COMET: A Cross-Layer Optimized Optical Phase-Change Main Memory Architecture

Febin Sunny*, Amin Shafiee*, Benoit Charbonnier†, Mahdi Nikdast*, and Sudeep Pasricha*
Electrical and Computer Engineering Department, Colorado State University, Fort Collins, USA*,
Universit´e Grenoble Alpes, CEA-Leti, Grenoble, France†,
febin.sunny@colostate.edu, amin.shafiee@colostate.edu, benoit.charbonnier@cea.fr, mahdi.nikdast@colostate.edu,
sudeep@colostate.edu

***Abstract*— Traditional DRAM-based main memory systems face several challenges with memory refresh overhead, high latency, and low throughput as the industry moves towards smaller DRAM cells. These issues have been exacerbated by the emergence of data-intensive applications in recent years. Memories based on phase change materials (PCMs) offer promising solutions to these challenges. PCMs store data in the material's phase, which can shift between amorphous and crystalline states when external thermal energy is supplied. This is often achieved using electrical pulses. Alternatively, using laser pulses and integration with silicon photonics offers a unique opportunity to realize high-bandwidth and low-latency photonic memories. But to realize photonic memories, several challenges that are unique to the photonic domain such as crosstalk, optical loss management, and laser power overhead must be addressed. In this work, we present *COMET*, the first cross-layer optimized optical main memory architecture that uses PCMs. In architecting *COMET*, we explore how to use silicon photonics technology and PCMs together to design a large-scale main memory system while exploring related challenges and proposing solutions at the PCM cell, photonic memory circuit, and memory architecture levels. Based on our evaluations, *COMET* offers 5.1× better bandwidth (BW), 12.9× lower energy-per-bit (EPB), and 65.8× better BW/EPB than the best-known prior work on photonic main memory architecture design.**

Keywords—phase change material, silicon photonics, main memory.

I. Introduction

In recent decades, the surge in big data and machine learning applications, including large language models [1], intrusion detection systems [2], and graph processing [3], [4], has led to a demand for terabyte range data storage and memory bandwidths that exceed terabytes-per-second (TB/s). Traditional electronic memory technologies, such as dynamic random-access memory (DRAM) are struggling to meet the increasing bandwidth and capacity demands in an energy-efficient manner [4]. Scaling DRAM technology down below the 12–14 nm nodes (e.g., Micron's 1α/1β nodes) creates charge retention issues, structural integrity concerns, and substantial delay and power penalties associated with bit lines [4]. Although 3D-stacking can help, the growing need for capacity, throughput, and energy efficiency is driving the search for new main memory technologies.

Non-volatile memories (NVMs) offer solutions to DRAM data retention challenges and eliminate the need for refresh cycles and related latency issues. However, NVMs based on ferroelectric (FRAM) [5] and resistive metal oxide (RRAM) [6] technologies often face reliability and write endurance problems. For improved reliability without sacrificing NVM advantages, phase change materials (PCMs) are a promising alternative [7]. PCM cells offer greater energy efficiency, bit density, and bandwidth compared to other NVM types [11], [8]. PCMs can switch between amorphous and crystalline states, providing high-resistance contrast for data storage. Electrically controlled PCM (EPCM) cells achieve state transitions using current pulses, with potential for multi-level cells (MLCs). However, PCM resistance depends non-linearly on write voltage [9], making intermediate states challenging to achieve and due to the resulting resistance drift, limits electrical PCM bit density to a few bits per cell [10].

One solution to these EPCM limitations is optically controlled PCM (OPCM) cells, where PCM layers are deposited on photonic waveguides (e.g., silicon-on-insulator (SOI)). In OPCM cells, laser pulses induce state transitions by heating the material, altering the refractive index contrast between amorphous and crystalline states, enabling optical data storage and retrieval. High contrast PCM candidates can help realize multi-level OPCM cells. Furthermore, OPCM memory can leverage high-bandwidth silicon photonic links for data transfer and support emerging optical computing [12], enabling high-speed, energy-efficient fully photonic computing systems with minimal electro-optic conversions.

In this paper, we introduce *COMET*, the first cross-layer optimized optical PCM-based main memory architecture. Our novel contributions include:

- We design a low-loss and energy-efficient silicon photonic PCM-based multi-level memory cell as a basic building block;
- We design and optimize a novel all-optical, loss-aware silicon photonic PCM-based photonic main memory architecture;
- We present detailed comparison of our designed photonic main memory architecture against state-of-the-art electronic and photonic main memory architectures;
- Lastly, we perform a case study aimed at showcasing the effectiveness of *COMET* within a photonic AI accelerator.

