

MIT Open Access Articles

DSENT - A Tool Connecting Emerging Photonics with Electronics for Opto-Electronic Networks-on-Chip Modeling

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Sun, Chen, Chia-Hsin Owen Chen, George Kurian, Lan Wei, Jason Miller, Anant Agarwal, Li-Shiuan Peh, and Vladimir Stojanovic. "DSENT - A Tool Connecting Emerging Photonics with Electronics for Opto-Electronic Networks-on-Chip Modeling." 2012 IEEE/ACM Sixth International Symposium on Networks-on-Chip (n.d.).

As Published: <http://dx.doi.org/10.1109/NOCS.2012.31>

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: <http://hdl.handle.net/1721.1/85863>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



DSENT – A Tool Connecting Emerging Photonics with Electronics for Opto-Electronic Networks-on-Chip Modeling

Chen Sun, Chia-Hsin Owen Chen, George Kurian, Lan Wei, Jason Miller, Anant Agarwal,
Li-Shiuan Peh and Vladimir Stojanovic
Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Cambridge, MA 02139

Email: {sunchen, owenhsin, gkurian, lanwei, jasonm, agarwal, peh, vlada}@mit.edu

Abstract—With the rise of many-core chips that require substantial bandwidth from the network on chip (NoC), integrated photonic links have been investigated as a promising alternative to traditional electrical interconnects. While numerous opto-electronic NoCs have been proposed, evaluations of photonic architectures have thus far had to use a number of simplifications, reflecting the need for a modeling tool that accurately captures the tradeoffs for the emerging technology and its impacts on the overall network. In this paper, we present DSENT, a NoC modeling tool for rapid design space exploration of electrical and opto-electrical networks. We explain our modeling framework and perform an energy-driven case study, focusing on electrical technology scaling, photonic parameters, and thermal tuning. Our results show the implications of different technology scenarios and, in particular, the need to reduce laser and thermal tuning power in a photonic network due to their non-data-dependent nature.

Index Terms—photonics; networks on chip; power

I. INTRODUCTION

As CMOS technology scales into the deep sub-100 nm regime, improvements in transistor density have resulted in greater processor parallelism as the means to improve processor performance, leading to ever-higher processor core counts. The rise of the many-core era, however, comes with the challenge of designing the on-die interconnect fabric to allow for efficient delivery of bits between an ever increasing number of processor cores, memories, and specialized IP blocks both on- and off-chip. Traditional approaches, such as the shared bus or global crossbars, scale poorly in either performance or cost for large numbers of network endpoints, driving the need for efficient Network-on-Chip (NoC) architectures to tackle the communication requirements of future many-core machines.

Recognizing the potential scaling limits of electrical interconnects, architects have recently proposed emerging nanophotonic technology as an option for both on-chip and off-chip interconnection networks [1, 2, 3, 4]. As optical links avoid the capacitive, resistive and signal integrity constraints imposed upon electronics, photonics allows for efficient realization of physical connectivity that is costly to accomplish electrically.

Photonics technology itself, however, remains immature and there remains a great deal of uncertainty in its capabilities. Whereas there has been significant prior work on electronic NoC modeling (see Section II-C), evaluations of photonic NoC architectures have not yet evolved past the use of fixed energy costs for photonic devices and interface circuitry [1, 3, 4, 5], whose values also vary from study to study. In order to gauge the true potential of this emerging technology, inherent interactions between electronic/photonic components and their impact on the NoC need to be quantified.

In this paper, we propose a unified framework for photonics and electronics, DSENT (Design Space Exploration of Networks Tool), that enables rapid cross-hierarchical area and power evaluation of opto-electronic on-chip interconnects¹. We design DSENT for two primary usage modes. When used standalone, DSENT functions as a fast design space exploration tool capable of rapid power/area evaluation of hundreds of different network configurations, allowing for impractical or inefficient networks to be quickly identified and pruned before more detailed evaluation. When integrated with an architectural simulator [6, 7], DSENT can be used to generate traffic-dependent power-traces and area estimations for the network [8].

Through DSENT, our paper makes the following contributions:

- Presents the first tool that is able to capture the interactions at electronic/photonic interface and their implications on a photonic NoC.
- Provides the first network-level modeling framework for electrical NoC components featuring integrated timing, area, and power models that are accurate (within 20%) in the deep sub-100 nm regime.
- Identifies the most profitable opportunities for photonic network optimization in the context of an entire opto-electronic network system. In particular, we focus on

¹We focus on the modeling of opto-electrical NoCs in this paper, though naturally, DSENT's electrical models can also be applied to pure electrical NoCs as well




Fig. 1: A typical opto-electronic NoC including electrical routers and links, and a wavelength division multiplexed intra-chip photonic link.

the impact of network utilization, technology scaling and thermal tuning.

The rest of the paper is organized as follows. Section II introduces the main building blocks of photonic NoCs and recaps existing work in photonic architectures and NoC modeling. We describe the DSENT framework in Section III and present its models for electrical and optical components in Sections IV and V, respectively. Validation of DSENT is shown in Section VI. Section VII presents an energy-efficiency-driven network case-study and Section VIII concludes the paper.

II. BACKGROUND

A. Silicon Photonics Technology

a) Waveguides, Couplers, and Lasers: Waveguides are the primary means of routing light within the confines of a chip. Vertical grating couplers [9] allow light to be directed both into and out-of the plane of the chip and provide the means to bring light from a fiber onto the chip or couple light from the chip into a fiber. In this paper, we assume commercially available off-chip continuous wave lasers, though we note that integrated on-chip laser sources are also possible [10, 11].

b) Ring Resonators: The optical ring resonator is the primary component that enables on-chip wavelength division multiplexing (WDM). When coupled to a waveguide, rings perform as notch filters; wavelengths at resonance are trapped in the ring and can be potentially *dropped* onto another waveguide while wavelengths not at resonance pass by unaffected. The resonant wavelength of each ring can be controlled by adjusting the device geometry or the index of refraction. As resonances are highly sensitive to process mismatches and temperature, ring resonators require active thermal tuning [12].

c) Ring Modulators and Detectors: Ring modulators modulate its resonant wavelength by electrically influencing the index of refraction [13]. By moving a ring's resonance in and out of the laser wavelength, the light is modulated (on-off keyed). A photodetector, made of pure germanium or SiGe, converts optical power into electrical current, which can then be sensed by a receiver [14] and resolved to electrical ones and zeros. Photodetectors standalone are generally wideband and require ring filters for wavelength selection in WDM operation.

d) Photonic Links: The dynamics of a wavelength-division-multiplexed (WDM) photonic architecture are shown in Figure 1. Wavelengths are provided by an external laser

source and coupled into an on-chip waveguide. Each wavelength is modulated by a resonant ring modulator dropped at the receiver by a matching ring filter. Using WDM, a single waveguide can support dozens of independent data-streams on different wavelengths.

