

Physical-Layer Modeling and System-Level Design of Chip-Scale Photonic Interconnection Networks

Johnnie Chan, *Student Member, IEEE*, Gilbert Hendry, *Student Member, IEEE*, Keren Bergman, *Fellow, IEEE*, and Luca P. Carloni, *Senior Member, IEEE*

Abstract—Photonic technology is becoming an increasingly attractive solution to the problems facing today’s electronic chip-scale interconnection networks. Recent progress in silicon photonics research has enabled the demonstration of all the necessary optical building blocks for creating extremely high-bandwidth density and energy-efficient links for on-chip and off-chip communications. From the feasibility and architecture perspective however, photonics represents a dramatic paradigm shift from traditional electronic network designs due to fundamental differences in how electronics and photonics function and behave. As a result of these differences, new modeling and analysis methods must be employed in order to properly realize a functional photonic chip-scale interconnect design. In this paper, we present a methodology for characterizing and modeling fundamental photonic building blocks which can subsequently be combined to form full photonic network architectures. We also describe a set of tools which can be utilized to assess the physical-layer and system-level performance properties of a photonic network. The models and tools are integrated in a novel open-source design and simulation environment. We present a case study of two different photonic networks-on-chip to demonstrate how our improved understanding and modeling of the physical-layer details of photonic communications can be used to better understand the system-level performance impact.

Index Terms—Optical communications, optical crosstalk, optical losses, photonic interconnection networks, simulation software, system analysis and design.

Manuscript received August 21, 2010; revised January 24, 2011; accepted April 21, 2011. Date of current version September 21, 2011. This research is supported in part by the Defense Advanced Research Projects Agency (DARPA) Microsystems Technology Office, under Grant ARL-W911NF-08-1-0127, the National Science Foundation, under Award 0811012, and the Interconnect Focus Center, one of the five research centers funded under the Focus Center Research Program by the Semiconductor Research Corporation and DARPA Program. The views, opinions, and findings contained in this paper are those of the authors and should not be interpreted as representing the official views or policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the Department of Defense. An earlier version of this paper was presented at the 2010 Design, Automation Test in Europe Conference (DATE), Dresden, Germany, March 2010. This paper was recommended by Associate Editor R. Marculescu.

J. Chan, G. Hendry, and K. Bergman are with the Department of Electrical Engineering, Columbia University, New York, NY 10027 USA (e-mail: johnnie@ee.columbia.edu; gilbert@ee.columbia.edu; bergman@ee.columbia.edu).

L. P. Carloni is with the Department of Computer Science, Columbia University, New York, NY 10027 USA (e-mail: luca@cs.columbia.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/TCAD.2011.2157157

I. INTRODUCTION

THE COMPUTING industry has been steadily increasing the number of cores on a single processor die to raise performance through parallel computation. In order for chip multiprocessor (CMP) systems to properly function and take advantage of the multiple cores, interconnection networks are required to provide both on-chip (e.g., core to core) and off-chip (e.g., core to main memory) communication links. Electronic interconnection networks have thus far been capable of coping with the communication demands of today’s applications, however further scaling of the number of cores and memory requirements of applications may not be able to be matched with equivalent communication improvements [1]. This inability of electronic interconnects to scale in performance is a consequence of the large amount of power that is dissipated by the electronic interconnects and the limited heat-dissipation capabilities of current packaging technology. A previous study has shown that over 50% of the dynamic power dissipated in some high-performance microprocessors comes from the interconnects exclusively [2]. These problems will be further exacerbated as the requirements of CMP systems continue to grow, exemplifying the need for a new interconnect technology that can deliver energy-efficient high-bandwidth communications.

Photonics technology has emerged as a promising chip-scale interconnect solution to the various challenges facing CMP scaling. Photonic signaling using wavelength division multiplexing (WDM) can enable orders of magnitude higher bandwidth density than electronics which is becoming increasingly constrained by the wire and pin densities that can be achieved [1]. The power dissipation of photonic signaling can be designed to be practically independent of distance and data rate. This allows for high-speed data to flow seamlessly between the on-chip and off-chip domains. All the necessary optical devices for creating chip-scale photonic interconnection networks have been demonstrated using complementary metal-oxide-semiconductor (CMOS)-compatible fabrication techniques, as described in Section IV. This compatibility allows them to be economically produced in existing fabrication lines. Moreover, CMOS compatibility allows these optical devices to be directly integrated with electronic digital circuits, providing a flexible and powerful means to create a high-performance interconnect fabric.

In light of these recognized advantages of optics, many challenges still exist in fathoming and creating a chip-scale

photonic network. Chip-scale silicon photonics currently provides no means of in-flight buffering and logical processing. The only way to accomplish these tasks is to use optical-electronic-optical (O-E-O) conversion, and do them in the electronic domain. O-E-O conversion is tolerated in switches for large-scale networks, however the additional power dissipation required would have a detrimental impact on chip-scale systems. Signal regeneration is also not easily accomplished in the CMOS-compatible photonic platform, therefore network architecture designs must carefully consider the optical losses to ensure signal integrity throughout the transmission path. Conventional network simulators are not well suited for photonic networks since they are incapable of capturing all the physical and functional details that differentiate photonic devices from their electronic counterparts.

In this paper, we present a methodology for designing, modeling, and analyzing the performance of photonic interconnection networks. Furthermore, this paper highlights several techniques to synergistically study a photonic architecture's system-level properties through physical-layer analysis. We have also developed the PhoenixSim environment which integrates the modeling and analysis aspects of our methodology and has been made publicly available [3]. PhoenixSim is implemented using OMNeT++, a C++-based event-driven simulation environment [4]. Our methodology and PhoenixSim represent a novel set of tools which system architects can use to see how integrated photonics can potentially impact the performance of a particular computing system.

In Section II, we review related work in the area of photonic network architectures and other simulators and tools that have been developed to model and design them. In Section III, we present our methodology and outline the design flow supported by PhoenixSim which guides designers from the modeling of the basic silicon photonics devices, through the composition of the devices into a complex interconnection network, to the analysis of the network performance and scalability. In Section IV, we describe our method for modeling photonic devices and overview a library of fundamental building blocks we have implemented. In Section V, we present a unique set of analysis tools for optimizing photonic interconnection networks. We then discuss two photonic network-on-chip architectures as case studies in Section VI to demonstrate how our methodology can be used to design and understand photonic interconnection networks. We finish with our concluding remarks in Section VII.

This paper expands upon work that has been published previously [5]. New contributions include: 1) a complete discussion of our proposed design methodology; 2) the modeling of Mach-Zehnder switches in our *Photonic Device Library*; 3) the expansion of our noise model to handle intra-message crosstalk and receiver noise; and 4) inclusion of serialization/de-serialization (SerDes) power within our power model. We also conduct our case study on new network topologies and include a more extensive set of results and analysis. Lastly, we include an expanded literature review of photonic interconnection network architectures

and photonic interconnect computer-aided design (CAD) tools.

II. RELATED WORK

A. Photonic Interconnection Networks

Many photonic interconnect architectures have been proposed in recent years for improving the performance and efficiency of computing systems. Photonics offers an entirely new design space to explore and optimize, potentially causing paradigm shifts in a variety of architectural issues such as memory hierarchies, programming models, and physical layouts. Many topologies have been proposed by the research community, consisting of designs that leverage wavelength-, spatial-, and time-based routing. Below, we review a sample of the various proposed topologies.

The Corona network is a proposed network that integrates photonic interconnects for use in both on-chip and off-chip communications [6]. The on-chip network is implemented with a serpentine crossbar and arbitration is handled with an optical token-ring based scheme and transmission is conducted using wavelength routing. Each memory controller can interface will multiple memory modules to create a memory solution with scalable capacity.

Joshi *et al.* proposed a mapping of the Clos topology for use as an on-chip photonic network [7]. This proposed network architecture connects a 64-core CMP through a global photonic crossbar. The network is source routed and arbitration is accomplished through the selection of an optical transmission wavelength and transmission waveguide from a wide optical bus. Center stages of the Clos are performed electronically, resulting in an O-E-O conversion. Other networks have also been proposed that use similar wavelength-routing scheme for arbitration [8], [9].

