

Pho\$: A Case for Shared Optical Cache Hierarchies

Haiyang Han*, Theoni Alexoudi†, Chris Vagionas‡, Nikos Pleros† and Nikos Hardavellas‡

* ECE, Northwestern University † Aristotle University of Thessaloniki, Greece ‡ CS & ECE, Northwestern University
haiyang.han@u.northwestern.edu {theonial, chvagion, npleros}@csd.auth.gr nikos@northwestern.edu

Abstract—Conventional electronic memory hierarchies are intrinsically limited in their ability to overcome the memory wall due to scaling constraints. Optical caches and interconnects can mitigate these constraints, and enable processors to reach performance and energy efficiency unattainable by purely electronic means. However, the promised benefits cannot be realized through a simple replacement process; to reach its full potential, the architecture needs to be holistically redesigned.

This paper proposes Pho\$, an opto-electronic memory hierarchy architecture for multicores. Pho\$ replaces conventional core-private electronic caches with a large shared optical L1 built with optical SRAMs. A novel optical NoC provides low-latency and high-bandwidth communication between the electronic cores and the shared optical L1 at low optical loss. Our results show that Pho\$ achieves on average 1.41× performance speedup (3.89× max) and 31% lower energy-delay product (90% max) against conventional designs. Moreover, the optical NoC for core-cache communication consumes 70% less power compared to directly applying previously-proposed optical NoC architectures.

I. INTRODUCTION

It has been nearly 25 years since the performance gap between CPUs and main memory, or the “Memory Wall”, was identified as the main obstacle in increasing the performance of computer systems [1]. To mitigate the memory wall, stemming from the high latency of electronic memories and the limited bandwidth of electronic off-chip memory interconnects, modern chip multiprocessors (CMPs) have resorted to deep cache hierarchies. However, on-chip caches can occupy as much as 40% of the die area [2] and 32% of the processor’s power [3].

Alternatively, optical interconnects and nanophotonic technologies have emerged as promising yet underdeveloped solutions to tackle the disparity between processor and memory speeds. Optical Networks on Chip (NoCs) demonstrate higher bandwidth and energy efficiency than the traditional electronic NoCs used in CMPs [4]. Optically-connected memory (OCM) raises the possibility to switch much of the data transports between the processor and DRAM chips to the optical domain [5]. Optical Flip-Flops (FFs) in photonic crystal nanocavities (PhC) [6], [7] can form the building blocks of all-optical memory cells [8], which have demonstrated both speed and energy benefits over their electronic counterparts by boasting read/write speeds up to 40 Gbps [6], [9]. We appear to have all the ingredients to design novel optical cache architectures.

However, the application of an optical cache is not a simple plug-and-play replacement of its conventional electronic counterpart. While prior works [10], [11] have tried to explore this

topic, the proposed designs are infeasible for capacities larger than a few kB due to unrealistically high power consumption, and do not consider the challenges of interconnecting the electronic and optical domains. In this paper, we address the issues that arise with the introduction of such optical cache devices, and bridge the gap between device- and architecture-level designs. More specifically, our contributions are:

- For the first time to our knowledge, we make optical caches practical. We employ a cascaded two-level row decoder to reduce laser power, active rather than passive components to reduce off-ring optical losses, and use new technology for the optical bit cells that dramatically lowers the static power consumption.
- We propose Pho\$, an opto-electronic memory hierarchy for CMPs. Pho\$ replaces all the core-private levels of a conventional electronic cache hierarchy with a single-level shared L1 optical cache (split I/D) that utilizes PhC-based optical memory cells [6] operating at 20 GHz. Pho\$ enables for the first time L1 caches to be **high capacity** (multiple MB), **fast** (2-processor-cycle access time), and **shared** (obviating cache coherence).
- We propose Pho\$Net, a novel hybrid MWSR/R-SWMR optical NoC to connect processor cores with optical cache banks in Pho\$. Pho\$Net disaggregates the request/reply paths to reduce laser power, and co-arbitrates both subnets simultaneously through a novel arbitration protocol.
- We perform comprehensive modeling and evaluation of Pho\$’s performance, power, and energy characteristics. Pho\$ is up to 3.89× faster (1.41× on average) over a traditional electronic cache hierarchy, while achieving up to 90% lower energy-delay product (31% on average). Under realistic assumptions, the Pho\$Net optical NoC achieves up to 70% power savings compared to directly applying previously-available optical NoC architectures.