B. Prior Photonic NoC Architectures

Many photonics-augmented architectures have been proposed to address the interconnect scalability issue posed by rapidly rising core-counts. The Corona [4] architecture uses a global 64x64 optical crossbar with shared optical buses employing multiple matching ring modulators on the same waveguide. Firefly [3] and ATAC [2] also feature global crossbars, but with multiple matching receive rings on the same waveguide in a multi-drop bus configurations. The photonic Clos network [5] replaces long electrical links characteristic of Clos topologies with optical point-to-point links (one set of matching modulator and receiver ring per waveguide) and performs all switching electrically. Phastlane [15] and Columbia [16] networks use optical switches in tile-able mesh-like topologies. While each of these prior works performs evaluations of their respective networks, we note that the analyses in these prior works all rely on fixed numbers for active photonic devices and electronic components, making it difficult to explore design tradeoffs and interactions between photonics and electronics.

C. Existing NoC Modeling Tools

Several modeling tools have been proposed to estimate the timing, power and area of NoCs. Chien proposed a timing and area model for router components [17] that is curve-fitted to only one specific process. Peh and Dally proposed a timing model for router components [18] based on logical effort that is technology independent; however, only one size of each logic gate and no wire model is considered in its analysis. These tools also only estimate timing and area, not power.

Among all the tools that provide power models for NoCs [19, 20, 21, 22], *Orion* [19, 22], which provides parametrized power and area models for routers and links, is the most widely used in the community. However, *Orion* lacks a delay model for router components, allowing router clock frequency to be set arbitrarily without impacting energy/cycle or area. Furthermore, *Orion* uses a fixed set of technology parameters and standard cell sizing, scaling the technology through a gate length scaling factor that does not reflect the effects of other technology parameters. For link components, *Orion* supports only limited delay-optimal repeated links. *Orion* does not model any optical components.

PhoenixSim [23] is the result of recent work in photonics modeling, improving the architectural visibility concerning the trade-offs of photonic networks. *PhoenixSim* provides parametrized models for photonic devices. However, *PhoenixSim* lacks electrical models, relying instead on *Orion* for all electrical routers and links. As a result, *PhoenixSim* uses fixed numbers for energy estimations for electrical interface

circuitry, such as modulator drivers, receivers, and thermal tuning, losing many of the interesting dynamics when transistor technology, data-rate, and tuning scenarios vary. *PhoenixSim* in particular does not capture trade-offs among photonic device and driver/receiver specifications that result in an area or power optimal configuration.

To address shortcomings of these existing tools, we propose DSENT to provide a unified electrical and optical framework that can be used to model system-scale aggressive electrical and opto-electronic NoCs in future technology nodes.

III. DSENT FRAMEWORK

In our development of the generalized DSENT modeling framework, we observe the constant trade-offs between the amount of required user input and overall modeling accuracy. All-encompassing technology parameter sets can enable precise models, at the cost of becoming too cumbersome for predictive technologies where only basic technology parameters are available. Overly simplistic input requirements, on the other hand, leaves significant room for inaccuracies. In light of this, we design a framework that allows for a high degree of modeling flexibility, using circuit- and logic-level techniques to simplify the set of input specifications without sacrificing modeling accuracy. In this section, we introduce the generalized DSENT framework and the key features of our approach.

A. Framework Overview

DSENT is written in C++ and utilizes the object-oriented approach and inheritance for hierarchical modeling. The DSENT framework, shown in Figure 2, can be separated into three distinct parts: user-defined models, support models, and tools. To ease development of user-defined models, much of the inherent modeling complexity is off-loaded onto support models and tools. As such, most user-defined models involve just simple instantiation of support models, relying on tools to perform analysis and optimization. Like an actual electrical chip design, DSENT models can leverage instancing and multiplicity to reduce the amount of repetitive work and speed up model evaluation, though we leave open the option to allow, for example, all one thousand tiles of a thousand core system to be evaluated and optimized individually. Overall, we strive to keep the run-time of a DSENT evaluation to a few seconds, though this will vary based upon model size and complexity.

B. Power, Energy, and Area Breakdowns

The typical power breakdown of an opto-electronic NoC can be formulated as Equation 1. The optical power is the wall-plug laser power (lost through non-ideal laser efficiency and optical device losses). The electrical power consists of the power consumed by electrical routers and links as well as electric-optical interface circuits (drivers and receivers) and

ring tuning.

$$P_{total} = P_{electrical} + P_{optical} \quad (1a)$$

$$P_{electrical} = P_{router} + P_{link} + P_{interface} + P_{tuning} \quad (1b)$$

$$P_{optical} = P_{laser} \quad (1c)$$

Power consumption can be split into *data-dependent* and *non-data-dependent* parts. Non-data-dependent power is defined as power consumed regardless of utilization or idle times, such as leakage and un-gated clock power. Data-dependent power is utilization-dependent and can be calculated given an energy per each event and frequency of the event. Crossbar traversal, buffer read and buffer write are examples of high-level events for a router. Power consumption of a component can thus be written as $P = P_{NDD} + \sum E_i \cdot f_i$, where P_{NDD} is the total non-data-dependent power of the module and E_i, f_i are the energy cost of an event and the frequency of an event, respectively.

Area estimates can be similarly broken down into their respective electrical (logic, wires, etc.) and optical (rings, waveguides, couplers, etc.) components. The total area is the sum of these components, with a further distinction made between active silicon area, per-layer wiring area, and photonic device area (if a separate photonic plane is used).

We note that while the area and non-data-dependent power can be estimated statically, the calculation for data-dependent power requires knowledge of the behavior and activities of the system. An architectural simulator can be used to supply the event counts at the network- or router-level, such as router or link traversals. Switching events at the gate- and transistor-level, however, are too low-level to be kept track of by these means, motivating a method to estimate transition probabilities (Section IV-D).