Batten *et al.* proposed an off-chip photonic interconnection network for connecting each core of a CMP to multiple memory banks [10]. The topology uses a crossbar matrix, implemented using ring filters for wavelength-based routing, to fully connect each on-chip computing node to all memory banks. The network is nonblocking and only requires a simple source-routing scheme for transmission. The Phastlane topology is another dedicated photonic network used for handling cache-coherency of a shared-memory CMP [11].

A circuit-switching style network was originally proposed by Shacham *et al.* that uses active switching for the spatial routing of photonic messages [12]. The network dedicates the entire optical spectrum for data parallelism to create high-bandwidth links using WDM, which contrasts with the aforementioned topologies that leverage the spectrum for routing purposes. Further studies of this routing scheme have also been conducted for varying topologies [13], [14] and applications [15].

Hendry *et al.* [16] proposed time-division multiplexing (TDM) as an alternative to circuit-switching to eliminate the latency overhead of the path-setup protocol. TDM routing is enabled by temporally dividing the transmission medium into a constant series of frames. Multiple time slots compose each frame and a unique network configuration consisting of some




Fig. 1. Design flow of modeling a network in the PhoenixSim environment.

network connectivity is assigned to each time slot. The set of all time slots within a frame will contain a set of connections that completely connect all nodes within the network.

B. CAD Tools

As the interest for using photonic interconnects continues to grow, so does the need for tools that can harness the potential of this new technology. In the realm of simulation, two levels exist which are of interest to photonic network designers: link-level and system-level. Simulation is an especially important predictive tool for gauging the performance of these photonic interconnect systems which are too complex for manufacturing in current fabrication technology. Beyond simulation, design tools will be needed to effectively and accurately design complex and efficient photonic interconnection networks. Most conventional simulation and design tools are not ideally suited for capturing the physical and performance characteristics of chip-scale photonic interconnection devices and networks. Therefore the development of photonically enabled tools is needed to fill the void.

As photonic interconnect topologies are becoming increasingly complex, layout tools and optimization techniques will be required for efficient and accurate design. Ding *et al.* have developed Optical Interconnect Library (OIL) a synthesis-like CAD tool for optimizing optical router designs in terms of insertion loss [17]. The methodology allows for constraint based optimization in terms of latency and insertion loss. Similarly, Minz *et al.* have devised a synthesis tool for timing-driven optimization of optical waveguide placement in an on-chip network [18]. VANDAL is a place-and-route tool for on-chip photonic architectures which uses a library of modeled and characterized components, and includes automation tools for rapid design and synthesis [19].

With link-level simulation, the primary concern is detailed physical modeling of all the end-to-end aspects of a photonic path to determine performance metrics such as signal integrity and link reliability. O'Connor *et al.* [20] proposed a link-level simulation environment for heterogeneous photonic integrated circuits which leverages detailed synthesizable models of building-block components for the purpose of determining interconnect density, area, link delay, and link power requirements. Similarly, De Wilde *et al.* [21] presented an approach for characterizing CMOS-to-CMOS links in terms

of timing, error rates, and noise sensitivity. The IBM optical link simulator was created to design and analyze telecom-scale and LAN-scale links through metrics such as failure rates, power penalties, and signal performance (e.g., eye diagrams) [22].

System-level simulation uses a higher-level of abstraction than link-level simulation and is primarily concerned with determining network performance metrics (e.g., bandwidth, application latency, and system power dissipation). Briere *et al.* [23] have developed the ONoC SystemC model which focuses on the simulation of optical networks-on-chip using the SystemC framework and primarily addressing high-level system concerns including device timing and network-level power dissipation. Their modeling is currently specific to topologies that leverage the *lambda router*, which routes optical traffic based on the wavelength of light that is being used by the source. Optisim is a system-level simulator for modeling optical interconnects in board-based and cluster-based computing [24].

PhoenixSim is primarily categorized as a system-level simulation environment that includes some aspects of link-level simulation. Our PhoenixSim environment closely resembles Optisim with respect to the use of a photonic building block library, and extractability of physical and system metrics. We differentiate our work from Optisim through combination of our focus on chip-scale systems, support for spatial and temporal based photonic chip-scale architectures, and synergistic study of physical-layer and system-level performance metrics.

III. METHODOLOGY AND DESIGN FLOW OVERVIEW

An overview of our design methodology is illustrated in Fig. 1. The sequence of design stages we employ for modeling photonic interconnection networks primarily consists of six design steps: 1) specification of the network building blocks; 2) specification of the target application; 3) modeling of the network architecture; 4) system-level performance analysis; 5) physical-layer characterization; and 6) iterative refinement of parameters and design.

Step 1 (as labeled in Fig. 1) involves the specification of the fundamental network building blocks that will be used for creating the interconnection network. The collection of




Fig. 2. Subset of the photonic devices in the *Interconnect Building Block Library*.

network building blocks is named the *Interconnect Building Block Library*. Within this library is a set of photonic devices that are characterized using the Basic Element Device Model (Fig. 2), described in further detail in Section IV. Users of this design methodology can choose to design a network based on the included library of devices, or extend the library themselves with other novel photonic building blocks.

The library for electronic building blocks consists of switch, arbitrator, and buffer blocks for creating standard pipelined routers. PhoenixSim leverages the Orion simulator [25] for deriving detailed values for electronic delay and energy dissipation. The electronic router model is highly configurable and includes parameters for clock rate, buffer size, channel width, and number of virtual channels. In addition to the standard router design, the electronic router model also includes additional methods for interfacing with photonic devices. Electro-optic photonic devices can take an electronic input to influence its optical behavior and are essential components for enabling the active types of switching used in some proposed networks [6], [12].

Next, Step 2 consists of specifying the target application. PhoenixSim currently supports the use of both synthetically generated traffic patterns and communication traces, with eventual plans for integration with a cycle-accurate microarchitecture simulator. A variety of synthetic patterns have already been created within the environment (e.g., random, hotspot, nearest neighbor, and tornado) and is extensible to others. Communication traces can be generated by monitoring the network traffic during the execution of a real application and used as an input into PhoenixSim. Performance results gained by using communication traces are useful in assessing the application-specific performance gains of photonic networks [15].

The design and modeling of the network occurs in Step 3 of the design flow. The devices from the *Interconnect Building Block Library* can be combined to create higher-order networking components and entire interconnection network topologies. By accounting for the target applications, a network architect can optimize the topology design to target specific requirements such as message size, latency, and/or throughput. For instance, Fig. 3 illustrates how a 4×4 nonblocking switch can be derived within PhoenixSim by connecting various devices from the *Building Block Library*.

Step 4 involves the characterization of the network architecture at the physical layer, which involves metrics such as the optical power budget, crosstalk, and power dissipation. The overall physical-layer performance of a




Fig. 3. (a) Schematic of a design for a 4×4 nonblocking photonic switch. (b) Screenshot of how PhoenixSim composes the switch by instantiating basic photonic devices.

derived photonic component or topology can be determined from the aggregate performance of the individual photonic devices. Although this is not as rigorous as a true link-level simulator, this hierarchical building process enables an accurate first-order physical characterization of an entire network through the characterization of a small number of foundational components.

Step 5 measures the system-level performance characteristics of the network architecture in terms of data throughput and latency. Many of the physical properties that are identified in Step 4 have an impact on network functionality and scalability and play a crucial role in determining overall system performance.

Finally, Step 6 forms the basis for an iterative process, where the performance results and analysis of the modeled network can be used to refine the topology design and device parameters to further optimize the overall performance. Previous work has demonstrated the effectiveness of this iterative step. The initial physical-layer characterizations showed the dramatic impact that waveguide crossing loss had on performance and a subsequent analysis of a system with improved crossings resulted in a dramatic improvement in overall performance [14].

