

Architecting a Computer with a Full Optical RAM

Jorge Gonzalez, Lois Orosa and Rodolfo Azevedo
Institute of Computing, University of Campinas, Brazil
{jorge.gonzalez,lois.rosa,rodolfo}@ic.unicamp.br

Abstract—On-chip photonics has gained attention in research for high-speed processor communication networks, and recent developments in optical fabrication techniques and data buffering has offered new opportunities for processor systems. In this work, we evaluate a processor with a full optical main memory system. We design it using recent optical devices that leverages the high-bandwidth optical capabilities to obtain low memory access latency, similar to those in state of the art L2 caches. This characteristic enables the possibility of eliminating the second level of caches, saving processor area. Experimental results show the average speedup is $\times 1.34$ with SPEC2006 and $\times 1.80$ with irregular applications.

I. INTRODUCTION

Photonics have been used successfully in long link communications for many years due to its intrinsic characteristics such as distance-independent low-power dissipation, high communication bandwidth and low crosstalk. With recent advance in device fabrication processes, the monolithic integration of silicon photonics modules is becoming feasible [1]. ITRS [2] roadmap highlights the potential of on-chip silicon photonics interconnections, and other recent roadmaps [3] point out the promising scalability expectations in the next two decades to achieve a full optical device integration in a single chip. Despite optical fabrication technology is still not mature, several small-scale integrated silicon photonic devices have been recently fabricated [4]–[6].

Design space exploration for optical on-chip devices is a growing research topic in computer architecture. The most popular research direction is Optical Network-on-Chip (ONoC), which uses wavelength routing mechanisms to communicate multiple IP-blocks within a chip under different topologies and arbitration schemes [7].

Another emerging photonic research topic is optical memory fabrication. Our work focuses on the co-design of processors and these memories. Optical memories leverage different techniques to achieve data buffering in optical memory cells. To date, a 1-bit memory cell can be implemented as a Semiconductor Optical Amplifier (SOA) based flip-flop [8], and a wider cell is feasible using emergent nanocavity technologies [5].

We propose a new architecture for a single core processor with an ONoC and an optical RAM (o-RAM) that could overcome the limitations of current electrical memory subsystems. Unlike other approaches using an ONoC and electrical memory [9], our proposal does not require optoelectrical conversion on the memory side, reducing its overhead. Furthermore, the o-RAM operation latency is in the order of picoseconds (the carrier is in THz).

With this approach, we tackle one of the greatest challenges in computer research, the memory wall. The disparity between the performance of current DRAM technology and the processor data access requirements is growing, causing the development of new memory devices, such as optical-based

memory cells, which is in early development stage, or PCM and memristor, which are relatively mature. This work seeks to address how to architect a reasonable sized o-RAM (256 MB) and its communication link based on state-of-the-art photonic devices such as transceivers, switches and memory cells. We envision a computer system with our memory proposal, and we evaluate it using microarchitectural execution-driven simulation. We define the characteristics of the optical memory system by inferring parameters from reported results, reports, and trends in photonic research.

Our exploration (Section II-D) indicates o-RAM low operation latencies between the values of a current first level caches. As a result, we obtain speedup compared to a conventional electric system due to the low latency for memory operations.

Contributions: 1) We explore key design considerations of o-RAM using state-of-the-art devices with a wavelength routing method. 2) We propose a novel architecture with a full optical memory subsystem able to get better performance than a traditional electric system, additionally reducing chip area.

II. BUILDING A FULL OPTICAL RAM

In our proposed architecture, the optical communication between processor and o-RAM uses the Wavelength Division Multiplexing technique (WDM), where a laser source can carry multiple wavelengths (λ). Although in this work we consider an external continuous wave light source, laser research shows promising energy efficient devices for on-chip integration [10].

We use modular transceivers, switches, and memory cells to architect our optical subsystem, as shown in Figure 1. All considered devices are silicon based for its potential low-cost integration with current CMOS devices [1]. The following subsections describe these devices.

A. Optical DWDM Transceiver

The processor outputs are electrical signals that need to be serialized and converted to the optical frequency via modulation of laser light. There are two distinct approaches based on SERDES (Serializer/Deserializer) for building optical transceivers: 1) to encode/decode all parallel bits into a single modulated wavelength. 2) a single modulated wavelength for each bit through an optical bus, such as DWDM (Dense Wavelength Division Multiplexing) [6].

