

EcoLaser: An Adaptive Laser Control for Energy-Efficient On-Chip Photonic Interconnects

Yigit Demir and Nikos Hardavellas

Northwestern University, Department of Electrical Engineering and Computer Science, Evanston, IL, USA
yigit@u.northwestern.edu, nikos@northwestern.edu

ABSTRACT

The high-speed and low-cost modulation of light make photonic interconnects an attractive solution for the communication demands of manycore processors. However, the high optical loss of many nanophotonic components results in high laser power consumption, most of which is wasted during periods of system inactivity. We propose EcoLaser, an adaptive laser control mechanism that saves between 24-77% of the laser power by turning off the laser when not needed. These power savings allow the cores to exploit a higher power budget and achieve speedups of 1.1-2x.

Categories and Subject Descriptions

B.4.3 [Hardware]: Interconnections; C.1.2 [Computer Systems Organization]: Multiprocessors— *Interconnection architectures*

Keywords

Nanophotonic Interconnection Networks; Adaptive Laser Control

1. INTRODUCTION

Silicon photonics have emerged as a promising solution to meet the growing demand for high-bandwidth, low-latency, and energy-efficient communication in manycore processors. Silicon waveguides can be manufactured alongside CMOS logic on the same die by adding a few new steps in the manufacturing process [5], and they are more efficient for long-distance on-chip communication than electrical signaling [18]. However, the high optical loss of typical silicon waveguides, optical couplers, and on-ring resonators dramatically increase the laser power consumption. While some optical interconnect topologies better balance power and performance [6,18,16] most of these costs are hard to avoid, and the laser power remains a considerable fraction of the total power budget. These costs together with the low efficiency of WDM-compatible lasers (in the range of 5-10% [27]), result in wall-plug laser power that is 10-20x higher than the required laser output power.

The majority of this power is typically wasted when activity is low because photonic interconnects are always on. By comparison, electrical interconnects stay idle until a packet traverses them. It is common for the interconnect to stay idle often for long periods of time, both in scientific computing (compute-intensive execution phases underutilize the interconnect), and in server computing (Google-scale datacenters have a typical utilization of 30% [1]).

Permission to make digital or hard copies of all or part 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 bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

ISLPED'14, August 11 - 13 2014, La Jolla, CA, USA

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2975-0/14/08...\$15.00.

<http://dx.doi.org/10.1145/2627369.2627620>

Motivated by these observations, we propose EcoLaser, a collection of static and adaptive laser control mechanisms that react to the demands of the aggregate workload and opportunistically turn the laser off during periods of low activity to save energy, and leave it on during periods of high activity in order to meet the high bandwidth demand. More specifically, our contributions are:

1. We propose laser control as a viable technique to save power and quantify the maximum opportunity.
2. We propose EcoLaser, a collection of static and dynamic laser control mechanisms and policies for SWMR and MWSR crossbars that approximate the maximum possible savings.
3. We evaluate the impact of EcoLaser on the performance and energy of a multicore running a range of synthetic and scientific workloads, under realistic physical constraints, and across a range of optical crossbar sizes.

Our results indicate that EcoLaser saves between 24-77% of the laser power for radix-16 and radix-64 SWMR and MWSR crossbars on real-world workloads. EcoLaser closely tracks (within 2-3% on average) a perfect controller with full knowledge of future interconnect requests. Thus, EcoLaser harvests the vast majority of the energy benefits that can be achieved by controlling the laser source. Moreover, the power savings of EcoLaser leave a higher power budget to the cores, which allows a multicore chip with EcoLaser to outperform a baseline scheme with no control by 1.1-2x.

2. LASER CONTROL SCHEMES

The objective of the laser control is to save laser energy by turning the lasers off whenever the bus (i.e., data channel) is idle. When the laser is off, the messages have to wait for the laser to turn on before transmission. EcoLaser capitalizes on recent advancements in Ge-based lasers [12,14], which enable DWDM-compatible on-chip laser sources that can be turned on within 1 ns, during which period the laser consumes the same power as when it is lasing. The laser control schemes we propose aim to maximize energy savings while minimizing the laser turn-on delay overhead.

2.1 Laser Control for SWMR Crossbar

A router in a Single-Reader-Multiple-Writer (SWMR) [11] crossbar (Figure 2) writes to its own dedicated bus, and reads from the other routers' busses. Reservation channels provide exclusion on the data bus [18]. The shaded components in the router microarchitecture in Figure 1 correspond to components added by EcoLaser. The laser controller turns the laser on if there is a message at any of the injection buffers, and it does not turn it off unless (a) there is no message at the injection buffers, and (b) the laser has stayed on for the minimum laser stay-on time "K". The laser controller keeps the switch allocator waiting while the laser turns on. When the laser is ready, the switch allocator moves messages to the modulators.