IV. DSENT MODELS AND TOOLS FOR ELECTRONICS

As the usage of standard cells is practically universal in modern digital design flows, detailed timing, leakage, and energy/op characterization at the standard-cell level can enable a high degree of modeling accuracy. Thus, given a set of technology parameters, DSENT constructs a standard cell library and uses this library to build models for the electrical network components, such as routers and repeated links.

A. Transistor Models

We strive to rely on only a minimal set of technology parameters (a sample of which is shown in Table I) that captures the major characteristics of deep sub-100 nm technologies without diving into transistor modeling. Both interconnect and transistor properties are paramount at these nodes, as interconnect parasitics play an ever larger role due to poor scaling trends [25]. These parameters can be obtained and/or calibrated using ITRS roadmap projection tables for predictive technologies or characterized from SPICE models and process design kits when available.

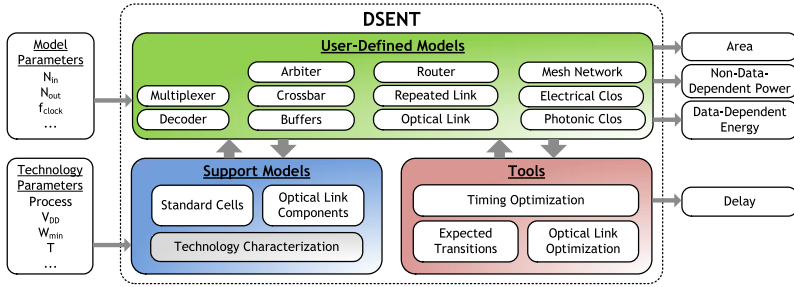


Fig. 2: The DSENT framework with examples of network-related user-defined models.

TABLE I: DSENT electrical parameters

Process Parameters	45 nm SOI	11 nm TG
Nominal Supply Voltage (V_{DD})	1.0V	0.6V
Minimum Gate Width	150 nm	40 nm
Contacted Gate Pitch	200 nm	44 nm
Gate Capacitance / Width	1.0 fF/ μ m	2.42 fF/ μ m
Drain Capacitance / Width	0.6 fF/ μ m	1.15 fF/ μ m
Effective On Current / Width [24]	650 μ A/ μ m	738 μ A/ μ m
Single-transistor Off Current	200 nA/ μ m	100 nA/ μ m
Subthreshold Swing	100 mV/dec	80 mV/dec
DIBL	150 mV/V	125 mV/V
Interconnect Parameters	45 nm SOI	11 nm TG
Minimum Wire Width	150 nm	120 nm
Minimum Wire Spacing	150 nm	120 nm
Wire Resistance (Min Pitch)	0.700 Ω / μ m	0.837 Ω / μ m
Wire Capacitance (Min Pitch)	0.150 fF/ μ m	0.167 fF/ μ m

Shown values are for NMOS transistors and the global wiring layer

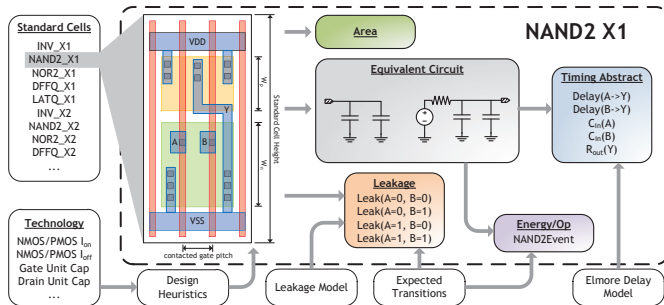


Fig. 3: Standard cell model generation and characterization. In this example, a NAND2 standard cell is generated.

Currently, DSENT supports the 45 nm, 32 nm, 22 nm, 14 nm and 11 nm technology nodes. Technology parameters for the 45 nm node are extracted using SPICE models. Models for the 32 nm node and below are projected [26] using the virtual-source transport of [27] and the parasitic capacitance model of [28]. A switch from planar (bulk/SOI) to tri-gate transistors is made for the 14 nm and 11 nm nodes.

B. Standard Cells

The standard-cell models (Figure 3) are portable across technologies, and the library is constructed at run-time based on design heuristics extrapolated from open-source libraries [29] and calibrated with commercial standard cells.

We begin by picking a global standard cell height, $H = H_{ex} + \alpha \cdot (1 + \beta) \cdot W_{min}$, where β represents the P-to-N ratio, W_{min} is the minimum transistor width, and H_{ex} is the extra height needed to fit in supply rails and diffusion separation. α is heuristically picked such that large (high driving strength) standard cells do not require an excessive number of transistor folds and small (low driving strength) cells do not waste too much active silicon area. For each standard cell, given a drive strength and function, we size transistors to match pull-up and pull-down strengths, folding if necessary. As lithography limitations at deep sub-100 nm force a fixed gate orientation and periodicity, the width of the cell is determined by the max of the number of NMOS or PMOS transistors multiplied by the contacted gate pitch, with an extra gate pitch added for separation between cells.

C. Delay Calculation and Timing Optimization


To allow models to scale with transistor performance and clock frequency targets, we apply a first-order delay estimation and timing optimization method. Using timing information in the standard cell models, chains of logic are mapped to stages of resistance-capacitance (RC) trees, shown in Figure 4a. An Elmore delay estimate [30, 31] between two points i and k can be formed by summing the product of each resistance and the total downstream capacitance it sees:

$$t_{d,i-k} = \ln(2) \cdot \sum_{n=i}^k \sum_{m=n}^k R_n \cdot C_m \quad (2)$$

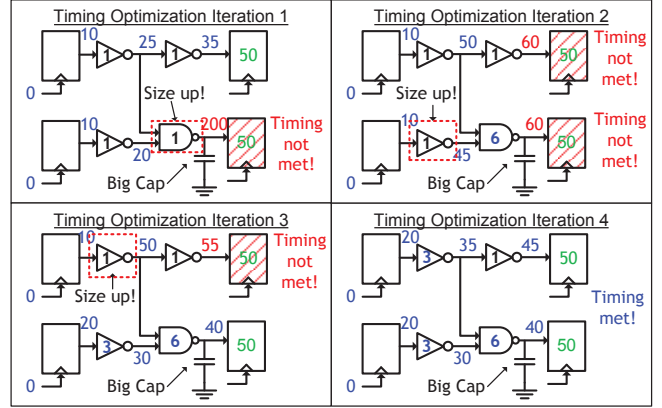
Note that any resistances or capacitances due to wiring parasitics is automatically factored along the way. If a register-to-register delay constraint, such as one imposed by the clock period, is not satisfied, timing optimization is required to meet the delay target. To this end, we employ a greedy incremental timing optimization algorithm. We start with the identification of a critical path. Next, we find a node to optimize to improve the delay on the path, namely, a small gate driving a large output load. Finally, we size up that node and repeat these three steps until the delay constraint is met or if we realize that it is not possible and give up. Our method optimizes for minimum energy given a delay requirement, as opposed to logical-effort based approaches employed by existing models [18, 32, 33], which optimize for minimum delay, oblivious to energy. Though lacking the rigorosity of timing optimization algorithms used by commercial hardware synthesis tools, our approach runs fast and performs well given its simplicity.

