LaC: Integrating Laser Control in a Photonic Interconnect

Yigit Demir and Nikos Hardavellas
Northwestern University, Department of Electrical Engineering and Computer Science, Evanston, IL 60208, USA
yigit@u.northwestern.edu, nikos@northwestern.edu

Abstract—Silicon photonics are an attractive technology for global interconnects in manycore processors. However, they come with significant laser power consumption. Turning off underutilized portions of the interconnect can save 62-92% of the laser power.

1. INTRODUCTION

Silicon photonics have emerged as a promising technology to meet the growing demand for high-bandwidth, low-latency, and energy-efficient communication in manycore processors. However, the high optical loss of typical silicon waveguides, optical couplers, and on-ring resonators, together with the low efficiency (5-10% [5]) of WDM-compatible lasers, dramatically increase the laser power consumption. Thus, the wall-plug laser power requirement 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 consuming only leakage power, until a packet traverses them. Idle times are quite common, as the interconnect stays idle often for long periods of time, both in scientific computing (compute-intensive execution phases underutilize the interconnect), and in server computing (servers in Google-scale datacenters have a typical utilization of less than 30% [1]).

Motivated by these observations, in our previous work we proposed EcoLaser [2], a laser control mechanism that reacts to the demands of the aggregate workload by opportunistically turning the laser off during periods of low activity to save energy. The laser gating mechanism capitalizes on recent advancements in Ge lasers [4], which enable energy-efficient on-chip laser sources that can be turned on or off within nanoseconds. In this paper we propose the laser control scheme LaC which improves upon EcoLaser [2] by keeping the majority of the data bus off while sending small (data-less) messages and provides better performance by turning the laser on proactively for large data messages.

Our results indicate that LaC saves between 62-92% (77% on average) of the laser energy while running real-world applications compared to traditional interconnects that keep the lasers always on (No-Ctrl). The energy savings and performance of LaC stay within 2-4% of a Perfect scheme [2] with future knowledge of interconnect accesses. More importantly, the power savings of LaC allow the cores to exploit a higher power budget and run between 2-2.2x faster (2.1x on average) on a 64-core processor.

2. LASER CONTROL SCHEME

The objective of the laser control is to save laser energy by turning off the lasers whenever the bus (i.e., data channel) or a portion of the bus is idle. However, the laser should be turned back on before the bus can be used again, which introduces a laser turn-on delay on messages. EcoLaser [2] exploits this trade-off by leaving the laser on for slightly longer every time it is used, to encourage opportunistic usage of the laser by senders without a turn-on delay. EcoLaser provides high throughput under heavy utilization, while maintaining high laser energy savings at low injection rates. In this paper we propose a new **La**ser **C**ontrol mechanism, **LaC**, which improves upon the prior art by keeping a portion of the optical data bus off when sending small messages (typically control messages from the cache coherence protocol), and activates the lasers ahead of time when long data messages are anticipated.

As photonic links provide high bandwidth, they offer wide busses which can send a data message in one cycle. 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. However, an optical bus does not need to be more than 300 bits wide (i.e., 300 wavelengths), because the optical links run at double the processor frequency. On the other hand, small coherence messages are transmitted in two 44-bit wide flits (64-bit address, 20-bit ID and 4-bit message type), which means 256 bits of this optical bus are used only when sending data messages and remain idle otherwise. LaC capitalizes on this by activating only 44 wavelengths of the optical bus (keeping the remaining 256 wavelengths off) when sending small coherence messages. The whole bus (300 wavelengths) will only be activated to send data messages.

LaC anticipates the early laser activation for data messages by correlating cache coherence request messages to replies, and turns the laser on proactively which reduces the latency overhead. In a directory-based cache coherence protocol, every data message is generated upon receipt of a read or write request. The request results in a lookup in the local L2 cache slice, which resolves quickly if it misses as only the tags are accessed first. However, if the tag hits, an additional long latency access to the data array is required to get the data. LaC detects that and turns on the laser 1ns (i.e., one laser turn-on delay) before the data become available.

3. EXPERIMENTAL RESULTS

We evaluate the energy savings and the performance improvement of LaC on a radix-64 SWMR crossbar, using the methodology presented in [2,3]. LaC attempts to hide the laser turn-on latency by activating the lasers ahead of time for upcoming messages. However, lasers with a high turn-on delay decrease the potential savings that LaC can achieve. We investigate LaC's sensitivity to the laser turn-on delay by modeling LaC schemes across a range of laser turn-on

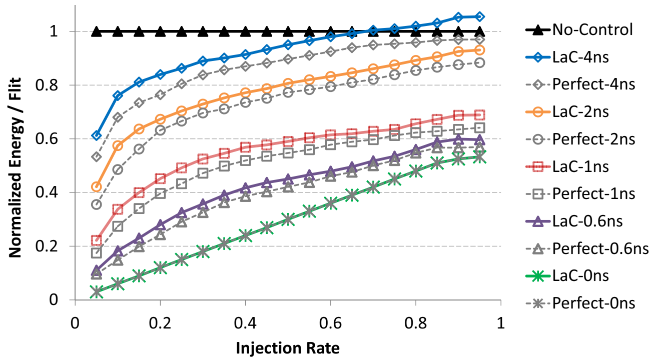

FIGURE 1. Energy savings' sensitivity to laser turn-on delay.

delays (Figure 1). The Ge-based laser [4] assumed in this work turns on in 1 *ns,* and allows LaC to save 32-88% of the laser energy (45% on average across injection rates). The energy savings of LaC decrease to 8% on average with lasers that have 4ns turn-on delay, indicating that laser control techniques are promising only when the lasers can turn on relatively fast. On the other end of the spectrum, employing LaC on a theoretical laser that can turn on and off instantaneously increases the energy savings up to 71%, providing an upper bound on what laser control schemes can achieve. Figure 1 shows that across the entire range of laser turn-on delays we evaluate, LaC closely tracks the energy savings of a theoretical Perfect scheme [2] which controls the lasers with perfect knowledge of future interconnect accesses (shown in gray dashed lines). Thus, LaC is near-optimal, and harvests the vast majority of the energy benefits that can be achieved.