2.2 Laser Control for MWSR Crossbar

In a Multiple-Writer-Single-Reader [24] (MWSR) crossbar every router reads from its own bus, and writes on the other routers' busses (Figure 2). Because lasers can be built within a waveguide [14],

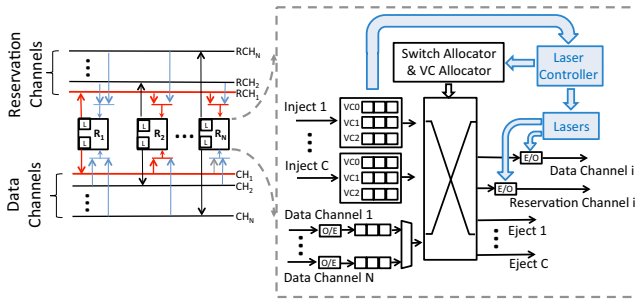


FIGURE 1. SWMR crossbar and router microarchitecture.

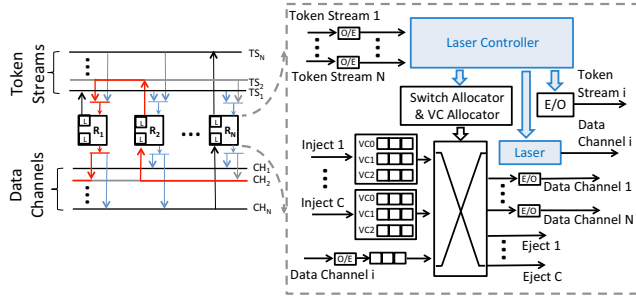


FIGURE 2. MWSR crossbar and router microarchitecture.

each receiver in EcoLaser holds the lasers for its own bus. This complicates laser control, as the receiver does not know that a sender wants to transmit. Contention in MWSR occurs when two routers try to transmit simultaneously to the same destination. Token-based arbitration [17,24] resolves the contention by using ring-shaped waveguides to move the tokens in the direction of data travel, and one cycle ahead of the data slot. The reader node collects back its returning tokens to control the input buffer utilization [17]. Any writer can exploit this and send a “Laser turn-on request” to the reader with a returning token.

We construct the tokens to perform three tasks: (a) maintain the time share on the bus, (b) indicate if there is light in the data bus hence the writer can write immediately, and (c) bring the laser turn-on requests back to the reader. Note that only the reader can inject tokens in his token stream, and any type of snooping of the token by writers is destructive. In order to meet all these needs, we design 3-bit tokens as shown at the top of Figure 3: the “ T ” bit provides mutual exclusion on the data bus; the “ L ” bit indicates if the laser was on when this token was released from the reader node (i.e., the subsequent slot in the data bus has light that can be modulated); the “ S ” signals the reader to turn on the laser. The photonic links operate at 2x the processor frequency, thus 2 wavelengths suffice (S and T are sent on the same wavelength in a single processor cycle).

Figure 3 shows the **Laser Controller Logic**. When the bus is idle, the data lasers stay at the “Laser OFF” state; they do not consume energy, and the released tokens indicate this with a clear L bit (note that the token-stream laser is always on). Writers send the laser turn-on signal by clearing the S bit of a token. When the laser controller receives the turn-on signal, the data channel lasers move to the “Laser Warm-up” state and the controller pushes a “1” into a 5-bit barrel shifter. During the warm-up the lasers consume full power preparing photons, but cannot emit any light yet. A “1” popping out from the other end of the 5-bit shifter indicates that the lasers are ready (1 ns is 5 cycles at 5 GHz). The lasers start emitting light into the data bus (emit data slots) and move to the “Laser Dedicated” state. The first emitted data slot is dedicated to the writer who requested the laser turn-on; the corresponding token has a

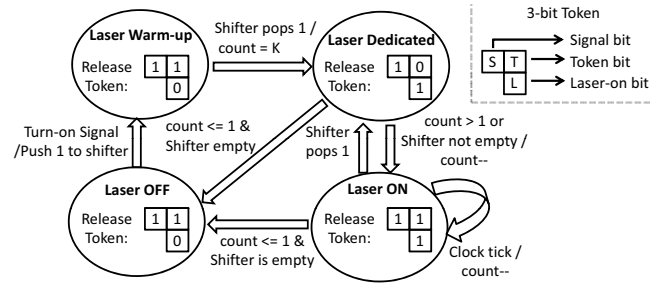


FIGURE 3. 3-bit Token and Laser Controller FSM.

