Exploring the Design of 64- and 256-Core Power Efficient Nanophotonic Interconnect

Randy Morris, *Student Member, IEEE*, and Avinash Karanth Kodi, *Member, IEEE*

***Abstract*—High-performance and low-power network-on-chips (NoCs) will be required to support the increasing number of cores in future chip multiprocessors. In this paper, we propose a scalable low-power 64-core NoC design called PROPEL that uses emerging nanophotonic technology. PROPEL strikes a balance between cheaper electronics and more expensive optics by facilitating nanophotonic interconnects for long distance interrouter communication and electrical switching for routing and flow control. In addition, PROPEL reduces the number of required components by facilitating communication in both the x- and y-directions. We also propose a 256-core scaled version of PROPEL called E-PROPEL that uses four separate PROPEL networks connected together by an optical crossbar. We also propose two different optical crossbar implementations using single and double microring resonators, where the single microring design has minimal optical losses (-4.32 dB) and the double microring design has minimal area overhead (0.0576 mm^2). We have simulated both PROPEL and E-PROPEL using synthetic and SPLASH-2 traffic, where our results indicate that PROPEL and E-PROPEL significantly reduce power (tenfold) and increase performance (twofold) over other well-known electrical networks.**

***Index Terms*—Interconnects, low-power architecture, network-on-chip (NoC), optoelectronic.**

I. Introduction

THE ERA of multi-cores is upon us and current technology trends have shown that future chip multiprocessors (CMPs) will be comprised of 10s–100s and even 1000s of cores. It is well known that bus-based on-chip networks will quickly become saturated with an increase in the number of cores. Scalability of on-chip networks combined with the wire delay problem (the number of transistors that can be reached in one clock cycle with technology scaling [1]) have forced chip designers to adopt a more modular network-on-chips (NoCs) paradigm [2]–[5]. With technology scaling, metallic interconnects will be limited by crosstalk, impedance mismatch, electromagnetic interference (EMI), and power dissipation, requiring alternate technology for future NoCs designs.

Nanophotonic technology is a potential solution and provides several significant advantages over metallic interconnects such as: 1) bit rates independent of distance; 2) higher bandwidth due to multiplexing of wavelengths; 3) larger bandwidth density by multiplexing wavelengths on the same waveguide/fiber; and 4) lower power by dissipating power only at the endpoints of the communication channel and many more [6]. Optics is the technology of choice at long distances (LAN, WAN) and short distances (board-to-board), as evidenced by industrial products such as Intel connects [7] and active cables [8]. However, the recent surge in photonic components and devices [9]–[13], such as silicon-on-insulator (SOI)-based microring resonators that offer extraordinary performance in terms of density, power efficiency, high bandwidth characteristics, and CMOS compatibility, are generating interest for even on-chip interconnects [3]–[5], [14].

In this paper, we propose PROPEL—a scalable NoC that uses emerging on-chip nanophotonic components to meet the power and bandwidth demands of future multicores with acceptable nanophotonic hardware complexity. PROPEL is designed for 64-core, and in addition, we also propose a 256-core version called extended (E)-PROPEL. Both PROPEL and E-PROPEL strike a balance between cheaper electronics and more expensive optics by facilitating nanophotonic interconnects for long-distance interrouter communication and electrical switching for routing and flow control. The adoption of nanophotonic interconnects allows for a reduction in network complexity and area overhead, as multiple communication channels can travel in one waveguide using wavelength-division multiplexing (WDM). For E-PROPEL, we implement an $N \times N$ optical crossbar using microring resonators and connect four, 64-core PROPELs to provide scalable interchip bandwidth with reduced power consumption. We propose optical crossbar designs using microring resonators and provide a detailed analysis. The proposed power-efficient optical crossbar designs can be used as a building block for many core (above 256) CMPs. We provide a thorough quantitative analysis of PROPEL and E-PROPEL in terms of throughput and power for both synthetic and SPLASH-2 benchmarks. Our results indicate that PROPEL has a tenfold reduction in power and an average of twofold speed up for SPLASH-2 applications when compared to mesh. Our results also indicate that E-PROPEL has a fivefold reduction in power and a 20% increase in performance when compared to mesh.

II. Related Work

With the recent surge in on-chip nanophotonic components, few on-chip network designs have been proposed. In a nanophotonic network proposed by HP [3], a 3-D stacked 64-cluster, 256-core optical crossbar that uses optical tokens for media access arbitration has been proposed. This design scales as $O(N^2)$, where N is the number of clusters, which increases the cost and complexity of the network. Another nanophotonic design proposed by Batten *et al.* [5] uses optical interconnects for direct access to dynamic RAM (DRAM). This design tackles the high

Manuscript received September 1, 2009; revised October 28, 2009 and November 17, 2009; accepted November 29, 2009. Date of publication January 8, 2010; date of current version October 6, 2010. This work was supported in part by the National Science Foundation under Grant ECCS-0725765 and Grant CCF-0953398.

The authors are with the Department of Electrical Engineering and Computer Science, Ohio University, Athens, OH 45701 USA (e-mail: kodi@ohio.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSTQE.2009.2038075

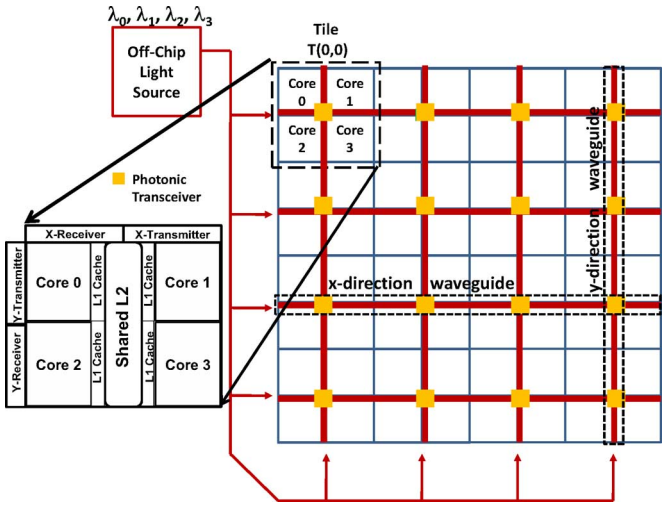

Fig. 1. Proposed layout of PROPEL architecture for 64-core architecture.