IV. PHOTONIC DEVICE LIBRARY

Our method for modeling photonic devices is designed to enable the assessment of the physical-layer performance at a first-order approximation while concurrently allowing for system-level analysis with a reasonable computational requirement. Many simulation packages use techniques such as finite-difference time-domain (FDTD) to accurately model an electromagnetic field according to Maxwell's equations. FDTD analysis, however, is usually limited to a single or small set of devices since it is computationally intensive and can have a large memory requirement. We use a more efficient level of abstraction by establishing a set of characteristic device parameters that are key to measuring the physical and system metrics which are important to our understanding of photonic interconnection networks. This simplified model enables PhoenixSim simulations to run on conventional computers in a period of minutes or hours. The device characteristics can be determined experimentally, through simulation, or projected. This set of modeled devices




Fig. 4. Parameters for characterizing a photonic device using the *Basic Element Model*.

composes the *Photonic Device Library*. While the descriptions included in this paper mostly highlight silicon ring-based topologies, the modeling methodology can easily be used to describe devices based on other technology domains such as Mach-Zehnders (also described in this section), photonic crystals, and MEMS.

The parameters used to describe basic photonic devices, called *Basic Elements*, are shown in Fig. 4. We refer to optical inputs and outputs as ports. Each port is physically bi-directional, therefore ports from which an optical signal can ingress into can also be used to egress from, and vice versa. Certain network topologies may still require uni-directional operation of the ports to facilitate simplicity or satisfy some other design requirement. Nonetheless, the bi-directional nature of each port is still represented for accuracy. The ports of the device are enumerated $0 \dots N - 1$ where N is the number of ports of a photonic device. N also determines the size of additional parameter matrices used in defining the photonic device behavior and characteristics.

We use a logical routing table to determine the path a message takes through the device. Fig. 4 shows how the routing table can be represented as a length- N vector, where the index represents the ingress port of an optical signal and the value at the index represents the egress port.

Additionally, we use two tables to represent the latency and the optical insertion loss properties of the device. Each property is represented as a $N \times N$ matrix where the row corresponds to the port through which the optical signal ingresses from (input) and the column represents the port from which the optical signal egresses from (output). Each entry in a matrix corresponds to the value used for the particular input/output combination. The latency for a particular input-output port combination is measured as the time between when optical signal enters the input port and when the same optical signal exits the output port. The insertion loss is a measure of the optical power attenuation an optical signal receives when traveling through a device and is useful in characterizing network-level insertion loss and crosstalk.

A. Static Elements

The *Basic Element Model* is most suitable for describing static optical devices that have characteristics that do not

change at runtime. The current library of devices focus on 2-D planar devices that are capable of being fabricated in a CMOS-compatible process. These static devices include waveguides, waveguide bends, waveguide crossings, and couplers.

1) *Waveguides*: Waveguides act as the optical wires used to link all the various devices, sources, and destinations. Optical signals that travel along a waveguide exhibit insertion loss in the form of propagation loss which is the attenuation experienced from traveling through the waveguide. Propagation loss is affected by a variety of parameters including waveguide dimensions, fabrication technique, and material properties. Waveguides are modeled as 2-port devices with parameters for length, group velocity per unit length, and insertion loss per unit length.

A waveguide's routing table is $[1, 0]$; which indicates that an optical signal ingressing on either end will egress on the opposite side. For a waveguide of length L_{wg} and propagation delay t_{wg} , the latency matrix will be

$$Latency_{wg} = \begin{bmatrix} - & L_{wg}t_{wg} \\ L_{wg}t_{wg} & - \end{bmatrix}. \quad (1)$$

Note that the elements along the diagonal represent the latency of a reflection. Since reflections are nonexistent in waveguides, the elements of the matrix that represent the latency of the reflection are marked as do not-care values. Similarly, the same waveguide with propagation loss of α_{wg} will have a insertion loss matrix of

$$Loss_{wg} = \begin{bmatrix} \infty & L_{wg}\alpha_{wg} \\ L_{wg}\alpha_{wg} & \infty \end{bmatrix}. \quad (2)$$

While reflections do not occur in the waveguide, it is useful to assign infinite insertion loss to the reflection path for crosstalk calculation purposes.

The production of low-loss on-chip waveguides for the CMOS platform is an important goal in realizing integrated photonic networks. Silicon waveguides with cross sectional areas of approximately $500 \text{ nm} \times 250 \text{ nm}$ have been demonstrated with a 1–2-dB/cm insertion loss using [26], [27]. Lower losses can be achieved using more exotic fabrication techniques such as with etchless silicon waveguides that have been shown to have losses of 0.3 dB/cm [28]. The freedom of parameter specification also enables the investigation of waveguides composed of nonsilicon materials such as silica fiber (losses on the order of tenths of a dB per kilometer) and silicon-nitride (losses of 0.1 dB/cm, [29]).

2) *Waveguide Bends*: Bends in waveguides are required to properly direct all the optical paths in the creation of switches and topologies. Waveguide bends contribute additional insertion loss to the waveguide's existing propagation loss, which we refer to as bending loss. Bends are modeled as 2-port devices and take parameters for loss per degree and angle of the bend. Silicon waveguide bending losses have been experimentally measured to be 0.005 dB per 90° turn with a bending radius of $6.5 \mu\text{m}$ [26].

3) *Waveguide Crossings*: Waveguide crossings are inherently required in silicon-based on-chip topologies due to the 2-D planar nature of the technology platform. Crossings occur whenever two waveguides intersect and can exhibit both




Fig. 5. Organization of building block element classes within PhonixSim.

insertion loss and crosstalk which can have an impact on system scalability and performance. This is in distinct contrast with electronic interconnects, which do not allow arbitrary crossings of two wires since this would cause a short circuit. The model for crossings are configured as 4-port devices with parameters for the loss and crosstalk.

Since many topologies require a large number of waveguide crossings, it is important for these devices to exhibit both low insertion loss and low crosstalk. A $6\ \mu\text{m} \times 6\ \mu\text{m}$ double-etched crossing design has been fabricated and tested, and was shown to have fairly low insertion loss at 0.16 dB and high crosstalk suppression at about $-40\ \text{dB}$ [30].

4) *Couplers*: The cross-boundary interface that separates the on-chip and off-chip domain presents a distinct situation where photonics can break through performance bottlenecks that are typically experienced by electronics. The capacitive effects of metal wires cause limitations in both the distance and rate at which data can be transmitted electronically, consequently causing problems when trying to scale I/O performance which can potentially require long wires that travel off-chip and across a board. Instead, optical signals are practically transparent to these issues and can be transmitted without penalty to extremely long distances and extremely high data rates. The optical I/O interface between the on-chip and off-chip world is a coupler, which is essentially a device for transferring light from one guiding medium to another (e.g., from an on-chip silicon waveguide to an off-chip silica fiber). This is modeled as a 2-port device with a single parameter for insertion loss.

There are currently two methods for implementing a coupler. Lateral coupling can be accomplished by building a nanotaper at the perimeter of a chip to couple into a fiber and has been calculated to have theoretical losses of under 1 dB [31]. Vertical couplers enable the I/O interface to be placed on the planar surface of the chip and can be accomplished with a Bragg grating for a ~ 1 -dB loss [32]. The advantage of vertical coupling over lateral coupling is the ability to position optical I/O ports anywhere on the chip. However, vertical coupling is fairly wavelength dependent due to the selectivity of Bragg structures while lateral coupling is spectrally more broadband.

B. Ring-Resonator Elements

Ring resonators are waveguides that form a closed loop which can be designed to manipulate the flow of light in a way that enables network functionality. Light interacts with the rings at specific periodically spaced wavelengths in the optical spectrum, called *resonant modes*. Light that enters the




Fig. 6. Propagation through a ring-resonator device depends on the signal wavelength and the resonant modes of the device. (a) Small rings with larger mode spacings (shown as periodic peaks) can be designed to interact with a single wavelength channel from a WDM signal (indicated by arrows). (b) Broadband switch have tightly spaced modes, enabling many WDM channels to couple into the device cohesively. (c) Path of propagation depends on whether the wavelength of the message is on or off-resonance with the ring.

ring is said to exhibit *on-resonance* behavior, whereas light that is transparent to the ring is said to be *off-resonance*. The *free spectral range* (FSR) specifies the spacing of the modes, which is inversely proportional to the optical length of the ring. A large-diameter ring will exhibit a small FSR (more spectrally dense), while a small-diameter ring will exhibit a large FSR (more selective). The FSR can be manipulated either by altering the physical length of the ring loop or by dynamically changing the refractive index of the device through electrical or thermal methods. Electrical manipulation can be accomplished by creating a p-i-n structure on the ring with the waveguide acting as the intrinsic region. Electrically biasing the p-i-n structure will cause a shift in refractive index due to the free-carrier plasma dispersion effect in silicon [33]. This contrasts with thermal manipulation which uses the thermo-optic properties of the material for index changes [34]. By using these techniques, the ring resonator can be engineered to perform a diverse range of network tasks [35]–[39].

