Energy-Proportional Photonic Interconnects

YIGIT DEMIR, Intel

NIKOS HARDAVELLAS, Northwestern University

54

Photonic interconnects have emerged as the prime candidate technology for efficient networks on chip at future process nodes. However, the high optical loss of many nanophotonic components coupled with the low efficiency of current laser sources results in exceedingly high total power requirements for the laser. As optical interconnects stay on even during periods of system inactivity, most of this power is wasted, which has prompted research on laser gating. Unfortunately, prior work has been complicated by the long laser turn-on delays and has failed to deliver the full savings. In this article, we propose ProLaser, a laser control mechanism that monitors the requests sent on the interconnect, the cache, and the coherence directory to detect highly correlated events and turn on proactively the lasers of a photonic interconnect. While ProLaser requires fast lasers with a turn-on delay of a few nanoseconds, a technology that is still experimental, several types of such lasers that are suitable for power gating have already been manufactured over the last decade. Overall, ProLaser saves 42% to 85% of the laser power, outperforms the current state of the art by $2\times$ on average, and closely tracks (within 2%–6%) a perfect prediction scheme with full knowledge of future interconnect requests. Moreover, the power savings of ProLaser allow the cores to exploit a higher-power budget and run faster, achieving speedups of 1.5 to $1.7\times$ ($1.6\times$ on average).

CCS Concepts: • **Hardware** → **Photonic and optical interconnect**; **Emerging optical and photonic technologies**; *Network on chip*; *Interconnect power issues*

Additional Key Words and Phrases: Nanophotonic interconnects, laser control, power efficiency, energy efficiency, energy proportionality

ACM Reference Format:

Yigit Demir and Nikos Hardavellas. 2016. Energy-proportional photonic interconnects. *ACM Trans. Archit. Code Optim.* 13, 4, Article 54 (December 2016), 26 pages.
DOI: <http://dx.doi.org/10.1145/3018110>

1. INTRODUCTION

The global on-chip communication among the processor cores and between cores and other on-chip components (e.g., caches, memory controllers) consumes an increasing fraction of the total chip power budget and imposes delays that impact the overall performance of a multicore system. Based on the evolutionary ITRS [ESIA et al. 2012] roadmap, and barring any revolutionary developments (e.g., interconnects based on carbon nanotubes), electrical interconnects seem unlikely to keep up with the combined requirements on bandwidth density, energy efficiency, and latency [Heck and Bowers 2014], indicating the need to revisit the on-chip interconnect architecture. The prime candidate technology today to realize an efficient network on chip (NoC) at future process nodes is silicon photonics. Silicon photonics can be integrated alongside CMOS

New Article, Not an Extension of a Conference Paper.

This work is supported by the National Science Foundation under CAREER award CCF-1453853.

Authors' addresses: Y. Demir, Intel, 2501 NW 229th Ave, Hillsboro, OR 97124; email: yigit.demir@intel.com; N. Hardavellas, Northwestern University, 2145 Sheridan Rd, Evanston, IL 60208; email: nikos@northwestern.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2016 ACM 1544-3566/2016/12-ART54 \$15.00

DOI: <http://dx.doi.org/10.1145/3018110>

logic by adding a few new steps in the manufacturing process [Chen et al. 2005]. While it is an ongoing discussion where such an optical NoC will access the process flow, front-end-of-line and back-end-of-line monolithic implementations [Orcutt et al. 2011] as well as heterogeneous implementations using 3D-chip stacking [Ahn et al. 2009] are all viable options. Already silicon photonics are shown to outperform electrical RC and transmission-line links [Heck and Bowers 2014; Li et al. 2014; Nitta et al. 2013; Krishnamoorthy et al. 2009; Pan et al. 2009; Pan et al. 2010; Vantrease et al. 2008], and even conservative studies of NoCs for high-performance applications argue that photonics will soon be required [Stucchi et al. 2013].

While much work has been done on the architecture, design, and analysis of optical interconnects, the laser is generally excluded from the analysis and assumed to be lumped into a static off-chip power overhead that we should not be concerned much about. Unfortunately, the wall-plug laser consumption constitutes a major component of the overall power consumption of an optical interconnect. Even if the laser is off chip and its power consumption does not impact the chip power budget, the laser remains a major component of the system's power and energy consumption.

Lasers consume a significant amount of power because their output optical power needs to be high enough to compensate for the optical loss of silicon waveguides, optical couplers, and on-ring resonators. Typical silicon waveguides exhibit optical loss between 0.1 and 0.3dB/cm [Cardenas et al. 2009], resulting in modest optical loss over short distances. However, replacing global wires with waveguides that traverse the entire chip in a serpentine form can drastically increase the required laser power. Firefly [Pan et al. 2009] on a 580mm² chip needs a 16cm waveguide, which increases the laser power by 1.5 to 3 $\times$. Aggressive technology can produce low-loss waveguides (0.05dB/cm [Krishnamoorthy et al. 2009]), which allow the routing of long optical channels, but these low-loss waveguides are much wider than conventional ones [Krishnamoorthy et al. 2009; Li et al. 2012]. Their high area occupancy may force the use of narrow data paths (e.g., 2-bit links for an 8 $\times$ 8 array in Oracle MacroChip [Krishnamoorthy et al. 2009]), which in turn impose significant serialization delays that degrade performance and ultimately increase power consumption.

To make matters worse, WDM-compatible lasers are highly inefficient, with typical efficiencies in the range of 5% to 8% [Tanaka et al. 2012] and up to 10% [Zilkie et al. 2012]. Thus, the wall-plug laser power requirement is 10 to 20 $\times$ higher than the required laser output power. Manufacturing variations (e.g., coupler misalignment, waveguide edge roughness due to process variations) impose additional losses, forcing designers to increase the laser power even higher to maintain a safety margin. Sharing the optical path with other senders or receivers, typically done to keep the hardware costs manageable, may also increase the laser power, as sharing requires additional components that accumulate optical loss. While some topologies strike a better balance between power and performance [Pan et al. 2009], most of these costs are hard to avoid. As all of these factors are multiplicative, the wall-plug laser power can easily grow by over an order of magnitude when all inefficiencies are factored in.

Unfortunately, the majority of this power is typically wasted. In real workloads, it is often the case that the interconnect stays idle for long periods of time: compute-intensive execution phases underutilize the interconnect (common in many scientific applications), and servers in the cloud often stay idle or exhibit load imbalances (Google-scale datacenters are typically less than 30% utilized [Barroso and Holzle 2007]). Figure 1 shows that 50% of the messages are more than 15 cycles apart from the previous message in a wide range of compute- and memory-intensive applications. While the full laser power is required to support periods of high interconnect activity, the laser is wasted during idle times because photonic interconnects are always on. In a typical setting, light is continuously injected into the waveguides even if no packets

Fig. 1. CDF of message interarrival times.

are sent. In contrast, electrical interconnects stay idle, consuming only a small amount of leakage power, until a packet attempts to traverse them.

Motivated by these observations, recent work proposes laser gating (turning the laser off during idle periods to save power and turning it on during periods of high activity to meet the high bandwidth demand) as an effective technique to conserve laser power [Demir and Hardavellas 2014a, 2014b; Heck and Bowers 2014]. However, turning a laser on can incur significant delays, which degrades the performance of the system and in turn increases the power consumption. To obtain the benefits promised by laser gating, predictive techniques are required to hide the laser turn-on delay. EcoLaser [Demir and Hardavellas 2014a] proposes an adaptive mechanism, which leaves the laser on longer than the time required for the current transmission, to allow other senders to opportunistically find the laser on and transmit without incurring a delay. EcoLaser was shown to outperform by a wide margin a traditional photonic interconnect without laser control. However, its design is complicated and still wastes significant laser power ($2\times$ on average).

We propose *Proactive Laser (ProLaser)*, a laser control mechanism that monitors the requests sent on the interconnect, the L2 cache, and the coherence directory, to predict with high accuracy when to turn on the lasers of a nanophotonic NoC. Pro-Laser is contingent on fast lasers with a turn-on delay of a few nanoseconds. This technology, while still not mainstream, is real and not speculative. Several types of fast lasers suitable for power gating have been manufactured by various research groups and industrial labs over the last decade [Paniccia and Bowers 2006; Fang et al. 2006; Liu et al. 2010; Kotelnikov et al. 2012; Camacho-Aguilera et al. 2012], and their turn-on delay has been measured on actual hardware. While the jury is still out, we believe that the prospect of nanophotonics revolutionizing interconnects, coupled with the significant savings of laser power gating and the emergence of fast lasers, justifies research in this area [Heck and Bowers 2014]. In particular, our contributions are:

- We perform a limited study for the benefits of controlling on-chip and off-chip laser sources over a wide range of turn-on delays.
- We propose ProLaser, a mechanism that (1) monitors the messages sent in an NoC and correlates them to cache coherence protocol events to predict future messages and turn on the laser proactively, (2) turns on the data portion of the NoC only when it predicts messages that carry data, and (3) employs a Bloom filter at the L2 cache slice of each node in the NoC to predict a cache hit or miss and provide the laser turn-on request sufficiently early to hide the entire laser turn-on delay.
- We evaluate the impact of ProLaser on the performance and energy of a multi-core running a range of synthetic and scientific workloads under realistic physical constraints. ProLaser saves 49% to 88% of the laser power, outperforms the current

state of the art by $2\times$ on average, and tracks within 2% to 6% a perfect prediction scheme.

- We show that the power savings of ProLaser allow for providing a higher-power budget to the cores, which enables them to run faster. ProLaser on R-SWMR crossbars (e.g., Firefly [Pan et al. 2009]) results in a 50% to 70% speedup (60% on average) and 35% to 52% lower energy consumption per instruction (40% on average).

