Chip Scale Photonic Interconnects for Energy-Performance Optimized Computing

Keren Bergman

Department of Electrical Engineering, Columbia University, New York, NY 10027, USA, bergman@ee.columbia.edu

***Abstract*— As computational performance continues to scale in chip multiprocessors with increasing numbers of cores, the gap between the available off-chip bandwidth and that which is required to appropriately feed the processors continues to widen under current architectures. For many high-performance computing applications, the bandwidth available for both on- and off-chip communications can play a vital role in efficient execution due to the use of data-parallel or data-centric algorithms. Electronic interconnected systems are increasingly bound by their communications infrastructure and the associated power dissipation of high-bandwidth data movement. Recent advances in chip-scale silicon photonic technologies have created the potential for developing optical interconnection networks that can offer highly energy efficient communications and significantly improve computing performance-per-Watt. This paper examines the design and performance of photonic networks-on-chip architectures that support both on-chip communication and off-chip memory access in an energy efficient manner.**

I. Introduction

The trend toward many-core high-performance processor architectures has illustrated the need for high bandwidth, scalable, low-power on-chip and off-chip interconnection networks [1]. Silicon nanophotonic devices, integrated at the chip-scale, offer many relevant advantages to the construction of these networks, including high bandwidth achieved using wavelength division multiplexing, data rate-transparent switching power, and compatibility with CMOS fabrication. In order to properly investigate the design of the photonic interconnection network between processing cores and components of the memory hierarchy, detailed analysis at both the physical layer as well as the architectural layer must be conducted across a wide design space.

We present a comprehensive photonic interconnection simulation platform that captures both physical and architectural details of future silicon photonics-based interconnection networks. The simulator includes a library of the photonic device models, which are largely composed of silicon nanophotonic micro-ring resonators and waveguides that have been experimentally characterized at the physical level. These device models can then be used as building blocks to create complete photonic interconnection networks with various topologies. The network model is functionally driven by a complete instruction set simulator and a detailed memory model, together allowing comprehensive simulation of the computing system performance while extracting the associated power dissipation and optical physical layer characterization.

Photonic interconnection networks have recently been proposed as a replacement to conventional electronic network-on-chip solutions in delivering the ever increasing communication requirements of future chip multiprocessors. While photonics offers superior bandwidth density, lower latencies, and improvements in energy efficiency over electronics, the photonic network designs that can leverage these benefits cannot be easily derived by simply mimicking electronic layouts. In fact, proper implementations of photonic interconnects will require the careful consideration of a variety new physical-layer metrics and design requirements that did not exist with electronics. Here, we review some of the currently proposed designs for chip-scale photonic interconnection networks, the design methodologies required to produce viable network topologies, and a simulation environment, called PhoenixSim, that we have developed to accurately model and study those metrics and designs [2]. We present an overview of how photonic interconnection networks are designed and can be used to create high-bandwidth and energy-efficient links for on- and off-chip communications.

I. Photonic Interconnection Networks

We envision the integration of a dedicated photonic networking plane to provide low-power global communications through the three-dimensional integration (3DI) of future chip multiprocessors (CMP) [3]. The 3DI approach enables the stacking of multiple layers that are each dedicated for a different purpose. The bottom plane would contain the many processing nodes of the CMP, while the next several layers would be dedicated to providing a large amount of local memory. Found at the top would be the photonic plane that globally connects the many cores and memory units.

Due to recent advancements in optical technologies, photonic NoCs have become attractive solutions for the problems facing traditional electronic implementations. All the optical devices that would be required to generate, route, and receive an optical signal have been demonstrated in silicon which renders them compatible with traditional CMOS fabrication technology. These devices includes waveguides [4], waveguide crossings [5], modulators [6], detectors [7], filters [8], and switches [9][10] that can be combined to form a diverse range of topologies, giving chip designers a large design space to explore the possibilities.

The silicon ring resonator is a fundamental device in enabling the control of signal propagation on a photonic network. Ring

978-1-4244-5369-6/10/$26.00 ©2010 IEEE

resonators can be used off-resonance which allows the signals to pass by unobstructed, or in an on-resonance configuration which shifts the signal onto another waveguide. These two states of the ring resonator form the basis for all types of chip-scale photonic switching.

One method of controlling the flow of an optical signal is to treat the ring as a selective filter and controlling the wavelength of the incoming signal to be either on- or off-resonance . If a signal's wavelength is aligned with a resonant peak (i.e. resonant mode) of the ring, the signal will be on-resonance and drop into the ring, otherwise the signal will pass by. We call this *wavelength-selective switching* since the wavelength of the signal is used to determine the route that a signal will take. A network employing this type of switching has previously been proposed for use as an off-chip memory-access network architecture.

