PhoenixSim: A Simulator for Physical-Layer Analysis of Chip-Scale Photonic Interconnection Networks

Johnnie Chan, Gilbert Hendry, Aleksandr Biberman, Keren Bergman
Department of Electrical Engineering
Columbia University, New York, NY
{johnnie,gilbert,biberman,bergman}@ee.columbia.edu

Luca P. Carloni
Department of Computer Science
Columbia University, New York, NY
luca@cs.columbia.edu

***Abstract*—Recent developments have shown the possibility of leveraging silicon nanophotonic technologies for chip-scale interconnection fabrics that deliver high bandwidth and power efficient communications both on- and off-chip. Since optical devices are fundamentally different from conventional electronic interconnect technologies, new design methodologies and tools are required to exploit the potential performance benefits in a manner that accurately incorporates the physically different behavior of photonics. We introduce PhoenixSim, a simulation environment for modeling computer systems that incorporates silicon nanophotonic devices as interconnection building blocks. PhoenixSim has been developed as a cross-discipline platform for studying photonic interconnects at both the physical-layer level and at the architectural and system levels. The broad scope at which modeled systems can be analyzed with PhoenixSim provides users with detailed information into the physical feasibility of the implementation, as well as the network and system performance. Here, we describe details about the implementation and methodology of the simulator, and present two case studies of silicon nanophotonic-based networks-on-chip.**

I. Introduction

The scaling of chip multiprocessor (CMP) systems is introducing an increasingly communication-limited performance bottleneck and creating a need for scalable and power efficient interconnection networks. While electronics has thus far been able to cope with the bandwidth and performance demands of today's systems, further scaling will be strained by power dissipation limits in the processor package. Electronic interconnects and the associated communication infrastructure are taking an increasing portion of the chip power budget, drawing over 50% of the consumed dynamic power in some high-performance processors [1]. As these power issues continue to challenge the scaling of CMP system, new technologies may be needed to deliver energy-efficient high-bandwidth communications.

Thanks to recent progress in optical device integration [2][3][4][5][6], silicon photonics has emerged as a promising technology platform for chip-scale interconnection networks. In comparison to electronics, photonics can potentially provide higher bandwidth through wavelength-division-multiplexed (WDM) transmission and better energy efficiency for global on- and off-chip communications. Since photonic devices are fundamentally different in how they function, exploiting these advantages would require a drastic change in how on- and off-chip interconnects are designed. In particular, optical signals cannot be buffered nor processed without being converted first to the electronic domain. Typically, optical-electronic-optical (O-E-O) conversions are performed at the terminals in large scale optical networks. At the chip- and board-scale, however, O-E-O conversions should be minimized since the power-dissipation penalty incurred can be significant. Many photonic interconnect designs have been proposed to avoid in-flight processing or buffering [7][8][9][10][11]. Further, since signal regeneration in optics cannot be economically accomplished on the CMOS-compatible silicon photonic platform, all photonic transmissions must propagate through the length of the transmission path without accumulating significant optical loss. Despite these major design constraints, the potential advantage in power and performance that photonic interconnects can offer makes them a solution worth pursuing for next generation CMPs.

We propose PHOENIXSIM, the Photonic and Electronic Network Integration and Execution Simulator, for modeling and analyzing the performance of multiprocessor systems that use electronic networks, photonic networks, and hybrid networks (ones that leverages a combination of both technology domains). In contrast with conventional network simulators, PHOENIXSIM can capture the physical characteristics and metrics of the photonic interconnection devices and network elements which have no electronic equivalent. We emphasize the physical-layer characterizations that play a fundamental role in determining system performance, which differs from other recent works in the modeling of optical networks-on-chip (NoCs) that focus on system-level behavior [12]. Our simulator was developed in the OMNeT++ discrete-event simulation environment [13] and relies on a library of electrical and photonic device models that are highly parameterized. This allows us to analyze and simulate both systems that are based on currently-realizable devices as well as those based on performance projections of future devices. With PHOENIXSIM the performance metrics of interconnection networks can be analyzed both at the physical level (*e.g.* optical insertion loss, crosstalk, energy dissipation) and system level (*e.g.* latency, performance, execution time). We present the design methodology that is enabled by PHOENIXSIM and demonstrate its capabilities through case studies of a photonic on-chip interconnection network [10] and a photonic off-chip memory-access interconnect [8].