D. Expected Transitions

The primary source of data-dependent energy consumption in CMOS devices comes from the charging and discharging of transistor gate and wiring capacitances. For every transition of a node with capacitance C to voltage V , we dissipate an energy of $E = \frac{1}{2} C \cdot V^2$. To calculate data-dependent power usage, we sum the energy dissipation of all such transitions multiplied by their frequency of occurrence, $P_{DD} = \sum C_i \cdot V_i^2 \cdot f_i$. Node capacitance C_i can be calculated for each model and, for digital logic, V_i is the supply voltage. The frequency of occurrence, f_i , however, is much more



(a) Mapping standard cells to RC delays



(b) Timing optimization

Fig. 4: DSENT’s delay calculation and timing optimization framework. In (a), delay is estimated by mapping standard cells to sets of input capacitances and output drive resistances. Using these delay calculations, timing optimization in (b) may begin by incrementally sizing up cells until all delay constraints are met.

difficult to estimate accurately as it depends on the pattern of bits flowing through the logic. As event counts and signal information at the logic gate level are generally not available except through structural netlist simulation, DSENT uses a simplified expected transition probability model [34] to estimate the average frequency of switching events. Probabilities derived using this model are also used with state-dependent leakage in the standard cells to form more accurate leakage calculations.

E. Summary

DSENT models a technology-portable set of standard cells from which larger electrical components such as routers and networks are constructed. Given a delay or frequency constraint, DSENT applies (1) timing optimization to size gates for energy-optimality and (2) expected transition propagation to accurately gauge the power consumption. These features allow DSENT to outpace *Orion* in estimating electrical components and in projecting trends for future technology nodes.

V. DSENT MODELS AND TOOLS FOR PHOTONICS

A complete on-chip photonic network consists of not only the photonic devices but also the electrical interface circuits and the tuning components, which are a significant fraction of the link energy cost. In this section we present how we model these components in DSENT.

A. Photonic Device Models

Similar to how it builds the electrical network model using standard cells, DSENT models a library of photonic devices necessary to build integrated photonic links. The library includes models for lasers, couplers, waveguides, ring resonators, modulators and detectors. The total laser power required at the laser source is the sum of the power needed by each photodetector after applying optical path losses:

$$P_{laser} = \sum P_{sense,i} \cdot 10^{loss_i/10} \quad (3)$$

where $P_{sense,i}$ is the laser power required at photodetector i and $loss_i$ is the loss to that photodetector, given in dB. Note that additional link signal integrity penalties (such as near-channel crosstalk) are lumped into $loss_i$ as well.

B. Interface Circuitry

The main interface circuits responsible for electrical-to-optical and optical-to-electrical conversion are the modulator drivers and receivers. The properties of these circuits affect not only their power consumption, but also the performance of the optical devices they control and hence the laser power [12].

1) *Modulator Driver*: We adopt the device models of [12] for a *carrier-depletion* modulator. We first find the amount of charge ΔQ that must be depleted to reach a target extinction ratio, insertion loss, and data-rate. Using equations for a reverse-biased junction, we map this charge to a required reverse-biased drive voltage (V_{RB}) and calculate the effective capacitance using charge and drive voltage $C_{eff} = \Delta Q/V_{RB}$. Based on the data-rate, we size a chain of buffers to drive C_{eff} . The overall energy cost for a modulator driver can be expressed as:

$$E_{driver} = \frac{1}{\gamma} \Delta Q \cdot \max(V_{DD}, V_{RB}) + E_{buf}(C_{eff}, f) \quad (4)$$

where γ is the efficiency of generating a supply voltage of V_{RB} and $E_{buf}(C_{eff}, f)$ is the energy consumed by the chain of buffers that are sized to drive C_{eff} at a data-rate f .

2) *Receiver*: We support both the TIA and integrating receiver topologies of [12]. For brevity, we focus the following discussion on the integrating receiver, which consists of a photodetector connected across the input terminals of a current sense-amplifier. Electrical power and area footprints of the sense-amplifier is calculated based on sense-amplifier sizing heuristics and scaled with technology, allowing calculation of switching power. To arrive at an expression for receiver sensitivity (P_{sense}), we begin with an abbreviated expression for the required voltage buildup necessary at the receiver sense

amp’s input terminal:

$$V_d = v_s + v_{os} + v_m + \Phi(BER) \cdot \sqrt{\sum \sigma_n^2} \quad (5)$$

which is the sum of the sense-amp minimum latching input swing (v_s), the sense-amp offset mismatch (v_{os}), a voltage margin (v_m), and all Gaussian noise sources multiplied by the number of standard deviations corresponding to the receiver bit error rate. The required input can then be mapped to a required laser power requirement, P_{sense} at the photodetector:

$$P_{sense} = \frac{1}{R_{pd}} \cdot \frac{ER}{ER-1} \cdot V_d \cdot C_{in} \cdot \frac{2 \cdot f}{1 - 2 \cdot f \cdot t_j} \quad (6)$$

where R_{pd} is the photodetector responsivity (in terms of amps/watt), ER is the extinction ratio provided by the modulator, C_{in} is the total parasitic capacitance present at the receiver input node, f is the data rate of the receiver, and t_j is the clock uncertainty. The factor of 2 stems from the assumption that the photodetector current is given only half the clock period to integrate; the sense-amp spends the other half in the precharge state.

3) *Serializer and Deserializer*: DSENT provides models for a standard-cell-based serializer and deserializer (SerDes) blocks, following a mux/de-mux-tree topology [35]. These blocks provide the flexibility to run links and cores at different data-rates, allowing for exploration of optimal data-rates for both electrical and optical links.