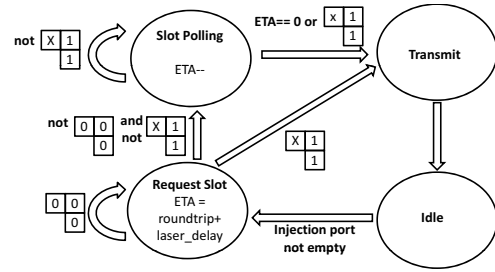


FIGURE 4. Writer Node FSM.

clear T bit, to prevent any other writer from grabbing the slot, thereby guaranteeing freedom from starvation. The lasers move to the “Laser ON” state and remain on for as long as there is a set bit in the shifter. The laser controller keeps track of the duration the laser has stayed on through the counter “ $count$ ”. When the laser emits the first (dedicated) slot, the count is assigned the value K . $Count$ decrements on every cycle, and the laser stays on and releases data slots which are available for any writer node to use (tokens indicate this availability with set T and L bits). When $count = 1$ the laser turns off, unless there is another set bit in the shifter, which indicates a new laser turn-on request. If there is, the laser remains on until the set bit pops out from the shifter, at which time the lasers move to the “Laser Dedicated” state like above.

Figure 4 shows the **Writer Node Logic**. When a writer has a message to send, it moves to the “Request Slot” state, and looks for an available data slot. The writer reads the T and L bits of the first token, and if they are both set (i.e., the data slot has light and is available), it modulates the message into the data slot. If T and L are not both set, the writer sends a laser turn-on signal by reading (clearing) the S bit of the token, and expects to receive a dedicated data slot at *Estimated Time of Arrival (ETA)* equal to round-trip plus laser-turn-on cycles later. If all token bits are clear, the token has been used to send out a laser turn-on signal already, so the writer stays in the “Request Slot” state and re-tries. After sending a laser turn-on signal, the writer moves to “Slot Polling”, in which the writer looks for an available slot (by reading both T and L bits) while waiting for its dedicated slot to arrive. The writer transmits when either an available slot with light arrives, or the writer’s dedicated slot arrives. By design the writer sends at most one laser turn-on signal, which avoids wasting laser energy. Also, writers can send a laser turn-on signal using a token that has been through another “Slot Polling” writer, which improves performance.

2.3 Adaptive Laser Control

The static laser control schemes discussed thus far leave the laser on for at least “ K ” cycles, where “ K ” is a fixed value. With lower laser stay-on time “ K ”, EcoLaser tends to turn off the laser quicker, which saves more laser energy when the crossbar is not heavily uti-

TABLE 1. Architectural Parameters.

CMP Size	64 cores, 480mm ²
Processing Cores	ULTRASPARC III ISA, up to 5Ghz, OoO, 4-wide dispatch/retirement, 96-entry ROB
L1 Cache	Split I/D, 64KB 2-way, 2-cycle load-to-use, 2 ports, 64-byte blocks, 32 MSHRs, 16-entry victim cache
L2 Cache	Shared, 512 KB per core, 16 way, 64-byte blocks, 14 cycle-hit, 32 MSHRs, 16-entry victim cache
Memory Controllers	One per 4 cores, 1 channel per Memory Controller Round-robin page interleaving
Main Memory	Optically connected memory [2], 10ns access
Networks	SWMR and MWSR crossbars, radix-16 and -64

lized. However, under heavier traffic, turning the laser off quickly results in lost opportunities to catch the laser on and send, and increases the number of times the laser has to be turned on anew. The frequent laser turn-on delays decrease performance. On the other hand, when K is high, the laser tends to stay on for longer, which increases performance under heavier traffic, but wastes more laser energy when the utilization is low. Thus, no static scheme is expected to perform best under all traffic conditions.

We propose an adaptive scheme that observes the amount of laser turn-on requests to adjust the laser stay-on time K at run time. Frequent laser turn-on requests hint to lost opportunities to transmit opportunistically, and the adaptive scheme increases K to keep the laser on for longer. A low number of laser turn-on requests hints at potentially wasted laser energy, so the adaptive scheme decreases K to save more laser energy, by turning the laser off more quickly.

To prevent oscillation or overshooting K from its ideal setting, we employ a hysteresis counter which robustly captures the laser turn-on request trends. The hysteresis counter decrements on every cycle on which there is no other counter activity. Upon sensing a laser turn-on signal, the counter increments by adding some value to it. Whenever the counter reaches its upper threshold, K increases by 1; whenever the counter reaches its lower threshold, K decreases by 1. The hysteresis counter controls the value of K in a stable manner, because increasing K results in a reduction of laser turn-on requests, as the likelihood of a writer finding an available data slot with light increases, and vice versa. The threshold settings and the increment and decrement values of the hysteresis counter change its reactive behavior (making it more lazy or aggressive). Through a design space exploration, we identified the settings that provide the highest energy savings for our workloads, and use these settings for the remainder of our study. Other than adapting K at run-time, the rest of the design of the adaptive laser control is the same as the designs described earlier for SWMR and MWSR crossbars.