DWDM allows inserting multiple wavelengths into a single fiber or waveguide. Each wavelength becomes an independent communication channel, and each bit line (n-bit link) from the processor could be directly modulated using a DWDM transceiver (n- λ link). Notice also that converting a multiple wire bus into a single waveguide bus coupled to the chip could minimize the electric pin problem [11].

In our design, we considered one of the first reports of integrated DWDM transceiver fabrication [6]. This device is based on $5\mu\text{m}$ radius microrings achieving 10 Gbps per chan-

nel supporting 24 channels. The fabricated device implements 5 channels and could be used as a voltage driven carrier-injection microring modulator for Tx or a bias-tuned microring coupled with a photodetector for Rx. Device channel outputs are combined using the DWDM technique.

To build our system, we require a 128 channel transceiver to communicate with the 128 bit memory cell (Section II-C). While optical fiber devices with more than 128 channels are well known in metro and long-haul communications, it is feasible that future integrated devices could achieve a higher number of channels.

B. Optical Switches (o-SW)

An optical switching device (o-SW) has analogous behavior to an electrical crossbar. Considering that an optical link has circuit switching routing, the target path must be established before signal propagation. Two of the most common approaches are: 1) cascaded micro-resonators as switches [12], [13], such as micro-rings. 2) cascaded electro-optical switches, such as Mach-Zehnder Interferometers (MZI) based switches [4] and plasmonic switches [14].

We selected electro-optical o-sw for our 128 bit wide design. Whereas ring resonators have lower power consumption and smaller footprint area, MZI and plasmonic based o-sw have larger bandwidth allowing more wavelengths (λ) to pass, and they are also less sensitive to temperature variations.

We use 2×2 switches that requires low driving voltage as our baseline, where it has two output ports: drop (ON) and through (OFF) and an electrical control input for selecting the output port. A 2×2 MZI o-SW [4] was fabricated using Silicon-on-insulator (SOI) wafers with an optical bandwidth of 110 nm and a measured switching delay of $< 4ns$. A plasmonic switch [14] was reported with a 100 nm bandwidth and expected ps switching delay (Section V).

C. Optical Memory Cells

We consider two optical buffering approaches, both are volatile: 1) optical flip-flop implementation composed by SOAs [8], and 2) photonic crystal nanocavity memory cell [5].

Unlike the flip-flop approach, that can only store one λ per cell (1 bit), the nanocavity approach can store multiple λ . Furthermore, the nanocavity cell control is simpler and easy to scale, the reason we choose it for our design.

A recent work [5] shows the implementation of a DWDM capable 128-bit memory cell in a *Si* photonic crystal, where the authors showed clear optical nonlinearity without overlapping in 105 of the 128 nanocavities.

Three operations are performed on the nanocavity cell varying the intensity of the optical bias input: read, write and reset. The write and read operations are performed in around ≈ 100 ps. This device has a waveguide that traverses its structure, where its both ends are the input and output ports. In the write case, λ is stored in the nanocavity cell, and in the read case, the content of the nanocavity cell is recovered in the output. In our specific case, where we are using DWDM with 128 λ (128 bits), the read-write operations are performed by optical pulses of the order of ps differing only in signal intensity. We designed the o-RAM bank as a set of these optical nanocavity cells.

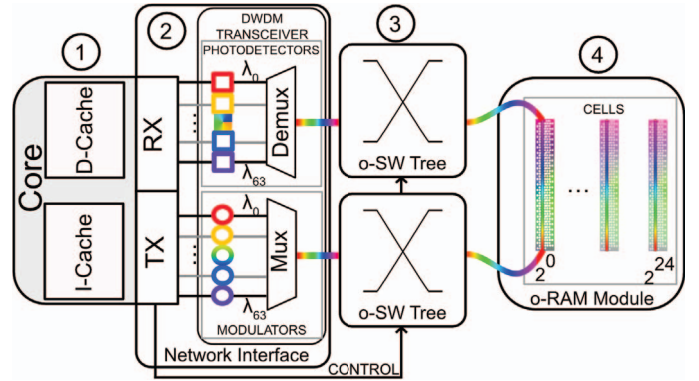


Fig. 1: Overview of a processor with an o-RAM system model: 1) a single core with L1 caches, 2) a DWDM transceiver based network interface, 3) an optical link based on o-SW and 4) an o-RAM bank.

D. Putting It All Together

In this section, we architect a system considering an electro-optical switch (o-SW) interconnection and an o-RAM with DWDM transceivers, and nanocavity cells.