II. Background and Related Work

A. OPCM: Fundamentals and Key Properties

PCMs undergo reversible phase transitions, shifting between amorphous and crystalline states by the application of thermal energy. These phase transitions lead to changes in electrical and optical properties. PCM states exhibit different electrical resistances, where the amorphous state represents binary 0 (high resistance) and the crystalline state denotes binary 1 (low

This research was supported by the National Science Foundation (NSF) under grants CNS-2046226 and CCF-1813370

979-8-3503-4859-0/DATE24/© 2024 EDAA

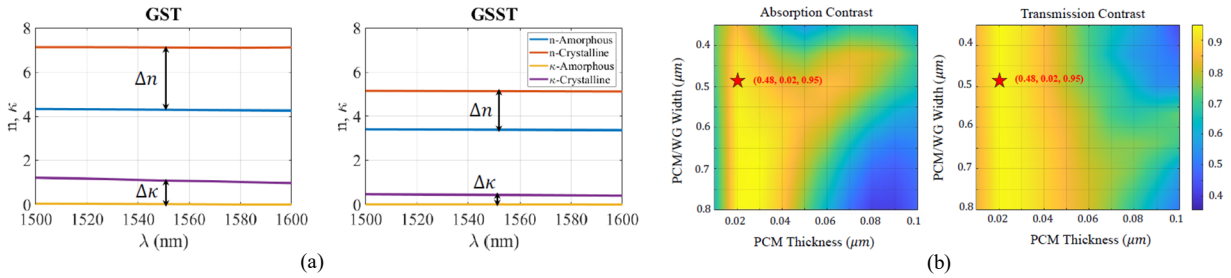

(a) (b)

Fig. 1. (a) Comparison of the refractive index (n) and extinction coefficient (κ) between GSST and GST in the optical C-band range. (b) Optical absorption contrast (left) and optical transmission contrast (right) of GST cell for different cell geometry (width and thickness) in the *COMET* architecture. The stars represent the geometric configuration selected with values for (width, thickness, absorption or transmission contrast ratio), based on our analysis.

resistance). This non-volatile resistance change enables 1T-1R EPCM memory configurations (e.g., [22]).

Optical PCM-based (OPCM) memories rely on refractive index changes between material phases, altering optical transmission to enable data storage and retrieval. Effective OPCM memory implementation requires an understanding of the PCM optical properties. High refractive index contrast (Δn) between amorphous and crystalline states is crucial for tolerating optical signal losses and noise, ensuring accurate readouts even with multiple intermediate phase-transition levels. Additionally, a high extinction coefficient contrast ($\Delta\kappa$), which measures optical power loss through a material, is also crucial in OPCM cells, allowing efficient laser power absorption for energy-efficient state transitions. The optical refractive index profile of $Ge_2Sb_2Te_5$ (GST) and $Ge_2Sb_2Se_4Te_1$ (GSST) has been depicted in Fig. 1(a). As GST exhibits higher Δn and $\Delta\kappa$ between amorphous and crystalline states than GSST, we have selected GST-based OPCM memory cells in our designs.

B. OPCM Main Memory

Many different approaches to OPCM cell design have been explored in the literature. One such design, introduced in [13], employs a straightforward crossbar-based cell with the PCM positioned atop waveguide crossings. This design formed the foundation for the *COSMOS* main memory architecture [13]. Access to this cell is facilitated through row and column access wavelength signals, both of which must be active simultaneously to enable write operations. The architecture also adopted a subtractive read method. Initially, the entire subarray is read, followed by a reset signal applied to the target row to erase its contents. Subsequently, the subarray is read again, and the memory controller (MC) performs subtraction between the two read values to derive the intended row values. This approach, combined with the assumption of 4-bits per cell, delivers high bit density within the architecture.

Fig. 2. Data corruption in crossbar-based OPCM memory from [13] due to crosstalk; (*left*) original image; (*right*) image after 4 writes to adjoining rows.

However, the cell design in the *COSMOS* architecture in [13] causes the rows to be susceptible to crosstalk from the write operations on the adjacent rows. The crosstalk signal, although small, can trigger changes in the OPCM cell due to the thermo-optic effect [14]. The energy from the write pulses can cause temperature changes, and hence refractive-index changes in the adjacent cells, causing severe data corruption. This renders the proposed read and write approaches in [13] prone to severe errors, as illustrated in Fig. 2.