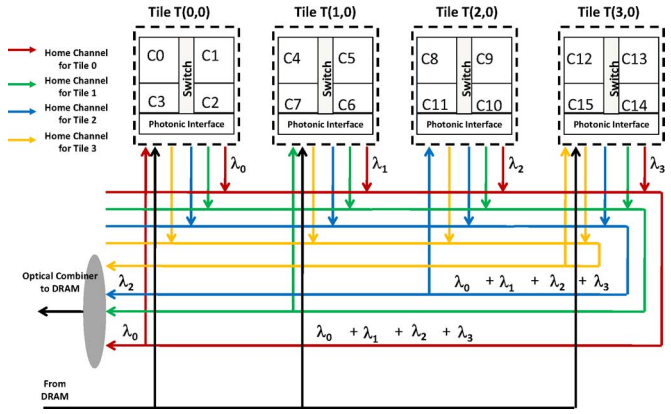

Fig. 2. RWA proposed for PROPEL for the x-dimension.

DRAM communication latency, but requires cores that are not in the same group to communicate with each other through slower memory. Shacham *et al.* [4] have proposed circuit-switched photonic interconnects, where electronic setup, photonic communication, and teardown are implemented. The disadvantage of this approach is the excess latency for path setup, which is performed using electrical interconnects. Firefly [15] is an optoelectronic multistage NoC that is comprised of an electrical network and an optical network. When packets enter the network, they first traverse the electrical network and then traverse the optical network or *vice versa*. Since Firefly requires packets to traverse an electrical network, this can result in high power dissipation. Recently, Phastlane [16], an optical mesh network, has been proposed that allows packets to move from router to router optically without requiring an initial setup circuit. Phastlane buffers incoming packets electrically if the output channel they need is blocked, thus resulting in higher packet latency and power dissipation.

III. PROPEL: Architecture and Implementation

A. Architecture

We choose 22 nm technology node for our paper, as prior research has shown that optics is advantageous compared to electronics in terms of power, delay, and area overhead at core-to-core distances. [3]–[5], [14]. It should be mentioned that even though PROPEL is implement in 22-nm technology, PROPEL could easily be implemented in newer technology nodes such as 17 nm. We simply focus on 22 nm technology, as we believe it will be the first technology node to allow nanophotonics to be more cost-effective than electronics [17]. In PROPEL, we combine four cores together and connect them through the use of a shared L2 cache, which we call a *tile*. This grouping reduces the cost of the interconnect as every core does not require lasers attached and more importantly, facilitates local communication through cheaper electronic switching [18]. Fig. 1 shows the layout of PROPEL, which consists of 16 tiles in a grid fashion with four tiles in x- and y-directions. Optical interconnects are used in two dimensions along the grid similar to an electronic 2-D mesh or torus, resulting in a maximum hop count of two. One hop count for traveling in the x-direction, and one hop count for traveling in the y-direction.

B. Implementation

Mach–Zehnder (MZ) modulators [19] and microring resonators [12] are two common devices used for indirect modulation of an optical signal (off-chip laser). When MZ modulators are compared to microring resonators, the latter are more favored due to their smaller footprint (10 μm) and lower power dissipation (0.1 mW) [12]. In addition, current microring resonators have been demonstrated with extinction ratios greater than 9 dB, optical losses as low as -0.12 dB/cm, and modulator insertion loss of 1 dB, which are sufficient for the receiver design used in PROPEL [5], [11], [12], [20]. Silicon waveguides are used over polymer waveguides for on-chip applications due to their bandwidth density, compatibility with CMOS, and lower waveguide pitch [3], [5], [14]. In addition, Ge-on-silicon-on-insulator (Ge-on-SOI) photodetector [11] provides high responsivity (0.56 A/W) and are sensitive to frequency ranges (850, 1350, and 1550 nm), thus making them ideal for on-chip applications. PROPEL is designed using microring resonators, silicon waveguides, and Ge-on-SOI photodetectors.

C. Intertile Routing and Wavelength Assignment

We adopt dimension-order routing (DOR) for intertile communication, by first routing packets in the x-direction and then in the y-direction [21]. We explain the routing in a single dimension (x) involving four tiles and a similar design can be extended to y-dimension. Fig. 2 shows tiles 0–3 arranged along the x-direction. Every tile modulates the same wavelength into a different waveguide. Each destination tile is associated with a waveguide called the *home channel*. For example, tile $T(0,0)$ has four modulators (ring resonators), all of which are resonant with the wavelength λ_0. Three λ_0 transmissions from tile $T(0,0)$ are used to communicate with the other three tiles $T(1,0)$, $T(2,0)$, and $T(3,0)$ on their home channel waveguides. The fourth resonant wavelength will be used to communicate with the memory bank. As shown in Fig. 2, the home channel for tile $T(0,0)$

consists of four wavelengths, $\Lambda = \lambda_0 + \lambda_1 + \lambda_2 + \lambda_3$ transmitted by tile $T(0,0)$, $T(1,0)$, $T(2,0)$, and $T(3,0)$, respectively. The wavelength-selective filters located at tile $T(0,0)$ will demultiplex all the wavelengths, except for λ_0, which originates from itself and is intended for the memory. Similarly, the wavelengths, λ_0 from tile $T(0,0)$, λ_1 from tile $T(1,0)$, λ_2 from tile $T(2,0)$, and λ_3 from tile $T(3,0)$ are combined and these are used to access the memory banks. These are also the same wavelengths at which the aforementioned tiles will receive data from the memory module. Our goal is to provide a scalable bandwidth to the memory similar to intertile communication. Similar wavelength assignment is replicated even in the y-direction for intertile communication.

The routing and wavelength assignment (RWA) algorithm designed for intertile communication involves selective merging of same wavelengths from source tiles into separate home channels for destination tiles. This design maximizes the bandwidth via WDM and reuses the same wavelengths on different waveguides via spaced division multiplexing (SDM). With recent research showing a potential of 64 wavelengths [5] traveling down one waveguide, the number of wavelengths a tile used to communicate with another tile is 16 wavelengths. This extends the number of wavelengths used by PROPEL to 64, resulting in a total waveguide bandwidth of 640 or 160 Gb/s between tiles. The effective bandwidth of a nanophotonic interconnect is given by $B = W_N \times W_{gN} \times B_R$, where W_N is the number of wavelengths, W_{gN} is the number of waveguides, and B_R is the effective bit rate of the channel. With $W_N = 64$, $W_{gN} = 1$, and $B_R = 10$ Gb/s, we obtain a bandwidth of 640 Gb/s. The bandwidth between tiles is 160 Gb/s with $W_N = 16$, $W_{gN} = 1$, and $B_R = 10$ Gb/s. The electronic switching performs localized arbitration for the output optical transmitters within each tile. As the wavelength for the destination tile is fixed, there is no more contention once the local electronic switching is completed.