978-3-9810801-6-2/DATE10 © 2010 EDAA

II. PHOTONIC DEVICE LIBRARY

At the foundation of PHOENIXSIM is the *Photonic Device Library* which consists of a set of elementary photonic devices that can be joined together to create photonic switches and topologies. A key goal of this modeling effort is to maintain a balanced level of detail and accuracy to enable the concurrent study of both physical-layer metrics and system-level performances with a reasonable amount of computing power. We have selected to build the environment initially with ring-resonator devices based on their versatility and wide use in photonic NoC design. Other photonic elements such as Mach-Zehnder modulators and tunable filters can also be readily incorporated into PHOENIXSIM.

Individual devices are described using our *Basic Element Model*, which abstracts the physical characteristics and behavior of the devices to create atomic blocks for building networks. Note that while the devices presented here are focused in the photonic domain, the modeling framework allows the generic implementation of many devices from other technology domains. In particular, PHOENIXSIM provides an interface through which a user can create device models simply by specifying the number of ports an optical signal can ingress or egress from, and the insertion loss and delay associated for every pair of ports. Basic Element devices are assumed to be broadband and therefore exhibit behavior that is independent of the wavelength of an incoming signal. Below we describe the key photonic devices, their interconnect functionality, and how they are modeled in the PHOENIXSIM environment.

A. Passive Elements

Waveguides. Waveguides provide the physical links between all sources and destinations and enables connectivity between all photonic devices. A photonic signal experiences insertion loss (*i.e.* attenuation) as it propagates through the waveguide due to light scattering at the waveguide sidewalls. This power loss due to propagation has been measured to be as low as 1.7 dB/cm [14], and is expected to improve with optimized fabrication techniques to less than 0.5 dB/cm. PHOENIXSIM models waveguides as 2-port devices with a single length parameter. Loss and delay are derived from the length and global variables that specify the material properties.

Waveguide Bends. Bending along waveguides is necessary to properly route optical paths and create compact switch and topology designs. Bends introduce additional insertion loss which has been experimentally measured to be 0.005dB per 90° [14]. In PHOENIXSIM, waveguide bends are modeled as 2-port devices with a radius of curvature of 2.5 μm and the calculation of loss and delay is based on the bend degrees.

Waveguide Crossings. Crossings are a byproduct of requiring planar topology fabrication for on-chip networks. Fortunately, photonic link crossings can be designed to mostly suppress insertion loss and crosstalk through the use of expanded double-etched crossing structures [15]. Insertion loss due to propagation through the waveguide crossing has been measured as low as 0.16 dB, and the crosstalk (light that leaks onto the waveguide that is orthoganal to the direction of propagation) has been measured to be about -40 dB [15]. In PHOENIXSIM, all cross elements have 4 ports and are assumed to be uniform in physical performance, deriving the insertion loss and delay from global variables.

Couplers. While traditional electronic system design is typically restrictive in cross-boundary data transmission (such as going from on-chip to off-chip), photonic interconnect-enabled systems possess the unique capability of crossing those boundaries with minimal impact on interconnect performance. Integrated optical I/O enables bandwidth transparency for off-chip signaling, and, unlike electrical I/O, the resulting signal integrity is practically impervious to propagation distance. Additionally, the power consumed in off-chip photonic communications is comparable to that of photonic on-chip message transfers, reducing the on- and off-chip bandwidth mismatch brought on by power limitations in current systems. The I/O interface can be accomplished through vertical coupling on the chip surface or lateral coupling at the chip edge, with theoretical losses of around 1 dB [16]. PHOENIXSIM models couplers with a single coupling loss parameter to account for the optical attenuation that is experienced at these interfaces.