C. Ring Tuning Models

An integrated WDM link relies upon ring resonators to perform channel selection. Sensitivity of ring resonances to ring dimensions and the index of refraction leaves them particularly vulnerable to process- and temperature-induced resonance mismatches [36, 37, 38], requiring active closed-loop tuning methods that add to system-wide power consumption [5]. In DSENT, we provide four models for four alternative ring tuning approaches [12]: *full-thermal tuning*, *bit-reshuffled tuning*, *electrically-assisted tuning*, and *athermal tuning*.

Full-thermal tuning is the conventional method of heating using resistive heaters to align their resonances to the desired wavelengths. Ring heating power is considered non-data-dependent, as thermal tune-in and tune-out times are too slow to be performed on a per-flit or per-packet basis and thus must remain always-on. Bit-reshufflers provide freedom in the bit-positions that each ring is responsible for, allowing rings to tune to its closest wavelength instead of a fixed absolute wavelength. This reduces ring heating power at the cost of additional multiplexing logic. Electrically-assisted tuning uses the resonance detuning principle of carrier-depletion modulators to shift ring resonances. Electrically-tuned rings do not consume non-data-dependent ring heating power, but is limited in tuning range and requires bit-reshufflers to make an impact. Note that tuning distances too large to be tuned electrically can still be bridged using heaters at the cost of non-data-dependent heating power. Athermal tuning represents an ideal scenario in which rings are not sensitive to temperature and all

process mismatches have been compensated for during post-processing.

D. Optical Link Optimization

Equations 4 and 6 suggest that both the modulator driver’s energy cost and the laser power required at the photodetector depend on the specification of extinction ratio and insertion loss of the modulator on the link. This specification can be used to tradeoff power consumption of the modulator driver circuit with that of the laser. This is an optimization degree of freedom that DSENT takes advantage of, looping through different combinations to find one that results in the lowest overall power consumption.

E. Summary

DSENT provides models not only for optical devices but also for the electrical backend circuitry including modulator driver, receiver and ring tuning circuits. These models enable link optimization and reveal tradeoffs between optical and electrical components that previous tools and analysis could not accomplish using fixed numbers.

VI. MODEL VALIDATION

We validate DSENT results against SPICE simulations for a few electrical and optical models. For the receiver and modulator models, we compare against a few early prototypes available in literature (fabricated at different technology nodes) to show that our results are numerically within the right range. We also compare our router models with a post-place-and-route SPICE simulation of a textbook virtual channel router and with the estimates produced by *Orion2.0* [22] at the 45 nm SOI technology node. To be fair, we also report the results obtained from a modified *Orion2.0* where we replaced *Orion2.0*’s original scaling factors with characterized parameters for the 45 nm SOI node and calibrated its standard cells with those used to calibrate DSENT. Overall, the DSENT results for electrical models are accurate (within 20%) compared to the SPICE simulation results. We note that the main source of inaccurate *Orion2.0* results is from the inaccurate technology parameters, scaling factors, and standard cell sizing. The recalibrated *Orion2.0* reports estimations at the same order of the SPICE results. The remaining discrepancy is partly due to insufficient modeling detail in its circuit models. For example, pipeline registers on the datapath and the multiplexers necessary for register-based buffers are not completely modeled by *Orion2.0*.

VII. EXAMPLE PHOTONIC NETWORK EVALUATION

Though photonic interconnects offer potential for improved network energy-efficiency, they are not without their drawbacks. In this section, we use DSENT to perform an energy-driven photonic network evaluation. We choose a 256-tile version of the 3-stage photonic cros network proposed by [5] as the network for these studies. Like [5], the core-to-ingress and egress-to-core links are electrical, whereas the ingress-to-middle and middle-to-egress links are photonic. The network configuration parameters are shown in Table III. While

TABLE II: DSENT validation points.

Model	Ref. Point	Orion2.0	Orion2.0 (re-calibrated)	DSENT	Configuration	
Ring Modulator Driver (fJ/bit)	[39]–50	N/A	N/A	60.87 (21.74%)	11 Gb/s, ER = 10 dB, IL = 6 dB	
Receiver (fJ/bit)	[14]–52	N/A	N/A	43.02 (-14.0%)	3.5 Gb/s, 45 nm SOI	
Router	Buffer (mW)	SPICE–6.93	34.4 (396%)	3.57 (-48.5%)	7.55 (8.94%)	<ul style="list-style-type: none"> • 6 input/output ports • 64-bit flit width • 8 VCs per port • 16 buffers per port • 1 GHz clock frequency • 0.16 flit injection rate
	Crossbar (mW)	SPICE–2.14	14.5 (578%)	1.26 (-41.1%)	2.06 (-3.74%)	
	Control (mW)	SPICE–0.75	1.39 (85.3%)	0.31 (-58.7%)	0.83 (10.7%)	
	Clock Dist. (mW)	SPICE–0.74	28.8 (3791%)	0.36 (-51.4%)	0.63 (-17.5%)	
	Total (mW)	SPICE–10.6	91.3 (761%)	5.56 (-47.5%)	11.2 (5.66%)	
	Total Area (mm ²)	Encounter–0.070	0.129 (84.3%)	0.067 (-4.29%)	0.062 (-11.2%)	

TABLE III: Network Configuration

Network Configuration	Values
Number of tiles	256
Chip area (divided equally amongst tiles)	400 mm ²
Packet length	80 Bytes
Flit width	128 bits
Core frequency	2 GHz
Clos configuration (m, n, r)	16, 16, 16
Link latency	2 cycles
Link throughput	128 bits/core-cycle
Router Configuration	Values
Number pipeline stages	3
Number virtual channels (VC)	4
Number buffers per VC	4

TABLE IV: Default Technology Parameters

Technology Parameters	Default Values
Process technology	11 nm TG
Optical link data-rate	2 Gb/s
Laser efficiency	0.25
Coupler loss	2 dB
Waveguide loss	1 dB/cm
Ring drop loss	1 dB
Ring through loss	0.01 dB
Modulator loss (optimized)	0.01–10.0 dB
Modulator extinction (optimized)	0.01–10.0 dB
Photodetector Capacitance	5 fF
Link bit error rate	10 ⁻¹⁵
Ring tuning model	Bit-Reshuffled [12, 41]
Ring heating efficiency	100 K/mW

DSENT includes a broader selection of network models, we choose this topology because there is an electrical network that is logically equivalent (an electrical clos) and carries a reasonable balance of photonic and electrical components. To obtain network-level event counts with which to animate DSENT’s physical models, we implement the clos network in Garnet [6] as part of the GEM5 [40] architecture simulator. Though the GEM5 simulator is primarily used to benchmark real applications, we assume a uniform random traffic pattern to capture network energy at specific loads. Given network event counts, DSENT takes a few seconds to generate an estimation.