To model the various ring resonator devices, we extend the *Basic Element Model* with subclasses for *Ring Elements* and *Dynamic Elements* (Fig. 5). The Dynamic Element Model is used to describe active devices which can exhibit changes in its routing table, latency matrix, and loss matrix during runtime. The properties of the active device during its operation is defined by *state* variables which can be changed and controlled. The Ring Element Model supports the definition of the resonant behavior of the devices. The behavior of ring-based devices is determined by the wavelength of the optical signal that interacts with the component. Also shown in Fig. 5 is how Dynamic-Ring Elements can be derived from the individual *Ring* and *Dynamic Element*. For instance, a ring-based broadband switch consists of a combination of ring resonators and electrical logic (described below) and can be electro-optically controlled to alter the optical flow of data.

1) *Filters*: Optical filters are useful in selectively extracting a subset of wavelengths from a WDM message. In the limiting case, an extremely small ring will have a large FSR and allow the filtering of a single wavelength channel. Filtering is accomplished by aligning the spectral mode of the ring

with the wavelength channel of interest [Fig. 6(a)]. Light at wavelengths that align with the mode of the ring (on resonance) will couple from the ingress waveguide, into the ring structure, and out onto a secondary waveguide; wavelengths of light that are not aligned (off resonance) will be unperturbed by the ring and continue down the injection waveguide [Fig. 6(c)]. We model ring filter devices as single-state 4-port *Ring Elements* with a parameter for the ring diameter (assuming a circle). Ring filters have been fabricated and demonstrated on SOI with 3- μm radius, corresponding to an FSR of 30 nm [39].

2) *Broadband Switches*: Ring resonators are also capable of controlling the flow of an entire WDM message by aligning each wavelength channel to a mode of the ring [Fig. 6(b)]. This can be accomplished in a limited spectral range by using a large ring with a correspondingly small FSR. When all the wavelength channels are on resonance, the entire WDM message will couple into the ring and onto a second waveguide, similar to the case of the filter. Additionally, if the FSR is manipulated electro-optically, all the modes can be shifted so that the wavelength channels are no longer on resonance, thus causing the entire WDM message to not couple into the ring. This functionality is illustrated in Fig. 6(c) for both a single-ring 1×2 photonic switching element (PSE) and a double-ring 2×2 PSE. These broadband switch elements are modeled as two-state 4-port devices. A 1×2 switch composed of a ring with a 100- μm radius and 0.8-nm FSR was shown to be capable of switching 20 wavelength channels simultaneously [38]. Elsewhere, a fifth-order switch was demonstrated being able to simultaneously route nine 40-Gb/s wavelength channels for an aggregate data rate of 360 Gb/s [37].

3) *Modulators*: Ring-based modulators are essentially high-speed switches. By electro-optically flipping the ring between an on and off-resonance state, a series of 0s and 1s can be encoded onto an optical stream of light. Light that couples into the ring will not egress into another waveguide like the filters and switches, but will eventually dissipate within the ring. A modulator array can be formed with multiple ring modulators so that several wavelength channels can be encoded in parallel, creating a WDM signal (Fig. 7). Modulators should have a small ring diameter to create a large FSR to ensure that the modulation does not interfere with other spectrally adjacent wavelength channels. The modulator device is modeled as a single-state device with parameters for energy dissipated per modulated bit and ring diameter. Ring-based modulation has been demonstrated at rates of 12.5 Gb/s in a 5- μm radius silicon ring resonator [36].

4) *Receivers (Photo-Detectors)*: Photodetectors are used for converting optical messages back into the electrical domain. While the detection element itself is not a ring resonator, photo-detectors still require rings to properly filter individual wavelength channels from an entire WDM message. Each ring filter will only allow the light from a single wavelength channel to be incident on the photo-detector it precedes, thereby allowing the receiver to convert a single wavelength channel's worth of data back into the electrical domain. Similar to modulators, filtering should be




Fig. 7. Schematic of the conversion process between the spatially parallel electronic domain and wavelength-parallel optical domain.

accomplished without disturbing other adjacent wavelength channels by using as high an FSR as possible. The *detector sensitivity* determines the minimum signal power that must be received at the photo-detector in order for data to be properly recovered from the optical domain and is an important parameter for determining the optical power budget (as discussed in Section V). This ring-based detection device takes parameters for energy dissipated per detected bit, sensitivity, and ring diameter. Integrated high-speed germanium detectors have been demonstrated operating at speeds of 40 Gb/s [40], [41].

C. Mach-Zehnder Elements

Switches and modulators can also be designed using the principle of Mach-Zehnder interferometry (MZI). Mach-Zehnder devices are designed to operate relatively uniformly over a large wavelength range and do not exhibit the sharp resonant peaks that ring resonators have. For instance, a MZI-based device can be used to modulate wavelengths of light that span a large continuous wavelength range while ring-resonator modulators are limited to specific resonance wavelengths. However, this operational difference between Mach-Zehnder devices and ring-resonator devices causes them to not be interchangeable. The ring-based network architectures analyzed in Section VI are not compatible with these devices and would require significant changes in the designs. Models for 1×2 and 2×2 Mach-Zehnder switches are currently included in the *Photonic Device Library*. A modulator and switch based on MZI has been demonstrated operating at up to 10 Gb/s [42].

V. PHYSICAL-LAYER PERFORMANCE ANALYSIS

The consideration of the photonic technology domain presents new design challenges that must be satisfied in order to produce feasible interconnect designs. Similar to electronics, it is important for photonic networks to consider power dissipation and system-level performance. Furthermore, photonic networks must also consider metrics that have no electronic equivalent such as insertion loss, the optical power budget, noise, and crosstalk. While a comprehensive analysis of a photonic interconnect design would involve the actual fabrication and operation of such a system, this is currently unrealistic since full-scale photonic on-chip networks are still in early stages of research. Therefore, the tools presented here




Fig. 8. Relationship of various parameters affecting the optical power budget. The difference in power of the total WDM signal (large arrow on the left) and the individual wavelength channels (five smaller arrows on the right) constrains the scalability of the system.

can give important insight into the physical feasibility of the designs and the performance that is expected.

A. Optical Power Budget

The optical power budget of a photonic network assesses the amount of WDM parallelism and insertion loss that can be tolerated. Many currently proposed photonic interconnection networks assume off-chip lasers to provide the optical sources, which are then coupled into the chip where they are modulated, routed, and received. Optical amplification in an on-chip environment is not easily accomplished in the CMOS platform. For this reason, the power that is received at the photodetectors must remain above a certain power threshold (labeled the detector sensitivity in Fig. 8) to ensure proper detection of data bit streams. This limitation can be partially compensated for by increasing the optical power that is injected into the chip. However, this also exhibits an upper limitation due to nonlinearities of the silicon material which will potentially distort the signal. Distortions are caused by nonlinearities within silicon which contribute additional insertion losses and can also cause unwanted shifts in the resonances of ring resonators. This limit is labeled as *nonlinear effects* in Fig. 8. The difference in the two thresholds is called the *optical power budget*.

As shown in Fig. 8, the optical power budget affects the design choices of a given network architecture by constraining the sum of the WDM factor and the network insertion loss. The *WDM factor* measures the power difference between an entire WDM signal and its constituent wavelength channels. This factor needs to be accounted for since the nonlinearity threshold is determined by the total power in the waveguide while the detector sensitivity depends on the power in the individual wavelengths. The remaining portion of the optical power budget must accommodate the worst-case insertion loss that an optical message could receive in the network. Fig. 9 shows an example of the calculation involved in determining the insertion loss for an optical signal being injected into a small network segment at 1 dBm. The signal is ejected at 0.24 dBm after propagating across a 0.1-cm distance, passing by two ring resonators, and entering four waveguide crossings. The total loss for this example is 0.76 dB. For a full-scale photonic network, all valid optical paths need to be examined to determine the highest-loss path.




Fig. 9. Calculation of insertion loss for a small network segment.