The crossbar cell design in [13], while attractive for its bit density, suffers from severe memory reliability issues. To ensure data integrity and retrieval, we need to isolate memory cells and implement access control mechanisms. One approach is to employ microring resonators (MRs) as access controllers, as detailed in [15] and [16]. Additionally, we can use electro-optic tuning [17], where MR resonance is controlled through carrier injection via a PN junction with nanosecond-scale latencies. It is important to note that this approach introduces higher optical losses, which we take into account in our design.

The next section describes our cell design as well as our photonic main memory architecture (created by tiling multiple cells to create subarray and bank structures) in detail.

III. *COMET* OPCM-BASED MAIN MEMORY DESIGN

A. OPCM Memory Cell Design

In our OPCM cell design, we consider GST deposited on a silicon-on-insulator (SOI) strip waveguide. As shown in Fig. 1(a), the extinction coefficient of the GST in the crystalline state (κ_c) is much higher than its amorphous state (κ_a). This leads to negligible optical transmission due to high absorption of the electric field in the PCM. The optical absorption contrast and the optical transmission contrast between fully crystalline and fully amorphous state for OPCM memory cells of different geometries and materials are shown in Fig. 1(b). Note that the optical transmission contrast is not only a function of optical absorption in the cells but also partially originates from the optical-refractive-index mismatch between the PCM and SOI waveguide due to high refractive-index contrast between silicon and GST. To avoid optical-refractive-index mismatch when designing GST-based OPCM memory cells, it is important to select a design where both optical transmission contrast and optical absorption contrast are maximized. This ensures that the optical transmission contrast stems from the optical power absorption which is controlled by crystallization of the GST.

For a 2-μm-long GST cell considered in our study (Fig. 1(b)), optical transmission and absorption are at 95% when the thickness of the cell is at 20 nm. Note that the impact of PCM

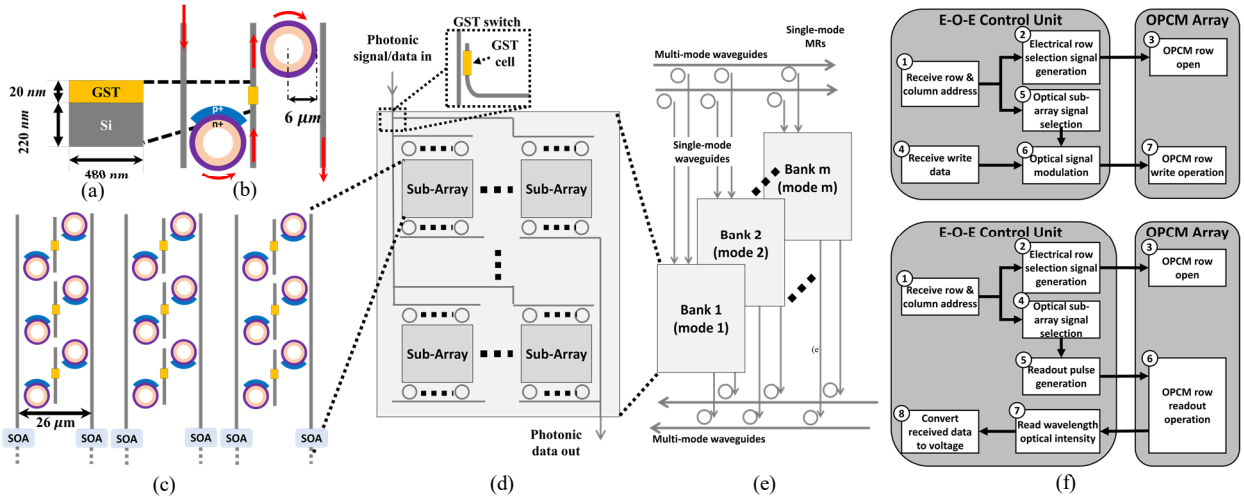

Fig. 3. (a) GST cell designed and simulated for this work. (b) Our proposed memory cell with MR-based access control. (*c*) OPCM memory array with intra-subarray semiconductor optical amplifiers (SOAs). (d) Single bank with multiple subarrays, with inset showing GST based signal switch. (e) Overall multi-bank architecture of *COMET*. (f) Steps during read (*top*) and write (*bottom*) operations in *COMET* architecture.

waveguide (WG) width on optical transmission and absorption is negligible. The SOI waveguide has a width of 480 nm (to ensure the single mode transmission of the light) and a thickness of 220 nm, where the GST deposited on it has the same width, but a thickness of 20 nm (see Fig. 3(a)). We have designed the GST cell with a reduced thickness, as a higher thickness makes heat transfer over the volume of the cell slower, leading to higher write and reset latencies.