II. OPTICAL SRAM OPERATION

Figure 1 shows the layout of an 8 B direct-mapped optical cache with a 2 B cache line, 2-bit index, and 5-bit tags [10]. The index and tag bits are each encoded with two wavelengths.

Read/write operations are controlled by the RW and \overline{RW} signals. During a write to the cache, a RW signal representing a logical “0” activates the Write Access Gates (WAG) ① and allows the incoming $data$ bits ②, the tag bits ③, and their complements \overline{data} and \overline{tag} to enter the optical RAM bank ④. At the same time, \overline{RW} represents a logical “1”, blocking the Read Access Gates (RAG) ⑤ and preventing a read operation. In the case of a read, the RW and \overline{RW} signals are set to logical “1” and “0”, respectively. This allows the data from the RAM

This work was partially funded by NSF award CCF-1453853, and HFRI and GSRT through the ORION (grant 585) and CAM-UP (grant 230) projects.

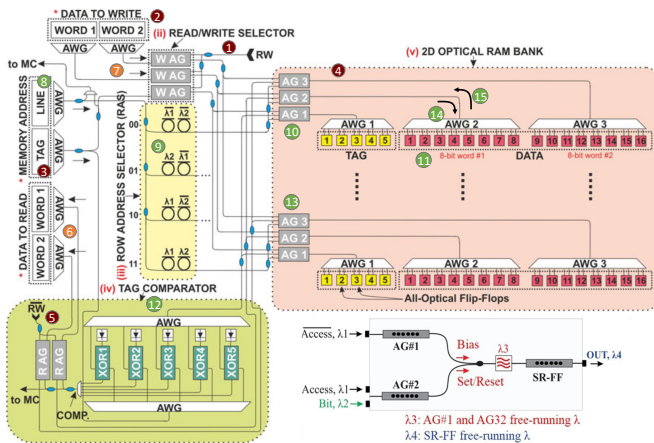


Fig. 1. 8B optical cache [10] and PhC nanocavity optical SRAM cell [7].




Fig. 2. (a) Packaged optical memory. (b) Monolithic InP Flip-Flop.

bank to propagate onto the data reply channel ⑥ and blocks the WAG to prevent any data from being overwritten ⑦.

The cache line to read or write is designated by the incoming *index* ⑧ and *index* bits which drive the passive Row Address Selector (RAS) ⑨. In Figure 1’s example, the RAS consists of 4 rows of two micro-rings (MRs) each. Each MR is tuned to a specific wavelength such that a pair of wavelengths λ_i and $\bar{\lambda}_i$ encode the logical “1” and “0” of the *i*-th bit of the *index*. The 2-bit *index* is encoded with 4 wavelengths: λ_1 , $\bar{\lambda}_1$, λ_2 , and $\bar{\lambda}_2$. The access gate (AG) ⑩ of the selected cache line now has a control signal of “0”, which allows either incoming data-to-write and tags to pass through to the optical Flip-Flops (FFs) for writing ⑪, or the contents of the FFs pass through to the tag comparator for reading ⑫. All other lines will have some wavelengths still propagating to their corresponding AGs, not activating them and blocking any data ⑬.

When the data and tag bits enter the optical RAM bank and propagate through the AGs in the row denoted by the index ⑩, the wavelengths are distributed to their corresponding optical FFs through Arrayed Waveguide Gratings (AWGs) ⑭. AWGs act as optical de-multiplexers that retrieve individual wavelengths from Dense Wavelength Division Multiplexing (DWDM) optical channels [12]. Each pair of wavelengths λ_i and $\bar{\lambda}_i$ drive the optical FF at the *i*-th bit in each 8-bit optical word. For a read, the AWGs multiplex the bits from the FFs into a single waveguide in the reverse direction ⑮.

We experimentally verified and characterized in our lab integrated photonic RAMs and optical FFs (Figure 2) which adopt the cross-coupled circuit-layout RAM cell architecture presented in Figure 1, and use technologies of optical gain elements integrated hybridly with InP PhC-on-SOI [7].