Our ONoC design relies on a tree topology to perform direct access to all memory cells by routing the light beam, where the tree branches are cascaded electro-optical o-SW (Section II-B). Figure 3 shows the o-SW topology as a complete binary tree graph where the nodes are the 1×2 o-SW, and its control relies on electrical signals. The number of memory cells is equal to the number of the tree leaves, and with this configuration, we can address a total of 256 MB. The binary tree has $height = 24$, and its total number of switches is $(2^{24} - 1)$, thus, processor address only requires 24 bits.

Figure 2 shows the timing diagram for the read/write operations assuming a 2 GHz processor clock, a memory operation latency below 1 processor cycle and a switching delay of 2 ns (4 cycles). Despite the state-of-the-art MZI switching delay is 4 ns, this is a conservative assumption, considering the ps switching latency implementations (Section II-B). Then, we consider three switching latencies: 8 (state-of-the-art MZI o-sw), 4 (conservative estimation) and 2 (optimistic future),

As result, the total o-RAM access latency is 7 cycles for the conservative case. In the cycle 1, the processor makes a memory request, and the network interface (NI) modifies the optical switch states to set the path to the requested data address (Addr). Cycle 2 contains the decoding address delay. The NI sets the path in 4 cycles (from cycle 3 to 6), as all the switches are activated simultaneously in a 2 ns electrical stimuli. In cycle 7, as result of picoscale operation of the optical cells: a) in the case of a store, the data is modulated with the transceiver and stored in the memory cell, or b) for memory reads, the memory cell outputs the data stored which moves to an electrical buffer.

III. SYSTEM LEVEL ARCHITECTURE

O-RAM effective read/write latencies are one order of magnitude lower than conventional electrical DRAMs and in the same order than on-chip caches. This allowed us to reassess the memory hierarchy of electrical caches with equal or higher latencies. We propose to only use L1 cache with the o-RAM, this could reduce the processor area since the L2 cache is




Fig. 2: Optical memory system read/write timing diagram.

approximately one-half of the die’s silicon area.

The processor communicates with the o-RAM using a Network Interface (NI) composed by the array of its electrical output (Tx) and input (Rx) pins, and the DWDM transceiver (Figure 1 ②). Each Tx’s and Rx’s pin is directly modulated or demodulated with the DWDM transceiver setting a high-bit-rate data link.

The physical link is composed by two o-SW trees as shown in Figure 1 ③, which are controlled to establish a direct access to the o-RAM by the NI. Due to the system circuit switching nature, the NI can only handle one in-flight request (serialized access). Therefore, o-RAM access could cause contention, increasing the memory latency. Then, read/write instructions would have to stall until previous o-RAM operations have finished. One way to minimize the contention is to increase the number of paths to each o-RAM cell, at the cost of higher area footprint. The set of paths to an o-RAM cell defines the total number of o-RAM ports. For example, in Figure 1 ④ an o-RAM cell has a single path. For an o-RAM implementation with two ports, we need to duplicate the number of NI and o-SW trees structures.

There are two o-SW trees in our single port design, one before the o-RAM cell to perform write operations, and the other after for read operations. When a write operation is performed a path is set from the processor NI Tx to the o-RAM, modulating the data on the initiator and storing it on the target o-RAM cell structure. However, for a read operation, a closed loop is established. This is because a path is set using both o-SW trees, one from processor NI Tx to o-RAM, and another from o-RAM to processor NI Rx. A Read operation is summarized as: 1) a modulated light beam with a defined intensity travels through the first path 2) traverses the o-RAM cell, and then the light beam carries the data, 3) the light beam travels from the cell output through the second path, 4) and finally reaches the processor NI Rx.

IV. EVALUATION

To assess the performance of our processor/o-RAM architecture we used a modified version of ZSIM simulator [15]. We use SPEC2006 benchmarks for evaluation, using the Simpoint [16] methodology. We also tested a set of irregular applications which includes a Page Rank algorithm implementation and a Random Memory Access application [17].

We considered two systems: an electrical two-level cache and an optoelectrical system as depicted in Section II-D. Both platforms have a single 2 GHz x86 core. The electrical system is a conventional electrical system with a 64 KB L1 data cache, a 64KB L1 instruction cache, and a 2MB L2 cache. The L1 and L2 caches have 4 and 7 latency cycles, respectively. Its main memory is a DDR3-1066-CL8 model.