To obtain the optical transmission characteristics of the GST cell we designed, we used FDTD simulations from Ansys Lumerical FDTD [18]. From these FDTD simulations, we were able to obtain the optical transmission levels of the intermediate states. To obtain the energy and latency of transitions for intermediate states, Ansys Lumerical HEAT [18] was used.

After finalizing the GST cell design, we designed the OPCM memory cell which integrates the GST cell and regulates signal access to the cell. This access control provides cell isolation between adjacent GST cells and is necessary to avoid optical crosstalk and associated thermo-optic-effect based data corruption. To ensure GST cell isolation, our memory cell has MR-based access control as shown in Fig. 3(b). MRs are switched in and out of resonance to enable and disable the signal access to the GST cell. We use MR designs from [27] with 6 μm radius, for low loss and efficient coupling during EO tuning to allow as much of the read/write signal to reach the GST cell from the Si waveguides. We opt for electro-optic tuning to tune the MRs in and out of resonance for the 2-ns access latencies that this mechanism provides [27].

B. *COMET Memory Bank Architecture Design*

The OPCM memory cells can form an array (see Fig. 3(c)) with columns dedicated to accessing GST cells. The cell architecture along with wavelength specific column access ensures cell isolation, significantly reducing crosstalk. Row access is achieved through electrical control of PN junctions in the MRs, allowing wavelength access to GST cells.

The memory bank thus formed has $N_r \times N_c$ OPCM cells, with a total capacity of $N_r \times N_c \times b$, where b is the bit capacity of the OPCM MLC. This bank is further divided into S subarrays, each containing $M_r \times M_c$ cells, such that $N_r = S_r \times M_r$ and $N_c = S_c \times M_c$, to enable energy-efficient access. The subarray access requires M_c wavelengths to be modulated, and $2 \times M_c$ MRs per bank have to be tuned to be in resonance. The memory bank requires N_c wavelengths to operate which are provided by the off-chip laser. Consequently, N_c MRs are needed to allow access to the columns, and another N_c MRs are needed to allow readout from the columns.

To access subarrays, we opt for GST-based waveguide switching [30], instead of passive splitters, to minimize laser power consumption. Note that using passive splitting will incur accumulative splitting losses, where each splitter warrants a 2× increase in laser power to ensure delivered power. The GST cell at the waveguide coupler switches from crystalline to amorphous state (Fig. 3(d)-top), incurring a 0.2 dB loss (for amorphous GST) and a 100 ns switching time, but significantly reduces losses and laser power requirements. Wavelength modulation and MR tuning signals originate from the electrical control unit, bridging OPCM memory banks with the processor.

To address data integrity issues, we integrate semiconductor optical amplifiers (SOA)-based gain tuning within subarrays and at the electrical interface (Fig. 3(c)-bottom). Row-wise loss-aware signal amplification compensates for optical losses, with SOAs within subarrays providing the needed power. The SOA gain required is static, as the losses and tolerances will be static. These SOAs counteract EO-tuned MR losses for read and write signals, and their power and latency overhead is considered in Section IV.

COMET is designed as a multi-bank OPCM memory (see Fig. 3(e)), using mode division multiplexing (MDM) and wavelength division multiplexing (WDM). Using MDM and WDM together allows COMET to limit the number of wavelengths but use their higher order modes as a multiplier to the communication channels. This reduces the overall laser power requirement. We enable parallel access across banks and interleaved cache lines via MDM with an MDM degree (B) of 4 to minimize overhead. As higher order modes are excited, the field will become leakier with higher losses. In addition, to support higher order modes, the waveguide width of the links and devices also needs to increase. Prior works have shown an MDM degree of 4 is achievable on a chip, without notable losses or area overhead [19].

C. Read and Write Operations in COMET

For reads (Fig. 3(f)-top), our isolated memory cells simplify the process. Row access is achieved through EO tuning, and column access is accomplished by sending readout pulses to the subarray. The row ID, obtained from the address, guides EO tuning signals to the row's MRs. Depending on the subarray ID from the physical address, relevant wavelengths are gain-tuned and directed to OPCM banks. Upon reaching the electrical interface, data is demodulated using an MR bank and passed to the processor for further processing. Likewise, for writes (Fig. 3(f)-bottom), the row ID from the address enables row access by tuning the MRs into resonance. Column IDs, also derived from the address, guide gain-tuning and modulation of corresponding wavelengths to represent the data to be written and transmit it to the OPCM subarray.