2.4 The Perfect Laser Control

A perfect laser control scheme has complete knowledge of future interconnect accesses. The perfect scheme saves the maximum laser energy without incurring any performance overhead by turning the laser on ahead of time, so the light reaches the writer at the exact time the writer attempts to transmit. After transmitting, the control deactivates the laser or leaves it on for an upcoming message, if deactivation could cause a delay. Thus, the perfect scheme presents the maximum energy savings for a given laser technology.

3. EXPERIMENTAL METHODOLOGY

3.1 Interconnect Performance and Energy

To evaluate the performance and energy consumption of EcoLaser in isolation from the interference of other system components or

TABLE 2. Nanophotonic Parameters and Laser Power.

	per Unit	Radix-16 Total	Radix-64 Total
DWDM		64	16
WG Loss	0.3 dB/cm[4]	3 dB	3 dB
Nonlinearity	1 dB	1 dB	1 dB
Modulator Ins.	0.5 dB	0.5 dB	0.5 dB
Ring Through	0.01 dB	10.24 dB	10.24 dB
Filter Drop	1.2 dB	1.2 dB	1.2 dB
Photodetector	0.1 dB	0.1 dB	0.1 dB
Total Loss		16.04 dB	16.04 dB
Detector		-20 dBm	-20 dBm
Laser Power	per Wavelength	0.401 mW	0.401 mW
Total Laser Power		20.1 W	78.1 W

application characteristics, we use a cycle-accurate network simulator based on Booksim 2.0 [7], which models radix-16 and radix-64 SWMR and MWSR crossbars servicing random uniform traffic (we use the notation $\langle type \rangle_XBAR_ \langle radix \rangle$ for the crossbars). We model a 480 mm² chip with single-cycle routers, 1-cycle E/O and O/E conversions, and a 10 cm waveguide that has a 5-cycle round-trip time. The link latency (1-5 cycles) is calculated based on the traversed waveguide length. The buffers are 20-flits deep, and flits are 300 bits. The maximum core frequency is 5 GHz, and the optical interconnect runs at 10 GHz. We evaluate the load-latency and energy-per-flit of EcoLaser, and compare it against a baseline without laser control (*No-Ctrl*), and against a perfect control scheme (*Perfect*). To demonstrate the merits of the adaptive mechanism, we compare adaptive laser control (*Adaptive*) with two static control mechanisms: *Static-1* and *Static-10*, with 1- and 10-cycle stay-on times respectively. *Static-1* is the quickest to turn the laser off; *Static-10* saves the most laser energy per packet among all static schemes on average across injection rates.

3.2 Multicore System Performance and Energy

To evaluate the impact of EcoLaser on a realistic multicore, we model a 64-core processor on a full-system cycle-accurate simulator based on Flexus 4.0 [9,25] integrated with Booksim 2.0 [7] and DRAMSim 2.0 [19]. We calculate the power consumption of the electrical interconnect using DSENT [21]. We target a 16 nm technology, and have updated our tool chain accordingly based on ITRS [8]. The simulated system executes a selection of SPLASH-2 and other scientific workloads. We model realistic multicores that use a throttling mechanism (without loss of generality, we use Dynamic Voltage and Frequency Scaling—DVFS) to keep the chip within safe operational temperatures (below 90°C).

We collect runtime statistics from full-system simulations, and use them to calculate the power consumption of the system using McPAT [13], and the power consumption of the optical networks using the analytical power model by Joshi *et al.* [10]. We estimate the temperature of the chip using HotSpot 5.0 [20]. We adjust the leakage power estimate based on the temperature and adjust DVFS based on the stable-state power and temperature estimates.

To put EcoLaser’s performance and energy consumption into perspective, we include in our evaluation a 2D-concentrated electrical mesh with express links (*CMesh*). For CMesh we model routers with 8 input and output ports and a 3-cycle routing delay. Routers are connected through 150-bit bi-directional links with 1-cycle local and 3-cycle global delay. To show the range of EcoLaser’s impact, we evaluate its application on two optical crossbars that are