In the following studies, we investigate the impact of different circuit and technology assumptions using energy cost per bit delivered by the network as our evaluation metric. Unless otherwise stated, the default parameters set in Table IV are used. The parameters we sweep are organized by section in Table V.

TABLE V: Sweep Parameters Organized by Section

Section	Sweep Parameter	Sweep Range
VII-A	Electrical Process	45 nm SOI, 11 nm Tri-Gate
VII-B	Waveguide Loss	0.0–2.5 dB
	Ring Heating Efficiency	10–400 K/mW
VII-C	Tuning Model	Full-Thermal, Bit-Reshuffled [12, 41], Electrically-Assisted [12]
	Link Data-Rate	2–32 Gb/s per λ

A. Scaling Electrical Technology and Utilization Tradeoff

We first compare the photonic clos network with an electrical equivalent, where all photonic links are replaced with electrical links of equal latency and throughput (128 wires, each at 2 GHz). We perform this comparison at the 45 nm SOI and 11 nm Tri-Gate technology nodes, representing present and future electrical technology scenarios, respectively. Energy per bit is plotted as a function of achieved network throughput (utilization) and a breakdown of the energy consumption at three specific throughputs is shown in Figure 5.

Note that in all configurations, the energy per bit rises sharply at low network utilizations, as non-data-dependent (NDD) power consumption (leakage, un-gated clocks, etc.) is amortized across fewer sent bits. This trend is more prominent in the photonic clos as opposed to the electrical clos due to a significantly higher NDD power stemming from the need to perform ring thermal tuning and to power the laser. As a result, the electrical clos becomes energy-optimal at low utilizations (Figure 5b). The photonic clos presents smaller data-dependent (DD) switching costs, however, and thus performs more efficiently at high utilization (Figure 5d).

Comparing 45 nm and 11 nm, it is apparent that both photonic and electrical clos networks benefit significantly from electrical scaling, as routers and logic become cheaper. Though wiring capacitance scales slowly with technology, link energies still scale due to a smaller supply voltage at 11 nm (0.6 V). Laser and thermal tuning cost, however, scale marginally, if at all, allowing the electrical clos implementation to benefit more. In the 11 nm scenario, the electrical clos is more efficient up to roughly half network of the saturation throughput. As networks are provisioned to not operate at high throughputs where contention delays are significant, energy efficiency at lower utilizations is critical.

B. Photonics Parameter Scaling

For photonics to remain competitive with electrical alternatives at the 11 nm node and beyond, photonic links must sim-




Fig. 5: Comparison of network energy per bit vs network throughput (a) and the energy per bit breakdown at various throughputs (b–d) for the electrical crosstalk (EClos) and photonic crosstalk (PClos) at both the 45nm and 11nm technology nodes. Utilization is plotted up to the point where the network saturates (defined as when the latency reaches $3 \times$ the zero-load latency).

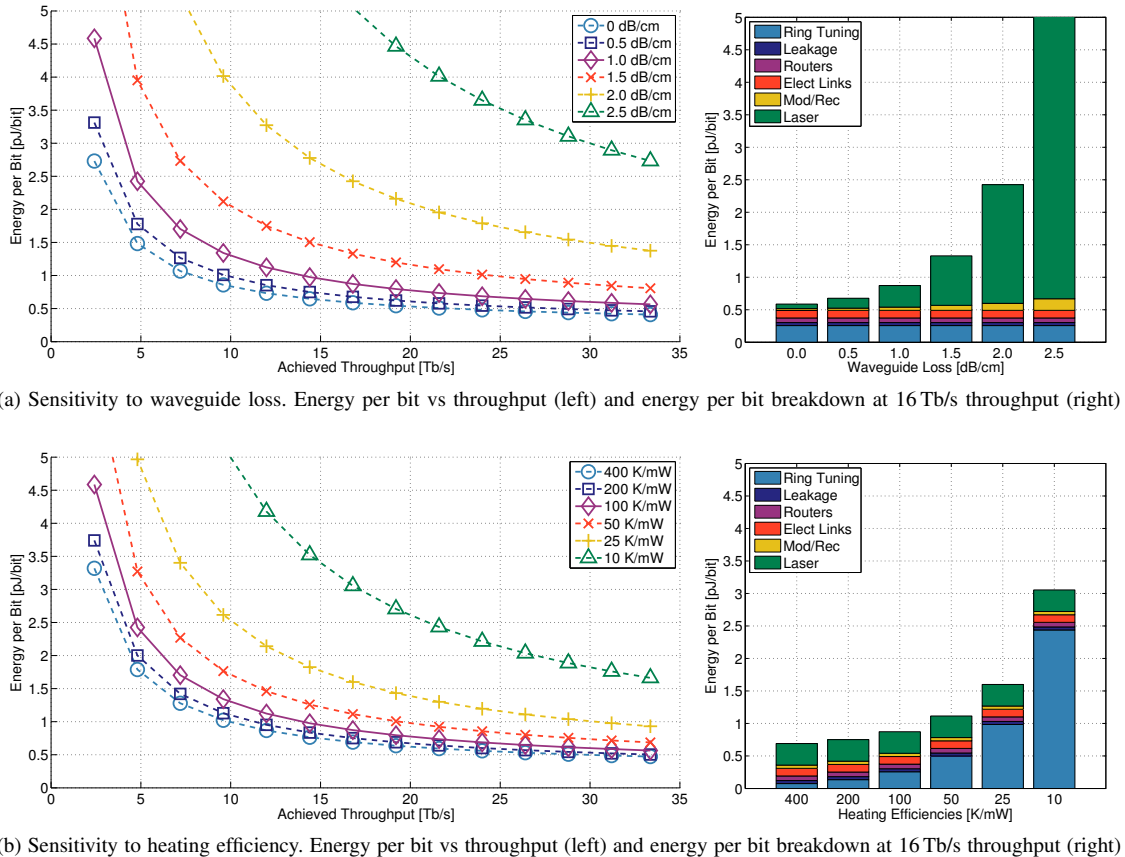


Fig. 6: Sensitivity of the photonic crosstalk network to a few select technology parameters. All plots assume the 11 nm electrical technology model.

ilarly scale. The non-data-dependent laser and tuning power as particularly problematic, as they are consumed even when links are used sporadically.