The relationship between the various device limitations and system-level metrics is summarized in the inequality

$$P - S \geq IL_{\max} + 10 \log_{10} n \quad (3)$$

where P is the power threshold we limit the optical power to and S is the detector sensitivity. The optical power budget is $P - S$. The worst-case optical path in terms of insertion loss is IL_{\max} and n specifies the number of wavelength channels being used. P , S , and IL_{\max} are expressed in decibel units.

While it may be desirable to maximize the number of wavelength channels used to increase bandwidth through parallelism, and to create scalable photonic networks at the cost of higher insertion losses, (3) shows the inherent limitation to this. From an architectural standpoint, P and S are fundamental design constraints imposed by the photonic devices. Therefore, a designer must strike a balance between the desired link bandwidth and the desired complexity of the network. In Section VI, we illustrate the evaluation of these tradeoffs which are made possible by PhoenixSim.

B. Data Integrity

A variety of interactions in a photonic interconnection network will work to degrade the integrity of transmitted data. Our current noise modeling methodology accounts for intensity noise generated at the laser sources, inter-message crosstalk, intra-message crosstalk, and electrical noise generated by the optical receivers (Fig. 10). The standard figure of merit for measuring the quality of signal is the signal-to-noise ratio (SNR) which is defined as the ratio between signal power and noise power. From a system perspective, the SNR can be used to determine the statistical likelihood that each bit of data is transmitted erroneously (e.g., a transmitted 0 is detected as a 1), also called a *bit error rate* (BER). An understanding of the potential noise in any interconnection network is critical to determining the effective throughput of the system since error detection and correction will invariably cause performance overheads.

The first source of noise is from the laser sources which inherently cause random fluctuations in an optical signal, called intensity noise. This noise is quantified as relative intensity noise (RIN), which is the ratio of the power variance of the optical signal to the mean optical power squared. Quantum cascade lasers have a measured RIN on the order of -150 dB Hz^{-1} with an output of 10-dBm mean optical power [43]. To convert to a SNR, we use the relation [44]

$$SNR_{laser} = \frac{m^2}{2B \cdot RIN} \quad (4)$$




Fig. 10. Sources of noise and crosstalk within a chip-scale photonic system.

where B is the noise bandwidth, assumed equal to the modulation rate, and m is the modulation index, equal to $1 - E$, where E is the extinction ratio of the modulator.

A second source of noise is *inter-message crosstalk* which occurs when multiple photonic messages concurrently propagate through a photonic device. In a waveguide crossing for example, the ideal situation is for two orthogonally propagating messages to be completely isolated from each other with no interaction. However, in reality a small amount of optical power from each message will leak onto the other message. A similar situation occurs in ring-resonator filters and switches due to imperfect coupling of each wavelength channel.

For the N -port device, the crosstalk power that a message on a particular port receives is given by the sum of the power that is leaked by any existing messages on the other $N - 1$ ports. If M is the set of all signals present in the device and the power of a signal k is given by the variable P_k , then the crosstalk power seen by signal s is given by

$$\sum_{k \in M, k \neq s} \frac{P_k}{IL(portin_k, portout_s)} \quad (5)$$

which aggregates the unwanted signal power that leaks into the output port being used by s . Function IL refers to the *insertion-loss matrix* (that was described in Section IV) of the device model with arguments for the input and output port. In (5), $portin_k$ denotes the input port of a message k , and $portout_s$ denotes the output port of s . This calculation is a first-order approximation that only considers crosstalk for messages that coexist in a device and not from leaked power that propagates across multiple devices before interfering with a foreign signal.

A third source of noise called *intra-message crosstalk* occurs due to imperfect filtering. For example, in order for a WDM message to be received and converted into an electrical signal, each wavelength channel must be individually filtered and fed into a photo-detector. Due to imperfect extinction, power from the adjacent wavelength channels will

leak through causing an additional source of noise. Intra-message crosstalk will also occur in any other location in a photonic network where filtering functionality is involved. The spectral response of a ring resonator mimics a periodic Lorentzian function. For simplicity we assume a periodic flat passband and constant extinction ratio for the stop bands. Lastly, our receiver model includes thermal and shot noise.

The combined effect of these multiple sources of noise can be used to compute an SNR for the final detected signal as follows:

$$SNR = \frac{P}{N_{laser} + N_{inter} + N_{intra} + N_{therm} + N_{shot}} \quad (6)$$

where P is the signal power and N corresponds to the noise power associated with the noise or crosstalk source indicated by the subscript.

C. Power Dissipation

To compute the power dissipation of the modeled networks, we add up the energy dissipation events from all devices. Our photonic device library tracks the power dissipation according to the type of model that is used, and can include both static (over a duration of time) and dynamic (instantaneous) power dissipation. *Dynamic Element* devices can have static power dissipation, which is determined by the occupied state. *Dynamic Element* devices can also have dynamic power dissipation, which is accumulated whenever there is a state transition. An additional source of power dissipation are *Ring Element* devices, which require constant thermal tuning to compensate for fabrication uncertainty and ambient temperature shifts. *Modulator* and *Detector Elements* also dissipate power during the transmission and detection of data, respectively.

Electronic routers are modeled as standard three-stage pipelines. The power modeling of the electronic routers is accomplished by leveraging the Orion simulator, which is currently capable of modeling down to the 32 nm technology node [25].

VI. CASE STUDY

In this section, we model two different photonic interconnection networks to demonstrate our methodology and the main capabilities of PhoenixSim. The tools presented in Section V are used to quantify the performance of the networks and to demonstrate design spaces that are allotted by our physical-layer analysis. This case study serves to demonstrate the various capabilities of the simulator. We will show that the two networks offer different advantages depending on the considered metric and traffic pattern. Therefore this analysis serves to give system architects recommendations based on their design objectives.

The first photonic network we model for this case study is the Photonic Mesh, a circuit-switching architecture originally proposed by Shacham *et al.* [12] for high-bandwidth optical communications on future CMPs. The network architecture consists of a photonic network plane and an electronic network plane. The electronic plane is used to transmit control




Fig. 11. Photonic Mesh topology. (a) High-level representation of a 4×4 Photonic Mesh. Parallel lines indicate two unidirectional waveguides, which are paired together to form bidirectional links. Boxes represent higher-order photonic components, which are labeled X for 4×4 nonblocking crossbar switch, I for injection gateway, and E for ejection gateway. Also shown are detail schematics of the (b) 4×4 nonblocking crossbar switch, (c) injection gateway, and (d) ejection gateway.

messages and provision optical resources for establishing circuits on the photonic plane where the data messages are actually transmitted. The Photonic Mesh [Fig. 11(a)] is similar to a typical electronic mesh since it is laid out in a matrix-like configuration of nodes, and has mechanisms for switching, entering the network, and exiting the network at each node. Although the mesh-based design presented here exhibits lower path diversity than previously proposed circuit-switching topologies [13], [14], the simpler architecture is beneficial to overall performance by lowering total insertion loss.

A 4×4 nonblocking crossbar switch [Fig. 11(b)] is found at each node of the network and is optimized for dimension-ordered routing (which is the case for the Photonic Mesh) by minimizing insertion losses along propagation paths that do not turn through the switch [45]. The injection gateway [Fig. 11(c)] and ejection gateway [Fig. 11(d)] designs, which are used by the underlying processing cores to transmit and receive optical data, are adapted from the TorusNX topology to help further reduce insertion loss overhead caused by more complex injection/ejection schemes [14]. Each switch and gateway is constructed using the devices previously described in Section IV.

The second photonic network we model for this case study is the Photonic Crossbar (Fig. 12). This design uses the crossbar concepts used previously in the Photonic Clos topology [7]. A set of waveguides are routed in a serpentine manner so that they intersect with all gateways in the network. Each individual waveguide is configured with two modulator banks and two receiver banks to connect a unique pair of gateways. For a topology with G gateways, a set of $G \cdot (G - 1)/2$ waveguides is required to fully connect the network. Since the required number of waveguides grows quadratically




Fig. 12. Photonic Crossbar topology. (a) High-level representation of a 2×4 Photonic Crossbar, connecting 16 cores. Boxes represent gateways with a concentration of two processing cores. (b) Detail schematic of the Photonic Crossbar gateway, showing 49 bypass waveguides and seven waveguides with modulator and receiver banks used to communicate to the other seven gateways.

with G , it can be advantageous to concentrate the traffic of a set of processing cores through a single photonic gateway. Each gateway exploits the bidirectionality of the waveguides and avoids receiving its own modulated signal by transmitting and receiving on different sets of wavelengths.