2. BACKGROUND

2.1. Laser Primer

Previous works [Batten et al. 2009; Batten et al. 2012; Kirman et al. 2006; Kurian et al. 2012; Pan et al. 2009] typically use off-chip lasers due to their ease of deployment and high energy efficiency (up to 30% for Gaussian comb lasers [Duan et al. 2009]). However, recent work [Heck and Bowers 2014] shows that output spectrum power variations and laser-to-fiber and fiber-to-chip coupling losses add 7 to 8dB optical loss, and thus off-chip lasers are in reality only 6% efficient. In comparison, on-chip laser sources [Koch et al. 2013] attain wall-plug efficiencies up to 15%, while enabling wavelength-division multiplexing (WDM). WDM can be implemented by feeding a set of wavelengths generated by an array of single-wavelength lasers into an optical bus. On-chip lasers offer energy efficiency and easy packaging, but their wall-plug power consumption counts against the processor's overall power budget. In either case, the laser power consumption remains a considerable overhead, especially when accounting for realistic optical loss parameters and laser efficiencies, emphasizing the need for power gating the laser source. Power gating on-chip lasers can increase the energy efficiency of a photonic interconnect by up to $4\times$ [Heck and Bowers 2014].

Laser power gating has been overlooked due to the high turn-on latency ($0.1\mu\text{s}$ [Heck and Bowers 2014]) of the traditional comb lasers that are widely assumed in photonic interconnects [Batten et al. 2009; Batten et al. 2012; Kirman et al. 2006; Kurian et al. 2012; Pan et al. 2009]. Comb lasers use diffraction grating to form the optical cavity. Temperature affects the diffraction grating pitch and the active region's refractive index, which alter the diffraction grating's wavelength selection, and hence the laser's emission wavelength. Thus, when comb lasers turn on, they need time to reach a set temperature and lock at the designated wavelength. This high delay hampers power gating. In contrast, Fabry-Perot (FP) lasers use two discrete mirrors to form the optical cavity, and their emission wavelength depends not on temperature but on the n-type doping level and the strain applied during the cavity development. Thus, when they are turned on (pumped to the lasing threshold), they lase at the designated wavelength without requiring time for temperature stabilization/locking and hence are suitable for power gating.

ProLaser, and laser power gating in general, strongly depends on fast lasers. While such technology is still experimental, it is important to note that fast lasers with ns-scale turn-on times have been manufactured and their turn-on delay has been characterized on real hardware prototypes [Kotelnikov et al. 2012; Paniccia and Bowers 2006; Fang et al. 2006; Camacho-Aguilera et al. 2012; Liu et al. 2010] and is in agreement with theoretically derived results. To turn the laser on, a supply current is applied to the laser. When the carrier density exceeds the threshold density, laser oscillation starts and light output increases drastically (laser turn-on). The time it takes from the current injection to the laser turn-on is the "laser turn-on delay," which is governed by the carrier lifetime and is in the order of nanoseconds [Petermann 1988, pp. 80–82]. The turn-on delay of Fabry-Perot lasers is highly tunable by design parameters, and nanosecond or subnanosecond laser turn-on delays are both theoretically predicted [Hisham et al. 2012a, 2012b; Petermann 1988, pp. 83] and achievable in real

implementations [Camacho-Aguilera et al. 2012; Liu et al. 2010; Kotelnikov et al. 2012; Petermann 1988].

For example, InP-based diode FP lasers [Kotelnikov et al. 2012] have been manufactured and shown to emit light with a 2ns-long electrical pulse excitation (so the laser turn-on latency is at most 2ns). InP-lasers have high peak power, and their emission wavelength is tunable in a wide range and highly stable with temperature, which makes them WDM compatible. Moreover, InP-lasers can be integrated on Si [Paniccia and Bowers 2006; Fang et al. 2006], so they can be used as an on-chip laser source. Similarly, Ge-based FP on-chip lasers have been manufactured [Liu et al. 2010] and the turn-on delay of real hardware prototypes was measured at 1.5ns at most for both optically and electrically pumped implementations [Camacho-Aguilera et al. 2012; Kimerling 2013; Liu et al. 2010]. We directly verified this claim for both optically and electrically pumped Ge-lasers with the lead author of Liu et al. [2010] in personal communication. Besides their fast turn-on time, Ge-lasers [Camacho-Aguilera et al. 2012] are suitable for on-chip photonic interconnects because they can be built within a standard-width ($1\mu\text{m}$) waveguide and occupy only $7.68 \times 10^{-3}\text{mm}^2$ area per laser, operate at room temperature, and are WDM compatible as they exhibit a gain spectrum over 200nm [Camacho-Aguilera et al. 2012].

2.1.1. Discussion. We want to emphasize that ProLaser does not depend on a singular laser technology. Any fast WDM-compatible continuous-wave laser that can be integrated on chip is suitable for laser power gating, including the InP and Ge lasers we assume in this work [Kotelnikov et al. 2012; Paniccia and Bowers 2006; Fang et al. 2006; Camacho-Aguilera et al. 2012; Liu et al. 2010]. Thus, there already exist technologies today that match the requirements of ProLaser. It is ultimately up to the processor architect to decide on the tradeoff between using ProLaser with a fast turn-on laser technology versus using another laser technology, potentially coupled with a different laser energy-saving technique. While exploring this tradeoff is beyond the scope of this article, it is an interesting avenue for future research.

If more efficient lasers were invented, the impact of ProLaser would drop, as would the impact of any laser energy-saving technique. Because ProLaser is within 2% to 6% of a perfect scheme, there is limited opportunity left beyond ProLaser for other techniques to save energy by turning the laser on/off. Thus, newer techniques would have to rely on a different mechanism (e.g., adjust laser power to the minimum level sufficient for reliable transmission of each flit). The exploration of such techniques is an open research question, but it is also beyond the scope of this work.

It is important to offer the interested reader an additional perspective on laser turn-on times: VCSELs can turn on with sub-100ps delay [Petermann 1988] and can be directly modulated over 35GHz [Wolf et al. 2013]. However, VCSELs are unsuitable for on-chip applications with WDM because they emit significant heat, and their operating wavelengths are defined by the epitaxial growth [Heck and Bowers 2014], which challenges the implementation of a multiwavelength VCSEL array on chip. Moreover, it is hard to protect the integrity of messages with direct laser modulation due to chirping and the pattern effect [Petermann 1988].

2.2. Nanophotonic Interconnect Topologies

In Single-Writer-Multiple-Reader (SWMR) [Kirman et al. 2006] crossbars, each router has its own dedicated data channel that delivers messages to all other routers (Figure 2(a)). R-SWMR [Kurian et al. 2012; Pan et al. 2009] crossbars add a reservation channel to SWMR. A sender in R-SWMR first broadcasts on its reservation channel a flit with the receiver's ID (in Figure 2(a), router R1 broadcasts on RCH1 a flit with ID = 2). Upon receiving a reservation flit, the receiver (R2) turns on its demodulators

Fig. 2. On-chip and off-chip laser configurations.

to receive the message from the sender's data channel (CH1), which is now dedicated to transfer data from the sender to the receiver. Reservation channels are narrow because reservation flits only carry the receiver ID and message type information. However, the laser power required to broadcast increases exponentially with the number of readers, making it impractical to broadcast at high-radix crossbars (e.g., radix-64). Instead of having a single broadcast link with many readers, slicing [Batten et al. 2012] spreads the readers across multiple waveguides and enables high-radix R-SWMR crossbars.

3. LASER CONTROL SCHEMES

A brief description the laser control schemes we model that can serve as reference to the reader appears in Table I. The table also includes a brief description of the no-control (No-Ctrl) and the power-equivalent (Power_Eq) networks for completeness, even though there is no laser control applied to these designs. We consider both on-chip and off-chip laser sources. Because the interconnects we model are based on R-SWMR buses, each sender employs its own private sets of lasers in both cases. When on-chip lasers are used, the lasers are placed next to each sender (Figure 2(a)). When off-chip lasers are used, each sender still has its own private sets of lasers that couple onto the chip through optical fibers. Each sender controls its lasers by employing an additional optical link that couples from the chip to the laser sources through an additional optical fiber (Figure 2(c)). This additional path increases the delay to control the lasers and is included in the evaluation. Section 3.6 describes off-chip laser control in detail, and Section 5.2 presents its hardware cost.

3.1. EcoLaser Control

Laser control schemes aim to save laser energy by turning the lasers off whenever the optical bus is idle. The energy savings come at the cost of a potential increase in

Table I. Brief Description of Modeled Laser Control Schemes

Name	Description
No-Ctrl	Baseline scheme; the lasers are not controlled and they are always on.
Power_Eq	An optical NoC with no laser control, similar to No-Ctrl. However, the network's width is scaled down to approximate the average power consumption of ProLaser.
EcoLaser	Proposed in Demir and Hardavellas [2014a]. It keeps the laser on for a few more cycles (adapting at runtime by monitoring the traffic) after the last transmission to allow senders to transmit opportunistically without incurring a laser turn-on delay.
Simple	EcoLaser with independent laser control for the data and the control bits of the optical bus. The data segment of the bus is turned on only when data are transmitted.
ATAC+	Proposed in Kurian et al. [2012]. It turns off the lasers when idle and turns them on with enough power for one receiver upon a transmit request (unicast mode).
ProLaser	The proposed scheme. It monitors coherence events and cache hits using Bloom filters to predict when communication is imminent and turns on the lasers just in time. The control and data portions of the optical bus are independently controlled.
Perfect	Oracle scheme with full knowledge of future interconnect requests. The lasers turn on and stabilize exactly when a packet is ready for transmission and turn off immediately afterward, unless the next request arrives before the laser can turn off and on.
xx-OffChip	Schemes with the suffix "-OffChip" refer to the corresponding laser control scheme for off-chip lasers. Schemes without a suffix imply their application to on-chip lasers.

message latency, because messages may have to wait for the laser to turn on. EcoLaser [Demir and Hardavellas 2014a] turns off the laser after a message transmission to save power, but it does not turn it off immediately; rather, it keeps the laser on for a few more cycles to allow subsequent senders to find the laser on and opportunistically transmit without incurring a delay. EcoLaser monitors the NoC traffic and dynamically adjusts the length of time the laser stays on (higher traffic benefits from longer laser-on times). While EcoLaser saves 50% of the laser power on average, it still wastes laser power because it always turns the entire optical bus on, even when sending small dataless coherence messages. A natural extension to EcoLaser would be the ability to control independently the data- and control-only portions of the optical bus.