In Figure 6, we evaluate the sensitivity of the photonic crosstalk to waveguide loss and ring heating efficiencies, which affect laser and tuning costs. We see that our initial loss assumption of 1 dB/cm brings the photonic crosstalk quite close to the ideal (0 dB/cm) and the network could tolerate up to around 1.5 dB/cm before laser power grows out of proportion. Ring tuning power will also fall with better heating efficiency. However, it is not clear whether a 400 K/mW efficiency is physically realizable and it is necessary to consider potential

alternatives.

C. Thermal Tuning and Data-Rate

Per wavelength data-rate of an optical link is a particularly interesting degree of freedom that network designers have control over. Given a fixed bandwidth that the link is responsible for, an increase in data-rate per wavelength means a decrease in the number of WDM wavelengths required to support the throughput. This affects the number of ring resonators and, as such, can impact the tuning power.

Under the more conservative full-thermal (no bit-resuffling) tuning scenario (Figure 7a), the energy spent

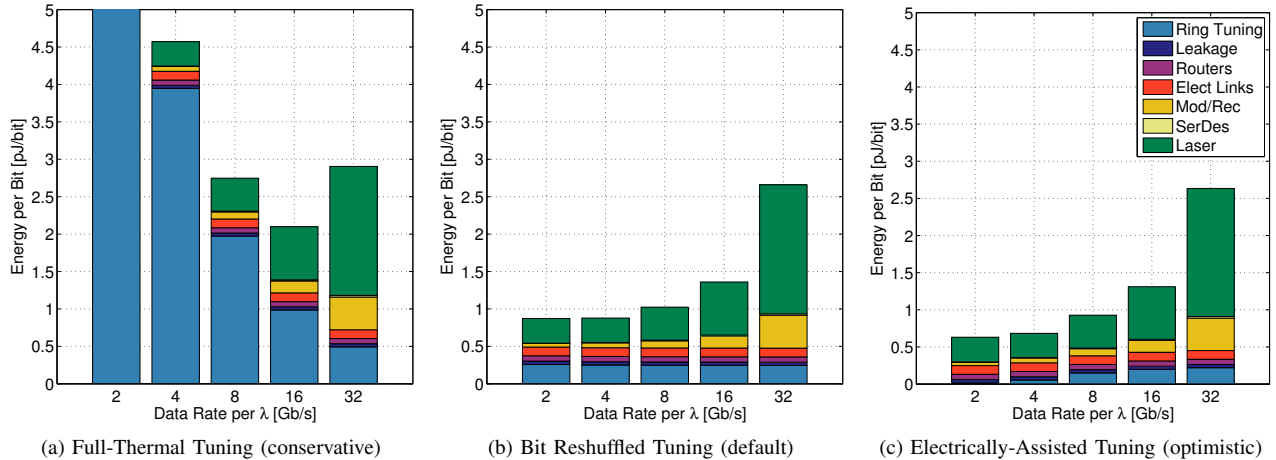


Fig. 7: A comparison of three thermal-tuning strategies discussed in Section V-C. Link data-rate is used as a degree of freedom to balance tuning power with other sources of power consumption. Since the throughput of each link is 128 bits/core-cycle at a 2GHz core clock, a data-rate of 2, 4, 8, 16, 32Gb/s per wavelength (λ) implies 128, 64, 32, 16, 8 wavelengths per link, respectively. All energy breakdowns are shown for half of saturation throughput (16.5 Tb/s).

on ring heating is dominant and will scale proportionally with the number of WDM channels (and thus inversely with per wavelength data-rate). Modulator and receiver energies, however, grow with data-rate as a result of more aggressive circuits. Laser energy cost per bit grows with data-rates due to a relaxation of modulator insertion loss/extinction ratios as well as clock uncertainty becoming a larger fraction of the receiver evaluation time. Routers and electrical links remain the same, though a small fraction of energy is consumed for serialization/deserialization (SerDes) at the optical link interface. These trends result in an optimal data-rate between 8-16 Gb/s, where ring tuning power is balanced with other sources of energy consumption, given the full-thermal tuning scenario.

This trend is no longer true once bit-reshuffling (the default scenario we assumed for Sections VII-A and VII-B) is considered, shown in Figure 7b. Following the discussion in V-C, a bit-reshuffler gives rings freedom in the channels they are allowed to tune to. At higher data-rates, there are fewer WDM channels and hence rings that require tuning. However, the channel-to-channel separation (in wavelength) is also greater. Given the presence of random process variations, sparser channels means each ring requires, on average, more heating in order to align its resonance to a channel. These two effects cancel each other out. Since the bit-reshuffler logic itself consumes very little power at the 11 nm node, ring tuning costs are small and remain relatively flat with data-rate.

If electrical-assistance is used (Figure 7c), tuning power favors high WDM channel counts (low data-rates). This is a consequence of the limited resonance shift range that carrier-depletion-based electrical tuners can achieve. At high WDM channel counts where channel spacing is dense, rings can align themselves to a channel by electrically biasing the depletion-based tuner without a need to power up expensive heaters. By contrast, when channels are sparse, ring resonances will often have to be moved a distance too far for the depletion

tuner to cover and costly heaters must be used to bridge the distance. As such, the lowest data-rate, 2 Gb/s per wavelength, is optimal under this scenario. A well-designed electrically-assisted tuning system could completely eliminate non-data-dependent tuning power. Hence, it is a promising alternative to aggressive optimization of ring heating efficiencies.

VIII. CONCLUSION

Integrated photonic interconnects is an attractive interconnect technology for future manycore architectures. Though it promises significant advantages over electrical technology, evaluation of photonics in existing proposals have relied upon significant simplifications. To bring additional insight into the dynamic behavior of these active components, we developed a new tool – DSENT – to capture the interactions between photonics and electronics. By introducing standard-cell-based electrical models and interface circuit models, we complete the connection between photonic devices and the rest of the optoelectrical network. Using our tool, we show that the energy-efficiency of a photonic NoC is poor at lower utilizations due to non-data-dependent laser and tuning power. These two components do not scale with electrical process technology and, in the case of thermal tuning, limited in photonics scaling potential. Using DSENT’s tuning models, we show that an electrically-assisted tuning scheme can eliminate non-data-dependent ring heating power for an NoC, significantly lowering the overhead of photonics and improve network energy efficiency. We will be releasing DSENT open-source [42].