III. THE PHO\$ ARCHITECTURE

The optical cache prototype presented in Section II achieves very low latency. The optical SRAM cells can perform reads

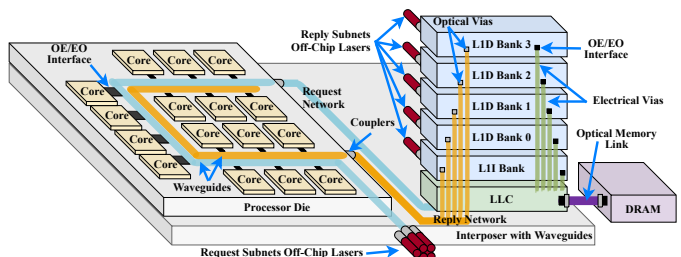


Fig. 3. Pho\$Net optical network topology.

and writes in under 50 ps, and the outside decoding processing time is 100 ps, resulting in 150 ps cache read and write latencies. As long as the core-to-cache optical bus takes no more than 50 ps, such an optical cache can perform single-cycle cache accesses for core frequencies up to 5 GHz. However, while the InP/Si PhC laser-based optical SRAM cells have fast on/off switching speeds, each cell requires a pump power of 103.5 μ W for writes [7]. Considering the number of components needed for a reasonably-sized cache, static power quickly reaches hundreds of Watts, which is unrealistic. Thus, prior designs [10], [11] are not implementable above 8 kB.

To avoid the additional pump power needed for biasing, Pho\$ instead utilizes the InGaAsP-based optical SRAM cells demonstrated by Nozaki *et al.* [6]. These cells require a static power of only 30 nW, and their switch-on latency of 44 ps is on par with the 50 ps latency of the InP/Si PhC laser, allowing cache reads to still be completed within one cycle at 5 GHz. Cache writes are slow at 7 ns, but this can be mostly mitigated by memory-level parallelism (MLP) and a modern core’s store queue. MLP allows for multiple concurrent memory requests, and store queues allow arithmetic operations and loads to bypass pending older writes. Thus, both MLP and store queues allow a core to overlap long write latencies with other work.

We propose Pho\$, an opto-electronic cache hierarchy architecture that replaces all the electronic L1D, L1I, and L2 caches in a traditional CMP with a single, shared, high capacity all-optical cache. We envision a shared optical L1D that employs 4 banks to provide high capacity and parallelism, and a shared optical L1I with one bank. The optical cache banks are fabricated on separate optical dies, while the processor cores remain on their original electronic die. The cores and optical caches are 2.5D-integrated on the same package and interconnected by an optical NoC, which handles arbitration and data transmission between the cores and the cache banks.

A. Pho\$Net Network Topology

Figure 3 shows a high-level view of Pho\$’s optical network topology, Pho\$Net. The electronic processor die on the left houses the cores and sits atop an interposer with photonic waveguides. The dies on the right are 3D-stacked. The L1D and L1I banks are on optical dies, while the Last Level Cache (LLC) is a traditional electronic cache with its own die. Each optical cache bank has one input and one output port.

Communication between the cores and caches is entirely in the optical domain. Two sets of optical waveguides are laid between the processor and L1 cache dies. Each waveguide line in the figure is abstracted to represent multiple sub-networks,




Fig. 4. Arbitration protocol. (a) Token circles the arbitration channel waiting to be grabbed. (b) Core 0 grabs the token and sends a request packet on the data channel. (c) Cache hit: reply packet sent, followed by a new token. (d) Cache miss: NACK sent, followed by a new token. (e) Cache has data following the miss, tries to grab token first. (f) Cache notifies core 0 with reservation channel, sends reply packet followed by a new token.

each comprising a bundle of waveguides with DWDM. The blue line depicts the subnets that carry requests from the cores to the cache banks (one subnet per bank). Within each request subnet, the cores are the writers and only one of the optical cache banks is the reader. Thus, each request subnet forms a Multiple-Writer Single-Reader (MWSR) crossbar [13] and uses token-based arbitration [14]. The orange line represents the reply subnets used by the cache banks to send data to the cores. For each reply subnet, one of the cache banks is the writer and the cores are the readers. Thus, the reply subnets are designed as Reservation-assisted Single-Writer Multiple-Reader (R-SWMR) crossbars [15]. In essence, Pho\$Net is a hybrid MWSR/R-SWMR optical network.

For a 16-core processor with 5 optical cache banks (as in Figure 3), there are in total 5 hybrid subnets, each comprising an MWSR request and an R-SWMR reply crossbar with arbitration and reservation channels, respectively. The request and reply subnets are powered by separate off-chip lasers to minimize laser power (Section III-D). Finally, the LLC can be connected to the DRAM through an optical interconnect [5] for low latency, high bandwidth DRAM accesses.