B. Ring Resonator Active Elements

One of the primary elements used by silicon photonic circuit designers is the micro-ring resonator, which has a large range of functionalities due to the flexibility of its design space [2][3][5][17][18]. Ring resonators are capable of guiding the path an optical signal will take through careful design of the dimensions and position of the resonator. Optical signals couple into ring resonators at specific regularly spaced wavelengths in the optical spectrum, called resonant modes. The modes are located at multiples of the *free spectral range* (FSR) which is inversely related to the circumference (optical length) of the ring. Larger rings have tighter mode spacing, while smaller rings have wider spacing. The FSR can be adjusted either by changing the physical dimensions or through changes in the index of refraction using electro-optic and thermal means. Consequently, the micro-ring resonator can be designed to perform many of the tasks required for photonic signal generation, routing, and reception.

In PHOENIXSIM ring-resonator devices are modeled using the *Ring Element Model*. Since their behavior is dependent on the light wavelength, the Ring Element Model must extend the Basic Element Model to include this dependency, which can be completely specified by the diameter of the ring. The model also includes a way to simulate electro-optic control through the specification of multiple *states*. Each state of a device can have entirely different properties, depending on the design.

Filters. For typical optical filters, it is only necessary to operate on a single wavelength channel at a time. This can be done by designing a ring resonator with as large a FSR as possible, so that only a single resonant mode appears within the spectrum of interest. Filtering can be accomplished by aligning a single wavelength channel (Fig. 1a) from a WDM signal with the mode of the ring filter. The on-resonance

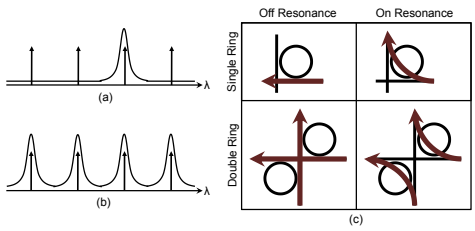

Fig. 1. Propagation through a ring-resonator device depends on the signal wavelength and the resonant modes of the device. (a) Small rings with larger mode spacings (shown as periodic peaks) can be designed to interact with a single wavelength channel from a WDM signal (indicated by arrows). (b) Broadband switch have tightly spaced modes, enabling many WDM channels to couple into the device cohesively. (c) The path of propagation depends on whether the wavelength of the message is on- or off-resonance with the ring.

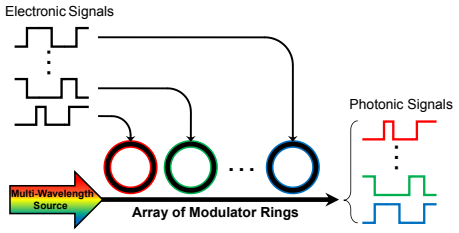

Fig. 2. Schematic of the conversion process between the spatially-parallel electronic domain and wavelength-parallel optical domain.

wavelength couples from the initial waveguide into the ring structure and then out to a second waveguide while off-resonance wavelengths pass by the ring uninterrupted (Fig. 1c). Ring filters have been demonstrated as small as 3 μm in diameter [18]. In PHOENIXSIM, filter elements are single-state 4-port devices with a single diameter parameter.

Modulators. Since ring-based modulators are designed to encode a data stream onto a single wavelength channel, they should have a minimal ring diameter to maximize FSR. By placing a series of uniquely-tuned rings on a common waveguide several individual wavelength channels can be modulated into a complete WDM signal (Fig. 2). Ring modulation has been demonstrated at rates of 12.5 Gbps [3]. In PHOENIXSIM modulator elements are modeled with parameters for energy-per-bit and ring diameter.

Broadband Switches. Micro-ring resonators can also be leveraged to route entire WDM messages between a source and destination. With *space routing* larger micro-ring resonators can be used to manipulate cohesively the entire WDM signal. This is accomplished by aligning all WDM channels into the periodic modes of the ring, which are more closely packed together due to the larger ring dimensions (Fig. 1b). The FSR of the ring can also be controlled through electro-optic means by shifting every mode away from the transmission channels and causing the entire signal to pass by the ring uninterrupted. This functionality is illustrated in Fig. 1c for both a single-ring 1×2 photonic switching element (PSE) and a double-ring 2×2 PSE: the entire WDM signal switches depending on whether the PSE is on- or off-resonance. In PHOENIXSIM broadband switch elements are modeled as 2-state 4-port devices (depending on the state either all or none of the channels of a WDM signal are extracted) [5][17].

