

# Integrated Photonics in Future Multicore Computing

## New Directions in Dependability and Power Efficiency

Sudeep Pasricha

Department of Electrical and Computer Engineering  
Colorado State University  
Fort Collins, CO – 80523-1373  
sudeep@colostate.edu

**Abstract**— Integrated silicon photonics represents one of the more promising solutions to overcome the challenge of worsening on-chip communication performance in emerging multicore platforms. This paper highlights recent innovations in architectures, protocols, and techniques to enhance performance, dependability, and power-efficiency for integrated photonics.

**Keywords**—integrated silicon photonics, photonic networks-on-chip, dependability, power-efficiency, high performance

### I. INTRODUCTION

With core counts on computing chips expected to reach into the hundreds or even thousands of cores in the near future, the need for high performance, dependable, and power-efficient communication among cores has never been more critical. Traditional electrical bus-based and network-on-chip (NoC) architectures that form the backbone of communication fabrics on computing chips today will not scale to meet the communication demands for such high core counts. This is because of the poor scaling of electrical wires with successive CMOS technology generations [1]. Already even for a few tens of cores on a chip, electrical wires are slow, power hungry, and unreliable, causing communication to become a major system-wide bottleneck [1]. A failure to respond to this challenge for on-chip communication will create a brick wall that will hinder advances in all forms of computing.

Fortunately, integrated silicon photonics represents a very promising solution to overcome this challenge. By transferring data using light signaling between cores and memory, orders of magnitude improvement in bandwidth, latency, reliability, and power-efficiency are possible. Integrated photonics also promises to pair well with existing board-to-board and chip-to-chip photonics offerings that are rapidly being adopted today. Not surprisingly, many semiconductor companies such as Intel and IBM are investing heavily in integrated photonics.

The challenge of designing integrated silicon photonic communication fabrics is actively being pursued from a variety of different perspectives [2]. New CMOS-compatible photonic devices (e.g., microring resonators, photodetectors, multi-mode lasers, waveguides) are emerging as a result of these concerted efforts from industry and academia, as are new on-chip and chip-to-chip network architectures, protocols, optimization and CAD techniques, and tools for rapid design and analysis. Successful solutions will likely adopt and encompass elements from all of these research thrusts.

### II. BACKGROUND: ON-CHIP PHOTONIC LINKS

A typical on-chip photonic link includes an off-chip (or on-chip) laser source that provides light, waveguide(s) that route these light signals, microring resonator (MR) couplers that modulate electrical signals to optical ones at the source, and MR filters to detect/drop light signals on to photo-detectors to recover the electrical signals at the destination. By sending multiple wavelengths of light simultaneously in a waveguide (i.e., using dense wavelength division multiplexing (DWDM)), high network bandwidth density can be achieved. Such DWDM-based photonic waveguides form the backbone of photonic network-on-chip (PNoC) architectures.

### III. PNOC ARCHITECTURES

Early PNoC architectures integrated simple ring-based photonic buses with electrical buses [3] and photonic mesh NoCs with electrical mesh NoCs [4]. These works showed how photonic links can provide viable alternatives to copper wires especially at the global interconnect level. To cope with higher core counts and support high data transfer bandwidths, recent architectures have favored a more complex photonic crossbar (e.g., Corona [5]) or a photonic crossbar/ring integrated with an electrical mesh NoC (e.g., Meteor [6], Firefly [7]).

To improve communication performance even further in the future, single-layer PNoC topologies will need to be modified by increasing the number of waveguides and MRs in the photonic layer. But this may not be practically feasible due to limitations imposed by physical design concerns, such as waveguide spacing requirements and higher waveguide crossing losses. For these reasons, recent efforts have focused on multi-layer PNoC implementations. Such architectures use multiple photonic layers, much like the multiple metal layers in contemporary CMOS chip designs. For example, OPAL [8] stacks multiple layers of photonic crossbars/rings with no optical waveguide crossover points, coupled to an electrical mesh topology. Such multi-layer (3D) PNoCs are likely to be more practical to fabricate, as they simplify intra-layer physical design over single-layer (2D) alternatives.

### IV. PNOC PROTOCOLS

PNoC architectures can suffer from network resource contention in shared photonic channels that can lead to significant performance degradation. This contention is not unlike that experienced in traditional shared buses and NoCs.

---

This research is supported by grants from SRC, NSF (CCF-1252500, CCF-1302693), and AFOSR (FA9550-13-1-0110).

Recent efforts have proposed protocols to resolve such contention issues. For example, [9] proposes a distributed concurrent token stream arbitration protocol that provides multiple simultaneous tokens for arbitration and increases photonic channel utilization over conventional token slot and token channel schemes [5], [10]. The scheme from [9] also supports the ability to dynamically transfer bandwidth between clusters of cores and re-prioritize multiple co-running applications to further improve channel utilization and adapt to time-varying application performance goals.

#### V. CAD OPTIMIZATION TOOLS FOR PNOCS

The successful adoption of integrated silicon photonics will also rely to a large extent on a robust ecosystem of CAD tools. Such tools will be indispensable to design, optimize, and test multi-billion transistor chips with disparate technologies such as silicon photonics. Such CAD tools are slowly emerging. For instance, [11] proposes VANDAL, a CAD tool for placing functional photonic devices, modifying their parameters, and routing waveguides between devices to create complex network structures. CAD frameworks at the system-level that allow exploration and analysis of PNoC architectures will also become increasingly essential for multi-objective optimization. For example, [12] proposes CAD optimization algorithms to balance power, bandwidth, and latency goals while customizing various facets of the METEOR PNoC [6], such as DWDM degree, waveguide count, photonic uplinks/downlinks, and modulation frequency. In [13], a similar system-level CAD framework is proposed, called HELIX, to optimize a hybrid PNoC architecture that combines free-space photonics with an electrical mesh NoC.