Core-private caches, as employed by traditional multicores, require core-to-core communication to maintain coherence, which in turn requires full-blown MWSR or R-SWMR crossbars with all-to-all connectivity. By employing an L1 cache that is shared among all cores, Pho\$ removes the need for cache coherency and inter-core traffic. Thus, it is no longer necessary to build physical links between cores. It suffices to implement separate networks for carrying either requests or reply packets directly to and from caches, and optimize each for their purpose. The hybrid Pho\$Net network capitalizes on this observation to shrink the network by avoiding full connectivity among all nodes, saving power, area, and cost.

B. Pho\$Net Arbitration Protocol

For each cache bank, all cores on the same request (or reply) subnet share the same channel, thus it is important to ensure that requests from different cores (or replies to different cores) do not conflict. As Pho\$Net is half MWSR and half R-SWMR (Section III-A), it requires a new way to arbitrate packets.

Arbitration in Pho\$Net is achieved through a protocol similar to optical token channel arbitration [14]. A single optical token circulates through each bank’s request-reply subnets. A core grabs the token, sends a request, and turns on its MRs to receive the reply. Upon an L1 hit, the cache injects the data to the reply network followed by a new token. Upon an L1

miss, the L1 injects the token back to allow future requests, along with a NACK reply so the requester can turn off its receiving MRs to minimize optical losses. Eventually, the L1 receives the data, waits for the token, signals the requesting core through the R-SWMR reservation channel to turn on its receiving MRs, and sends the data. Figure 4 shows an example arbitration in a simplified 3-core 1-cache-bank setup.

C. Pho\$ Optical Cache Architecture

Assuming a 64 B cache line, each of Pho\$’s five 1 MB direct-mapped cache banks has 16384 lines. Row decoding with an MR-based matrix, as in prior work [10], is impractical: the number of MRs needed for each line increases as the matrix scales up, consuming inordinate amounts of power. Instead, Pho\$ uses a two-level cascaded row decoding process. The first-stage de-multiplexing uses an active 9-to-512 tree global row selector, implemented with PhC nanocavity-based resonant switches [16], which activates only one of the 512 5-to-32 passive MR-based row decoders in the second stage. In this way, we build a 16384-line row selector with only 5 MRs per line instead of 14, drastically lowering laser power.

For the column decoding optical circuit, we use 8 1-to-128 AWGs to de-multiplex the wavelengths in the incoming light into their respective optical FFs. Each 1-to-128 AWG serves 64 bits, with 2 complementary channels per FF, so a total of 8 AWG-based column decoders are needed for a 64 B cache line. For each AWG, an AG controls the direction of data when switching between writing and reading the FFs. The AGs are controlled by 8 WAGs acting as read/write selectors. Data are fed into the reply waveguides through 8 RAGs (Section II).

D. Laser Power Sources and Optical NoC Parameters

The request subnet and the optical cells and reply subnet are powered by separate laser sources. The laser used to power the request subnet also powers the row decoders, column decoders, read/write selectors, and AGs before the optical FFs, because additional lasers along the path can overwrite any data already traveling on the waveguide. The token arbitration and reservation channels are also powered by the same laser. The FFs in the optical cache cells need a continuous power source to store data using photons, and the same laser can be used to power the tag comparators as well as the reply network.

We consider a comprehensive range of parameters for optical components by grouping the parameters of several seminal optical NoC designs from recent years [13], [17]–[22] into two groups, conservative and aggressive (Table I).

TABLE I
NANOPHOTONIC PARAMETERS FOR OPTICAL NOC.

Component	Conservative	Aggressive	Component	Conservative	Aggressive
Waveguide	1 dB/cm	0.05 dB/cm	Waveguide bending	0.005 dB	0 dB
Coupler	2 dB	1 dB	Waveguide crossing	0.12 dB	0.05 dB
Nonlinearity	1 dB	1 dB	Photodetector	0.1 dB	0.1 dB
Ring through	0.01 dB	0.001 dB	Modulator insertion	1 dB	0.001 dB
Filter drop	1.5 dB	0.5 dB	Detector sensitivity	-16 dBm	-28 dBm
Splitter	0.2 dB	0.1 dB	Laser Efficiency	30%	30%
Trimming	20 μ W/ring	5 μ W/ring	Modulation / Demod.	150 fJ/bit	20 fJ/bit

IV. METHODOLOGY

We evaluate Pho\$ using the Sniper simulator [23] running workloads from SPEC CPU2017 [24] (SPECspeed, ref inputs) and Parsec 3.0 [25] (simlarge inputs) benchmark suites. We compare our results with a baseline electronic multicore whose configuration is similar to a 16-core Intel Skylake (Table II).