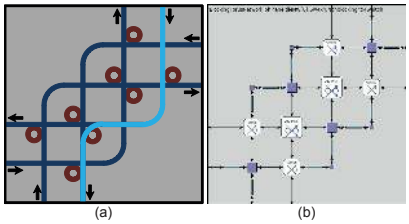

Fig. 3. (a) Schematic of a design for a 4×4 non-blocking photonic switch. (b) A screenshot of how PHOENIXSIM composes the switch by instancing basic photonic devices.

C. Receivers

Photo-Detector. A high-speed photo-detector translates the photonic message back into the electrical domain. Germanium is a CMOS-compatible material that absorbs in the wavelength range of interest. Significant advancements have been made in the performance of germanium and silicon-germanium photo-receivers [6]. Detector elements require the specification of an energy-per-bit and ring diameter parameter.

III. PHOTONIC INTERCONNECTION NETWORK MODELING

PHOENIXSIM combines the photonic modeling capabilities of the Photonic Device Library with additional models for electronic routers and traffic generators to produce a variety of switch and network fabrics.

Higher-Order Networking Components. Models from the Photonic Device Library can be combined to derive more complex photonic components and interconnect network structures. The overall performance of such complex aggregate components is determined by the performance of its individual elements. This form of encapsulation allows the designer to create large, complex, and physically-accurate networks spanning an entire system, while only requiring the characterization of the physical details of a few elementary devices. Examples of higher-order component are large-radix broadband switches (beyond the 1×2 and 2×2 PSEs) which are of particular importance in a variety of networks to support scalability and connectivity. The 4×4 switch shown in Fig. 3 is a critical component for efficient routing of traffic through an on-chip photonic network topology [4][11]. In PHOENIXSIM, this 4×4 switch can be composed from the building blocks, including the use of 1×2 and 2×2 switches, waveguides, waveguide bends, and waveguide crossings (Fig. 3b).

Electronic Routers. PHOENIXSIM uses a standard pipeline electronic router model, containing building blocks for buffering, arbitration, and switching. Electronic delay and energy dissipation leverage the ORION simulator [19]. The router model is highly configurable including parameters for buffer size, flit size, channel width, clock rate, and number of virtual channels. Additionally, each router model also supports the control and arbitration of actively-switched electro-optic photonic devices. Active network arbitration has been proposed using either electronic signaling [20] or optical signaling [7]. Electrical arbitration is enabled through a separate electronic control plane that signals the photonic circuit-switched network. In order for the control plane to properly arbitrate the photonic plane, an electronic router

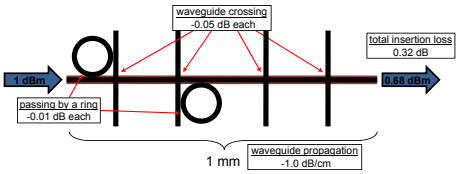

Fig. 4. Calculation of insertion loss for a small network segment.

must be placed at each photonic switch, effectively creating a mirrored topology in the electronic domain. The network uses a circuit-switching protocol to allocate paths and prevent message collisions on the photonic network. Control messages travel through the packet-switched control plane to enable the necessary ring resonators within each switch to trace out a complete path from source to destination.

Traffic Generators. PHOENIXSIM uses a processor model to generate synthetic traffic patterns. Currently the supported synthetic traffic patterns include random, hot-spot, nearest-neighbor, and tornado. Each pattern is parameterized for such variables as interarrival time and message size. Trace files generated by monitoring communication traffic of real applications can be read into the simulator.

IV. PERFORMANCE ANALYSIS TOOLS

In this section we describe a set of unique physical metrics that are important in characterizing the performance of photonic network designs and how they are evaluated in PHOENIXSIM. Since photonic NoCs are still in early stages of research, full-scale analysis must be done in simulation, and the tools presented here will give valuable information into the physical feasibility of the design.