Fig. 12(a) shows an 8-gateway network with two processing cores connected to each gateway. The gateway design is illustrated in Fig. 12(b). The gateway contains 49 bypass waveguides which are ignored, and is connected to the remaining seven waveguides through a set of seven modulator banks and seven receiver banks. Each connected waveguide will transmit to and receive from one of the other seven gateways in the network. Attached to each photonic gateway is a 9-port electronic router which must transport messages to and from the group of cores to the appropriate photonic transmitter or receiver bank.

A. Optical Power Budget Analysis

First, we used PhoenixSim to model both photonic topologies and analyze the worst-case insertion loss for network radices from 2×2 (four nodes) to 10×10 (100 nodes). The insertion loss parameters used in this paper are derived from experimentally demonstrated results and are listed in Table I. Fig. 13 shows the maximum total loss exhibited within each network and the breakdown according to type of loss. All network sizes assumed total chip dimensions of $2 \text{ cm} \times 2 \text{ cm}$ and the size of the network is designed to span the entire chip. Hence the spacing between nodes will decrease with larger radices. Crossing loss and propagation loss are the most significant contributors to total loss in the Photonic Mesh and Photonic Crossbar, respectively. The 10×10 Photonic Mesh has 18.1 dB of crossing loss caused by the existence of a network path with 113 waveguide crossings, accounting for

TABLE I
INSERTION LOSS PARAMETERS

Parameter	Value	Reference
Propagation loss (silicon)	1.7 dB/cm	[26]
Waveguide crossing	0.16 dB	[30]
Waveguide bend	0.005 dB/90°	[26]
Drop into a ring	0.6 dB	[38]
Pass by a ring	0.005 dB	[38]




Fig. 13. Insertion loss results for the (a) Photonic Mesh and (b) Photonic Crossbar of varying sizes. Labeled values that overlay the columns indicate the worst-case total network-level loss values. Columns illustrate the worst-case loss associated with the individually labeled loss component which does not necessarily occur in the network path with the worst total loss.

approximately 63% of the total network-level insertion loss. The serpentine waveguide design of the Photonic Crossbar causes repeated traversals of the chip, therefore causing high propagation loss. This analysis is an important indicator for device researchers who may seek to focus on improving the performance of a specific type of network architecture.

By taking the insertion loss results and applying (3), we can derive the allowed number of wavelength channels for varying radices and optical power budgets (Fig. 14). For the specified optical power budgets, points below and to the left of the plotted curve indicate physically realizable combinations of network size and number of wavelength channels. For example, both networks are realizable as a 4×4 network using 32 wavelength channels with devices that stay above a 30-dB optical budget, however the fabrication of an 8×8 network with 32 wavelength channels and a more aggressive 40-dB budget will only be possible for the Photonic Mesh. Furthermore, the plot indicates that the Photonic Crossbar is not capable of operating at sizes of 10×10 or greater with a 30-dB optical power budget.

B. Network Performance

The performance and power dissipation of the on-chip network are both important considerations for future scaling




Fig. 14. Wavelength channel allotment in the Photonic Mesh and Photonic Crossbar for varying network sizes and optical power budgets.

of CMPs. For this analysis, we assume a 64-core processor and compare the performance of the Photonic Mesh, the Photonic Crossbar, and a traditional electronic mesh. In each of the three networks, we assume a 2.5-GHz operating frequency for both electrical and optical signaling. Both photonic networks assume the use of 128-wavelength channels (the Photonic Crossbar will have two bi-directional 64-wavelength channels). Electronic routers for the Photonic Mesh are modeled with a 32-bit channel width and buffer size of 128 bit, which equates to a buffer depth of four control messages. Electronic routers for the Photonic Crossbar and electronic mesh have a 64-bit channel width and a 1024-bit buffer size. Additionally, for the Photonic Crossbar we assume a concentration of eight cores per gateway. All simulations are based on uniform random traffic.

Fig. 15 plots the network-level bandwidth and latency of the three networks under consideration. For 1-kbit messages, we see that the Photonic Crossbar exhibits the highest throughput. The Photonic Mesh performs the worst as a result of the costly overhead associated with circuit switching. In the case of 100-kbit messages, the Photonic Mesh now achieves the best performance since the latency overhead of circuit switching is now amortized over the duration of the message transmission. This indicates that the most suitable network design will be dependent on the type of traffic exhibited by the system.

C. Data Integrity Analysis

Whereas the insertion loss has an impact on physical size and bandwidth of the network, the noise has an impact on the quality of the data stream. Given the same network configuration used in the Network Performance results, the average noise power for each wavelength channel for all WDM transmissions under saturated random-traffic load is plotted in Fig. 16. These noise power results are based on the crosstalk and noise parameters listed in Table II. In this network, laser intensity noise, thermal noise, and shot noise are negligible quantities in comparison to inter-message and intra-message crosstalk.

In both networks intra-message crosstalk predominately occurs at the ejection gateway where filters are used to




Fig. 15. Bandwidth and latency performance on the Electronic Mesh, Photonic Crossbar, and Photonic Mesh for (a) 1-kbit and (b) 100-kbit message sizes.

select individual wavelength channels. The amount of intra-message crosstalk power exhibited by each optical message is predominately dependent on the number of co-propagating wavelengths. Therefore, it is practically independent of both network load and message size. We see that across the two different message sizes, the amount of intra-message crosstalk power remains approximately constant.

The trend in inter-message crosstalk reflects the probability that two WDM messages will intersect in the network. The Photonic Crossbar exhibits zero inter-message crosstalk since it contains no crossings or switches where a message intersection could occur. A longer duration optical packet from using fewer wavelength channels or large message sizes will create a scenario where the photonic message will occupy the network for a longer period of time, thereby increasing the likelihood that another message will be instanced in the network and interfere. In Fig. 16, we can see that indeed larger messages in the Photonic Mesh do produce a non-negligible amount of inter-message crosstalk.

Lastly, PhoenixSim also determines the signals SNR when the message is finally received. For 1-kbit message sizes, the average electrical SNR of the Photonic Mesh and Photonic Crossbar optical link is 6.4 dB and 3.5 dB, respectively. For 100-kbit message sizes, the average SNR for the Photonic Mesh and Photonic Crossbar is 6.5 dB and 2.9 dB, respectively. These results indicate that the Photonic Crossbar relatively outperforms the Photonic Mesh with respect to signal integrity. However these values also conclude that the optical link integrity of both networks will be detrimentally compromised. This performance penalty can be rectified by improved filter performance or through the use of fewer wavelength channels.

D. Power Dissipation Analysis

Last, we compare the power dissipation of the Electronic Mesh, Photonic Mesh, and Photonic Crossbar, assuming the

TABLE II
CROSSTALK AND NOISE PARAMETERS

Parameter	Value	Reference
Laser (relative intensity noise)	-150 dB/Hz	[43]
Modulator (extinction ratio)	16 dB	[46]
PSE through-port (extinction ratio)	25 dB	[38]
PSE drop-port (extinction ratio)	20 dB	[38]
Waveguide crossing (crosstalk)	-40 dB	[30]




Fig. 16. Average total noise power accumulated by each optical message in the Photonic Mesh and Photonic Crossbar for saturated network load. Laser noise, thermal noise, and shot noises are negligible quantities and are not listed.

same system configuration as before and the power parameters listed in Table III. The total power dissipation of each network, operating with maximum load, is plotted in Fig. 17. Each column is broken down into categories of photonic-related dissipation from ring modulators, photodetectors, optical switches, and thermal feedback tuning, and electronic-related dissipation from router logic, router buffers, and wires. While SerDes would be required in many proposed photonic interconnect architectures for every ring modulator and photodetector to up and down convert to the photonic transmission clock, in this case study we assume the same 2.5-GHz clock for both electronic and photonic domains.