Alternatively, control can be enabled by holding the wavelength of the signal constant and actively tuning the ring's resonant profile through electro-optic means. This type of active control would result in a concurrent shift of the entire resonant profile. We refer to ring resonators that use this method as broadband switches since the many wavelength channels of a WDM signal can each be aligned to a unique resonant mode and any shift in the resonant profile will simultaneously affect all wavelength channels. We label this form of control as *space switching* for its nature of shifting the entire WDM packet.

When comparing the two switching methods, spacing switching fully utilizes the optical spectrum to maximize link bandwidth by leveraging WDM, whereas wavelength-selective switching sacrifices some of the spectrum in order to use it as a controlling mechanism. Conversely, wavelength-selective switching can achieve much better latencies since it does not incur the delay overhead that is required by the circuit-switching protocol of the space-switching technique. Note that the two switching techniques described here can be considered distinctive techniques, but combinations of the two can also be used.

II. Photonic Metrics and Simulation Methodology

Design metrics are useful in determining the usefulness and feasibility of a NoC design. Electronics and photonics share many system-level metrics such as throughput and latency as well as some physical-layer metrics such as power efficiency. However, in order to fully understand the capabilities and limitations of a photonic NoC, a designer must also consider additional metrics unique to photonics, namely insertion loss and crosstalk. In fact, some of these new metrics directly impact others metrics, clearly making them a critical consideration for creating a high-performance energy-efficient design.

Since photonics is physically different from electronics, it renders typical network simulators incapable of clearly determining accurate performance characteristics of such networks. Also, current fabrication technology is incapable of producing entire photonic network topologies which further stresses the need for physically-accurate network simulation. The PhoenixSim simulation environment was developed with the goal of incorporating detail models of photonic components while also maintaining a level of customizability to allow a variety of topology designs to be explored. The toolset is utilized in a hierarchical manner, enabling the creation of a library of highly-parameterizable device models to be used as building blocks to form larger networking components as well as entire topologies.

The network-level power dissipation is a major component in limiting performance scaling of chip-scale systems. Photonic on-chip networks have been shown to drastically outperform electronic networks in both performance and energy, especially in the case of traffic patterns that require large data transmissions. Power and energy analysis has also been integrated into PhoenixSim.

III. Conclusions

Photonic technology offers a unique solution to mitigating many performance limitations facing today's electronic NoCs. Because of the fundamentally different nature of photonics, new strategies for designing photonic topologies must be devised to ensure that such a design can fully leverage what photonics has to offer. Novel tools that understand the unique characteristics of photonics, like PhoenixSim, are needed so that researchers can accurately model, analyze, and create optimized designs for them.

References

[1] N. Magen, A. Kolodny, U. Weiser, and N. Shamir, "Interconnect-power dissipation in a microprocessor," in *Proc. of the 2004 International Workshop on System Level Interconnect Prediction*, pp. 7–13, 2004.

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

[3] A. W. Topol, et al., "Three-dimensional integrated circuits," *IBM Journal of Research and Development*, vol. 50, no. 4/5, pp. 491–506, Jul/Sep 2006.

[4] F. Xia, L. Sekaric, and Y. Vlasov, "Ultracompact optical buffers on a silicon chip," *Nature Photonics*, vol. 1, pp. 65–71, 2006.

[5] W. Bogaerts, P. Dumon, D. V. Thourhout, and R. Baets, "Low-loss, low-cross-talk crossings for silicon-on-insulator nanophotonic waveguides," *Opt. Lett.*, vol. 32, no. 19, pp. 2801–2803, 2007.

[6] M. R. Watts, D. C. Trotter, R. W. Young, and A. L. Lentine, "Ultralow power silicon microdisk modulators and switches," in *Proc. IEEE International Conference on Group IV Photonics*, Sorrento, Italy, paper WA2, Sept. 2008.

[7] S. J. Koester, et al., "Ge-on-SOI-detector/Si-CMOS-amplifier receivers for high-performance optical-communication applications," *IEEE/OSA Journal Lightwave Technology*, vol. 25, no. 1, pp. 46–57, Jan. 2007.

[8] B. E. Little, et al., "Ultra-compact Si-SiO_2 microring resonator optical channel dropping filters," *IEEE Photon. Technol. Lett.*, vol. 10, no. 4, Apr. 1998.

[9] B. G. Lee, A. Biberman, P. Dong, M. Lipson, and K. Bergman, "All-optical comb switch for multiwavelength message routing in silicon photonic networks," *IEEE Photon. Technol. Lett.*, vol. 20, pp. 767–769, 2008.

[10] B. G. Lee, et al., "High-speed 2×2 switch for multi-wavelength silicon photonic networks on-chip," *IEEE/OSA Journal Lightwave Technology*, vol. 27, no. 14, pp. 2900–2907, Jul. 2009.