3.2. Simple Laser Control

To provide high bandwidth, photonic links offer wide buses, which can send a data message in one cycle. In a typical configuration, a data message is 600 bits wide and contains a 64-byte cache block, 64-bit address, 20-bit ID, and 4-bit message type. Messages carrying data are transmitted in one cycle on a 300-bit (or 300-wavelength) optical bus [Joshi et al. 2009; Pan et al. 2009; Vantrease et al. 2008], because the optical links transmit at both edges of the clock. Small coherence messages are transmitted in two 44-bit dataless flits, which means that 256 bits of this optical bus remain idle. Thus, 44 bits of the optical bus are activated for all messages (common bits), and the remaining 256 bits are used only for data. As the data bits are not always used, a laser gating scheme would benefit from controlling them separately from the control bits. Figure 2(b) illustrates the separation of the data bus into two independent sections: control bits and data-only bits.

To evaluate the impact of segregating the optical bus into control and data sections, we implement a simple scheme (Simple). It extends EcoLaser with the ability to control the two sections independently and activates the data portion of the bus only when sending data. Simple requires an additional set of lasers for independently controlling the data-only bus. This does not increase the total laser power, as the optical link loss and the total number of wavelengths remain the same. On the downside, the Simple scheme keeps the data-only portion of the bus switched off most of the time, which

lowers the data messages' likelihood of finding the data bus on. Thus, data messages may suffer from higher message latencies, which degrades performance. The hardware overhead is detailed in Section 5.2.

3.3. ATAC+ Laser Control

The laser control scheme proposed in Kurian et al. [2012] turns off the lasers when idle and turns them on with enough power for one receiver when in unicast mode. The unicast mode of ATAC+ is identical in operation to R-SWMM [Pan et al. 2009], which is the optical bus technology assumed and modeled in this work. ATAC+ does not employ a predictive mechanism to turn lasers on/off, and hence may expose some of the laser turn-on delay. We do not model the architecture of the ATAC+ network in its entirety. We only model the laser control scheme of ATAC+ and employ it on the same optical networks we use for all other control schemes (a radix-16 R-SWMM and a topology with four optical R-SWMM crossbars similar to Firefly [Pan et al. 2009]).

3.4. Proactive Laser Control

Proactive Laser (ProLaser) predicts when a message transmission on the NoC is imminent and turns on the corresponding lasers just enough cycles early to hide the turn-on delay, and also predicts if the data portion of the optical bus is required. ProLaser makes these predictions both for the cache miss events that initiate a coherence transaction and for the intermediate messages generated by the cache coherence protocol (e.g., invalidations, forwarded requests, acks, etc.).

ProLaser accurately predicts the initial cache miss events that initiate a coherence transaction by employing a small Bloom filter [Safi et al. 2008] in front of the L2 cache slice of each node in the NoC. A request to the L2 cache is sent in parallel to the Bloom filter. A lookup in the Bloom filter takes one cycle when an L2 hit takes up to 14 cycles. A Bloom filter lookup consumes 1.09pJ at 16nm technology (overestimated). When the Bloom filter predicts an L2 miss, ProLaser does not turn on the whole bus, thereby avoiding energy waste, but turns on only the common bits. When the Bloom filter predicts an L2 hit, the whole data bus is turned on 1.5ns before the L2 hit latency, so that the data bus will be ready when the data is ready to be sent out. ProLaser implements a 1KB counting Bloom filter for each L2 cache slice, with 4K 2-bit counters and a hash function that takes the lowest 22 bits of the address above the block offset. We faithfully model the Bloom filter in our simulations. We observe false positives between 1.64% (barnes) and 1.98% (em3d). False positives do not degrade performance but waste laser energy. One can implement a ProLaser scheme that predicts an L2 miss using the result of the L2 tag lookup (10–11 cycles) rather than using a one-cycle Bloom filter lookup. However, in that case, the small time window (three to four cycles) between the L2 tag lookup and L2 hit latency is not sufficient to completely hide the laser turn-on delay. Thus, such a scheme would be susceptible to higher laser turn-on delays (Section 5.4). Also, ProLaser aims to keep the design simple by introducing modifications only on the interconnect, whereas a tag-lookup scheme requires modifying the cache. For these reasons, ProLaser employs Bloom filters and not an early tag check.

ProLaser accurately predicts future messages generated by the cache coherence protocol by monitoring the message types. For example, in a directory-based cache coherence protocol, every data message is generated upon receipt of a read, write, or directory-forwarded request. When a node receives a forwarded request (i.e., the node is the owner or a sharer), ProLaser turns on the whole data bus, anticipating that this node will send out a data reply. When a node receives a read or write request, a lookup in the local L2 cache slice's Bloom filter decides the type of the reply: an L2 miss generates another read request (to the tile with the memory controller), while an L2 hit generates a data reply. ProLaser turns the lasers on proactively for both of these

Table II. Bloom Filter False Positives

fmm	1.6%	barnes	1.64%	appbt	1.74%	em3d	1.98%
moldyn	1.61%	tomcatv	1.76%	ocean	1.91%	bodytrack	1.92%

messages, but it turns on the data-only portion only upon predicting an L2 hit. ProLaser also sends out acknowledgment messages quickly by turning on the control bits right after a reply or an invalidation. ProLaser avoids the formation of longer queues at the output buffers, which improves the throughput of the network. ProLaser does not predict L1 misses. However, these requests do not incur high latency overhead, as they only use the common bits, which are frequently active. Other noncritical messages, such as write-backs, are also not predicted by ProLaser, as they do not have a significant impact on the overall performance.

Figure 2(a) shows the microarchitecture of ProLaser. ProLaser adds a Bloom filter to the L2 cache slice and a laser controller, which makes predictions by monitoring the Bloom filter and the message types in the injection buffers. ProLaser keeps the lasers on until all of the messages queued in the injection buffers are transmitted. Unlike prior schemes, ProLaser is applicable to off-chip laser sources as well.

3.5. Perfect Laser Control

The perfect control scheme has complete knowledge of future interconnect accesses and establishes the limit of energy savings with the given laser technology. The perfect scheme saves the maximum laser energy without any performance overhead by turning the laser on sufficiently ahead of time, so the light reaches the sender at the exact time he or she attempts to transmit. Also, Perfect keeps the laser on when a message will need it in the immediate future, if the energy consumed by keeping the laser on is lower than the energy consumed by turning the laser off and on again. Similarly to ProLaser, Perfect controls the data-only portion of the bus independently.

3.6. Controlling an Off-Chip Laser Source

Off-chip lasers are less efficient than on-chip ones (Section 2.1), but their power consumption is not counted against the processor's power budget. In contrast, on-chip lasers do not suffer from additional coupling inefficiencies, but their power consumption reduces the power budget available to cores and caches, may cause overheating, and may degrade the overall performance. To achieve the best of both worlds, Heck and Bowers [2014] propose to implement an off-chip laser source using an array of single-wavelength lasers similar to the on-chip lasers assumed in this work. This arrangement still requires a fiber-to-chip coupling and higher packaging costs but avoids comb laser losses and thermal concerns. With such an implementation, a feedback signal from the laser control in the processor die can turn on/off the lasers in the array, thereby improving energy efficiency.

Capitalizing on this observation, we extend ProLaser to off-chip WDM laser arrays [Heck and Bowers 2014]. We send a laser turn-on signal to the off-chip laser by redirecting the laser signals back to the off-chip laser array (Figure 2(c)). The green and red wavelengths in Figure 2(c) are dedicated to the laser control and they are always on. When the node wants to turn on the laser, it redirects these wavelengths back to the off-chip source using the microrings at the cost of one cycle. Signaling the laser source takes 0.4ns (2cm waveguide plus 4cm fiber travel), that is, two cycles at 5GHz. Two wavelengths control the data-only and common bits separately. Upon receiving a signal, the laser turns on after the turn-on delay, the emitted light travels back to the node within two cycles, and the message is sent out. Using light to signal the off-chip laser source requires minimal additional hardware (a waveguide, a fiber, and a few

Table III. Architectural Parameters

CMP Size	64 cores, 480mm ²
Cores	ULTRASPARC III ISA, OoO, 4-wide dispatch/retirement, 96-entry ROB
Clock Freq.	1.6–3.8GHz for chips under realistic physical constraints (90°C max temperature). DVFS sets the core frequency. The average across all applications is 3GHz. Designs not subject to power and thermal constraints (Figure 4(a) only) assume a 5GHz clock.
L1-I/D Caches	64KB, 2-way, 64-byte block, 2-cycle hit, 2 ports, 32 MSHRs, 16-entry victim cache
L2 Cache	Shared, 512KB/core, 16 way, 64-byte block, 14 cycles, 32 MSHRs, 16-entry vic. cache
Mem. Control.	One per 4 cores, 1 channel per memory controller, round-robin page interleaving
Main Memory	Optically connected memory [Batten et al. 2009], 10ns access
Networks	Radix-16 R-SWMR and Firefly

microrings) and provides low-energy and low-latency signaling, which ProLaser hides. Our evaluation accounts for the power and energy cost of the additional optical link for the off-chip laser control, as well as all the associated latencies.

4. EXPERIMENTAL METHODOLOGY

4.1. Standalone Interconnect Modeling with Synthetic Traffic

To evaluate the performance and energy consumption of ProLaser in isolation from the interference of other system components or application characteristics, we employ a cycle-accurate network simulator based on Booksim 2.0 [Dally and Towles 2004], which models a radix-16 R-SWMR crossbar servicing random uniform traffic. As this is a standalone simulation of the network, no cores are simulated; rather, the simulator simply injects traffic into the interconnect ports at a predetermined rate to evaluate the network properties in isolation from the impact of other components or application characteristics. Thus, the clock frequency for the optical interconnect has no bearing on the clock frequency of the cores modeled in the rest of this article. We model a 5GHz clock frequency for the optical interconnect. Because the optical links transmit at both edges of the clock [Joshi et al. 2009; Pan et al. 2009; Vantrease et al. 2008], a 5GHz clock results in the optical interconnect running at 10GHz.