The optical system has $M \times$ (o-RAM banks) with P ports each, 64 KB L1 data cache and a 64KB L1 instruction




Fig. 3: Optical switch (o-SW) topology for an 256 MB o-RAM.

cache. We evaluate three different o-RAM access latencies: 11 (state-of-the-art), 7 (conservative) and 5 (optimistic) cycles, as previously discussed in Section II-D. We use an interleaved configuration with $M = 1, 2$ and 4 memory banks. Each memory bank was modeled with $P = 1, 2$ and 4 ports.

Figure 4 reports the geometric mean speedup with SPEC2006 and irregular applications. Results are grouped by the number of modeled ports and memory banks and normalized to the electrical baseline platform detailed previously in this section. Each bar has 3 levels for the 11, 7 and 5 cycles access latencies evaluated, where the first bar shows the results with SPEC2006, and the other bar was obtained with irregular applications. Bars (A) and (B), are the results of a system with a L1 I-cache and D-cache with an o-RAM.

All cases obtained better performance than the electrical case, because its lower access latency. The speedup is up to %38 and %84 with SPEC2006 and irregular applications respectively. Furthermore, our experiments show promising results with irregular applications. As detailed in Fig. 4, the o-RAM system obtained and speedup up to 84% (B).

The speedup is higher when the grade of memory level parallelism increases due to a higher number of o-RAM banks or the number of ports. O-RAM system presents contention on its ports, as discussed in Section III. Fig. 5 shows the percentage of times that a port P is busy when an operation is required, in the case of an o-RAM with 7 cycles access latency with SPEC2006. With $(2P \times 2M)$ and $(4P \times 1M)$ the contention is 48.6% and 42.1% respectively. Both obtained an average data latency of 7 cycles. Port contention has a direct effect on the latency data access. The $2P \times 2M$ and $4P \times 1M$ configurations have a good balance between area and performance, where both cases obtain an $\approx \times 1.36$ speedup.

V. DISCUSSION

Implementation feasibility of our proposed architecture is related to the development of integrated photonics devices (DWDM transceiver, o-sw and o-RAM cell). In our design, we propose to use millions of o-sw for processor communication with a 256 MB o-RAM. Reported MZI and plasmonic o-sw have a low switching time that allows for low access time to o-RAM. On the other hand, other electro-optical devices [18] focus on low rise and fall times (in the ps scale), at the cost of a narrower bandwidth. Although we are not considering switching delays in the order of picoseconds in our design, this last work is a good indicator that lower switching delays could be achieved.

Recent work using plasmonic materials show promising results for area reduction, due to its high-integration and compatibility with CMOS processes. An MZI [4] o-sw has an area of $\approx 0.02 \text{ mm}^2$, then the o-sw tree area will be 0.3 m^2 . Using plasmonic o-sw [14], each with $4.8 \text{ } \mu\text{m}^2$, the total required area will be 80.5 mm^2 . This calculation also represents control circuitry area because of the stacked electro-




Fig. 4: Overall Speedup of processor with a L1 cache and o-RAM. On each set, the first bar (A) represent SPEC2006 and the second one (B) irregular applications.

optical fabrication. In [19] the authors fabricated an atomic scale plasmonic switch, this enables research for future low footprint o-sw devices.

Moreover, our design relies on cascaded electro-optical switches which cause loss in the transmitted signal. For a 24 height o-sw tree (section II-D), using plasmonic o-sw [14] the losses will be up to 57.6 dB, and using MZI [4] o-sw up to 69.6 dB. For our system implementation, we need intermediate stages of amplification between the levels of the o-sw tree which could significantly increase the final area. However, techniques and methods for obtaining reasonable energy and losses with plasmonic materials are open issues that have gained considerable attention in the optical community [20]. While there are technology challenges, we expect that future device fabrication techniques will alleviate this trade-off. Our goal is to show the potential impact of this technology and stimulate further studies at the physical level.

VI. RELATED WORK

In [21] the authors present a 1.48 cm^2 full optical SRAM chip, based on previous work on SOA based FF o-RAM cells [8] with a row/column wavelength selector. Despite the capacity of the chip is still a few bytes, in [22] they show great progress in high integration of optical device systems envisioned for processors. Unlike our work, the authors proposed an optical L1 cache replacement, whereas our proposal architect an o-RAM as main memory.

In [23] a PCM-optical memory cell was introduced. This device could perform photonic-multilevel bit store, up to 8 levels because of the material intermediate states. This work proved the feasibility of a non-volatile photonic memory cell implementation with a large potential for high integrability and high operation frequency.