D. Scaling PROPEL

PROPEL can be scaled directly to 256 cores by connecting 64 tiles in a 2-D grid manner, but this would result in a large 15×15 crossbar at each tile. In order to reduce the size of the crossbar, we propose an alternate design, called (extended) E-PROPEL, with which we can increase the communication bandwidth without significantly increasing the cost of the network. We utilize an optical crossbar that can provide an $N \times N$ switching functionality. An optical crossbar allows incoming light to be switched from one waveguide to another depending on the wavelength and the input waveguide. However, optical crossbars have dimensions in centimeter scale, making them unfavorable for on-chip applications [22]. Using single and double microring resonators and unique waveguide routing, we design an optical crossbar device with dimensions suitable for on-chip applications.

The proposed E-PROPEL is shown in Fig. 3. We combine four 64-core PROPELs (called a cluster) using optical crossbars to design a 256-core E-PROPEL. This creates a fat tree topology with multiple roots to provide scalable intercluster bandwidth. Every tile with similar coordinates $T(x, y)$ on different clusters

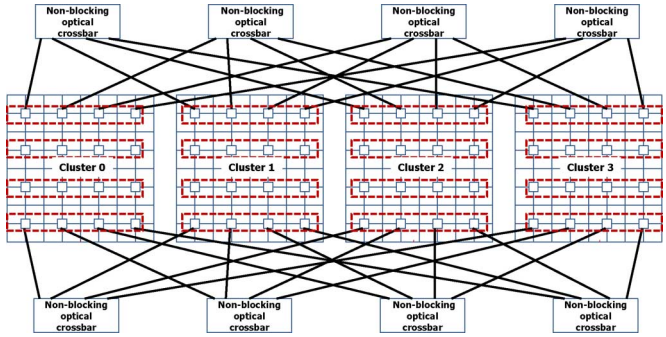

Fig. 3. Proposed E-PROPEL architecture. Each of the cluster is the original PROPEL architecture designed for 64 cores. Intercluster connectivity is established using four-input 64-wavelength optical crossbar implemented using microring resonators. All combination are not shown for clarity.

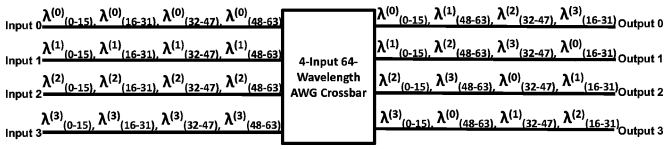

Fig. 4. Four-input, four-output, 64-wavelength optical crossbar functionality.

are connected together with the 4×4 optical crossbars previously designed. For example, tiles $T(0,0)$ on clusters 0, 1, 2, and 3 are connected together with the optical crossbar. Similarly, tiles $T(0,1)$ on clusters 0, 1, 2, and 3 are connected together with another optical crossbar and so on. Therefore, we will require 16 optical crossbars to connect all the tiles from different clusters. This scaling increases the crossbar size to 10×10 with three more connections between clusters. These three additional crossbar inputs and outputs allow packets coming from the L2 cache, x-direction, and y-direction to traverse across the clusters. This reduces the diameter of the network to three; one hop across clusters and two hops within the cluster.

As optical crossbar design occupies considerable area, we evaluate a reduced version of the proposed E-PROPEL, called RE-PROPEL in which we retain the optical crossbars only at the top and bottom of the cluster (identical to what is shown in Fig. 3). This alternate design trades off performance with area and increases the diameter of the network to four: one intracluster hop to get to an optical crossbar, one hop to traverse the optical crossbar, and two hops in the destination cluster. While the proposed E-PROPEL provides scalable bandwidth, there are other technological challenges in implementing the proposed architecture, such as routing signals to and from the chip and area overhead of optical crossbars. These are beyond the scope of the paper.

IV. Crossbar Implementation

A. Optical Crossbar Functionality

Fig. 4 shows an example of a 4×4 64-wavelength optical crossbar that is used in the construction of E-PROPEL. Here, the wavelengths are indicated as $\lambda^{(c)}_{(a-b)}$, where $a - b$ indicate the wavelength range and c indicates the input port. For the 4×4 64-wavelength optical crossbar, consider input port 0. All

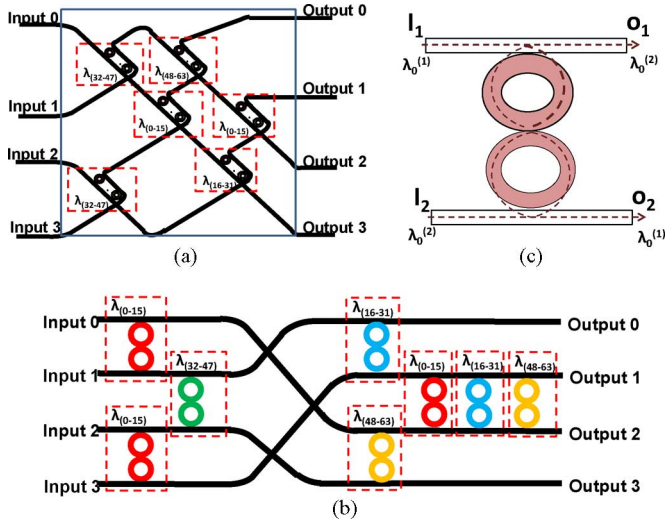

Fig. 5. Proposed four-input 64-wavelength AWG implementations using micro-ring resonator with (a) single-ring resonator and (b) double-ring resonators. (c) Double microring resonator switching two optical light beams.