We model one-cycle routers with one-cycle E/O and O/E bit conversions on a 480mm² chip. The chip employs a 10cm waveguide with a five-cycle round-trip time at 5GHz. The link latency (one to five cycles) is calculated based on the traversed waveguide length. The buffers are 20 flits deep, with a flit size of 300 bits. We measure latency as the time required for the network to process a sample of injected packets. We evaluate the load latency and energy per flit of ProLaser, Simple, and EcoLaser with 1.5ns laser turn-on delay and compare them against a baseline with no laser control (No-Ctrl) and a perfect control scheme with full knowledge of future messages (Perfect).

4.2. Modeling a Multicore with Optical and Electrical NoC

To evaluate the impact of laser gating on a realistic multicore, we model a 64-core processor on a full-system cycle-accurate simulator based on Flexus 4.0 [Hardavellas et al. 2004; Wenisch et al. 2006] integrated with Booksim 2.0 [Dally and Towles 2004] and DRAMSim 2.0 [Rosenfeld et al. 2011]. Table III details the architectural modeling parameters. We model a shared physically distributed L2 cache and directories. The memory controllers are uniformly distributed on the chip, and they use the same physical interconnect with VCs to avoid deadlock. All messages below L1 traverse the interconnect. We calculate the power consumption of the electrical interconnect using DSENT [Sun et al. 2012]. We target a 16nm node and have updated our tool chain according to ITRS projections [ESIA et al. 2012]. The simulated system executes a selection of workloads from PARSEC, SPLASH-2, and other scientific applications.

All systems we model employ a throttling mechanism to keep the chip within a safe operational temperature below 90°C. We assume a DVFS (dynamic voltage and frequency scaling) mechanism that aggressively scales voltage and frequency to maximize performance while staying within 90°C. Thus, core frequencies stay within a reasonable 1.6 to 3.8GHz range for all applications (3–3.2GHz on average). To decouple the impact of ProLaser from the exact design of DVFS, we assume a very fine-grained scheme that can switch at 200MHz steps within that range. Multicore designs not subject to power and thermal constraints (Figure 4(a)) do not have a need for DVFS, as they are not thermally constrained. These hypothetical designs allow the cores to run at a maximum speed that is only limited by core design and implementation. In that case, which applies only to Figure 4(a), we model 5GHz cores.

We collect runtime statistics from full-system simulations and use them to calculate the power consumption of the system using McPAT [Li et al. 2009], and the power consumption of the optical networks using the analytical power model by Joshi et al. [2009]. We provide these power calculations to HotSpot 5.0 [Skadron et al. 2003] to estimate the temperature of the chip. The estimated temperature is then used to refine the leakage power estimate through McPAT. The new power calculation is fed back to HotSpot to refine the temperature profile, and we iterate this process until temperature stabilizes. We adjust DVFS based on the stable-state power and temperature estimates.

To put ProLaser's performance and energy consumption into perspective, we include in our evaluation a high-performance all-electrical on-chip interconnect: a 4-ary 4-flat flattened butterfly (Flat-Butterfly) [Kim et al. 2007] derived by combining a four-level butterfly network with concentration 4. For Flat-Butterfly, we model routers with 10 input and output ports and a three-cycle routing delay. Routers are connected through 88-bit bidirectional links (one-flit control, seven-flit data messages) with one-cycle local, two-cycle semilocal, and three-cycle global link delays. To show the range of ProLaser's impact, we evaluate its application on two optical network topologies that are at the opposite ends of the spectrum, a radix-16 R-SWMR, and a topology with four optical R-SWMR crossbars (Firefly [Pan et al. 2009]).

The radix-16 R-SWMR approximates a worst-case scenario for ProLaser. It has low power consumption (similar to Flat-Butterfly's) and its high concentration factor (4) creates heavier traffic. The low power consumption and heavy traffic limit ProLaser's opportunity. Firefly connects 16 local clusters (four routers each) with four R-SWMR crossbars. The local clusters use an electrical ring to route packets within the destination cluster, and each of the routers in a local cluster is connected to a different R-SWMR crossbar. The local electrical ring has 150-bit bidirectional links with one-cycle delay. Firefly has high laser power consumption and a low concentration (one), which results in light traffic, thus giving ample opportunity to ProLaser to conserve laser power.

Section 4.1 describes the modeling of the optical R-SWMR crossbars. We also contrast ProLaser to a power-equivalent optical interconnect with no laser control (Power_Eq). Power_Eq is similar to No-Ctrl, but its interconnect width is scaled down to approximate ProLaser's average energy savings. We compare the performance (user instructions per sec) and energy per instruction (EPI) of Flat-Butterfly, No-Ctrl, Power_Eq, EcoLaser, Simple (Section 3.2), ProLaser, and Perfect. As our design impacts only the on-chip energy consumption, we calculate EPI accounting for the energy consumed on the processor chip only, without accounting for the energy to access main memory.

4.3. Laser Power Modeling

The photonic devices we assume include on-chip or off-chip lasers, ring modulators to modulate the light (electrical-to-optical conversion), waveguides to route optical signals to their destination, and resonant demodulators to demodulate the optical signal

Table IV. Nanophotonic Parameters and Laser Power

Radix-16 R-SWMR	Per Unit	On-Chip Laser Total	Off-Chip Laser Total
DWDM		32	32
Splitter [Sun et al. 2015]	1dB	3dB	3dB
WG Loss [Epping et al. 2015]	0.4dB/cm	4dB	4dB
Nonlinearity [Sun et al. 2015]	1dB	1dB	3dB
Modulator Ins. [Sun et al. 2015]	3dB	3dB	5.12dB
Ring Through [Sun et al. 2013]	0.01dB	5.12dB	1.5dB
Filter Drop [Sun et al. 2013]	1.5dB	1.5dB	1.5dB
Coupler [Sun et al. 2015]	1.2dB		2.4dB
Total Loss		17.62dB	20.02dB
Detector [Masini et al. 2012]		−20dBm	−20dBm
Laser Power Per Wavelength		0.578mW	1.01mW
Total Laser Power	15% Efficiency	18.56W	32.36W

(optical-to-electrical conversion). By employing Dense Wavelength Division Multiplexing (DWDM), lasers of different wavelengths can be guided in the same waveguide without interfering with each other, which increases the bandwidth density. There are several waveguide technologies that offer low optical loss in the 0.272 to 0.6dB/cm range [Bogaerts and Selvaraja 2011; Cardenas et al. 2009; Epping et al. 2015; Selvaraja et al. 2010; Takei et al. 2014]. For ProLaser, we choose to employ silicon nitride (Si₃N₄) waveguides because they have high light confinement, offer low intrinsic optical loss in the C-band (0.4dB/cm), and can achieve superior reproducibility in a CMOS-compatible platform [Epping et al. 2015]. Similarly, DWDM-compatible modulators and demodulators using resonant rings have been manufactured and characterized in Sun et al. [2015] and can handle energy-efficient signal conversions at speeds higher than 10GHz.

To keep our design grounded to reality, we model devices with parameters that have been experimentally demonstrated in recent manufactured prototypes. Table IV presents the optical parameters for the nanophotonic devices we model along with the corresponding parameter sources. The modulation and demodulation energy is 317fJ/bit at 10GHz [Sun et al. 2015]. The laser power per wavelength and total laser power are calculated in Table IV using the analytical models introduced in Joshi et al. [2009] (we use the analytical models to provide the breakdown; DSENT [Sun et al. 2012] calculates similar results). The total laser power in Table IV is the wall-plug laser power and accounts for both the data and reservation channels, plus the laser efficiency of 15%.

4.4. Sensitivity to Optical Parameters

There is little consensus on the optical loss parameters used or projected in the literature. In some cases, parameters exhibit a variance of over 10× across publications. However, we observe that the design of an optical interconnect highly depends on the losses of the optical components used. For example, if the off-ring through loss on the radix-16 crossbar was 10× higher (i.e., 0.1dB), the interconnect would not employ 32-way DWDM, as this would increase the laser power to unsustainable levels. Rather, the interconnect would be optimized with a lower six-way DWDM and it would employ more waveguides, resulting in a total optical loss (and hence laser power) similar to the interconnect modeled in our work. In the extreme case where the off-ring loss were to increase by 10×, and on top of that the modulator insertion, drop loss, detection, and nonlinearity losses were to double, a four-way DWDM would accommodate the increased losses and keep the total laser power at a similar level. In either case, the fraction of laser energy that ProLaser saves depends on the network utilization, not on the optical loss parameters. Moreover, the higher the total optical loss, the more

power in absolute terms ProLaser would save, which would have a higher impact on the performance of the processor if this power is given back to the cores. Thus, in this work, we remain conservative in our estimates of optical losses.

Previous work [Kimerling 2013; Kurian et al. 2012] estimates that lasers suitable for power gating can be turned on and off within 1ns at the technology node we are targeting. We remain conservative and assume lasers with eight-cycle (1.5ns) laser turn-on delay, in line with existing manufactured lasers [Kotelnikov et al. 2012; Paniccia and Bowers 2006; Fang et al. 2006; Camacho-Aguilera et al. 2012; Liu et al. 2010]. We analyze the sensitivity of ProLaser on the laser turn-on delay in Section 5.4.

4.5. Resonant Ring Heater Modeling

To calculate the total ring heating power, we extend the method by Nitta et al. [2011] by additionally accounting for the heating of the photonic die by the operation of the cores. We model the thermal characteristics of a 3D-stacked architecture where the photonic die sits underneath the logic die. We use the 3D-chip extension of HotSpot [Skadron et al. 2003] to model the transient temperature changes in the optical die. After we execute a workload and collect transient temperature traces, we calculate the ring heating power required to maintain the entire photonic die at the constant microring trimming temperature during the entire execution. In addition, we account for the individual ring trimming power required to overcome process variations, as described in Joshi et al. [2009]. The ring trimming power is less significant when using smaller-radix crossbars.

