth saturates, and the photonic networks diverge from the electronic ones. Latency, as expected, reaches a similar threshold where the circuit-switched networks become superior at larger message sizes under saturation. Zero-load latency is considerably better for the circuit-switched networks even at small message sizes.

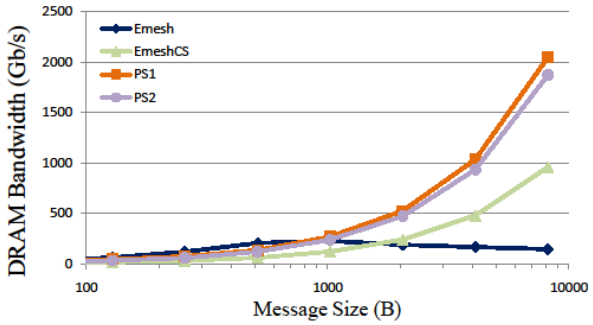

Figure 2: DRAM bandwidth under random traffic

In addition to modeling random memory access, we evaluate the proposed network architectures using the application modeling framework, Mapping and Optimization Runtime Environment (MORE) to collect traces from the execution of a high-performance embedded application component, the matrix projective transform, an important component in the SAR image processing chain. We varied the image size of the projective transform computed, collecting average network power and computational performance, measured in giga-operations per second (GOPS). These results are averaged across image sizes (with little variance), illustrated in Table 1. The circuit switched networks yield about 10% performance improvement over the packet switched network, with over 20% less power. Combining performance and power, we see that the photonic networks are both around 500 times more efficient than the Emesh. For this application, the nonblocking nature of the PS-2 switch makes it more efficient under our definition, though considering that both photonic networks dissipate less than one Watt, it is likely that the performance gain using the PS-1 switch is preferable.

Table 1: Projective Transform performance summary

Network	*Avg. Network Power (W)*	*Avg. Perf. (GOPS)*	*Power-Perf. Improvement over Emesh*
Emesh	21.69	2.35	-
EmeshCS	1.33	14.38	99.7
PS-1	0.75	**39.19**	482
PS-2	**0.54**	30.82	**526**

Conclusion

We have investigated a circuit-switched memory access methodology for use in embedded high-performance CMPs. We compared our design against a traditional electronic packet switched network with respect to bandwidth and latency using random traffic. Accessing memory using circuit-switching was found to have significant advantages for the projective transform algorithm, namely increased bandwidth through long burst lengths and decreased power by eliminating performance dependent buffers. Silicon nanophotonic technology adds to these benefits with low-energy transmission and higher bandwidth density which will allow scaling up of off-chip bandwidths and memory capacities.

References

[11] M. Duranton. The challenges for high performance embedded systems. In DSD '06: Proceedings of the 9th EUROMICRO Conference on Digital System Design, pages 3–7, Washington, DC, USA, 2006. IEEE Computer Society.

[2] W. Wolf. Embedded computer architectures in the mpsoc age. In WCAE '05: Proceedings of the 2005 workshop on Computer architecture education, page 1, New York, NY, USA, 2005. ACM.

[3] L. Chen, K. Preston, S. Manipatruni, and M. Lipson. Integrated GHz silicon photonic interconnect with micrometer-scale modulators and detectors. Optics Express, 17(17), August 2009.

[4] B. G. Lee et al. High-speed 2_2 switch for multi-wavelength message routing in on-chip silicon photonic networks. In European Conference on Optical Communication (ECOC), Sept. 2008.

[5] V. R. Almeida, R. R. Panepucci, and M. Lipson. Nanotaper for compact mode conversion. Optics Letters, 28(15):1302–1304, August 2003.

[6] J. Chan, G. Hendry, A. Biberman, K. Bergman, and L. P. Carloni. Phoenixsim: A simulator for physical-layer analysis of chip-scale photonic interconnection networks. In DATE: Design, Automation, and Test in Europe., Mar. 2010.