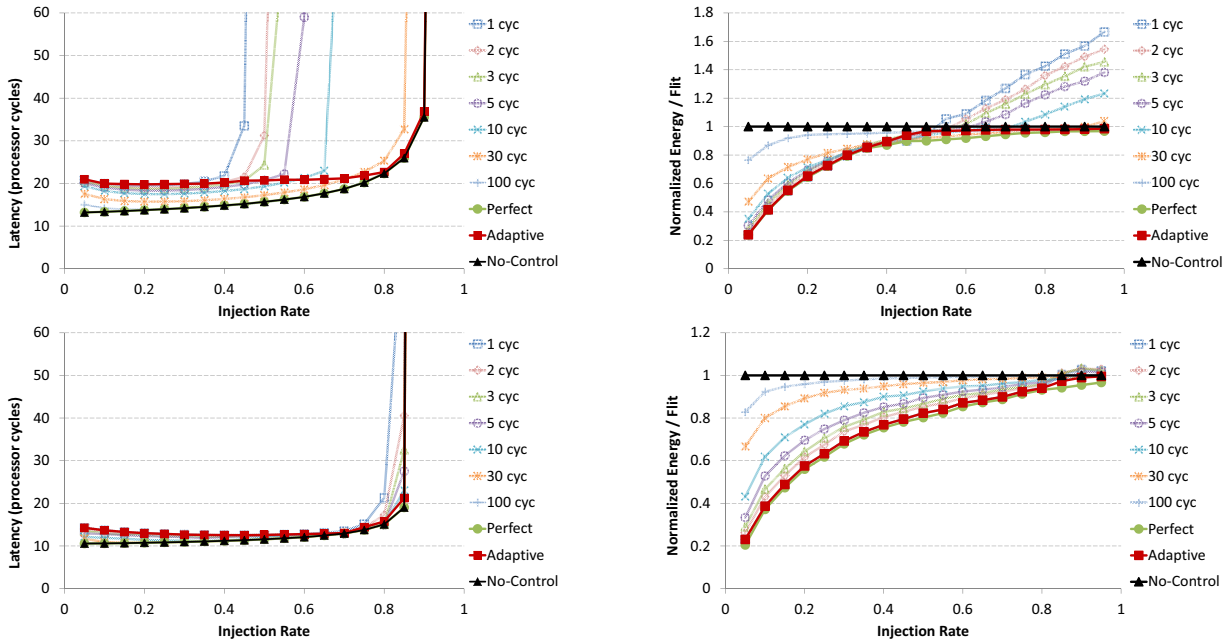


FIGURE 5. Load-Latency (left) and Energy-per-Flit (right) for radix-16 MWSR (top row) and SWMR (bottom row) crossbars.

at the opposite ends of the spectrum, a radix-16 and a radix-64 optical crossbar. The radix-16 crossbar approximates a worst case scenario for EcoLaser. It has low power consumption (similar to the power consumption of CMesh) and its high concentration factor (4) creates heavier traffic. The low power consumption and heavy traffic limit EcoLaser’s opportunity. The radix-64 crossbar corresponds to a better case for EcoLaser. It has high laser power consumption and a low concentration (1), which results in light traffic, thus giving ample opportunity to EcoLaser to conserve laser power. For each one of the radix-16 and radix-64 crossbars, we evaluate both SWMR and MWSR designs. The modeling of the optical interconnects is described in Section 3.1. Finally, we contrast EcoLaser to a power-equivalent optical interconnect design similar to No-Ctrl, but with its interconnect width scaled down to approximate EcoLaser’s average energy savings ($Power_Eq$).

3.3 Optical Network Power Consumption

Table 2 shows the assumed optical loss parameters for the modulators, demodulators, drop filters, and detectors, as introduced in [2]. The modulation and demodulation energy is 150 fJ/bit at 10 GHz [2]. The laser power per wavelength and total laser power are calculated in Table 2 using the analytical models introduced in [10]. Because the number of turned-off rings on a single optical path is high for a radix-64 crossbar, we limit the network to 16 DWDM. The total laser power in Table 2 includes the laser power for both data and reservation channels, plus the laser efficiency of 10%, so it is the wall plug power for the laser. The data bus is 300-bits wide.

To calculate the total ring heating power we extend the method by Nitta *et al.* [15] 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 using the 3D-chip extension of Hot-Spot [20]. When a workload executes, we calculate the ring heating power required to maintain the entire photonic die at the micro-ring trimming temperature during the entire execution. In addition, we account for the individual ring trimming power required to overcome process variations, as described in [10].

Unfortunately, there is little consensus on the optical loss parameters used or projected in literature, as parameters exhibit a variance over 10x across publications. However, the optical components typically drive the design of an optical interconnect (e.g., lower DWDM minimizes off-ring through loss, and low-radix crossbars can be used to implement high-radix ones [18]). In either case, the fraction of laser energy that EcoLaser saves depends on the network utilization, not on the optical loss parameters, and the higher the total optical loss, the higher EcoLaser’s impact would be. Thus, we remain conservative in our estimates of optical losses.