To get an insight into the optical NoC’s power consumption, we compare our hybrid optical NoC, Pho\$Net, against three network configurations. The first is a fully-connected MWSR crossbar with 21 21-to-1 MWSR links (16 cores and 5 cache banks, a total of 21 nodes) with a token arbitration protocol. The second is a fully-connected R-SWMMR crossbar with 21 1-to-21 reservation-assisted SWMMR links. Finally, we also compare against a “one channel” network where requests and replies share the waveguides as a single data channel, while all other characteristics are the same as in Pho\$Net. For this comparison, we ignore the static power needed for optical FFs to operate as this depends on the number of cache components and not the network configuration. We model a 64- λ DWDM.

We estimate the energy consumption of cores, electronic caches, electronic on-chip interconnects, and DRAM using McPAT [26]. The energy consumption of the optical caches and Pho\$Net are calculated using in-house tools. As the request and reply subnet lasers power the passive optical cache components, the optical cache dynamic energy is categorized as part of the NoC. The overall optical cache static power is calculated by multiplying the number of active components with the static power of each component. We use the 30 nW reported by Nozaki *et al.* [6] as the static power needed for every optical FF. For Pho\$Net we model the best configuration determined by our design-space exploration (Figure 6). The NoC dynamic power accounts for the modulation/demodulation during the EO/OE conversions at the cores and LLC.

We model both a conventional DRAM for Pho\$, as well as an optically-connected one (Pho\$_OCM; see Table II).

V. EXPERIMENTAL RESULTS

A. Benchmark Performance

Figures 5a and 5b summarize the speedup of Pho\$ and Pho\$_OCM over the baseline running SPEC CPU2017 and Parsec 3.0. Figures 5c and 5d show the normalized CPI stacks [27], respectively. Each bar shows the relative values of cycles per instruction that are spent waiting for a particular component in the system. The “busy” sub-bar denotes the fraction of time spent within the core itself. For each application, the left, middle, and right bars represent the normalized CPI

TABLE II
SIMULATED SYSTEM PARAMETERS.

	Baseline	Pho\$	Pho\$_OCM
Cores	16 cores, x86 ISA, 3.2 GHz, OoO, 4 wide dispatch/commit 224-entry ROB, 72-entry load queue, 56-entry store queue		
L1 ICache	electronic, private, 64 B line, 32 kB/core, 8-way, 4 cycles	optical, shared, 64 B line, 1 MB direct-mapped, 2-cycle read, 23-cycle write	
L1 DCache	electronic, private, 64 B line, 32 kB/core, 8-way, 4 cycles	optical, shared, 4 banks, 64 B line, 4 MB direct-mapped, 2-cycle read, 23-cycle write	
L2	electronic, private, 64 B line, 256 kB/core, 4-way, 14 cycles	N/A	
LLC	electronic, shared, non-inclusive, 64 B line, 22 MB, 11-way, 50 cycles		
Core-L1 Netw.	electronic, point-to-point	hybrid optical	
LLC Network	electronic, 4 \times 4 mesh (NUCA)		
Memory	electrically connected, 49.37 ns		optically connected, 41.61 ns

stacks of baseline, Pho\$, and Pho\$_OCM, respectively. Pho\$ achieves an average speedup of 1.34 \times and 1.41 \times without and with OCM, respectively. For CPU2017, we see an improved execution time across all applications, with *cactuBSSN* having a maximum of 3.89 \times speedup. Pho\$ is able to significantly decrease instruction fetch delays because of its fast L1 read latency and large L1I capacity. Similarly, most applications enjoy a decrease in total L1D and L2 delay, like *leela* and *gcc_1*. The increased L1 capacity also means there are fewer misses that must visit the much slower LLC, and this is indicated by a reduced CPI for *mem-llc* in applications like *gcc*, *mcf*, and *xz*. The slow 7 ns L1 write time does not seem to have much adverse effect. OCM-enabled Pho\$ makes an impact in applications like *fotonik3d* and *lbm*, providing on average an additional 5% speedup across the suite.