Regardless of the message size the Electronic Mesh dissipates approximately 8 W of power and the Photonic Mesh dissipates approximately 5 W. This is a result of both networks relying on some electronic routers to route data. Data on the Photonic Mesh is only transmitted optically, which provides a significant savings in power when the circuit-switching overhead can be amortized. In terms of energy efficiency when transmitting 1-kbit messages, the Photonic Crossbar outperforms the other networks at 2.9 pJ/bit, while the Photonic Mesh performs the worst at 55.9 fJ/bit. Nonetheless, with the larger 100-kbit messages, the Photonic Mesh achieves the highest efficiency with 3.2 pJ/bit as a result of the efficient optical transmission. This message-size/efficiency relationship of the circuit-switched Photonic Mesh design is a useful indicator as to which photonic design may be ideally suited for various application traffic patterns.

TABLE III
POWER DISSIPATION PARAMETERS

Parameter	Value
Modulators (dynamic energy)	85 fJ/bit
Modulators (static energy)	30 μ W
Photodetectors	50 fJ/bit
PSEs (dynamic energy)	375 fJ/bit
PSEs (static energy)	400 μ W
Thermal ring tuning	100 μ W/ring




Fig. 17. Network-level power dissipation breakdown of the Electronic Mesh, Photonic Mesh, and Photonic Crossbar for transmission of 1-kbit and 100-kbit messages. Values overlaying each column indicate the energy efficiency of the network in units of pJ/bit.

For instance, photonic circuits have been shown to be ideally suited for many classes of scientific applications that require long data messages [15].

VII. CONCLUSION

We have described a methodology for modeling, designing, and analyzing photonic interconnection networks at both the physical-layer and system-level. A *Photonic Device Library* has been devised to describe any type of fundamental photonic elements, which can then be combined and used to model large-scale photonic components and network topologies. We developed a set of physical-layer tools to accurately determine physical properties of the photonic networks and examine how they impact the network architectures in terms of system performance. Our PhoenixSim environment implements this methodology, which we have made open source and publicly available. We illustrated the capabilities of PhoenixSim through the analysis of two photonic networks and showed how various system-level design tradeoffs are made possible through an understanding of the physical-layer characteristics. The device library, analysis tools, and simulation environment form a comprehensive design flow for understanding and designing photonically enabled computing systems.

REFERENCES

- [1] D. Miller, "Device requirements for optical interconnects to silicon chips," *Proc. IEEE*, vol. 97, no. 7, pp. 1166–1185, Jul. 2009.
- [2] N. Magen, A. Kolodny, U. Weiser, and N. Shamir, "Interconnect-power dissipation in a microprocessor," in *Proc. Int. Workshop SLIP*, Feb. 2004, pp. 7–13.
- [3] *Photonic and Electronic Network Integration and Execution Simulator (PhoenixSim)* [Online]. Available: <http://lightwave.ee.columbia.edu/phoenixsim>
- [4] A. Varga. *OMNeT++ Discrete Event Simulation System* [Online]. Available: <http://www.omnetpp.org>
- [5] J. Chan, G. Hendry, A. Biberman, K. Bergman, and L. Carloni, "PhoenixSim: A simulator for physical-layer analysis of chip-scale photonic interconnection networks," in *Proc. DATE*, Mar. 2010, pp. 691–696.
- [6] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. P. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. G. Beausoleil, and J. H. Ahn, "Corona: System implications of emerging nanophotonic technology," in *Proc. 35th Annu. ISCA*, Jun. 2008, pp. 153–164.
- [7] A. Joshi, C. Batten, Y.-J. Kwon, S. Beamer, I. Shamim, K. Asanovic, and V. Stojanovic, "Silicon-photonic Clos networks for global on-chip communication," in *Proc. 3rd ACM/IEEE Int. NOCS*, May 2009, pp. 124–133.
- [8] 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. ISCA*, Jun. 2009, pp. 429–440.
- [9] N. Kirman, M. Kirman, R. K. Dokania, J. F. Martinez, A. B. Apsel, M. A. Watkins, and D. H. Albonese, "Leveraging optical technology in future bus-based chip multiprocessors," in *Proc. 39th Annu. IEEE/ACM Int. Symp. Microarch. (MICRO)*, Dec. 2006, pp. 492–503.
- [10] C. Batten, A. Joshi, J. Orcutt, A. Khilo, B. Moss, C. W. Holzwarth, M. A. Popovic, H. Li, H. I. Smith, J. L. Hoyt, F. X. Kartner, R. J. Ram, V. Stojanovic, and K. Asanovic, "Building many-core processor-to-DRAM networks with monolithic CMOS silicon photonics," *IEEE Micro*, vol. 29, no. 4, pp. 8–21, Jul.–Aug. 2009.
- [11] M. J. Cianchetti, J. C. Kerekes, and D. H. Albonese, "Phastlane: A rapid transit optical routing network," in *Proc. 36th Annu. ISCA*, Jun. 2009, pp. 441–450.
- [12] 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.
- [13] M. Petracca, B. G. Lee, K. Bergman, and L. P. Carloni, "Photonic NoCs: System-level design exploration," *IEEE Micro*, vol. 29, no. 4, pp. 74–85, Aug. 2009.
- [14] J. Chan, G. Hendry, A. Biberman, and K. Bergman, "Architectural exploration of chip-scale photonic interconnection network designs using physical-layer analysis," *J. Lightw. Technol.*, vol. 28, no. 9, pp. 1305–1315, May 2010.
- [15] G. Hendry, S. Kamil, A. Biberman, J. Chan, B. Lee, M. Mohiyuddin, A. Jain, K. Bergman, L. Carloni, J. Kubiatowicz, L. Oliker, and J. Shalf, "Analysis of photonic networks for a chip multiprocessor using scientific applications," in *Proc. 3rd ACM/IEEE Int. NOCS*, May 2009, pp. 104–113.
- [16] G. Hendry, J. Chan, S. Kamil, L. Oliker, J. Shalf, L. Carloni, and K. Bergman, "Silicon nanophotonic network-on-chip using TDM arbitration," in *Proc. 18th Annu. IEEE Symp. High Perform. Interconnects (HOTI)*, Aug. 2010, pp. 88–95.
- [17] D. Ding and D. Z. Pan, "OIL: A nano-photonics optical interconnect library for a new photonic networks-on-chip architecture," in *Proc. 11th Int. Workshop SLIP*, Jul. 2009, pp. 11–18.
- [18] J. Minz, S. Thyagaraja, and S. K. Lim, "Optical routing for 3D system-on-package," in *Proc. DATE*, vol. 1, Mar. 2006, pp. 1–2.
- [19] G. Hendry, J. Chan, L. P. Carloni, and K. Bergman, "VANDAL: A tool for the design specification of nanophotonic networks," in *Proc. DATE*, Mar. 2011, pp. 1–6.
- [20] I. O'Connor, F. Tissafi-Drissi, F. Gaffiot, J. Dambre, M. D. Wilde, J. V. Campenhout, D. V. Thourhout, J. V. Campenhout, and D. Stroobandt, "Systematic simulation-based predictive synthesis of integrated optical interconnect," *IEEE Trans. Very Large Scale Integr. Syst.*, vol. 15, no. 8, pp. 927–940, Aug. 2007.
- [21] M. De Wilde, O. Rits, W. Meeus, H. Lambrecht, and J. V. Campenhout, "Integration of modeling tools for parallel optical interconnects in a standard EDA design environment," in *Proc. DATE*, Feb. 2004 [Online]. Available: <http://citeseerx.ist.psu.edu/viewdoc/summary>