5. EXPERIMENTAL RESULTS

5.1. Standalone Network Performance with Synthetic Traffic

Laser control saves energy by turning off the lasers whenever the data bus is idle. However, the energy savings come at the potential cost of increased message latency, because messages may have to wait for the laser to turn back on. We investigate the tradeoff between the laser energy savings and the network performance on a radix-16 R-SWMR using random traffic (Figure 3) and compare ProLaser against No-Ctrl, EcoLaser, Simple, and Perfect. We extend this evaluation for both on-chip and off-chip laser sources. Injection rate is defined as the number of flits per router per cycle that are injected into the network. Thus, an injection rate of 1 implies that every router sends a flit in the network on every clock cycle (100% network utilization). We model one-cycle routers with one-cycle E/O and O/E bit conversions on a 480mm² chip. The chip employs a 10cm waveguide with a five-cycle round-trip time at 5GHz. The link latency (one to five cycles) is calculated based on the traversed waveguide length. The buffers are 20 flits deep, with a flit size of 300 bits. We measure latency as the time required for the network to process a sample of injected packets. We evaluate the load latency and energy per flit of ProLaser, Simple, and EcoLaser with a laser turn-on delay of 1.5ns and compare them against a baseline with no laser control (No-Ctrl) and a perfect control scheme with full knowledge of future messages (Perfect).

Messages in EcoLaser exhibit a six-cycle average delay at low injection rates instead of eight cycles, which is the laser turn-on delay, because some of the messages find the laser active and transmit immediately. This overhead decreases slightly for higher injection rates as more messages find the laser active. The Simple scheme incurs a slightly higher message delay because data messages cannot catch the laser actively as the data-only portion of the bus is turned off more frequently. In comparison, ProLaser incurs only a one-cycle overhead at low injection rates, because it foresees the majority of the messages and activates the laser ahead of time. EcoLaser and Simple saturate

Fig. 3. Load latency (top) and energy per flit (bottom) for a radix-16 R-SWMM crossbar.

early, providing 12% and 20% lower throughput than ProLaser, respectively, because they are more susceptible to exposing the laser turn-on delay.

Controlling an off-chip laser requires sending control signals back to it, after which the light emitted by the laser travels to the sender, who sees this as additional delay. We estimate an overhead of two cycles each way (2cm waveguide plus 4cm fiber travel). While EcoLaser and Simple expose this extra overhead, ProLaser hides most of it and shows only slight performance degradation when controlling off-chip lasers.

EcoLaser's energy savings disappear as the injection rate grows, as it keeps the laser on most of the time. Simple and ProLaser continue to exhibit high energy savings at high injection rates because they can keep the data-only portion of the data bus inactive (Figure 3, bottom). ProLaser consumes lower energy per flit than Simple because it provides higher throughput with lower message latency. On average over injection rates, ProLaser consumes 34% lower laser energy than EcoLaser and is within 4% of Perfect.

When using an off-chip laser, the couplers to carry the light to the chip add 2.4dB loss and increase the laser power consumption by $1.74\times$. The additional coupling losses and control latency make EcoLaser, Simple, and ProLaser consume higher laser energy with an off-chip laser source than an on-chip one. On average, over injection rates, ProLaser consumes 35% lower laser energy per flit than EcoLaser and achieves 7% lower energy savings than Perfect.

5.2. The Hardware, Energy, and Performance Cost of Laser Control

Each R-SWMM bus is powered by two sets of 32 lasers. Each set has one laser per wavelength in 32-way DWDM. The lasers of each set are first muxed and then split into

five waveguides; thus, 64 lasers create a 300-bit optical bus spread over 10 waveguides. We choose this architecture instead of 32 lasers split into 10 waveguides as the latter requires an additional level of splitters and incurs 26% higher optical loss.

ProLaser introduces an additional set of 44 lasers to control the common bits of the bus independently. Thus, ProLaser for radix-16 employs 704 more lasers than a baseline radix-16 crossbar. These additional lasers occupy 5.4mm^2 [Liu et al. 2010], increasing the overall area of the photonic devices by 6.2%. Overall, ProLaser on radix-16 R-SWMR crossbars employs a total of 1,728 lasers, which occupy 13.25mm^2 (i.e., 2.8% of the chip area of a 480mm^2 chip, of which 1.1% of the chip area is ProLaser's overhead). The area overhead of ProLaser is very small and it can be easily accommodated, especially when the photonic devices are integrated in a separate die that is 3D-stacked with the logic die, as assumed in this work (Section 4.5).

In the case of Firefly, the additional lasers are 2,816 and occupy 21.6mm^2 . ProLaser in the Firefly topology requires a total of 6,912 lasers, which occupy a total of 53mm^2 in a 480mm^2 chip (11% of the chip area, of which 4.5% is ProLaser's overhead). The area overhead of ProLaser is again small enough to be practical and leaves enough area for the other nanophotonic components.

When an off-chip laser is used, 64 wavelengths are muxed into one optical fiber to couple onto the chip. With ProLaser for the off-chip source, each bus requires three optical fibers (one for the data bits, one for the common bits, and one for off-chip laser control, Figure 2(c)). Thus, ProLaser for radix-16 uses 48 fibers attached to the chip, and 192 fibers for Firefly; 192 fibers need 48mm at a generous $250\mu\text{m}$ coupling pitch (tapered couplers need just $25\mu\text{m}$), occupying only 54% of the chip circumference.

Controlling the off-chip lasers requires an additional optical link powered by an always-on laser. The additional link increases static laser power by 0.6%, which corresponds to only 0.07% overhead over the chip's peak power because this is only a 2-bit link. On top of that static power consumption, each time the link is used, the system goes through an E/O conversion at the L2 tile and an O/E conversion at the off-chip laser source. These operations consume a total of 634fJ. Compared to an L2 access, which consumes 78.35pJ, each laser control command imposes a dynamic energy overhead of 0.8% over an L2 cache access, which corresponds to only 0.009% overhead over the average dynamic energy consumption of the chip across workloads. Our evaluation faithfully models the additional link and all these effects. In addition, ProLaser introduces a 1KB counting Bloom filter at each L2 cache slice (Section 3.4). The L2 cache slice is 512KB; thus, the Bloom filters impose a 0.2% area overhead on the L2 cache (i.e., less than 0.1% area overhead over the total 480mm^2 of the chip area). The Bloom filter is a very efficient L-CBF [Safi et al. 2008] structure of 4K 2-bit entries and consumes 1.09pJ per lookup (overestimated, Section 3.4). For comparison, an access to the 512KB 16-way L2 cache slice consumes 78.35pJ. Thus, the dynamic energy overhead of a Bloom filter is 1.4% of the L2 dynamic energy. All the overheads of incorporating Bloom filters are faithfully modeled in our evaluation.

Laser gating trades off message latency for energy savings, and thus it is expected to achieve lower performance than No-Ctrl. However, the power saved by laser gating of on-chip lasers may reduce the thermal emergencies and the need for core throttling, and thus increase performance. We aim to analyze separately the effects of increasing the effective message latency and reducing the need for core throttling.

First, we analyze the performance cost of laser gating mechanisms by evaluating them on a multicore that is not subject to thermal constraints; thus, cores run at their maximum frequency (5GHz). Our benchmark suite includes both memory-intensive workloads that generate high traffic and are sensitive to interconnect latency (bodytrack, em3d, ocean, appbt, tomcatv) and compute-intensive workloads that have low injection rates and are less sensitive to message latency (fmm, moldyn, barnes).

Figure 4(a) summarizes our findings. The injection rate of each application appears below its name. Injection rate is the number of flits per router per cycle that are injected into the network. For reference, we also present the performance of a multicore with an electrical-only network (Flattened-Butterfly).

Simple exposes most of the laser turn-on delay, and thus it underperforms other designs on memory-intensive workloads. Compared to No-Ctrl, Simple saves 46%, ATAC+ 31%, and EcoLaser 30% of the on-chip laser energy on average, at the cost of 17%, 16%, and 10% slowdown, respectively. In contrast, ProLaser saves 61% (85% max) with only 1.7% slowdown. Moreover, the laser energy savings of ProLaser are within 3% to 4% of Perfect's when running real workloads. ProLaser achieves higher energy savings on real workloads than on synthetic random traffic because real workloads have bursty and sparse memory accesses (Figure 1 shows that 70% of the messages are more than eight cycles apart, providing ample laser gating opportunities).

Controlling off-chip lasers achieves similar energy savings, but at a higher performance penalty due to the higher turn-on delay. ProLaser-OffChip successfully hides this overhead and causes only 4.5% slowdown on average, while saving 58% of the laser energy. Power_Eq approximates ProLaser's energy savings by scaling down its width (100-bit flits instead of 300-bit flits). However, Power_Eq suffers from high serialization delays and underperforms ProLaser. Thus, saving laser energy by reducing the width of the interconnect is not a good alternative to laser control.

5.3. Impact of ProLaser on a Realistic Multicore

The wall-plug power consumption of on-chip lasers can be a major contributor to a multicore's power budget due to the lasers' low efficiency (15% [Koch et al. 2013]). Under realistic thermal and power constraints, DVFS scales voltage and frequency to keep the chip within 90°C. Thus, core frequencies stay at a reasonable 1.6 to 3.8GHz range for all applications (3–3.2GHz on average). DVFS in No-Ctrl throttles the cores frequently to keep them below 90°C. ProLaser, however, reduces the laser power and leads to a cooler chip, with less core throttling and higher performance.

The impact of ProLaser depends on the total laser power consumption of the photonic network. Off-chip lasers incur additional coupling losses, and thus consume $2.5\times$ higher power than on-chip lasers. The majority of this power is dissipated away from the multicore chip, decreasing ProLaser's potential savings. However, the power consumption of the off-chip lasers can be as high as the power budget of the multicore itself (due to high coupling losses and low efficiency), so the impact of ProLaser's energy savings on the total *system energy* remains significant.