For the multi-threaded workloads in Parsec, Pho\$ is able to speed up the execution of most applications, obtaining on average 1.37 \times speedup. Instruction fetch delays are greatly reduced, which is most prominent in *bodytrack* and *x264*. We find that Pho\$ does not suffer from high contention from a shared L1I cache. This is due to Pho\$ combining the aggregate capacity of the individual L1Is in baseline into a larger shared L1I, allowing more of the instruction stream to be L1-resident. Each fetched cache line also includes multiple instructions, eliminating the need for fetching on every cycle. The CPI component for L1D in Pho\$ and Pho\$_OCM is 44% lower on average than the CPI contribution of L1D+L2 in the baseline. The benefits of a low read latency and large capacity outweigh the disadvantage of a high write latency. Like in CPU2017, the large capacity of Pho\$’s L1 cache also results in fewer visits to the LLC and thus fewer stalls. For example, Pho\$ in *blackscholes* almost eliminates the CPI contribution of LLC and in *streamcluster* reduces it by about 4 \times . On average, Pho\$ decreases LLC delays by 2.5 \times . Adding OCM to Pho\$ reduces the average CPI spent waiting for DRAM by 2 \times and increases the overall speedup to 1.48 \times .

B. Optical NoC Power Analysis

Figure 6 shows the normalized optical power consumption of Full MWSR, Full R-SWMMR, Pho\$Net, and One Channel normalized to the Full MWSR configuration (normalized separately for the conservative and aggressive nanophotonic technologies). Our estimates include the power consumption of the off-chip laser, heating for MRs, and modulation/demodulation.

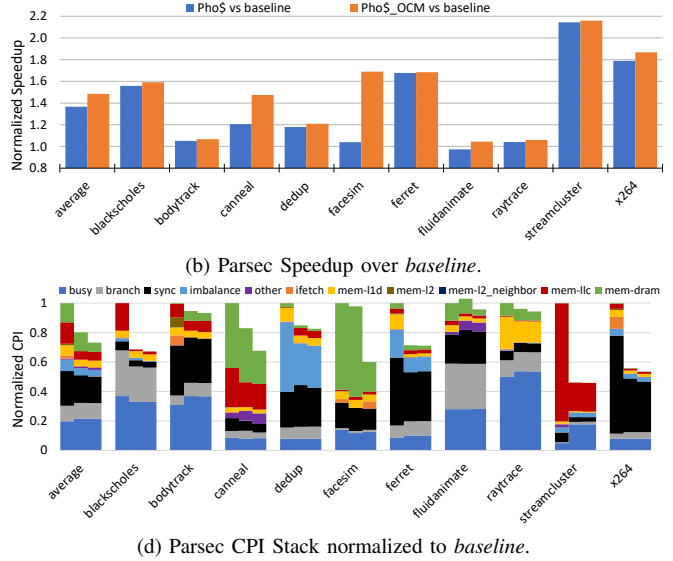
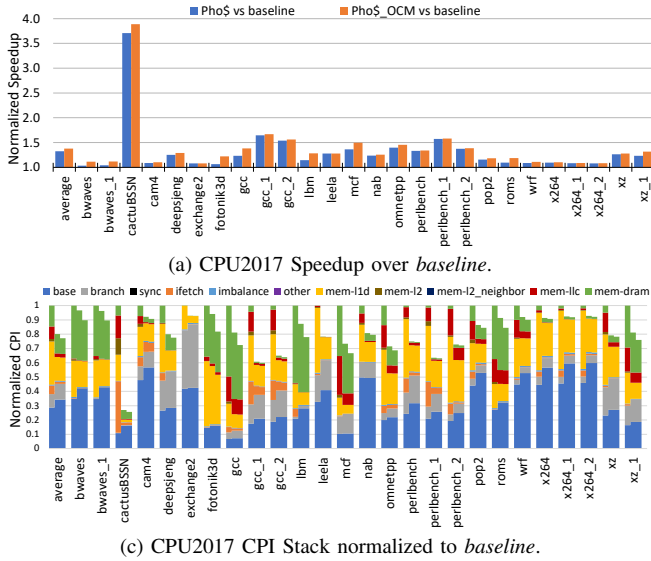


Fig. 5. Speedup and CPI Stacks for CPU2017 and Parsec; the three bars per benchmark correspond to baseline, Pho\$, and Pho\$_OCM.