input wavelengths are indicated as $\lambda^{(0)}_{(0-15)}$, $\lambda^{(0)}_{(16-31)}$, $\lambda^{(0)}_{(32-47)}$, and $\lambda^{(0)}_{(48-63)}$. After traversing the series of ring resonators, the wavelengths $\lambda^{(0)}_{(0-15)}$ arrive at output port 0, $\lambda^{(0)}_{(16-31)}$ arrive at output port 1, $\lambda^{(0)}_{(32-47)}$ arrive at output port 2, and $\lambda^{(0)}_{(48-63)}$ arrive at output port 3. This enables a $1 \times N$ switching functionality per waveguide. Now consider input port 1. All input wavelengths are indicated as $\lambda^{(1)}_{(0-15)}$, $\lambda^{(1)}_{(16-31)}$, $\lambda^{(1)}_{(32-47)}$, and $\lambda^{(1)}_{(48-63)}$. After traversing the series of ring resonators, the wavelengths $\lambda^{(1)}_{(0-15)}$ arrive at output port 1, $\lambda^{(1)}_{(16-31)}$ arrive at output port 2, $\lambda^{(1)}_{(32-47)}$ arrive at output port 3 and $\lambda^{(1)}_{(48-63)}$ arrive at output port 0. In a similar manner, the input wavelengths for input port 2 are $\lambda^{(2)}_{(0-15)}$, $\lambda^{(2)}_{(16-31)}$, $\lambda^{(2)}_{(32-47)}$, and $\lambda^{(2)}_{(48-63)}$. After traversing the series of microring resonators, $\lambda^{(2)}_{(0-15)}$ arrives at output 2, $\lambda^{(2)}_{(16-31)}$ arrive at output 3, $\lambda^{(2)}_{(32-47)}$ arrives at output 0, and $\lambda^{(2)}_{(48-63)}$ arrives at output 1. Lastly, input four wavelengths are given as $\lambda^{(3)}_{(0-15)}$, $\lambda^{(3)}_{(16-31)}$, $\lambda^{(3)}_{(32-47)}$, and $\lambda^{(3)}_{(48-63)}$. After traversing the series of microring resonators, $\lambda^{(3)}_{(0-15)}$ arrives at output 3, $\lambda^{(3)}_{(16-31)}$ arrives at output 0, $\lambda^{(3)}_{(32-47)}$ arrives at output 1, and $\lambda^{(3)}_{(48-63)}$ arrives at output 2. This creates a 4×4 switching functionality device. It should be mentioned that the the aforementioned process can be directly applied to create any size switching devices.

B. Single-Ring Optical Crossbar

This section discusses the construction of an optical crossbar using single microring resonators. Fig. 5(a) shows the single microring resonator implementation of a 64-wavelength 4×4 optical crossbar. It should be noted that this optical crossbar implementation is an extended version of the optical crossbar proposed by Zhou *et al.* [23]. As shown, each waveguide is routed in a manner that allows it to come in close proximity to the other three waveguides. At these close proximity points, a selected range of wavelengths are switched between waveguides. This switching of light between two different waveguides allows light coming from one waveguide to be switched to another waveguide. In terms of functionality, this allows a single tile to have the ability to communicate with multiple other tiles using only one input waveguide. In Fig. 5(a), $\lambda_{(32-47)}$ is switched at the intersection of waveguide 0 and waveguide 1 and also at the intersection of waveguide 2 and waveguide 3, $\lambda_{(0-15)}$ is switched at the intersection of waveguide 0 and waveguide 3 and also at the intersection of waveguide 1 and waveguide 2, $\lambda_{(16-31)}$ is switched at the intersection of waveguide 0 and waveguide 2, and lastly, $\lambda_{(48-63)}$ is switched at the intersection of waveguide 1 and waveguide 3. The aforementioned switching of selected wavelengths at unique waveguide intersections results in a 4×4 64-wavelength optical crossbar.

For a further understanding of the single-ring optical crossbar, we will show how light from waveguide 0 is switched and arrives on the four output waveguides. As light travels down the waveguide 0, it first encounters the intersection with waveguide 1. At this point, $\lambda^{(0)}_{(32-47)}$ is placed on waveguide 1, thus allowing input 0 to communicate with output 2 using $\lambda^{(0)}_{(32-47)}$. Then light traveling down waveguide 0 encounters the intersection with waveguide 3. At this point, $\lambda^{(0)}_{(0-15)}$ is placed on waveguide 3, thus allowing input 0 to communicate with output 0 using $\lambda^{(0)}_{(0-15)}$. As the light continues traveling down, it encounters the intersection with waveguide 2. At this point, $\lambda^{(0)}_{(16-31)}$ is placed on waveguide 2, thus allowing input 0 to communicate with output 1 using $\lambda^{(0)}_{(32-47)}$. Lastly, $\lambda^{(0)}_{(48-63)}$ arrives at output 3, as it was the only light that is not switched. This allows input 0 to communicated with output 0 using $\lambda^{(0)}_{(48-63)}$. This concept is expanded for other inputs which create a 4×4 64-wavelength optical crossbar.

C. Double Rings AWG

This section discusses the construction of an AWG using double microring resonators. A double microring resonator consists of two microring resonators that are placed in between two waveguide to retain the same direction of light from input to output port. Fig. 5(c) shows the operating principle of a double microring resonator allowing light of the same wavelength to be switched between the two waveguides. $\lambda^{(1)}_0$ is switched to the bottom waveguide, and $\lambda^{(2)}_0$ is switched to the top waveguide. Fig. 5(b) shows the double microring resonator implementation of a 64-wavelength 4×4 optical crossbar. It should be mentioned that each microring resonator in the figure represents 16 microring resonators, which are not shown for clarity and would be placed adjacent to each other, thus allowing the 15 additional wavelengths to be switched between waveguides. In the double rings optical crossbar, wavelengths are switched at locations, where two waveguides are running parallel to each other. At this point, light is switched from one waveguide to the other. This switching enables an input port to be connected

TABLE I
COMPARISON BETWEEN THE SINGLE AND DOUBLE MICRORING OPTICAL CROSSBAR DESIGNS

	Single	Double
Micro-ring resonators	96	256
Area (mm^2)	1.38	0.0576
Optical loss (-dB)	4.32	6.23

to all output ports. The operating principle and functionality of the double rings optical crossbar is identical to the single-ring optical crossbar, where the input and output wavelengths are the same. The difference between the two designs is how each one will switch wavelengths. In Fig. 5(b), $\lambda_{(0-15)}$ is switched between waveguide 0 and waveguide 1, waveguide 2 and waveguide 3, and also between waveguide 0 and waveguide 3, $\lambda_{(16-31)}$ is switched between of waveguide 1 and waveguide 2, $\lambda_{(32-47)}$ is switched between waveguide 0 and waveguide 3 and also between waveguide 1 and waveguide 3, and lastly, $\lambda_{(48-63)}$ is between of waveguide 0 and waveguide 2 and also between waveguide 0 and waveguide 3. This creates the functionality identical to the single ring optical crossbar implementation.