We evaluate ProLaser through two opposing case studies: a radix-16 crossbar and a Firefly topology. The radix-16 crossbar approximates a worst-case scenario. It has low power consumption and its high concentration (4) creates heavy traffic. These two effects limit ProLaser's potential. Firefly [Pan et al. 2009] has high laser power consumption ($4\times$ higher than radix-16's) and a low concentration (1), which results in light traffic, thus giving ample opportunity to ProLaser to conserve laser power.

2.1.1. Case Study 1: Radix-16 R-SWMR. The wall-plug power consumption for ProLaser on a radix-16 crossbar with on-chip lasers is 18.56W. All laser gating schemes save a significant fraction of this power with minimal slowdown when running compute-intensive workloads, and thus they outperform No-Ctrl (Figure 4(b)). However, Simple, ATAC+, and EcoLaser underperform No-Ctrl on memory-intensive workloads due to their higher effective message latency and low laser energy savings. Power_Eq achieves similar savings as ProLaser but incurs serialization delays, while ProLaser achieves both high performance and high energy savings. ProLaser outperforms No-Ctrl by 5%,

Fig. 4. (a) Speedup over No-trl for radix-16 R-SWMMR on a hypothetical multicore without thermal constraints. The clock speed is 5GHz as the absence of physical constraints allow the system to reach maximum speed. (b) Speedup over Flat-Butterfly for radix-16 R-SWMMR on a multicore under realistic thermal constraints. Because the chip is under realistic power and thermal constraints, voltage and frequency scaling limit clock speeds. The average clock speed across all applications is 3GHz.

ATAC+ by 15%, and EcoLaser by 10.6% on average, and has 14% and 9.3% lower EPI than ATAC+ and EcoLaser, respectively (Figure 5).

ProLaser on a radix-16 crossbar with off-chip lasers consumes 32.4W, of which only 2.6W are dissipated on chip. All laser control schemes are slower than No-Ctrl, as they can no longer increase performance by reducing core throttling. However, ProLaser saves a significant portion of the laser energy, which reduces the total energy consumption of the chip. ProLaser-OffChip is only 2% slower than No-Ctrl-OffChip (Figure 4(b)) but has 12% lower EPI (Figure 5). In contrast, EcoLaser both is slower and consumes more energy than No-Ctrl for off-chip lasers. Thus, segregation of the data bus and proactive laser turn-on are essential for off-chip lasers. ProLaser outperforms ATAC+ by 19% and EcoLaser by 12%, and has 21% and 17% lower EPI, respectively. ProLaser's performance and EPI are always within 4% to 5% of Perfect's.

2.1.1. Case Study 2: Firefly. Firefly consists of four radix-16 R-SWMMR crossbars with slightly longer waveguides; thus, its power consumption is slightly more than $4\times$ that of a radix-16 R-SWMMR crossbar. The wall-plug power consumption for Firefly with on-chip lasers is 76.7W. All laser control schemes outperform No-Ctrl for all workloads (Figure 6(a)), because the laser energy savings are a considerable fraction of the chip's power budget. ProLaser outperforms No-Ctrl by 60% on average and has 39% lower EPI, and outperforms EcoLaser by 15.5% and has 12.4% lower EPI (Figure 6). ProLaser outperforms ATAC+ by 24% while providing 18% lower EPI. Firefly with No-Ctrl cannot deliver high performance on a realistic multicore as over half the chip's power budget is consumed by the interconnect. ProLaser enables Firefly to reach high performance by making the interconnect energy proportional.

Off-chip lasers consume 134.06W for Firefly, of which only 10.96W are dissipated on the multicore chip. On average, ProLaser-OffChip is slightly slower than No-Ctrl-OffChip, but has 44% lower EPI (52% max, Figure 6(b)). ProLaser outperforms ATAC+ by 23% and EcoLaser by 16%, and has 29% and 24% lower EPI, respectively. In all cases, ProLaser's performance and EPI are within 5% to 6% of Perfect's, which indicates that ProLaser is harvesting the majority of the possible laser energy savings. It is important to note that the off-chip lasers for Firefly consume 28% more power than the processor chip itself (No-Ctrl-OffChip, Figure 6(b)). ProLaser reduces the laser energy consumption by $5.8\times$ on average, allowing the lasers to use only 21% of the total processor energy, and thus rendering this topology practical.

5.4. Laser Turn-On Latency Tolerance

ProLaser foresees the arrival of messages and activates the lasers proactively; thus, it can tolerate higher laser turn-on delays. In contrast, EcoLaser lacks a proactive laser turn-on mechanism and is susceptible to high laser turn-on delays. Figure 7 reports the average message latency and the laser energy savings as a function of laser turn-on delay, under uniform traffic with an injection rate similar to the average injection rate of our benchmark suite (0.11 packets/router/cycle, Section 5.3). We model a 2.5GHz clock to reflect the lowest speed of the power-limited radix-16 R-SWMMR multicores (Section 5.3). We make this decision because we need to select a single clock frequency for this experiment, and we wanted to obtain conservative estimates. ProLaser is very accurate at predicting future interconnect requests, and thus it rarely exposes the laser turn-on delay, but the remaining schemes incur this penalty often. At a 3.8GHz clock, the penalty would be 50% more processor cycles than at the much lower 2.5GHz range. Thus, we conservatively run this experiment at the lowest clock speed to help all the competing schemes gain an edge against ProLaser.

Figure 7 shows that the network performance and the laser energy savings decrease with increasing laser turn-on delay, which emphasizes the need for fast lasers. Laser

Fig. 5. Energy per instruction for radix-16 R-SWMMR under realistic thermal constraints. The evaluated designs are from left to right: Flat-B. (F), No-Ctrl (N), Power_Eq (e), Simple (S), Ecolaser (E), ATAC+ (A), ProLaser (P), and Perfect (I), and their off-chip implementations. Because the chip is under realistic power and thermal constraints, voltage and frequency scaling limit clock speeds. The average clock speed across all applications is 3GHz.

Fig. 6. (a) Speedup over Flat-Butterfly for Firefly topology under realistic thermal constraints. Because the chip is under realistic power and thermal constraints, voltage and frequency scaling limit clock speeds. The average clock speed across all applications is 3.2GHz. (b) Energy per instruction for Firefly topology under realistic thermal constraints. The evaluated designs are from left to right: Flat-B (F), No-Ctrl (N), Power_Eq (e), Simple (S), Ecolaser (E), ATAC+ (A), ProLaser (P), and Perfect (I), and their off-chip implementations. Because the chip is under realistic power and thermal constraints, voltage and frequency scaling limit clock speeds. The average clock speed across all applications is 3.2GHz.

Fig. 7. Laser turn-on latency tolerance. The core speed is set to 2.5GHz.

control schemes become inefficient with high laser turn-on delays, because they slow down the system and end up consuming more laser energy to send messages. ATAC+ cannot tolerate a laser turn-on delay beyond 3ns, and its laser energy savings are lower than EcoLaser's. EcoLaser can tolerate up to only a 3ns laser turn-on delay, where it saves 18% of the laser energy and increases message latency by 50%, and becomes impractical beyond that. A ProLaser scheme without Bloom filters that relies on early L2-tag lookup has three to four cycles between the L2 tag lookup (laser turn-on) and the L2 hit, so it can tolerate higher laser turn-on delays than EcoLaser (up to 5ns). ProLaser with Bloom filters hides the increased laser turn-on delay even more, because the laser controller has 14 cycles to turn on the lasers between the Bloom filter lookup and L2 cache hit. ProLaser can tolerate up to a 7ns laser turn-on delay, where it still saves 23% of the laser energy compared to No-Ctrl.

ProLaser has limited impact for laser turn-on delays higher than 7ns and barely saves energy for delays longer than 9ns. The 0.11 packets/router/cycle injection rate for uniform traffic results in a nine-cycle average packet interarrival time. Thus, lasers with a turn-on delay longer than nine cycles have very little chance of saving power. In conclusion, early laser turn-on prediction with Bloom filters allows ProLaser to withstand $2.3\times$ higher turn-on delays than the state of the art (EcoLaser). This allows ProLaser with Bloom filters to remain an effective laser control scheme even under relatively high laser turn-on delays, when competing schemes fail.

6. RELATED WORK

Several NoCs exploit CMOS-compatible photonics for future multicore processors. Firefly [Pan et al. 2009] partitions optical R-SWMMR crossbars to connect clusters of electrical mesh networks. Batten et al. [2009, 2012] present energy-efficient and scalable implementations of R-SWMMR crossbars. All R-SWMMR-based NoCs can utilize ProLaser to lower laser energy while maintaining their performance.

Previous work has explored segregating the interconnect used for core communication from the interconnect used for communication with the cache [Huh et al. 2005; Lotfi-Kamran et al. 2012] to lower the network cost or to optimize for data placement and partitioning. However, such designs have not been proposed or evaluated in the context of photonic interconnects. ProLaser segregates the data portion of the photonic interconnect from the control portion and manages them separately, to maximize power savings without hurting performance.

EcoLaser [Demir and Hardavellas 2014a] keeps the laser on for a few cycles after the last transmission to increase the likelihood that a sender will opportunistically

transmit without incurring a laser turn-on delay. However, it does not directly predict upcoming traffic and exposes some of the laser turn-on delay, or adopt a conservative configuration that misses opportunities to save laser energy. Kurian et al. [2012] propose an optical R-SWMMR crossbar and electrical hybrid interconnection network, and improve performance by utilizing the coherence protocol and by turning lasers on/off. However, unlike ProLaser, the laser control scheme in Kurian et al. [2012] turns lasers on/off on demand without employing a predictive mechanism.

Thonnart et al. [2008] power down the unused units of an electrical interconnect to reduce static power. Zhou et al. [2013] control active splitters to tune channel bandwidth on a binary tree network and increase channel utilization, which leads to higher energy efficiency. Chen et al. [2013] distribute laser power across multiple busses in a multi-bus NoC based on the changes in the bandwidth demand. Flexishare [Pan et al. 2010] minimizes static power consumption by fully sharing a reduced number of channels across the network. Neel et al. [2015] employ runtime power management techniques to reduce the laser power by adapting the width of the network, that is, scaling the number of channels available for communication. Chen et al. [2015] dynamically activate/deactivate L2 cache banks and switch on/off the corresponding silicon-photonics links in the NoC. Nitta et al. [2012] show the energy inefficiency of photonic interconnects under low utilization and propose to improve efficiency by recapturing the energy of photons, which are not used for communication. ProLaser can complement these techniques and enhance laser energy savings.