We analyze the impact of LaC on the performance of a multicore system by running real-world workloads on a simulated processor that employs LaC on a radix-64 SWMR crossbar. Figure 2-left shows the performance of a multicore that is not subject to thermal constraints (always runs at 5 GHz), normalized against the performance of a topology that keeps the lasers always on (No-Ctrl). In this thermally-unconstraint case, LaC is only 6% slower on average than No-Ctrl, because LaC hides the laser turn-on delay by proactively turning the lasers on. At the same time, however, LaC saves between 62%-92% of the laser energy (77% on average). As the total power budget of the chip is a limited resource, the power that LaC saves can be used by the cores to increase their speed and performance. Thus, in a realistic power- and thermally-limited processor, LaC can enable performance improvements. Figure 2-right shows the performance of a realistic simulated processor that is subject to thermal constraints, with cores throttled to keep the chip within safe operating temperatures (under 90C). LaC improves the system's performance by 2.1x on average (running at 3.5 GHz) over No-Ctrl (running at 1.5 GHz). Moreover, LaC's laser energy savings are within 4% of the Perfect scheme's [2]. Overall, LaC improves the energy efficiency of the SWMR crossbar, making it a feasible and attractive alternative to other power-equivalent optical interconnects with narrower data paths (Power_Eq) or traditional concentrated-mesh electrical interconnects (CMesh).

4. CONCLUSION

In this paper we propose LaC, a laser-control mechanism that turns the laser off during periods of inactivity to save energy. LaC improve upon the EcoLaser [2] design by keeping the majority of the data bus off while sending small (data-less) messages and turning the laser on proactively when it anticipates long data messages. Compared to keeping the lasers always on (No-Ctrl), LaC saves 62-92% of the laser energy (77% on average) when running real-world applications. The energy savings and performance of LaC stay within 2-4% of a Perfect scheme [2] with future knowledge of interconnect accesses. More importantly, the power savings of LaC allow the cores to exploit a higher power budget and run 2-2.2x faster (2.1x on average) on a 64-core processor.

5. ACKNOWLEDGEMENTS

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

6. REFERENCES

[1] L. A. Barroso and U. Holzle. The case for energy-proportional computing. *IEEE Computer*, 40(12):33- 37, 2007.

[2] Y. Demir and N. Hardavellas. Ecolaser: An adaptive laser control for energy efficient on-chip photonic interconnects. In *Proceedings of the IEEE International Symposium on Low Power Electronics and Design*, ISLPED'14, August 2014. Pre-print available as Technical Report NU-EECS-14-02 at http://www.eecs.northwestern.edu/images/eecs-faces/2014_2_Eco-Laser.pdf.

[3] Y. Demir, Y. Pan, S. Song, N. Hardavellas, J. Kim, and G. Memik. Galaxy: A high-performance energy-efficient multi-chip architecture using photonic interconnects. In *Proceedings of the 28th ACM International Conference on Supercomputing*, ICS'14, June 2014. Pre-print available as Northwestern University Technical Report NU-EECS-13-08 at http://www.eecs.northwestern.edu/images/docs/techreports/2013_TR/NU-EECS-13-08.pdf

[4] 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):679-681, Mar 2010.

[5] 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)*, pages 317- 319, 2012.

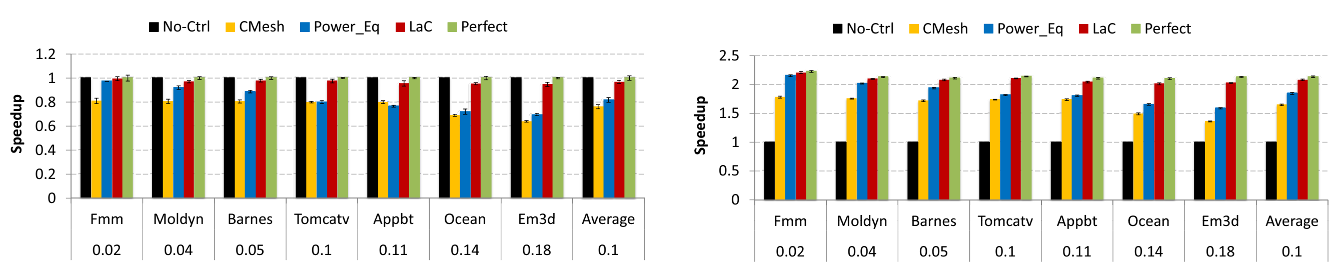

FIGURE 2. Speedup of No-Ctrl, CMesh, Power_Eq, LaC and Perfect at 5GHz (left) and on a power-limited multicore (right).