ACKNOWLEDGMENT

The authors would like to thank the Integrated Photonics teams at both University of Colorado, Boulder and MIT. This work was supported in part by DARPA, NSF, FCRP, MARCO IFC, SMART LEES, Trusted Foundry, Intel, APIC, MIT CICS, and NSERC.

REFERENCES

- [1] S. Beamer *et al.*, “Re-architecting DRAM memory systems with monolithically integrated silicon photonics,” in *ISCA*, 2010.
- [2] G. Kurian *et al.*, “ATAC: A 1000-core cache-coherent processor with on-chip optical network,” in *PACT*, 2010.
- [3] Y. Pan *et al.*, “Firefly: Illuminating on-chip networks with nanophotonics,” *ISCA*, 2009.
- [4] D. Vantrease *et al.*, “Corona: System implications of emerging nanophotonic technology,” in *ISCA*, 2008.
- [5] A. Joshi *et al.*, “Silicon-photonics crosstalk networks for global on-chip communication,” in *NOCS*, 2009.
- [6] N. Agarwal *et al.*, “GARNET: A detailed on-chip network model inside a full-system simulator,” in *ISPASS*, 2009, pp. 33–42.
- [7] J. E. Miller *et al.*, “Graphite: A distributed parallel simulator for multicores,” in *HPCA*, 2010.
- [8] G. Kurian *et al.*, “Cross-layer energy and performance evaluation of a nanophotonic manycore processor system using real application workloads,” in *IPDPS*, 2012.
- [9] D. Taillaert *et al.*, “Compact efficient broadband grating coupler for silicon-on-insulator waveguides,” *OL*, vol. 29, no. 23, pp. 2749–2751, 2004.
- [10] J. Liu *et al.*, “Ge-on-si laser operating at room temperature,” *OL*, vol. 35, no. 5, pp. 679–681, 2010.
- [11] “Intel hybrid silicon laser,” Website, <http://techresearch.intel.com/ProjectDetails.aspx?Id=149>.
- [12] M. Georgas *et al.*, “Addressing link-level design tradeoffs for integrated photonic interconnects,” in *CICC*, 2011.
- [13] C. Pollock and M. Lipson, *Integrated Optics*. Springer, 2003.
- [14] M. Georgas *et al.*, “A monolithically-integrated optical receiver in standard 45-nm soi,” in *ESSCIRC*, 2011.
- [15] M. J. Cianchetti *et al.*, “Phastlane: a rapid transit optical routing network,” in *ISCA*, 2009.
- [16] G. Hendry *et al.*, “Circuit-switched memory access in photonic interconnection networks for high-performance embedded computing,” in *ICS*, 2010.
- [17] A. A. Chien, “A cost and speed model for k-any n-cube wormhole routers,” in *HOTI*, 1993.
- [18] L.-S. Peh and W. J. Dally, “A delay model and speculative architecture for pipelined routers,” in *HPCA*, 2001.
- [19] H. Wang *et al.*, “Orion: A power-performance simulator for interconnection networks,” in *MICRO*, 2002.
- [20] N. Banerjee *et al.*, “A power and performance model for network-on-chip architectures,” in *DATE*, 2004.
- [21] J. Balfour and W. J. Dally, “Design tradeoffs for tiled CMP on-chip networks,” in *ICS*, 2006.
- [22] A. Kahng *et al.*, “ORION 2.0: A fast and accurate NoC power and area model for early-stage design space exploration,” in *DATE*, 2009.
- [23] J. Chan *et al.*, “PhoenixSim: a simulator for physical-layer analysis of chip-scale photonic interconnection networks,” in *DATE*, 2010.
- [24] M. H. Na *et al.*, “The effective drive current in CMOS inverters,” in *IEDM*, 2002.
- [25] D. C. Pham *et al.*, “Overview of the architecture, circuit design, and physical implementation of a first-generation cell processor,” *JSSC*, vol. 41, no. 1, pp. 179–196, 2006.
- [26] A. Khakifirooz and D. Antoniadis, “MOSFET performance scaling - part II: Future directions,” *ITED*, vol. 55, no. 6, pp. 1401–1408, 2008.
- [27] A. Khakifirooz *et al.*, “A simple semiempirical short-channel MOSFET current-voltage model continuous across all regions of operation and employing only physical parameters,” *ITED*, vol. 56, no. 8, pp. 1674–1680, 2009.
- [28] L. Wei *et al.*, “Parasitic capacitances: Analytical models and impact on circuit-level performance,” *ITED*, vol. 58, no. 5, pp. 1361–1370, 2011.
- [29] “NCSU FreePDK45,” Online Website, <http://www.eda.ncsu.edu/wiki/FreePDK>.
- [30] R. Gupta *et al.*, “The elmore delay as a bound for rc trees with generalized input signals,” *IEEE TCAD*, vol. 16, no. 1, pp. 95–104, 1997.
- [31] J. M. Rabaey *et al.*, *Digital Integrated Circuits: A Design Perspective, second edition*. Prentice Hall, 2003.
- [32] “CACTI6.5,” Online Website, <http://www.hpl.hp.com/research/cacti>.
- [33] S. Li *et al.*, “McPAT: An integrated power, area, and timing modeling framework for multicore and manycore architectures,” in *MICRO*, 2009.
- [34] R. Marculescu *et al.*, “Probabilistic modeling of dependencies during switching activity analysis,” *IEEE TCAD*, vol. 17, no. 2, pp. 73–83, 1998.
- [35] H. Hatamkhani *et al.*, “A 10-mW 3.6-Gbps I/O transmitter,” in *VLSIC*, 2003.
- [36] J. S. Orcutt *et al.*, “Nanophotonic integration in state-of-the-art CMOS foundries,” *OpEx*, vol. 19, no. 3, pp. 2335–2346, 2011.
- [37] S. Selvaraja *et al.*, “Fabrication of uniform photonic devices using 193nm optical lithography in silicon-on-insulator,” *ECIO*, 2008.
- [38] C. Nitta *et al.*, “Addressing system-level trimming issues in on-chip nanophotonic networks,” in *HPCA*, 2011.
- [39] P. Dong *et al.*, “High speed silicon microring modulator based on carrier depletion,” in *Natl. Fiber Optic Engineers Conf.*, 2010.
- [40] N. Binkert *et al.*, “The GEM5 simulator,” *CAN*, vol. 39, pp. 1–7, 2011.
- [41] ———, “The role of optics in future high radix switch design,” *ISCA*, 2011.
- [42] “Download link,” Website, <http://www.rle.mit.edu/isg/technology.htm>.