Insertion Loss. Photonic transmission at the intra- and inter-chip scale must be accomplished without signal regeneration due to the difficulties in creating silicon-based optical amplifiers. For this reason, it is critical for a photonic network design to minimize *insertion loss*, which is the power attenuation incurred by an optical signal along its path of propagation. Everywhere along the transmission path, beginning at the laser source and ending at a receiving optical detector, the optical message accrues insertion loss as it interacts with all the photonic devices. Fig. 4 shows a simple example of a signal injected into a network segment at 1 dBm and being received at 0.68 dBm after a propagation distance of 1 mm, passing by two micro-ring resonators, and going through four waveguide crossings. The insertion loss in this case is 0.32 dB. The complexity and size of a network is ultimately limited by the insertion loss since a photonic link can only exhibit a certain amount of loss before the signal becomes too weak to be received properly. It is critical for system designers to account for insertion loss since it plays a direct role in scalability and reliability of the network.

Closely related to the insertion loss is the *optical loss budget*, which has implications on the design, scalability, and performance of the entire photonic network. This parameter is assessed from the difference of the maximum injectable laser power into the network and the minimum detectable power at the receivers. The light source injection power is limited by the threshold of undesirable nonlinear optical effects in silicon, which deteriorate the signal integrity when the signal power is too high. While WDM transmission enables data signals to be transmitted in parallel across different wavelength channels, the total optical power (sum across all present wavelength channels) must still remain below this nonlinear limit, reducing the allowed injected power for each wavelength channel. The relationship between the device limitations and system-level metrics is summarized in the inequality $P - S \geq IL_{max} + 10log_{10}n$, where P is the power threshold we limit the optical power to and S is the sensitivity of the photodetector. The optical loss 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. This relationship shows that network designers should create smaller and less complex networks if they desire high-bandwidth connections. High-radix networks can also be supported by sacrificing bandwidth.

Crosstalk. An optical message typically leaks a small amount of optical power onto intersecting waveguides along its path of propagation. If another signal is present on this perpendicular waveguide, then each message will interfere with the other in the form of crosstalk. Similarly, crosstalk also occurs at ring-resonator filters and switches due to imperfect coupling of the wavelength channels. The crosstalk a message receives in a device depends on the power levels of all other signals present within the device. If a device is modeled as having N ports from which an optical signal can ingress or egress, then the message can receive crosstalk from up to $N-1$ foreign messages. If M is the set of signals present in the device and P_k is the power of signal k, then the crosstalk seen by signal s is given by

$$\sum_{k \in M, k \neq S} \frac{P_k}{IL(p_{k.in}, p_{s.out})}$$

which aggregates the unwanted signal power that leaks into the output port assumed by s. Function IL specifies the insertion loss between any two ports on the device, where $p_{k.in}$ denotes the input port of a message other than s, and $p_{s.out}$ denotes the output port of s. This calculation is a simple approximation that only considers crosstalk for messages that coexist in a device and not from leakage power that propagates a certain distance before leaking into a foreign signal. Improving this approximation is planned as future work.

The signal-to-noise ratio (SNR) captures the integrity of a message. From a system performance standpoint, the SNR measures how likely a signal can be received without errors. The signal power is simply the mean power at which the message is received at the detector while the noise power is accumulated from crosstalk, laser noise, and detector noise. Laser noise contributions can be measured in terms of relative intensity noise (RIN), which is the ratio of variance of the optical power to the mean optical power squared. Quantum cascade lasers have been measured to have RIN on the order of -150 dB Hz^{-1} with an output of 10 dBm mean optical power [21]. To convert to a more meaningful value (SNR), we use the theoretical relation $SNR = m^2/(2B \cdot RIN)$ [22], where

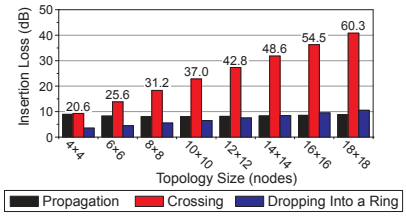

Fig. 5. Insertion loss results for varying size photonic tori. Displayed values represent the total loss for the worst-case path.