7. THE POTENTIAL IMPACT OF PROLASER ON THE DATACENTER

Our benchmark suite covers a wide range of workload categories from scientific and parallel computation but does not include scale-out server workloads. Datacenter workloads exhibit bursty traffic behavior, which requires a high-bandwidth interconnect to maximize performance. At the same time, this burstiness leaves interconnect links idle for long periods. For example, links on a university datacenter network remain idle for 54% to 62% of the time [Demir and Hardavellas 2016]. This behavior provides ample opportunities for laser power gating. Thus, we believe that the trends we observe for ProLaser would still hold for datacenter workloads.

While some scale-out workloads may be challenging, we expect others to provide ample opportunity. Latency-sensitive workloads (e.g., OLTP) may suffer disproportionate performance penalties if the laser turn-on delay is exposed. ProLaser in that case may have to be conservative and sacrifice some of the energy savings to avoid exposing the laser turn-on delay. Other workloads (e.g., decision support systems and analytics) may be tolerant to the occasional delay as they can hide it through thread- and instruction-level parallelism, and may exhibit compute-heavy execution phases often enough to provide significant opportunities for turning off the lasers. As bursty behavior is prevalent in large-scale datacenter installations [Kandula et al. 2009], we believe that ProLaser will continue to provide significant laser energy savings.

There are many opportunities left for power savings from the optical part of the datacenter. Datacenters typically employ interconnects with path diversity. A laser control scheme could capitalize on that diversity to turn off enough of the datacenter network to minimize power consumption and still provide full connectivity and fail-safe redundancy, and grow or shrink the interconnect bandwidth following the shape of traffic [Demir and Hardavellas 2016]. Also, significant savings may be obtained by employing disintegrated processors (e.g., Galaxy [Demir et al. 2014b]). These architectures allow chip-to-chip communication entirely in the optical domain, without the need for multiple O/E/O conversions or metallic interconnects.

8. CONCLUSION

In this work, we propose ProLaser, a mechanism that (1) monitors the messages sent in an NoC and correlates them to cache coherence protocol events to predict future messages and turn on the laser proactively, (2) turns on the data portion of the NoC only when it predicts messages that carry data, and (3) employs a Bloom filter in front of the L2 cache slice of each node in the NoC to predict a cache hit or miss and provide the laser turn-on request sufficiently early to hide the entire laser turn-on delay. Our results indicate that ProLaser saves between 42% and 85% of the laser power, outperforms the current state of the art by $2\times$ on average, and closely tracks (within 2%–6%) a perfect prediction scheme with full knowledge of future interconnect requests. Moreover, the power savings of ProLaser allow the cores to exploit a higher power budget and run faster, achieving speedups of 50% to 70% (60% on average).

REFERENCES

- J. Ahn, M. Fiorentino, R. Beausoleil, N. Binkert, A. Davis, D. Fattal, N. Jouppi, M. McLaren, C. Santori, R. Schreiber, S. Spillane, D. Vantrease, and Q. Xu. 2009. Devices and architectures for photonic chip-scale integration. *Applied Physics A*, 95, 4, 989–997.
- L. A. Barroso and U. Holzle. 2007. The case for energy-proportional computing. *IEEE Computer* 40, 12, 33–37.
- 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. 2009. Building many-core processor-to-DRAM networks with monolithic CMOS silicon photonics. *IEEE Micro* 29, 4, 8–21.
- C. Batten, A. Joshi, V. Stojanovic, and K. Asanovic. 2012. Designing chip-level nanophotonic interconnection networks. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems* 2, 2, 137–153.
- W. Bogaerts and S. K. Selvaraja. 2011. Compact single-mode silicon hybrid rib/strip waveguide with adiabatic bends. *IEEE Photonics Journal* 3, 3, 422–432.
- R. E. Camacho-Aguilera, Y. Cai, N. Patel, J. T. Bessette, M. Romagnoli, L. C. Kimerling, and J. Michel. 2012. An electrically pumped germanium laser. *Optics Express* 20, 10, 11316–11320.
- J. Cardenas, C. Poitras, J. Robinson, K. Preston, L. Chen, and M. Lipson. 2009. Low loss etchless silicon photonic waveguides. *Optics Express* 17, 6, 4752–4757.
- C. Chen, J. L. Abellan, and A. Joshi. 2015. Managing laser power in silicon-photonic NoC through cache and NoC reconfiguration. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 34, 6, 972–985.
- C. Chen and A. Joshi. 2013. Runtime management of laser power in silicon-photonic multibus NoC architecture. *IEEE Journal of Selected Topics in Quantum Electronics* 19, 2, 3700713–3700713.
- G. Chen, H. Chen, M. Haurylau, N. Nelson, P. M. Fauchet, E. Friedman, and D. Albonesi. 2005. Predictions of CMOS compatible on-chip optical interconnect. In *Proceedings of the 7th International Workshop on System-Level Interconnect Prediction*. 13–20.
- W. J. Dally and T. Bowles. 2004. *Principles and Practices of Interconnection Networks*. Morgan Kaufmann.
- Y. Demir and N. Hardavellas. 2014. Ecolaser: An adaptive laser control for energy efficient on-chip photonic interconnects. In *Proceedings of the International Symposium on Low-Power Electronics and Design*. 3–8.
- Y. Demir and N. Hardavellas. 2014. LaC: Integrating laser control in a photonic interconnect. In *Proceedings of the 2014 IEEE Photonics Conference*. 28–29.
- Y. Demir and N. Hardavellas. 2016. SLac: Stage laser control for a flattened butterfly network. In *Proceedings of the IEEE International Symposium on High Performance Computer Architecture*. 321–332.
- Y. Demir, Y. Pan, S. Song, N. Hardavellas, J. Kim, and G. Memik. 2014. Galaxy: A high-performance energy-efficient multi-chip architecture using photonic interconnects. In *Proceedings of the 28th ACM International Conference on Supercomputing*. 303–312.
- G. H. Duan, A. Shen, A. Akrouit, F. V. Dijk, F. Lelarge, F. Pommereau, O. LeGouezigou, J. G. Provost, H. Gariah, F. Blache, F. Mallecot, K. Merghem, A. Martinez, and A. Ramdane. 2009. High performance InP-based quantum dash semiconductor mode-locked lasers for optical communications. *Bell Labs Technical Journal* 14, 3, 63–84.
- J. P. Epping, M. Hoekman, R. Mateman, A. Leinse, R. G. Heideman, A. van Rees, P. J. van der Slot, C. J. Lee, and K. J. Boller. 2015. High confinement, high yield Si₃N₄ waveguides for nonlinear optical applications. *Optics Express* 23, 2, 642–648.

- European, Japan, Korean, Taiwan, and United States Semiconductor Industry Associations. 2012. *The International Technology Roadmap for Semiconductors (ITRS)*. Retrieved from <http://www.itrs.net/>.
- A. W. Fang, H. Park, O. Cohen, R. Jones, M. J. Paniccia, and J. E. Bowers. 2006. Electrically pumped hybrid AlGaInAs-silicon evanescent laser. *Optics Express* 14, 20, 9203–9210.
- N. Hardavellas, S. Somogyi, T. F. Wenisch, R. E. Wunderlich, S. Chen, J. Kim, B. Falsafi, J. C. Hoe, and A. G. Nowatzky. 2004. SimFlex: A fast, accurate, flexible full-system simulation framework for performance evaluation of server architecture. *SIGMETRICS Performance Evaluation Review, Special Issue on Tools for Computer Architecture Research* 31, 4, 31–35.
- M. Heck and J. Bowers. 2014. Energy efficient and energy proportional optical interconnects for multi-core processors: Driving the need for on-chip sources. *IEEE Journal of Selected Topics in Quantum Electronics* 20, 4, 1–12.
- H. Hisham, G. Mahdiraji, A. Abas, M. Mahdi, and F. Adikan. 2012. Characterization of transient response in fiber grating fabry-perot lasers. *IEEE Photonics Journal*, 4, 6, 2353–2371.
- H. Hisham, G. Mahdiraji, A. Abas, M. Mahdi, and F. Adikan. 2012. Characterization of turn-on time delay in fiber grating Fabry-Perot lasers. *IEEE Photonics Journal* 4, 5, 1662–1678.
- J. Huh, C. Kim, H. Shafi, L. Zhang, D. Burger, and S. W. Keckler. 2005. A NUCA substrate for flexible CMP cache sharing. In *Proceedings of the Annual International Conference on Supercomputing*. 31–40.
- A. Joshi, C. Batten, Y. J. Kwon, S. Beamer, I. Shamim, K. Asanovic, and V. Stojanovic. 2009. Silicon-photonics networks for global on-chip communication. In *Proceedings of the IEEE International Symposium on Networks-on-Chip*. 124–133.
- S. Kandula, S. Sengupta, A. Greenberg, P. Patel, and R. Chaiken. 2009. The nature of data center traffic: Measurements and analysis. In *Proceedings of the 9th ACM SIGCOMM Internet Measurement Conference*, 202–208.
- J. Kim, W. Dally, and D. Abts. 2007. Flattened butterfly: A cost-efficient topology for high-radix networks. In *Proceedings of the 34th Annual International Symposium on Computer Architecture*. 126–137.
- L. C. Kimerling. 2013. Scaling functionality with silicon photonics: Achievement and potential. In *UK Silicon Photonics Showcase Event*. Retrieved January 1, 2015, from http://www.orc.soton.ac.uk/fileadmin/seminar_pdf/UKSP_Showcase_-_Lionel_Kimerling.pdf.
- N. Kirman, M. Kirman, R. K. Dokania, J. F. Martinez, A. B. Apsel, M. A. Watkins, and D. H. Albonesi. 2006. Leveraging optical technology in future bus-based chip multiprocessors. In *Proceedings of the 39th IEEE/ACM Annual International Symposium on Microarchitecture*. 492–503.
- B. R. Koch, E. J. Norberg, B. Kim, J. Hutchinson, J. H. Shin, G. Fish, and A. Fang. 2013. Integrated silicon photonic laser sources for telecom and datacom. In *Proceedings of the Optical Fiber Communication Conference/National Fiber Optic Engineers Conference 2013*. 1–3.
- E. Kotelnikov, A. Katsnelson, K. Patel, and I. Kudryashov. 2012. High-power single-mode InGaAsP/InP laser diodes for pulsed operation. In *Proceedings of SPIE*. 8277:827715–1–827715–6.
- A. Krishnamoorthy, R. Ho, X. Zheng, H. Schwetman, J. Lexau, P. Koka, G. Li, I. Shubin, and J. Cunningham. 2009. Computer systems based on silicon photonic interconnects. *Proceedings of the IEEE* 97, 7, 1337–1361.
- G. Kurian, C. Sun, C. H. Chen, J. Miller, J. Michel, L. Wei, D. Antoniadis, L. S. Peh, L. Kimerling, V. Stojanovic, and A. Agarwal. 2012. Cross-layer energy and performance evaluation of a nanophotonic manycore processor system using real application workloads. In *Proceedings of the 26th IEEE International Parallel Distributed Processing Symposium*. 1117–1130.
- C. Li, M. Browning, P. V. Gratz, and S. Palermo. 2014. LumiNoC: A power-efficient, high-performance, photonic network-on-chip. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 33, 6, 826–838.
- G. Li, J. Yao, H. Thacker, A. Mekis, X. Zheng, I. Shubin, Y. Luo, J. H. Lee, K. Raj, J. E. Cunningham, and A. V. Krishnamoorthy. 2012. Ultralow-loss, high-density SOI optical waveguide routing for macrochip interconnects. *Optics Express*, 20, 11, 12035–12039.
- S. Li, J. H. Ahn, R. D. Strong, J. B. Brockman, D. M. Tullsen, and N. P. Jouppi. 2009. McPAT: An integrated power, area, and timing modeling framework for multicore and manycore architectures. In *Proceedings of the 42nd IEEE/ACM Annual International Symposium on Microarchitecture*. 469–480.
- J. Liu, X. Sun, R. Camacho-Aguilera, L. C. Kimerling, and J. Michel. 2010. Ge-on-si laser operating at room temperature. *Optics Letters* 35, 5, 679–681.
- P. Lotfi-Kamran, B. Grot, and B. Falsafi. 2012. NoC-out: Microarchitecting a scale-out processor. In *Proceedings of the 45th Annual IEEE/ACM International Symposium on Microarchitecture*. 177–187.
- G. Masini, A. Narasimha, A. Mekis, B. Welch, C. Ogden, C. Bradbury, C. Sohn, D. Song, D. Martinez, D. Foltz, D. Guckenberger, J. Eicher, J. Dong, J. Schramm, J. White, J. Redman, K. Yokoyama, M. Harrison, M.