4. EXPERIMENTAL RESULTS

4.1 Network Performance

EcoLaser turns off the lasers and saves energy at the cost of higher message latency. At low injection rates, EcoLaser on SWMR has a 4-cycle latency overhead, which is lower than the 5-cycle laser turn-on delay, as some messages catch the laser on (Figure 5). The overhead decreases for higher injection rates as more messages catch the laser on. Similarly, EcoLaser on MWSR exhibits 8 cycles latency overhead, instead of the full 11-cycle laser turn-on delay, as the token design allows senders to transmit immediately when they find the laser on. Static schemes with high laser stay-on time “ K ” (keeping the laser on at least K cycles), increase the likelihood of finding the laser on, so they have lower latency overhead and provide higher throughput. However, they don’t save much energy at low injection rates, as they may needlessly leave the lasers on. Static schemes with lower K turn off the lasers quickly, saving significant laser energy at low injection rates. However, they don’t provide enough throughput under heavy utilization, increasing the overall energy consumption. Adaptive outperforms all Static schemes because it adjusts K at runtime, thus it achieves high energy savings at low injection rates, and high throughput at high injection rates. Adaptive’s performance improvement over Static schemes is higher for MWSR, because it sends turn-on requests through the token stream (which takes longer), while SWMR can turn on or keep the laser on much quicker. Overall, Adaptive’s energy consumption is within 2-3% of the Perfect scheme.





FIGURE 6. Speedup for radix-16 (left) and radix-64 (right) MWSR on a hypothetical multicore without thermal constraints.





FIGURE 7. Speedup over CMesh for radix-16 (left) and radix-64 (right) MWSR crossbars under realistic thermal constraints.

4.2 The Performance Cost of Laser Control

EcoLaser is expected to degrade performance compared to No-Ctrl, as sometimes transmission is delayed while the laser turns on. In reality, however, EcoLaser recoups the losses and even increases performance by minimizing thermal emergencies and core throttling that DVFS employs to keep a chip within safe operating temperatures. Controlling the laser lowers the power consumption by a significant margin compared to No-Ctrl, which allows for a cooler chip, reduces core throttling, and increases performance. Thus, even though EcoLaser trades off network latency for energy savings, a realistic power-limited system may exhibit higher performance with EcoLaser because the cores will be throttled less often.

We analyze the two effects (increasing the network latency, and reducing core throttling) separately. We analyze the performance cost of EcoLaser by evaluating it on a multicore that is not subject to thermal constraints, thus cores are not throttled and run at maximum frequency (5 GHz). Our workload suite includes both memory-intensive workloads that generate high traffic and are sensitive to interconnect latency (em3d, ocean, appbt, tomcatv), as well as compute-intensive workloads that are less sensitive to message latency (fmm, moldyn, barnes). Figure 6 summarizes our findings. The injection rate of each application appears below its name. Overall, laser control saves more energy on real-world workloads than on synthetic random traffic patterns, because real-world workloads typically have bursty (and sparse) memory access patterns.

In radix-16 MWSR, Static-1 saves the most laser energy (49% on average) at the expense of slowing down the memory intensive workloads. Static-10 achieves high throughput, but it wastes laser energy at compute-intensive workloads (saves 32% on average).





FIGURE 8. Energy x Delay Product in radix-16 and radix-64 MWSR crossbar. The evaluated designs are from left to right: No-Ctrl (N), Power_Eq (E), Static-1 (1), Static-10 (10), Adaptive (A), and Perfect (P).

Adaptive combines the benefits of both: it saves 45% of the laser energy on average for radix-16 and 68% for radix-64 MWSR crossbars, at the cost of 4.8% and 7.5% slowdown respectively, while on SWMR it saves 53% and 72% of the laser energy for radix-16 and radix-64 respectively, with only 4% slowdown.

Power_Eq is a scaled-down version of No-Ctrl (150-bit flits for radix-16, and 100-bit flits for radix-64) to approximate Adaptive's laser energy consumption. While it achieves similar energy savings, Power_Eq suffers from high serialization delays and underperforms EcoLaser. Thus, saving laser energy by reducing the width of the interconnect is not a good alternative to laser control.

4.3 EcoLaser on a Realistic Multicore

Under realistic thermal (power) constraints, DVFS in No-Ctrl throttles the cores to keep the chip within a safe temperature. EcoLaser, however, reduces the laser power and results in a cooler chip, less core throttling, and higher performance. The static schemes typically work well at only one end of the spectrum. Static-1 speeds up workloads with low injection rates, as it saves the most power and reduces throttling, but slows down memory-intensive workloads due to frequent laser turn-on delays (Figure 7-left). Static-10 speeds up workloads with high injection rates, as it increases the likelihood that a sender finds the laser on and transmits without delay, but wastes power when the injection rate is low and leads to more core throttling. Power_Eq achieves low laser power, but at the expense of serialization delays due to its limited width. Overall, the performance and energy-delay product (EDP, Figure 8) of the static schemes is much worse than that of Perfect's. Thus, static laser control or reduced width often lead to slow and energy-inefficient systems.