- [22] P. K. Pepeljugoski and D. M. Kuchta, "Design of optical communications data links," *IBM J. Res. Dev.*, vol. 47, nos. 2–3, pp. 223–237, Mar. 2003.
- [23] M. Briere, E. Drouard, F. Mieleveille, D. Navarro, I. O'Connor, and F. Gaffiot, "Heterogeneous modelling of an optical network-on-chip with SystemC," in *Proc. 16th IEEE Int. Workshop RSP*, Jun. 2005, pp. 10–16.
- [24] A. Kodi and A. Louri, "Optisim: A system simulation methodology for optically interconnected HPC systems," *IEEE Micro*, vol. 28, no. 5, pp. 22–36, Sep.–Oct. 2008.
- [25] 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. Microarch. (MICRO)*, Nov. 2002, pp. 294–305.
- [26] F. Xia, L. Sekaric, and Y. Vlasov, "Ultracompact optical buffers on a silicon chip," *Nature Photonics*, vol. 1, pp. 65–71, Dec. 2006.
- [27] M. Gnan, S. Thorns, D. Macintyre, R. D. L. Rue, and M. Sorel, "Fabrication of low-loss photonic wires in silicon-on-insulator using hydrogen silsesquioxane electron-beam resist," *IET Electron. Lett.*, vol. 44, no. 2, pp. 115–116, Jan. 2008.
- [28] J. Cardenas, C. B. Poitras, J. T. Robinson, K. Preston, L. Chen, and M. Lipson, "Low loss etchless silicon photonic waveguides," *OSA Optics Express*, vol. 17, no. 6, pp. 4752–4757, 2009.
- [29] M. J. Shaw, J. Guo, G. A. Vawter, S. Habermehl, and C. T. Sullivan, "Fabrication techniques for low-loss silicon nitride waveguides," *Micromach. Technol. Micro-Optics Nano-Optics III*, vol. 5720, pp. 109–118, 2005.
- [30] W. Bogaerts, P. Dumon, D. V. Thourhout, and R. Baets, "Low-loss, low-cross-talk crossings for silicon-on-insulator nanophotonic waveguides," *OSA Optics Lett.*, vol. 32, no. 19, pp. 2801–2803, 2007.
- [31] V. R. Almeida, R. R. Panepucci, and M. Lipson, "Nanotaper for compact mode conversion," *OSA Optics Lett.*, vol. 28, no. 15, pp. 1302–1304, 2003.
- [32] G. Roelkens, D. V. Thourhout, and R. Baets, "High efficiency grating coupler between silicon-on-insulator waveguides and perfectly vertical optical fibers," *OSA Optics Lett.*, vol. 32, no. 11, pp. 1495–1497, 2007.
- [33] M. Lipson, "Compact electro-optic modulators on a silicon chip," *IEEE J. Selected Topics Quantum Electron.*, vol. 12, no. 6, pp. 1520–1526, Nov.–Dec. 2006.
- [34] W. Green, H. Hamann, L. Sekaric, M. Rooks, and Y. Vlasov, "Ultracompact reconfigurable silicon optical devices using micron-scale localized thermal heating," in *Proc. OFC*, Mar. 2007, p. OTuM3.
- [35] M. Watts, D. Trotter, R. Young, and A. Lentine, "Ultralow power silicon microdisk modulators and switches," in *Proc. 5th IEEE Int. Conf. Group IV Photonics*, Sep. 2008, pp. 4–6.
- [36] Q. Xu, S. Manipatruni, B. Schmidt, J. Shakya, and M. Lipson, "12.5 Gbit/s carrier-injection-based silicon micro-ring injection modulators," *OSA Optics Express*, vol. 15, no. 2, pp. 430–436, 2007.
- [37] Y. Vlasov, W. M. J. Green, and F. Xia, "High-throughput silicon nanophotonic wavelength-insensitive switch for on-chip optical networks," *Nat. Photonics*, vol. 2, pp. 242–246, Apr. 2008.
- [38] 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 Photonics Technol. Lett.*, vol. 20, no. 10, pp. 767–769, May 2008.
- [39] B. Little, J. Foresi, G. Steinmeyer, E. Thoen, S. Chu, H. Haus, E. Ippen, L. Kimerling, and W. Greene, "Ultra-compact Si-SiO₂ microring resonator optical channel dropping filters," *IEEE Photonics Technol. Lett.*, vol. 10, no. 4, pp. 549–551, Apr. 1998.
- [40] S. Assefa, F. Xia, S. W. Bedell, Y. Zhang, T. Topuria, P. M. Rice, and Y. A. Vlasov, "CMOS-integrated 40 GHz germanium waveguide photodetector for on-chip optical interconnects," in *Proc. OFC*, Mar. 2009, p. OMR4.
- [41] L. Vivien, J. Osmond, J.-M. Fédéli, D. Marris-Morini, P. Crozat, J.-F. Damlencourt, E. Cassan, Y. Lecunff, and S. Laval, "42 GHz p.i.n germanium photodetector integrated in a silicon-on-insulator waveguide," *OSA Optics Express*, vol. 17, no. 8, pp. 6252–6257, 2009.
- [42] H.-W. Chen, Y.-H. Kuo, and J. E. Bowers, "High speed hybrid silicon evanescent Mach-Zehnder modulator and switch," *OSA Optics Express*, vol. 16, no. 25, pp. 20571–20576, 2008.
- [43] T. Gensty, W. Elsässer, and C. Mann, "Intensity noise properties of quantum cascade lasers," *OSA Optics Express*, vol. 13, no. 6, pp. 2032–2039, 2005.
- [44] C. Miller, *Fiber Optic Test and Measurement*. Englewood Cliffs, NJ: Prentice-Hall, 1998.

- [45] J. Chan, A. Biberman, B. Lee, and K. Bergman, "Insertion loss analysis in a photonic interconnection network for on-chip and off-chip communications," in *Proc. 21st Annu. Meeting IEEE LEOS*, Nov. 2008, pp. 300–301.
- [46] K. Preston, S. Manipatruni, A. Gondarenko, C. B. Poitras, and M. Lipson, "Deposited silicon high-speed integrated electro-optic modulator," *OSA Optics Express*, vol. 17, no. 7, pp. 5118–5124, 2009.



Johnnie Chan (S'08) received the B.S. degree (high distinction) in computer and electrical engineering and the M.S. degree in electrical engineering from the University of Virginia, Charlottesville, in 2005 and 2007, respectively. He is currently pursuing the Ph.D. degree with the Department of Electrical Engineering, Columbia University, New York, NY.

His current research interests include silicon photonic devices for chip-scale interconnection networks and optical networks for high-performance computing.



Gilbert Hendry (S'08) received the B.S. and M.S. degrees in computer engineering from the Rochester Institute of Technology, Rochester, NY, in 2007, and the Ph.D. degree from Columbia University, New York, NY, in 2011.

He is currently a Post-Doctoral Researcher with the Department of Scalable Modeling and Analysis, Sandia National Laboratories, Livermore, CA. His current research interests include modeling, simulation, layout, and design automation of chip-scale networks, particularly those using

nanophotonics.



Keren Bergman (S'87–M'93–SM'07–F'09) received the B.S. degree from Bucknell University, Lewisburg, PA, in 1988, and the M.S. and Ph.D. degrees from the Massachusetts Institute of Technology, Cambridge, in 1991 and 1994, respectively, all in electrical engineering.

She is currently a Professor and Chair of the Department of Electrical Engineering, Columbia University, New York, NY, where she also directs the Lightwave Research Laboratory. She leads multiple research programs on optical interconnection networks for advanced computing systems, data centers, optical packet switched routers, and chip multiprocessor nanophotonic networks-on-chip.

Dr. Bergman is a fellow of the OSA and currently serves as the Co-Editor-in-Chief of the *IEEE/OSA Journal of Optical Communications and Networking*.



Luca P. Carloni (S'95–M'04–SM'09) received the Laurea degree (summa cum laude) in electrical engineering from the Università di Bologna, Bologna, Italy, in 1995, and the M.S. and Ph.D. degrees in electrical engineering and computer sciences from the University of California, Berkeley, in 1997 and 2004, respectively.

He is currently an Associate Professor with the Department of Computer Science, Columbia University, New York, NY. He has authored over 80 publications and is the holds one patent. His current research interests include design methodologies and tools for integrated circuits and systems, distributed embedded systems design, and design of high-performance computer systems.

Dr. Carloni received the Faculty Early Career Development (CAREER) Award from the National Science Foundation in 2006, was selected as an Alfred P. Sloan Research Fellow in 2008, and received the ONR Young Investigator Award in 2010. He is the recipient of the 2002 Demetri Angelakos Memorial Achievement Award for "recognition of altruistic attitude toward fellow graduate students." In 2002, one of his papers was selected for the "Best of ICCAD," a collection of the best IEEE International Conference on Computer-Aided Design papers of the past 20 years. He is a Senior Member of the Association for Computing Machinery.