TABLE I
INSERTION LOSS COMPONENT PARAMETERS

Component	Loss Parameter
Propagation Loss (Silicon)	1.5 dB/cm
Waveguide Crossing	0.15 dB
Waveguide Bend	0.005 dB/90°
Drop Into a Ring	0.5 dB
Pass By a Ring	0.005 dB

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. Detector noise is modeled with standard thermal noise and shot noise equations.

Energy Dissipation. PHOENIXSIM calculates the total energy dissipated accounting from all individual devices found in the network model. For power modeling of electronic routers we leverage ORION, which outputs values for both dynamic and static power dissipation at various technology nodes including 32 nm and 22 nm [19]. Modulators and broadband switches both require driver circuitry for the electro-optic control which also leaks power both dynamically and statically. In addition, we assume that all ring devices require real-time thermal tuning to compensate for thermal fluctuations in the system and fabrication imperfections. Currently PHOENIXSIM calculates thermal tuning by modeling it as a static dissipation cost.

V. CASE STUDIES

To demonstrate the diverse capabilities of PHOENIXSIM we analyze two fundamentally different photonic interconnection networks, each with a different PHOENIXSIM analysis tool.

A. On-Chip Photonic Torus Network

The photonic torus network proposed by Shacham *et al.* is a circuit-switched NoC that uses photonic broadband switches for high-bandwidth communications on a CMP [10]. We implemented a model of this NoC in PHOENIXSIM and produced worst-case insertion loss values for network sizes of 4×4 to 18×18 using the conservative loss parameters listed in Table I. The results are shown in Fig. 5. Losses due to bending and passing by a ring-resonator were found to be negligible in this network. The plot indicates that the losses due to waveguide crossings dominate the overall insertion loss as the network scales up in size. These are clearly a major impediment to the realization and scaling of this photonic torus design without drastic improvements in device losses.

Fig. 6 shows the number of wavelengths (solid lines) that can be achieved by varying the topology size, while continuing to assume the insertion loss results based on

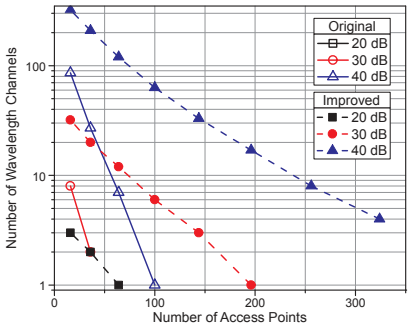

Fig. 6. Photonic torus NoC size vs. possible number of wavelengths for given optical power budgets.

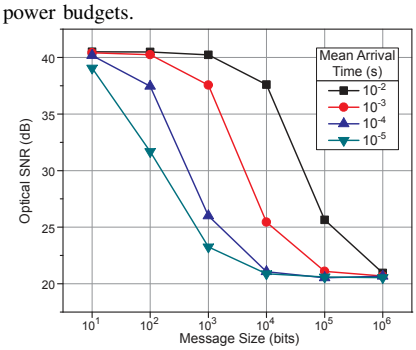

Fig. 7. Optical SNR values for a 4×4 photonic torus NoC for varying message sizes and mean arrival times.

the conservative parameters. Also shown are results for a hypothetical improvement in crossing loss from 0.15 dB to 0.05 dB (dashed lines). The largest torus network possible with a 40-dB optical loss budget before improvement is a 10×10 (100 access points), whereas the network is able to reach a size of 18×18 (324 access points, limit of this set of simulations) with the improvement in crossing loss.

Optical SNR (OSNR) is the measure of the SNR just before the detector, which accounts for only the laser intensity noise and crosstalk but not the noise component due to the detector circuit. The OSNR of a 4×4 photonic torus is plotted in Fig. 7. Regardless of message size or mean arrival time, the network exhibits a maximum OSNR of approximately 40.5 dB, the limiting case when the noise power is exclusively due to the laser intensity noise, and a minimum OSNR of approximately 20.5 dB, which occurs when the network is fully saturated.