Adaptive EcoLaser tracks the workload's needs, and provides both low power and high throughput. The impact of EcoLaser is more pronounced on 64-radix crossbars, because their energy savings are a significant fraction of the total chip power, and hence allow the cores to run faster. For example, Perfect runs fmm at 3.25 GHz, Adaptive at 3.2 GHz, and No-Ctrl at only 1.5 GHz. For the same reason, No-Ctrl is 1.7x slower than CMesh even though it has higher bandwidth and lower latency. Compared to No-Ctrl, adaptive EcoLaser on radix-64 MWSR and SWMR crossbars is 2x faster and has 74-77% lower EDP on average (10% faster and 20% lower EDP for radix-16). In all cases, Adaptive's performance and EDP are within 2-6% of Perfect's.

5. RELATED WORK

Corona [24] and many others [23,17,16], implement an MWSR crossbar topology for on-chip communication. Firefly [18] uses partitioned SWMR optical crossbars to connect clusters of electrically-connected mesh networks. Batten *et al.* [2,3] connects a many-core processor to DRAM memory using SWMR crossbars. These works can exploit EcoLaser for laser energy efficiency.

Thonnart *et al.* [22] propose techniques to reduce the static power consumption in electrical interconnects by powering down unused units. Zhou *et al.* [26] identify the constant laser power consumption as an inefficiency, and propose a mechanism to increase average channel utilization, by controlling splitters to tune bandwidth on a binary tree network. Kurian *et al.* [12] propose an optical SWMR crossbar and electrical hybrid network, and mention that a Ge-based laser can be controlled to improve the laser energy efficiency, but don't present nor evaluate a laser control scheme.

6. CONCLUSION

In this paper we propose EcoLaser, a laser-control mechanism that turns the laser off during periods of inactivity to save energy, and meets high bandwidth demands by turning the laser on for as long as necessary. EcoLaser saves between 24-77% of the laser power for radix-16 and radix-64 SWMR and MWSR crossbars on real-world workloads. EcoLaser harvests the vast majority of the energy benefits, as it closely tracks (within 2-3% on average) a perfect controller with full knowledge of future interconnect requests. Moreover, the power savings of EcoLaser allow for providing a higher power budget to the cores, which enables a multicore chip with EcoLaser to achieve speedups of 1.1-2x over a multicore with a traditional optical interconnect that keeps the lasers always on.

7. ACKNOWLEDGEMENTS

This work was partially supported by National Science Foundation award CCF-1218768, and ISEN booster award, and the June and Donald Brewer Chair in EECS at Northwestern University.

8. REFERENCES

- [1] L. A. Barroso and U. Holzle. The case for energy-proportional computing. *IEEE Computer*, 40(12):33- 37, 2007.
- [2] C. Batten, A. Joshi, J. Orcutt, A. Khilo, B. Moss, C. W. Holzwarth, M. A. Popovic, H. Li, H. I. Smith, J. L. Hoyt, F. X. Kartner, R. J. Ram, V. Stojanovic, and K. Asanovic. Building many-core processor-to-dram networks with monolithic CMOS silicon photonics. *IEEE Micro*, 29(4):8- 21, 2009.
- [3] C. Batten, A. Joshi, V. Stojanovic, and K. Asanovic. Designing chip-level nanophotonic interconnection networks. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 2(2), 2012.
- [4] J. Cardenas, C. Poitras, J. Robinson, K. Preston, L. Chen, and M. Lipson. Low loss etchless silicon photonic waveguides. *Optics Express*, 17(6):4752- 4757, 2009.
- [5] G. Chen, H. Chen, M. Haurylau, N. Nelson, P. M. Fauchet, E. Friedman, and D. Albonese. Predictions of cmos compatible on-