- Peterson, M. Saberi, M. Mack, M. Sharp, P. D. Dobbelaere, R. LeBlanc, S. Leap, S. Abdalla, S. Gloeckner, S. Hovey, S. Jackson, S. Sahni, S. Yu, T. Pinguet, W. Xu, and Y. Liang. 2012. CMOS photonics for optical engines and interconnects. In *Proceedings of the National Fiber Optic Engineers Conference and Optical Fiber Communication Conference and Exposition 2012*. 1–3.
- B. Neel, M. Kennedy, and A. Kodi. 2015. Dynamic power reduction techniques in on-chip photonic interconnects. In *Proceedings of the 25th Edition on Great Lakes Symposium on VLSI*. 249–252.
- C. Nitta, M. Farrens, and V. Akella. 2011. Addressing system-level trimming issues in on-chip nanophotonic networks. In *Proceedings of the IEEE International Symposium on High Performance Computer Architecture*. 122–131.
- C. Nitta, M. Farrens, and V. Akella. 2012. DCOF: An arbitration free directly connected optical fabric. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 2, 2, 169–182.
- C. J. Nitta, M. K. Farrens, and V. Akella. 2013. *On-Chip Photonic Interconnects: A Computer Architect's Perspective*. Morgan & Claypool Publishers.
- J. S. Orcutt, A. Khilo, C. W. Holzwarth, M. A. Popovic, H. Li, J. Sun, T. Bonifield, R. Hollingsworth, F. X. Kartner, H. I. Smith, V. Stojanovic, and R. J. Ram. 2011. Nanophotonic integration in state-of-the-art CMOS foundries. *Optics Express* 19, 3, 2335–2346.
- Y. Pan, J. Kim, and G. Memik. 2010. Flexishare: Channel sharing for an energy-efficient nanophotonic crossbar. In *Proceedings of the IEEE International Symposium on High-Performance Computer Architecture*. 1–12.
- Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, and A. Choudhary. 2009. Firefly: Illuminating future network-on-chip with nanophotonics. In *Proceedings of the Annual International Symposium on Computer Architecture*. 429–440.
- M. Paniccia and J. Bowers. 2006. First electrically pumped hybrid silicon laser. Retrieved May 2014, from <http://www.intel.com/content/dam/www/public/us/en/documents/technology-briefs/intel-labs-hybrid-silicon-laser-announcement.pdf>.
- K. Petermann. 1988. *Laser Diode Modulation and Noise*. Advances in Optoelectronics, Vol. 3. Springer.
- P. Rosenfeld, E. Cooper-Balis, and B. Jacob. 2011. DRAMsim2: A cycle accurate memory system simulator. *Computer Architecture Letters* 10, 1, 16–19.
- E. Safi, A. Moshovos, and A. Veneris. 2008. L-CBF: A low-power, fast counting bloom filter architecture. *IEEE Transactions on Very Large Scale Integration Systems* 16, 6, 628–638.
- S. K. Selvaraja, W. Bogaerts, D. V. Thourhout, and R. Baets. 2010. Record low-loss hybrid rib/wire waveguides for silicon photonic circuits. In *Proceedings of the 7th International Conference on Group IV Photonics*. 1–3.
- K. Skadron, M. R. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan. 2003. Temperature-aware microarchitecture. In *Proceedings of the Annual International Symposium on Computer Architecture*. 2–13.
- M. Stucchi, S. Cosemans, J. Van Campenhout, Z. Tkei, and G. Beyer. 2013. On-chip optical interconnects versus electrical interconnects for high-performance applications. *Microelectronic Engineering* 112, 84–91.
- C. Sun, C. H. O. Chen, G. Kurian, L. Wei, J. Miller, A. Agarwal, L. S. Peh, and V. Stojanovic. 2012. Dsent - a tool connecting emerging photonics with electronics for opto-electronic networks-on-chip modeling. In *Proceedings of the 6th IEEE/ACM International Symposium on Networks-on-Chip*. 201–210.
- C. Sun, Y. H. Chen, and V. Stojanovic. 2013. Designing processor-memory interfaces with monolithically integrated silicon-photonics. In *Proceedings of the Conference on Lasers and Electro-Optics Pacific Rim*. 1–2.
- C. Sun, M. T. Wade, Y. Lee, J. S. Orcutt, L. Alloatti, M. S. Georgas, A. S. Waterman, J. M. Shainline, R. R. Avizienis, S. Lin, B. R. Moss, R. Kumar, F. Pavanello, A. H. Atabaki, H. M. Cook, A. J. Ou, J. C. Leu, Y.-H. Chen, K. Asanovic, R. J. Ram, M. A. Popovic, and V. M. Stojanovic. 2015. Single-chip microprocessor that communicates directly using light. *Nature* 528, 7583, 534–538.
- R. Takei, S. Manako, E. Omoda, Y. Sakakibara, M. Mori, and T. Kamei. 2014. Sub-1 dB/cm submicrometer-scale amorphous silicon waveguide for backend on-chip optical interconnect. *Optics Express* 22, 4, 4779–4788.
- S. Tanaka, S. H. Jeong, S. Sekiguchi, T. Kurahashi, Y. Tanaka, and K. Morito. 2012. Highly-efficient, low-noise Si hybrid laser using flip-chip bonded SOA. In *Proceedings of the IEEE Optical Interconnects Conference*. 12–13.
- Y. Thonnart, E. Beigne, A. Valentian, and P. Vivet. 2008. Automatic power regulation based on an asynchronous activity detection and its application to ANOC node leakage reduction. In *Proceedings of the 14th IEEE International Symposium on Asynchronous Circuits and Systems*. 48–57.

- D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. P. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. G. Beausoleil, and J. H. Ahn. 2008. Corona: System implications of emerging nanophotonic technology. In *Proceedings of the Annual International Symposium on Computer Architecture*. 153–164.
- T. F. Wenisch, R. E. Wunderlich, M. Ferdman, A. Ailamaki, B. Falsafi, and J. C. Hoe. 2006. SimFlex: Statistical sampling of computer system simulation. *IEEE Micro* 26, 4, 18–31.
- P. Wolf, P. Moser, G. Larisch, W. Hofmann, H. Li, J. Lott, C. Y. Lu, S. Chuang, and D. Bimberg. 2013. Energy-efficient and temperature-stable high-speed VCSELs for optical interconnects. In *Proceedings of the 15th International Conference on Transparent Optical Networks*. 1–5.
- L. Zhou and A. Kodi. 2013. Probe: Prediction-based optical bandwidth scaling for energy-efficient NoCs. In *Proceedings of the 7th IEEE/ACM International Symposium on Networks on Chip*. 1–8.
- A. Zilkie, B. Bijlani, P. Seddighian, D. C. Lee, W. Qian, J. Fong, R. Shafiha, D. Feng, B. Luff, X. Zheng, J. Cunningham, A. V. Krishnamoorthy, and M. Asghari. 2012. High-efficiency hybrid III-V/Si external cavity DBR laser for 3um SOI waveguides. In *Proceedings of the 9th International Conference on Group IV Photonics*. 317–319.

Received June 2016; revised October 2016; accepted November 2016