D. Comparison

In this section, we compare the two optical crossbar implementations in terms of area, optical loss, and number of microring resonators. For each design, we take into account the achievable loss, bandwidth, and crosstalk of each optical crossconnect (optical switch).

1) Single-ring Optical Crossbar Analysis: The single-ring optical crossbar consists of three and four sets of microring resonators in the vertical and horizontal directions, respectively. The estimated area overhead for each set is 60 μm $\times$ 90 μm. As the four-input 64 wavelengths optical crossbar consists of 16 microrings resonators at each set, the horizontal and vertical lengths will increase by 16-fold. This causes the four-input 64 wavelength optical crossbar area to be 1440 μm $\times$ 960 μm. In addition, each waveguide has three waveguide crossings and traverses a maximum distance of about 1.7 mm. Moreover, the optical crossbar is constructed with 96 microring resonators.

2) Double Rings Optical Crossbar Analysis: The double rings optical crossbar consists of a total of eight sets of double-ring resonators. Each set is comprised of 32 microring resonators or a total of 256 microring resonators are used to construct the double rings optical crossbar. The optical crossbar contains 96 double microring resonators, where each double microring resonator has dimensions of 15 μm $\times$ 20 μm. This includes a 5-μm spacing between each double microring resonator that is used to prevent cross coupling between adjacent microring resonators. This results in the optical crossbar having a height of 80 μm (three double microring resonator sets and four waveguides) and a width of 720 μm (three double microring sets). The double rings optical crossbar has two waveguide crossings with a maximum distance of about 1 mm.

Table I shows the calculated values for each optical crossbar design. In calculating the optical loss, we used a waveguide loss of -1.3 dB/cm, a microring traversal loss of -1dB, waveguide crossover loss of -0.05 dB, and a bending loss of -1 dB. The single-ring optical crossbar has less optical power loss than the double rings optical crossbar, mainly due to the fact that the double rings incur a traversal loss of -2 dB. The major contribution of optical loss for the single-ring optical crossbar is bending losses, as multiple bends are required to accommodate the 16-wavelength switching. The double rings optical crossbar occupy less area, which may come as a surprise, as it uses 160 more microring resonators. The reason that the double rings optical crossbar has less area overhead than the single-ring optical crossbar is because multiple ring resonators can be stacked on top of each other due to the different wavelengths being switched. Moreover, the single-ring optical crossbar requires more bending than the double rings optical crossbar, which increases the area overhead. In comparison, the single-ring optical crossbar design should be used if optical loss needs to be minimized, and the double rings optical crossbar design should be used if area overhead needs to be minimized.

V. PERFORMANCE EVALUATION

In this section, we compare PROPEL to mesh and flattened-butterfly [24] in terms of throughput, latency, and power for synthetic and SPLASH-2 traffic traces. In addition, we provide a comparison between each optical network in terms of optical components and the minimum optical power required for each network.

A. Optical Hardware Complexity and Power Analysis

We analytically compare the optical hardware complexity in terms of wavelengths, optical components (waveguides, ring resonators), and total optical power required. For all networks, we assume an off-chip laser source and the following losses consistent across all networks [3]–[5], [14]: a star splitter loss (L_S) of $-3(\log_2 N)$, where N is the number of times the waveguide is split, a splitter/coupler loss (L_C) of -3 dB (50% loss of signal), off-chip laser-to-fiber coupling loss ($L_{\rm LF}$) of -0.5 dB, off-chip to on-chip fiber-to-waveguide coupling loss ($L_{\rm FW}$) of -2 dB, waveguide loss (L_W) of -1.3 dB/cm, bending loss (L_B) of -1 dB, a modulator traversal loss (L_M) of -1 dB, a waveguide crossover loss (L_X) of -0.05 dB, and a waveguide-to-receiver loss ($L_{\rm WR}$) of -0.5 dB. The aforementioned component losses includes both scattering and crosstalk losses to account for the maximum potential optical loss of a component. It should be mentioned that there is sufficient spacing between components to minimize crosstalk.

PROPEL uses a total of 3072 ring resonators (192 per tile, 96 each for x- and y-directions), 32 silicon waveguides (16 each for x and y-directions), and 1536 photodetectors (96 per tile), which results in PROPEL having a total optical area of 64.6 mm^2. In addition, PROPEL is comprised of 16 electrical routers and 1536 optical receivers (96 per tile), resulting in a total electrical area of 50 mm^2. The maximum power loss in PROPEL is given by $L_S + L_{\rm LF} + L_{\rm FW} + 2 \times L_M + L_{\rm WR} + 4 \times L_B + 32 \times L_X + L_W$, where L_S will be -15 dB ($=-3\log_2 32$) and L_W will be -6.5 dB. This makes the total power loss to be -32.1 dB.

Table II shows various optical components and losses of various photonic interconnects for 256 cores. We compare the

TABLE II
OPTICAL HARDWARE COMPLEXITY COMPARISON BETWEEN VARIOUS ON-CHIP PHOTONIC ARCHITECTURES FOR 256 CORES

	Proc. DRAM [5]	CORONA [3]	PROPEL	E-PROPEL
Wavelengths	64	64	64	64
Waveguides	128	387	272	208
Ring Resonators	24,576	1,081,344	28,672	19,968
Power Loss (dB)	-36	-49	-44	-42
Optical Area mm^2	22	337	181	96
Electrical Area mm^2	186	860	395	280
Photodetectors	4,096	32,768	14,336	9,216

Processor-DRAM architecture [5], as this was designed for 256 cores. It should be noted that we did not consider the electrical area overhead for the mesh within the Processor-DRAM architecture that is used for intercore communication within a group. PROPEL and E-PROPEL are designed for core-to-core communication and utilize comparable components and devices. CORONA requires almost twofold and 3.5-fold of optical area when compared to PROPEL and E-PROPEL. PROPEL and E-PROPEL require almost threefold lesser electrical area than CORONA architecture. E-PROPEL requires lesser nanophotonic components and occupies lesser area than PROPEL for 256 cores. The proposed architectures PROPEL and E-PROPEL provide a balanced architecture that reduces the number of nanophotonic and electronic components to design an area-efficient and cost-efficient on-chip architecture.