- chip optical interconnect. In *7th International Workshop on System-Level Interconnect Prediction (SLIP)*, pp. 13- 20, 2005.
- [6] M. J. Cianchetti, J. C. Kerekes, and D. H. Albonese. Phastlane: a rapid transit optical routing network. In *36th Annual International Symposium on Computer Architecture*, 2009.
- [7] W. J. Dally and T. B. Principles and Practices of Interconnection Networks. Morgan Kaufmann Publishing Inc., 2004.
- [8] European (ESIA), Japan (JEITA), Korean (KSIA), Taiwan (TSIA), and United States (SIA) Semiconductor Industry Associations. The international technology roadmap for semiconductors (ITRS), 2012.
- [9] N. Hardavellas, S. Somogyi, T. F. Wenisch, R. E. Wunderlich, S. Chen, J. Kim, B. Falsafi, J. C. Hoe, and A. G. Nowatzky. SimFlex: a fast, accurate, flexible full-system simulation framework for performance evaluation of server architecture. *SIGMETRICS Performance Evaluation Review, Tools for Comp. Arch. Research*, 2004.
- [10] A. Joshi, C. Batten, Y.-J. Kwon, S. Beamer, I. Shamim, K. Asanovic, and V. Stojanovic. Silicon-photonics networks for global on-chip communication. In *IEEE International Symposium on Networks-on-Chip (NOCS)*, pp. 124- 133, 2009.
- [11] N. Kirman, M. Kirman, R. K. Dokania, J. F. Martinez, A. B. Apse, M. A. Watkins, and D. H. Albonese. Leveraging optical technology in future bus-based chip multiprocessors. In *39th IEEE/ACM Annual International Symposium on Microarchitecture*, pp. 492- 503, 2006.
- [12] 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. Cross-layer energy and performance evaluation of a nanophotonic manycore processor system using real application workloads. In *26th International Parallel Distributed Processing Symposium*, 2012.
- [13] S. Li, J. H. Ahn, R. D. Strong, J. B. Brockman, D. M. Tullsen, and N. P. Jouppi. Mpcat: an integrated power, area, and timing modeling framework for multicore and manycore architectures. In *42nd IEEE/ACM Annual International Symposium on Microarchitecture*, 2009.
- [14] J. Liu, X. Sun, R. Camacho-Aguilera, L. C. Kimerling, and J. Michel. Ge-on-Si laser operating at room temperature. *Opt. Lett.*, 35(5), 2010.
- [15] C. Nitta, M. Farrens, and V. Akella. Addressing system-level trimming issues in on-chip nanophotonic networks. In *17th IEEE International Symposium on High Performance Computer Architecture*, 2011.
- [16] Y. Pan, J. Kim, and G. Memik. Flexishare: Channel sharing for an energy-efficient nanophotonic crossbar. In *16th IEEE International Symposium on High-Performance Computer Architecture*, 2010.
- [17] Y. Pan, J. Kim, and G. Memik. Featherweight: low-cost optical arbitration with qos support. In *44th IEEE/ACM Annual International Symposium on Microarchitecture*, pp. 105-116, 2011.
- [18] Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, and A. Choudhary. Firefly: Illuminating future network-on-chip with nanophotonics. In *36th Annual International Symposium on Computer Architecture*, 2009.
- [19] P. Rosenfeld, E. Cooper-Balis, and B. Jacob. Dramsim2: A cycle accurate memory system simulator. *Comp. Arch. Letters*, 10(1), 2011.
- [20] K. Skadron, M. R. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan. Temperature-aware microarchitecture. In *30th Annual International Symposium on Computer Architecture*, 2003.
- [21] C. Sun, C.-H. O. Chen, G. Kurian, L. Wei, J. Miller, A. Agarwal, L.-S. Peh, and V. Stojanovic. Dsent - a tool connecting emerging photonics with electronics for opto-electronic networks-on-chip modeling. In *6th IEEE International Symposium on Networks-on-Chip*, 2012.
- [22] Y. Thonnart, E. Beigne, A. Valentian, and P. Vivet. Automatic power regulation based on an asynchronous activity detection and its application to anoc node leakage reduction. In *14th IEEE International Symposium on Asynchronous Circuits and Systems*, pp. 48-57, 2008.
- [23] D. Vantrease, N. L. Binkert, R. Schreiber, and M. H. Lipasti. Light speed arbitration and flow control for nanophotonic interconnects. In *42nd Annual International Symposium on Microarchitecture*, 2009.
- [24] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. P. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. G. Beausoleil, and J. H. Ahn. Corona: System implications of emerging nanophotonic technology. In *35th Annual International Symposium on Computer Architecture*, 2008.
- [25] T. F. Wenisch, R. E. Wunderlich, M. Ferdman, A. Ailamaki, B. Falsafi, and J. C. Hoe. SimFlex: statistical sampling of computer system simulation. *IEEE Micro*, 26(4):18- 31, Jul-Aug 2006.
- [26] L. Zhou and A. Kodi. Probe: Prediction-based optical bandwidth scaling for energy-efficient NoCs. In *7th IEEE/ACM International Symposium on Networks on Chip (NoCS)*, pages 1- 8, 2013.
- [27] A. Zilkie, B. Bijlani, P. Seddighian, D. C. Lee, W. Qian, J. Fong, R. Shafiiha, D. Feng, B. Luff, X. Zheng, J. Cunningham, A. V. Krishnamoorthy, and M. Asghari. High-efficiency hybrid III-V/Si external cavity DBR laser for 3um SOI waveguides. In *9th IEEE International Conference on Group IV Photonics (GFP)*, 2012.