Pho\$Net shows the lowest power consumption among alternatives. Under conservative nanophotonic parameters, laser power constitutes over 99% of optical power for all configurations. This is due to the high optical loss accumulated along the data path. For each waveguide with a DWDM of 64 wavelengths, 64 MRs need to be placed at each node for Full MWSR, Full R-SWMR, and Pho\$Net topologies (128 for One Channel). Pho\$Net gains an advantage over the other three topologies because it does not need to keep all nodes fully connected, requiring the fewest MRs along each datapath as well as the fewest data channels, thus reducing its total off-ring losses. However, the high optical loss per device under conservative technology parameters still results in unrealistically high power requirements. Even the most power-efficient Pho\$Net configuration under the highly-conservative nanophotonic parameters consumes 511 W for the network, requiring a 506 W laser power and 38 mW per wavelength.

We perform the same analysis using the aggressive photonic parameters. The optical loss for off-resonance rings decreases from 0.01 dB to 0.001 dB. As a result, the total laser power can be lowered to a reasonable level. Pho\$Net achieves the lowest optical power of 6.52 W, requiring 5.43 W for the laser, 0.94 W for ring heating, and 0.15 W for modulation/demodulation. Compared to the other designs, Pho\$Net still benefits from stripping off unnecessary links in the network and fewer rings on each waveguide. Having fewer MRs also means that the MR heating and modulation/demodulation power are reduced. As a result, Pho\$Net saves 70% of power compared to the two fully connected topologies and 16% compared to One Channel.




Fig. 6. Optical NoC power for a range of nanophotonic parameters.


Overall, the study using aggressive nanophotonic parameters gives us a very promising power consumption outlook with the lowest power consumption being under 7 W.

C. Energy Evaluation


Figure 7 shows Pho\$’s normalized energy per instruction (EPI, $J/insn$) and normalized energy \times delay product (EDP, $J \times s$). The three bars for each workload represent baseline, Pho\$, and Pho\$_OCM. Pho\$’s L1 static energy is considered to be the total pump energy needed for optical FF operations to be stable, and it is mostly on the same level with the combined L1 and L2 electrical static energy in the baseline. Note that Pho\$’s L1 dynamic energy is considered part of the NoC, as the NoC lasers also power the optical cache’s passive components. Pho\$ also has lower core and LLC energy consumption as there are less frequent core stalls and fewer LLC accesses. Overall, Pho\$_OCM saves on average 12% EPI and 31% EDP, and is most energy-efficient in applications such as *blackscholes*, *streamcluster*, and *cactuBSSN*.

D. Comparison with Previous Optical Cache Designs


A number of practical problems exist in previous optical cache designs from Maniotis *et al.* [10]. First, it relies on set-associative optical caches, but no optical cache designs are capable of set-associative replacement. Second, its high static power due to all-passive decoder and power-inefficient PhC cells [7] makes it impractical. Finally, its TDM-based optical bus requires the entire optical system to operate at 50–80 GHz, as 1 CPU cycle needs to correspond to 16 optical cycles. To the best of our knowledge this is currently unattainable for optical interconnects and optical memory [4], [8]. Figure 8 shows the performance (speedup) and energy comparison (*log scale*) between Pho\$ and Maniotis *et al.* [10], even under the assumption that the associativity and TDM challenges are resolved. Pho\$ is able to achieve a performance increase despite a slower writing speed, while maintaining a two orders-of-magnitude lower energy consumption.




(a) CPU2017 energy per instruction normalized to *baseline*.



(b) Parsec energy per instruction normalized to *baseline*.



(c) CPU2017 energy \times delay product normalized to *baseline*.



(d) Parsec energy \times delay product normalized to *baseline*.

Fig. 7. Normalized energy per instruction and energy \times delay product for CPU2017 and Parsec. For each benchmark, the three bars from left to right correspond to baseline, Pho\$, and Pho\$_{OCM}\$, respectively.




Fig. 8. Normalized speedup, energy per instruction (nJ/insn), and energy \times delay product (J \times us) of Maniotis *et al.*'s optical cache and Pho\$

VI. CONCLUSIONS

Recent discoveries of new materials and research on optical SRAM cells enable us to build fast, low-power optical cache architectures. In this paper we propose Pho\$, an optoelectronic memory hierarchy architecture for multicores. Pho\$ replaces private electronic L1 and L2 caches with a large shared optical cache, and on-chip electronic mesh networks with a novel optical NoC that uses a unique network arbitration protocol. We estimate that Pho\$ is on average $1.41 \times$ faster and 31% more energy-efficient (in terms of EDP) over purely-electronic designs. Pho\$'s network design, Pho\$Net, consumes 70% less power than previously-proposed optical NoCs. We also solve a number of problems that make previous optical cache designs impractical, achieving a performance lead and two orders-of-magnitude lower energy consumption.