B. Off-Chip Memory Access Network

Then we implemented in PHOENIXSIM a model for the off-chip photonic memory-access network proposed by Batten *et al.* for manycore processors [8]. This network leverages wavelength selective routing to enable contention-free photonic traffic for up to 256 cores and 16 memory banks. The network integrates passive ring-resonator filters in a centralized photonic crossbar to route core-to-memory and memory-to-core data transmissions. To facilitate the sharing of network resources, the chip layout is arranged into groups of 16 cores and 16 photonic access points (one to each memory bank) connected via a local electronic mesh network. The memory crossbar network traffic is routed through a combination of source routing and wavelength-selective routing, which does not require active switching

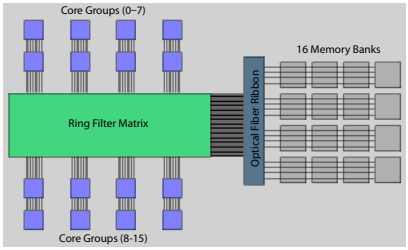

Fig. 8. PHOENIXSIM model of a many-core processor-to-DRAM system [8].

TABLE II
ENERGY DISSIPATION PARAMETERS FOR PHOTONIC COMPONENTS

Component	Energy Parameter
Thermal Ring Tuning	100 μW
Modulator	85 fJ/bit
Detector	50 fJ/bit

of the ring resonators and, therefore, does not incur the arbitration overhead that is required in the photonic torus. Upon reaching the edge of the chip, the optical signal couples into an off-chip optical silica-fiber ribbon where it is guided to a remote memory bank. The top-level view of the memory-access network simulation model is shown in Fig. 8. Note that while the model is faithful to the proposed design, additional assumptions were made about layout and device performances to fully populate the parameter list required by PHOENIXSIM.

Fig. 9 shows the power dissipation breakdown for this photonic memory-access network, based on the energy parameters per component given in Table II and the ORION model for the 32-nm technology with normal voltage-threshold transistors and a 2.5 GHz clock rate. Router buffers and logic account for a dominant portion of the total power dissipation while the contribution of the electronic wires grows with the injection rate. In contrast, the power dissipated by the photonic components remains limited across the various injection rates. Saturation occurs at an injection rate of 10^9 packets per second for a total network throughput of approximately 22 Tbps, which matches the results from [8].

VI. CONCLUSIONS

Photonic network technology, once relegated exclusively to large-scale telecommunication networks, has in recent years gradually been penetrating into smaller scale networking domains, with the potential to eventually become a viable architectural solution for on-board and chip-scale systems. We have introduced PHOENIXSIM as a simulation environment for the design, analysis, and optimization of these high-performance interconnection networks in a manner that accurately captures the physical-layer aspects of the devices while enabling system performance evaluation. This combination of models and tools in a single integrated environment provides a unique resource for the design exploration of the new systems enabled by photonic networks.

ACKNOWLEDGMENTS

This research is partially supported by DARPA MTO under grant ARL-W911NF-08-1-0127, the NSF (Award #: 0811012), and the FCRP Interconnect Focus Center (IFC).

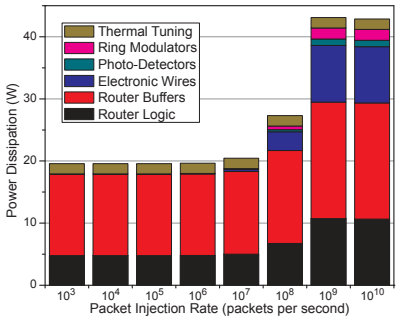

Fig. 9. Power dissipation breakdown of the memory-access network for varying injection rates per core. Injection rate is the inverse of the average time between the transmission of a message and the arrival of the next message.