#### VI. CROSSTALK MITIGATION IN PNOCS

MR modulators are used extensively in PNoCs. One challenge with these devices is that they can suffer from intrinsic crosstalk-noise and power-loss due to design imperfections. The crosstalk noise severely impacts PNoCs with high MR counts, where the generated crosstalk is intensified, leading to transmission errors. For example, the Corona [5] crossbar PNoC has worst-case SNR of 14dB in its data channels, which is insufficient for reliable data communication, as its corresponding bit-error-rates (BER) are very high, in the order of  $10^{-3}$ . In [14], it was shown that when transmitting data in PNoCs, crosstalk noise in MRs depends on the characteristics of data values propagating in the photonic waveguide. The work went on to propose techniques to reduce undesirable data value occurrences in a photonic waveguide. These techniques enabled trade-offs between reliability, performance, and energy overhead for the Corona and Firefly PNoCs, and reduced worst-case SNR by up to 18%.

#### VII. THERMAL-AWARE DESIGN FOR PNOCS

MR modulators are also highly sensitive to temperature fluctuations. With increase in temperature, the refractive index of an MR device changes, causing a change in its resonance wavelength that has been statically assigned at design-time. As a result of this variation in resonance wavelength, an MR may be unable to write or read data in the waveguide at higher temperatures. A cross-layer framework for thermal reliability-aware run-time management was proposed in [15] that

integrates adaptive MR assignment at the device-level and dynamic thread migration at the system-level. The framework reduced total power by up to 61% for the Corona PNoC.

#### VIII. PHOTONICS FOR CHIP-TO-CHIP COMMUNICATION

Realizing the full potential of integrated silicon photonics will require a tight integration of intra-chip PNoCs with inter-chip photonic links. Perhaps the most significant impact that photonics can have beyond the chip-level is to allow high-bandwidth and low-latency transfers between processing chips and main-memory, to alleviate main-memory bottlenecks. For example, [16] proposed a photonics-based communication link between the CPU and DRAM. It was shown that energy-efficiency of the proposed photonic interface significantly outperformed several conventional electrical interfaces such as DDR3, LPDDR3, Wide-I/O, and differential serial interface. In fact, without photonics, it was shown that 3D DRAM architectures could not provide high bandwidth transfers [16].

#### REFERENCES

- [1] S. Pasricha, N. Dutt. "On-Chip Communication Architectures", *Morgan Kaufman, ISBN 978-0-12-373892-9, Apr 2008.*
- [2] Y. Xu, S. Pasricha, "Silicon Nanophotonics for Future Multicore Architectures: Opportunities and Challenges", *IEEE Design and Test (D&T), Sep/Oct, pp. 9-17, 2014.*
- [3] S. Pasricha, N. Dutt, "ORB: An On-chip Optical Ring Bus Communication Architecture for Multi-Processor Systems-on-Chip", in *IEEE ASPDAC, Seoul, Korea, Jan 2008*
- [4] A. Shacham, K. Bergman, L. Carloni, "Photonic Networks-on-Chip for Future Generations of Chip Multiprocessors," *IEEE Trans. Comput., 57(9):1246-1260, 2008.*
- [5] D. Vantrease, et al., "Corona: System Implications of Emerging Nanophotonic Technology," in *ISCA, 2008.*
- [6] S. Bahirat, S. Pasricha, "METEOR: Hybrid Photonic Ring-Mesh Network-on-Chip for Multicore Architectures", in *ACM Transactions on Embedded Computing Systems (TECS), 13(3):116:1-116:33, Mar 2014.*
- [7] Y. Pan, et al., "Firefly: illuminating future network-on-chip with nanophotonics," in *ISCA, 2009.*
- [8] S. Pasricha, S. Bahirat, "OPAL: A Multi-Layer Hybrid Photonic NoC for 3D ICs," in *IEEE/ACM ASPDAC, 2011.*
- [9] S. V. R. Chittamuru, S. Desai, S. Pasricha, "A Reconfigurable Silicon-Photonic Network with Improved Channel Sharing for Multicore Architectures," in *ACM GLSVLSI, May 2015.*
- [10] D. Vantrease et al., "Light speed arbitration and flow control for nanophotonic interconnects," in *IEEE/ACM MICRO, Dec. 2009.*
- [11] G. Hendry, J. Chan, L. Carloni, K. Bergman, "VANDAL: A tool for the design specification of nanophotonic networks," in *DATE, 2011.*
- [12] S. Bahirat, S. Pasricha, "A Particle Swarm Optimization Approach for Synthesizing Application-specific Hybrid Photonic Networks-on-Chip," in *IEEE ISQED, 2012.*
- [13] S. Bahirat, S. Pasricha, "HELIX: Design and Synthesis of Hybrid Nanophotonic Application-Specific Network-On-Chip Architectures ", in *IEEE ISQED, Mar. 2014.*
- [14] S. V. R. Chittamuru, S. Pasricha, "Crosstalk Mitigation for High-Radix and Low-Diameter Photonic NoC Architectures ", *IEEE Design and Test (D&T), May/June 2015.*
- [15] S. V. R. Chittamuru, S. Pasricha, "SPECTRA: A Framework for Thermal Reliability Management in Silicon-Photonic Networks-on-Chip," in *IEEE VLSI, Jan 2016.*
- [16] I. Thakkar, S. Pasricha, "3D-ProWiz: An Energy-Efficient and Optically-Interfaced 3D DRAM Architecture with Reduced Data Access Overhead", in *IEEE Trans. on Multi-Scale Comp. Sys. (TMSCS), 2015.*