REFERENCES

- W. A. Wulf and S. A. McKee, "Hitting the memory wall: implications of the obvious," *ACM SIGARCH Comp. Arch. News*, vol. 23, 1995.
- S. Borkar and A. A. Chien, "The future of microprocessors," *CACM*, vol. 54, 2011.
- R. Sen and D. A. Wood, "Cache power budgeting for performance," Univ. of Wisconsin-Madison, Computer Science Dept., Tech. Rep., 2013.
- S. Werner, J. Navaridas, and M. Luján, "A survey on optical network-on-chip architectures," *ACM Computing Surveys*, vol. 50, 2017.
- C. Batten *et al.*, "Building many-core processor-to-dram networks with monolithic cmos silicon photonics," *IEEE Micro*, vol. 29, 2009.
- K. Nozaki *et al.*, "Ultralow-power all-optical ram based on nanocavities," *Nature Photonics*, vol. 6, 2012.
- T. Alexoudi *et al.*, "Iii-v-on-si photonic crystal nanocavity laser technology for optical static random access memories," *IEEE J. Sel. Topics Quantum Electron.*, vol. 22, 2016.
- T. Alexoudi, G. T. Kanellos, and N. Pleros, "Optical ram and integrated optical memories: a survey," *Light: Science & Applications*, vol. 9, 2020.
- C. Vagionas *et al.*, "XPM- and XGM-based optical RAM memories: Frequency and time domain theoretical analysis," *IEEE J. Quantum Electron.*, vol. 50, 2014.
- P. Maniotis *et al.*, "Optical buffering for chip multiprocessors: a 16ghz optical cache memory architecture," *J. Lightw. Tech.*, vol. 31, 2013.
- P. Maniotis *et al.*, "An optically-enabled chip-multiprocessor architecture using a single-level shared optical cache memory," *Optical Switching and Networking*, vol. 22, 2016.
- C. Vagionas *et al.*, "Optical ram row access and column decoding for wdm-formatted optical words," in *Natl. Fiber Opt. Eng. Conf.*, 2013.
- D. Vantrease *et al.*, "Corona: System implications of emerging nanophotonic technology," in *ISCA*, 2008, pp. 153–164.
- D. Vantrease *et al.*, "Light speed arbitration and flow control for nanophotonic interconnects," in *MICRO*, 2009.
- Y. Pan *et al.*, "Firefly: Illuminating future network-on-chip with nanophotonics," in *ISCA*, 2009.
- K. Nozaki *et al.*, "Ultralow-energy and high-contrast all-optical switch involving fano resonance based on coupled photonic crystal nanocavities," *Optics express*, vol. 21, 2013.
- Y. Pan, J. Kim, and G. Memik, "FeatherWeight: low-cost optical arbitration with qos support," in *MICRO*, 2011.
- Y. Demir *et al.*, "Galaxy: A high-performance energy-efficient multi-chip architecture using photonic interconnects," in *ICS*, 2014.
- V. K. Narayana *et al.*, "Morphonoc: Exploring the design space of a configurable hybrid noc using nanophotonics," *Microprocessors and Microsystems*, vol. 50, 2017.
- P. K. Hamedani, N. E. Jerger, and S. Hessabi, "Qut: A low-power optical network-on-chip," in *NOCS*, 2014.
- S. Werner, J. Navaridas, and M. Luján, "Efficient sharing of optical resources in low-power optical networks-on-chip," *IEEE J. Opt. Comm. Netw.*, vol. 9, 2017.
- C. A. Thraskias *et al.*, "Survey of photonic and plasmonic interconnect technologies for intra-datacenter and high-performance computing communications," *IEEE Commun. Surveys Tuts.*, vol. 20, 2018.
- T. E. Carlson *et al.*, "An evaluation of high-level mechanistic core models," *ACM Trans. Arch. Code Opt.*, 2014.
- "Spec cpu 2017." [Online]. Available: <https://www.spec.org/cpu2017/>
- C. Bienia, "Benchmarking modern multiprocessors," Ph.D. dissertation, Princeton University, January 2011.
- S. Li *et al.*, "McPAT: An integrated power, area, and timing modeling framework for multicore and manycore architectures," in *MICRO*, 2009.
- W. Heirman *et al.*, "Using cycle stacks to understand scaling bottlenecks in multi-threaded workloads," in *IISWC*, 2011.