B. Power Estimates

For electrical interconnects, we consider wires implemented in semiglobal metal layers for interrouter links. The wire capacitances, resistances, and device parameters were obtained from International Roadmap for Semiconductors [25] and Berkeley Predictive Technology models [26]. At 22 nm with a flit size of 128 bits, the power dissipation will be 198 mW considering a 9-GHz link. A flit, or flow control digit, is the smallest data unit of a packet that can be individually routed [21]. To reduce the power dissipation at future technology nodes, we reduce the network frequency to 2 GHz and reduce the power consumption to 44 mW, which is comparable to power values from [15]. Pan *et al.* [15] used an energy of 19 pJ per flit per hop or 38 mW per flit per hop given a 2 GHz clock.

At 22 nm, we estimate the buffer power to be 8.06 mW and occupies an area of 185 μm^2, which is similar to the power value estimated from [27]. Grot *et al.* [27] use a energy of 61.7 pJ for a 567-bit flit at 45 nm technology. If a 128-bit flit is scaled to 22 nm, the buffer power would be 6.84 mW. A 5×5 matrix crossbar with tristate buffer connectors [28] is considered for the regular NoC design. The area of the crossbar is estimated by the number of input/output signals that it should accommodate. At 22 nm, we estimate the power value for a 5×5 crossbar to be 8.66 mW and this value is similar to the power value estimated from [27]. We believe that the values obtained here for electrical components (buffers, links, and crossbars) [15], [27] closely matches to other network designs, giving us confidence in our calculations. For optical links, we assume a power dissipation of 1.1 mW/(Gb·s) per link [11] and a power dissipation of 0.1 mW/(Gb·s) per modulator [12].

TABLE III
CORE AND CACHE PARAMETERS USED FOR SPLASH-2 SUITE SIMULATION

Parameter	Value
L1/L2 coherence	MOESI
L2 cache size/accoc	4MB/16-way
L2 cache line size	64
L1 cache/accoc	64KB/4-way
L1 cache line size	64
Core Frequency(GHz)	2.5
Threads(core)	2
Issue policy	In-order
Memory Size(GB)	4
Memory Controllers	16

C. 64-core SPLASH-2 Traffic Results

We evaluated PROPEL and compared to mesh and flattened-butterfly [24] using real SPLASH-2 application traces that were collected using Simics, a full system simulator [29]. Table III shows the tile parameters used to evaluate PROPEL. During each Simics simulation, the multifacet general execution-driven multiprocessor simulator (GEMS) [30] package was enabled for correct cache coherence requests. Once the traces were collected, each trace was run on the cycle accurate simulator called OPTISIM [31]. We use the following SPLASH-2 application benchmarks: fast Fourier transform (FFT) (16K particles), LU (512×512), radiosity (largeroom), ocean (258×258), raytrace (teapot), radix (1 M integers), water (512 molecules), FMM (16K particles), and barnes (16K particles).

Fig. 6(a) shows the speed up relative to mesh for the select SPLASH-2 benchmarks. In the ocean and radix benchmarks, PROPEL has the highest speed up factor of more than 2.5. This results because ocean and radix applications have a higher percentage of nonlocal traffic that requires multiple hops in a mesh network. Since PROPEL has a maximum hop count of two, nonlocal traffic would see a higher speed up than mesh. The LU, water, FFM, and barnes applications are comprised of both local and nonlocal traffic. This results in a speed up of around 2–2.5 for these applications. Lastly, FFT and radiosity are comprised of mainly local traffic, which is why PROPEL has the least speed up for these benchmarks.

Fig. 6(b) shows the power dissipation relative to mesh. From the figure, PROPEL dissipates about tenfold less power than mesh and about sixfold less power than flattened-butterfly for each application. PROPEL uses less power than mesh because PROPEL has an average hop count of two and mesh has an average hop count of eight. During each hop in PROPEL, the power dissipation is the addition of one crossbar traversal and one buffer storage (15 mW per flit per router traversal). In should be mentioned that the power dissipation for a flit to traverse

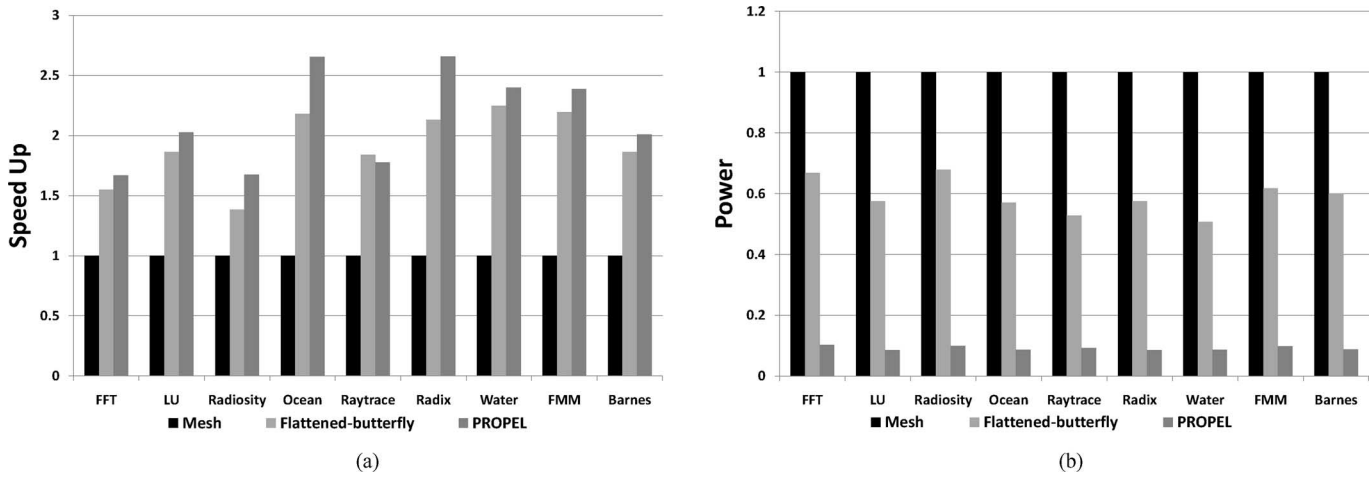

Fig. 6. Simulation results showing (a) speed up and (b) power dissipation for select SPLASH-2 traffic traces.