A electrical RISC-V processor and 1 MB electrical SRAM were integrated on a single chip with an optical link for communication between them [24]. The authors used standard microelectronics fabrication processes, becoming one of the first works of the new electronic-photonic on-chip era.

VII. CONCLUSIONS

This work presented the design of a full optical memory subsystem (o-RAM) as an alternative for main memory in conventional processors. The proposed architecture exploits the high bandwidth and low-latency characteristics of photonic devices, specifically DWDM transceivers, electro-optical switches and nanocavity memory cells. We design it under a tree topology for direct serialized access interconnection obtaining an average speedup of 34% over a conventional electrical system with SPEC2006.




Fig. 5: O-RAM port contention with SPEC2006 caused by access serialization on 7-cycle access latency

On-chip photonic fabricated devices had notorious development in recent years. Despite this technology is still immature and area is not yet reasonable for computer systems, it brings new opportunities for custom o-RAM system devices and explore the trade-off between area, energy and performance.

REFERENCES

- [1] X. Chen *et al.*, "Device engineering for silicon photonics," *NPG Asia Mater.*, 2011.
- [2] ITRS Interconnect. (2007). [Online]. Available: {<http://www.itrs.net/>}
- [3] MIT Communication Technology Roadmap. [Online]. Available: {<https://mphotronics.mit.edu>}
- [4] J. Van Campenhout *et al.*, "Low-power, 2x2 silicon electro-optic switch with 110-nm bandwidth for broadband reconfigurable optical networks," *Opt. Express, OE*, 2009.
- [5] E. Kuramochi *et al.*, "Large-scale integration of wavelength-addressable all-optical memories on a photonic crystal chip," *Nature*, 2014.
- [6] T.-C. Huang *et al.*, "DWDM nanophotonic interconnects: toward terabit/s chip-scale serial link," in *2015 IEEE 58th MWCAS*, 2015.
- [7] C. Batten *et al.*, "Designing Chip-Level Nanophotonic Interconnection Networks." *IEEE J Emerg Sel Topics Circuits Syst*, 2012.
- [8] N. Pleros *et al.*, "Optical Static RAM Cell," *IEEE Photon. Technol. Lett.*, 2008.
- [9] P. Grani, "From hybrid electro-photonic to all-optical on-chip interconnections for future CMPs." *HPCS*, 2014.
- [10] J. Pu *et al.*, "Heterogeneously integrated III-V laser on thin SOI with compact optical vertical interconnect access," *Opt. Lett., OL*, 2015.
- [11] Z. Wang *et al.*, "Improve Chip Pin Performance Using Optical Interconnects," *IEEE Trans VLSI Syst*, 2015.
- [12] A. W. Poon *et al.*, "Cascaded Microresonator-Based Matrix Switch for Silicon On-Chip Optical Interconnection," *Proc. IEEE*, 2009.
- [13] N. Ophir *et al.*, "Silicon Photonic Microring Links for High-Bandwidth-Density, Low-Power Chip I/O," *Micro, IEEE*, 2013.
- [14] C. Ye *et al.*, "A compact plasmonic MOS-based 2x2 electro-optic switch," *arXiv.org*, 2015.
- [15] C. K. Daniel Sanchez, "ZSim: Fast and Accurate Microarchitectural Simulation of Thousand-Core Systems," *ISCA*, 2013.
- [16] E. Perelman *et al.*, "Picking statistically valid and early simulation points," *PACT-03*.
- [17] S. Lloyd and M. Gokhale, "In-Memory Data Rearrangement for Irregular, Data-Intensive Computing," *Computer*, 2015.
- [18] L. Lu *et al.*, "Low-power 2x2 silicon electro-optic switches based on double-ring assisted Mach-Zehnder interferometers," *OL*, 2014.
- [19] A. Emboras *et al.*, "Atomic Scale Plasmonic Switch," *NL*, 2016.
- [20] J. B. Khurgin, "How to deal with the loss in plasmonics and metamaterials," *Nature Publishing Group*, 2015.
- [21] T. Alexoudi *et al.*, "WDM-enabled optical RAM and optical cache memory architectures for Chip Multiprocessors," 2015.
- [22] N. Pleros *et al.*, "Optical interconnect and memory technologies for next generation computing," in *2016 18th ICTON*, 2016.
- [23] C. Rios *et al.*, "Integrated all-photonic non-volatile multi-level memory," *Nature Photonics*, 2015.
- [24] C. Sun *et al.*, "Single-chip microprocessor that communicates directly using light," *Nature*, 2015.