REFERENCES

[1] N. Magen *et al.*, "Interconnect-power dissipation in a microprocessor," in *Intl. Workshop on Syst.-Level Interconnect Prediction*, 2004, pp. 7–13.
[2] M. R. Watts *et al.*, "Ultralow power silicon microdisk modulators and switches," in *IEEE Int. Conf. on Group IV Photonics*, Sep. 2008.
[3] Q. Xu *et al.*, "12.5 Gbit/s carrier-injection-based silicon micro-ring silicon modulators," *Opt. Exp.*, vol. 15, no. 2, pp. 430–436, 2007.
[4] B. G. Lee *et al.*, "Multi-wavelength message routing in a non-blocking four-port bidirectional switch fabric for silicon photonic networks-on-chip," in *Optical Fiber Commun. Conf.*, 2009.
[5] Y. Vlasov, W. M. J. Green, and F. Xia, "High-throughput silicon nanophotonic wavelength-insensitive switch for on-chip optical networks," *Nature Photonics*, vol. 2, pp. 242–246, Apr. 2008.
[6] S. Koester *et al.*, "Ge-on-SOI-detector/Si-CMOS-amplifier receivers for high-performance optical-communication applications," *J. of Lightwave Technology*, vol. 25, no. 1, pp. 46–57, Jan. 2007.
[7] D. Vantrease *et al.*, "Corona: System implications of emerging nanophotonic technology," in *Intl. Symp. on Comput. Architecture*, Jun. 2008, pp. 153–164.
[8] C. Batten *et al.*, "Building many-core processor-to-DRAM networks with monolithic CMOS silicon photonics," *IEEE Micro*, vol. 29, no. 4, pp. 8–21, Jul.-Aug. 2009.
[9] N. Kirman *et al.*, "On-chip optical technology in future bus-based multicore designs," *IEEE Micro*, vol. 27, no. 1, pp. 56–66, 2007.
[10] A. Shacham, K. Bergman, and L. P. Carloni, "Photonic networks-on-chip for future generations of chip multiprocessors," *IEEE Trans. on Comput.*, vol. 57, no. 9, pp. 1246–1260, Sep. 2008.
[11] 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, Jul.-Aug. 2009.
[12] M. Briere *et al.*, "Heterogeneous modelling of an optical network-on-chip with SystemC," in *IEEE Intl. Workshop on Rapid Syst. Prototyping*, 2005, pp. 10–16.
[13] A. Varga, "OMNeT++ discrete event simulation system," http://www.omnetpp.org.
[14] F. Xia, L. Sekaric, and Y. Vlasov, "Ultracompact optical buffers on a silicon chip," *Nature Photonics*, vol. 1, pp. 65–71, 2006.
[15] W. Bogaerts *et al.*, "Low-loss, low-cross-talk crossings for silicon-on-insulator nanophotonic waveguides," *Opt. Lett.*, vol. 32, no. 19, pp. 2801–2803, 2007.
[16] V. R. Almeida, R. R. Panepucci, and M. Lipson, "Nanotaper for compact mode conversion," *Opt. Lett.*, vol. 28, pp. 1302–1304, 2003.
[17] B. G. Lee *et al.*, "High-speed 2×2 switch for multi-wavelength message routing in on-chip silicon photonic networks," in *Eur. Conf. on Optical Commun.*, vol. 2, no. 96, Sep. 2008.
[18] B. E. Little *et al.*, "Ultra-compact Si-SiO_2 microring resonator optical channel dropping filters," *IEEE Photonics Technology Lett.*, vol. 10, no. 4, pp. 549–551, Apr. 1998.
[19] H. Wang *et al.*, "ORION: A power-performance simulator for interconnection networks," in *Int. Symp. on Microarchitecture*, 2002.
[20] A. Shacham, K. Bergman, and L. P. Carloni, "On the design of a photonic network-on-chip," in *Symp. on Networks-on-Chip*, May 2007, pp. 53–64.
[21] T. Gensty, W. Elsäßer, and C. Mann, "Intensity noise properties of quantum cascade lasers," *Opt. Exp.*, vol. 13, no. 6, pp. 2032–2039, 2005.
[22] C. Miller, *Fiber Optic Test and Measurement*. Prentice Hall, 1998.