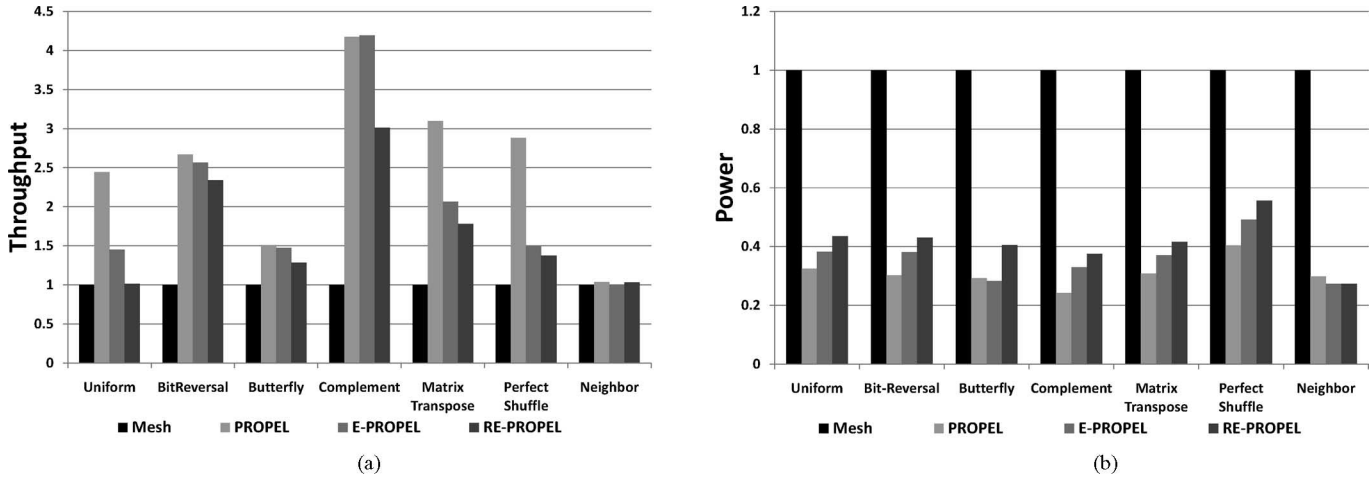

Fig. 7. Simulation results showing (a) saturation throughput and (b) power dissipation for different synthetic traffic patterns for 256 cores. With $M = 8$ and $N = 8$, PROPEL is designed for 256 cores, E-PROPEL is designed by using four 64-core PROPELs with 16 optical crossbars, and RE-PROPEL is designed with four 64-core PROPELs and eight optical crossbars.

across an optical link is not included in PROPEL's average power dissipation, as power dissipation of an optical link is much less than an electrical link. This results in PROPEL having an average power dissipation per flit of 45 mW for two hops. As for mesh, the power dissipation is the addition of one crossbar traversal, one buffer storage, and one link traversal (59 mW per flit per router traversal). This results in mesh having an average power dissipation per flit of 487 mW if we assume an average hop count of 8. The average power dissipation difference between mesh and PROPEL is about 10.8, which is what is seen in Fig. 6(b). Flattened-butterfly dissipates lower power than mesh, as it requires lesser number of routers.

D. 256-core Synthetic Traffic Results

In evaluating E-PROPEL, we use synthetic traffic for emulation due to the difficulty in obtaining SPLASH-2 traces for 256 cores. Fig. 7(a) shows the normalized throughput and Fig. 7(b) shows the normalized power dissipation for mesh, PROPEL, E-PROPEL, and RE-PROPEL. As shown, PROPEL outperforms all other networks for 256 cores, but consumes more area and increases the complexity of the switch. E-PROPEL reduces the size of the crossbar with minimal loss in performance. RE-PROPEL reduces the number of optical crossbars from 16 to 8 and this results in some loss in performance, as packets have to go over one extra hop to reach the edge tiles. From Fig. 7(b), we can see that PROPEL, E-PROPEL and RE-PROPEL consume power in increasing order, as they require more hops. PROPEL topologies dissipate 70% less power than mesh topology for 256 cores.

VI. Conclusion

In this paper, we proposed an on-chip scalable NoC called PROPEL that uses emerging nanophotonic components to overcome the limited bandwidth and high power dissipation bottlenecks found metallic based NoCs. Our analysis clearly shows significant saving (tenfold) in power and an increase in performance (twofold), when PROPEL is compared to both mesh and flattened-butterfly for Splash-2 suite. Moreover, this architecture

has the desirable features identical to a mesh architecture, which can be scaled in two dimensions, and provides fault-tolerance due to multipath connectivity. We also analyzed the scalability of PROPEL architecture and developed an extended E-PROPEL architecture. E-PROPEL and RE-PROPEL reduces the crossbar radix while delivering scalable performance for 256 cores. We also propose two different optical crossbar implementations using single and double microring resonators.

REFERENCES

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

[2] L. Benini and G. D. Micheli, "Networks on chips: A new soc paradigm," *IEEE Comput.*, vol. 35, no. 1, pp. 70–78, Jan. 2002.

[3] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. Jouppi, M. Fiorentino, A. Davis, N. Binker, R. Beausoleil, and J. H. Ahn, "Corona: System implications of emerging nanophotonic technology," in *Proc. 35th Int. Symp. Comput. Archit.*, Jun. 2008, pp. 153–164.

[4] A. Shacham, K. Bergman, and L. P. Carloni, "Photonic networks-on-chip for future generations of chip multiprocessors," *IEEE Trans. Comput.*, vol. 57, no. 9, pp. 1246–1260, Sep. 2008.

[5] C. Batten, A. Joshi, J. Orcutt, A. Khilo, B. Moss, C. Holzwarth, M. Popovic, H. Li, H. Smith, J. Hoyt, F. Kartner, R. Ram, V. Stojanovi, and K. Asanovic, "Building manycore processor-to-dram networks with monolithic silicon photonics," in *Proc. 16th Annu. Symp. High-Perform. Interconnects*, Stanford, CA, Aug. 27–28, 2008.

[6] D. A. B. Miller, "Device requirements for optical interconnects to silicon chips," *Proc. IEEE*, vol. 97, no. 7, pp. 1166–1185, Jun. 10, 2009.

[7] (2007). [Online]. Available: http://www.intel.com/pressroom/kits/hpc/index.htm

[8] [Online]. Available: http://www.hpcwire.com/topic/networks/finisar-intros-150-gbps-parallel-active-optical-cable-48806432.html

[9] S. Manipatruni, Q. Xu, B. Schmidt, J. Shakya, and M. Lipson, "High speed carrier injection 18 gb/s silicon micro-ring electro-optic modulator," in *Proc. IEEE/LEOS 2007*, Oct. 21–25, pp. 537–538.

[10] I. O'Connor, "Optical solutions for system-level interconnect," in *Proc. 2004 Int. Workshop Syst. Level Interconnect Prediction*. ACM, New York, NY, 2004, pp. 79–88.

[11] S. J. Koester, C. L. Schow, L. Schares, and G. Dehlinger, "Ge-on-soi-detector/si-cmos-amplifier receivers for high-performance optical-communication applications," *J. Lightw. Technol.*, vol. 25, no. 1, pp. 46–57, Jan. 2007.

[12] Q. Xu, S. Manipatruni, B. Schmidt, J. Shakya, and M. Lipson, "12.5 gbit/s carrier-injection-based silicon micro-ring silicon modulators," *Opt. Exp.: Int. Electron. J. Opt.*, vol. 15, no. 2, pp. 430–436, Jan. 2007.

[13] B. Jalali and S. Fathpur, "Silicon photonics," *J. Lightw. Technol.*, vol. 24, pp. 4600–4615, 2006.

[14] N. Kirman, M. Kirman, R. Dokania, J. Martinez, A. Apsel, M. Watkins, and D. Albonesi, "Leveraging optical technology in future bus-based chip multiprocessors," in *Proc. 39th Int. Symp. Microarchitecture*, Dec. 2006, pp. 492–503.

[15] Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, and A. Choudhary, "Firefly: Illuminating future network-on-chip with nanophotonics," in *Proc. 36th Annu. Int. Symp. Comput. Archit.*, 2009, pp. 429–440.

[16] M. Cianchetti, J. Kerekes, and D. Albonesi, "Phastlane: A rapid transit optical routing network," presented at the 36th Int. Symp. Comput. Archit., Austin, TX, USA, Jun. 2009.

[17] M. Haurylau, G. Chen, H. Chen, J. Zhang, N. Nelso, D. Albonesi, E. Friedman, and P. Fauchet, "On-chip optical interconnect roadmap: Challenges and critical direction," *IEEE J. Sel. Topics Quantum Electron.*, vol. 12, no. 6, pp. 1699–1705, Nov./Dec. 2006.

[18] J. Balfour and W. J. Dally, "Design tradeoffs for tiled cmp on-chip networks," in *Proc. 20th ACM Int. Conf. Supercomputing (ICS)*, Cairns, Australia, Jun. 28–30, 2006, pp. 187–198.

[19] W. M. J. Green, M. J. Rooks, L. Sekarie, and Y. Vlasov, "Ultra-compact, low rf power, 10 gb/s silicon mach-zehnder modulator," *Opt. Exp.*, no. 25, pp. 17 106–17 113, Dec. 2007.

[20] A. Gondarenko, J. Levy, and M. Lipson, "High confinment micro-scale silicon nitride high q ring resonator," *Opt. Exp.*, vol. 17, no. 14, pp. 11 366–11 370, Jul. 2009.

[21] W. J. Dally and B. Towles, *Principles and Practices of Interconnection Networks*. San Fransisco, CA: Morgan Kaufmann, 2004.

[22] Y. Lin, N. Rahmanian, S. Kim, G. N. an C. Topping, D. Smith, and J. Ballato, "Ultracompact awg using air-trench bends with perfluorocyclobutyl polymer waveguides," *J. Lightw. Technol.*, vol. 26, no. 17, pp. 3062–3070, Sep. 2008.

[23] L. Zhou, S. Djordjevic, R. Proietti, D. Ding, S. Yoo, R. Amirtharajah, and V. Akella, "Design and evaluation of an arbitration-free passive optical crossbar for on-chip interconnection networks," *Appl. Phys. A: Mat. Sci. Process.*, vol. 95, no. 4, pp. 1111–1118, Jun. 2009.

[24] J. Kim, W. J. Dally, and D. Abts, "Flattened butterfly: Cost-efficient topology for high-radix networks," in *Proc. 34th Annu. Int. Symp. Comput. Archit. (ISCA)*, Jun. 2007, pp. 126–137.

[25] I. T. R. for Semiconductors. (2007). [Online]. Available: http://www.itrs.org

[26] Predictive Technology Model. (2008). [Online]. Available: http://www.eas.asu.edu/∼ptm/

[27] B. Grot, J. Hestness, S. W. Keckler, and O. Mutlu, "Express cube topologies for on-chip interconnects," in *Proc. 15th Int. Symp. High Perform. Comput. Archit.*, Feb. 2009, pp. 163–174.

[28] H. S. Wang, X. Zhu, L. S. Peh, and S. Malik, "Orion: A power-performance simulator for interconnection networks," in *Proc. 35th Annu. ACM/IEEE Int. Symp. Microarchitecture*, Istanbul, Turkey, Nov. 18–22, 2002, pp. 294–305.

[29] S. Woo, M. Ohara, E. Torrie, J. Singh, and A. Gupta, "The splash-2 program: Characterization and methodological considerations," in *Proc. Int. Symp. Comput. Archit.*, 1995, pp. 24–36.

[30] M. Martin, D. Sorin, B. Beckmann, M. Marty, M. Xu, A. Alameldeen, K. Moore, M. Hill, and D. Wood, "Multifacet's genreal execution-driven multiprocessor simulator (gems) toolset," *ACM SIGARCH Comput. Archit. News*, vol. 33, no. 4, pp. 92–99, Nov. 2005.

[31] A. Kodi and A. Louri, "A system simulation methodology of optical interconnects for high-performance computing (hpc) systems," *J. Opt. Netw.*, vol. 6, pp. 1282–1300, Dec. 2007.

Randy Morris (S'09) received the B.S. and M.S. degrees in electrical engineering and computer science from Ohio University, Athens in 2007 and 2009, respectively, from where he is currently working toward the Ph.D. degree. His research interests include optical interconnects, network-on-chips, and computer architecture.

Avinash Karanth Kodi (M'07) received the Ph.D. and M.S. degrees in electrical and computer engineering from the University of Arizona, Tucson, in 2006 and 2003, respectively.

He is currently an Assistant Professor of electrical engineering and computer science with Ohio University, Athens. His research interests include computer architecture, optical interconnects, chip multiprocessors, and network-on-chips.