D. COMET Power Consumption

The *COMET* memory architecture can achieve a capacity of $(B \times N_r \times N_c \times b)$ bits. With our SOA-based loss mitigation strategy, we assume $M_c = N_c$, making, $S_c = 1$. The subarrays can be arranged in an array of $\sqrt{S_r} \times \sqrt{S_r}$ for layout and addressing. The M_r value would depend on the b value to ensure the cache line readout across the B banks. Our intra-subarray SOAs provide a gain of 15.2 dB to their input signal [20]. Given that the EO-tuned MRs have a through loss of 0.33 dB, there needs to be an SOA array at every 46 rows. This necessitates a total SOA count of $\frac{B \times N_r \times N_c}{46}$. To minimize power overhead, we only enable the SOAs within the accessed subarrays. Assuming 1.4 mW power consumption for 0 dBm i.e., 1 mW output [20], total SOA power consumption at any instance during *COMET* operation will be $\left(\frac{B \times M_r \times M_c}{46} \times 1.4\right)$ mW.

Apart from these power considerations, we must consider power consumption of the photonic links. Our WDM-MDM link requires N_c wavelengths and $2 \times B \times N_c$ MRs to access the OPCM arrays. These MRs are completely passive as they need not perform any switching operations. The laser power consumption can be calculated based on the various losses the signal will experience on its way to and from the OPCM arrays, as will be discussed in Section IV. Lastly, we need to consider the power consumption by the EO tuning mechanism within the OPCM arrays. Given that the MRs need to be tuned only within the row of the subarray being accessed, the tuning mechanism will contribute to $(B \times 2 \times M_c \times P_{EO})$ W of power, where P_{EO} is power consumption for EO tuning a single MR.

E. Address Mapping in COMET

To access the data within the memory banks, we need to perform an address mapping to the cells in our architecture. *COMET* can consider cache lines of various sizes (e.g., 32, 64, and 128 bytes), to reflect popular last level cache line sizes, interleaved across the B banks. The mapping process should perform the following mapping:

$$\{Channel_{ID}, Row_{ID}, Bank_{ID}, Column_{ID}\} \rightarrow \{Channel_{ID}, Subarray_{ID}, Subarray_{ROW}, Bank_{ID}, Subarray_{COL}\} \quad (1)$$

For our architecture, as described in Section III.B, the channel and bank IDs can remain the same. Row_{ID} must be mapped to $Subarray_{ID}$ and $Subarray_{ROW}$ and $Column_{ID}$ to $Subarray_{ID}$ and $Subarray_{COL}$. The values for these parameters can be calculated as follows:

$$ID_1 = int\left(\frac{Row_{ID}}{M_r}\right) \quad (2)$$

$$ID_2 = int\left(\frac{Column_{ID}}{M_C}\right) \quad (3)$$

$$Subarray_{ID} = ID_2 \times \sqrt{S_r} + ID_1 \quad (4)$$

$$Subarray_{ROW} = (Row_{ID} \% M_r) \quad (5)$$

$$Subarray_{COL} = (Column \% M_C) \quad (6)$$

IV. Experiments and Evaluation

We employ a customized version of NVMain 2.0 [21], a main memory simulator that we heavily modified to accommodate 4-bit/cell MLC operation, our photonic memory configurations, and the addressing scheme. Our experiments and evaluations are based on an 8 GB main memory chip capacity. Performance metrics encompass application latency, bandwidth, and energy-per-bit (EPB). We benchmark *COMET* against *COSMOS* [13], EPCM-MM [22], a proposed EPCM main memory architecture, and 2D and 3D configurations of DDR3 and DDR4 DRAMs (labeled as 2D_DDR3, 3D_DDR3, 2D_DDR4, and 3D_DDR4). Memory traces from the SPEC benchmark suite [23] are utilized for architecture evaluation. Table I outlines the parameters considered for power modeling in the *COMET* architecture.

TABLE I: Optical losses and power parameters considered for *COMET* power modeling.

Loss parameters	Values
Coupling loss	1 dB [24]
MR drop loss	0.5 dB [25]
MR through loss	0.02 dB [26]
EO tuned MR drop loss	1.6 dB [27]
EO tuned MR through loss	0.33 dB [27]
Propagation loss	0.1 dB/cm [28]
Bending loss	0.01 dB/90° [29]

Power parameters	Values
EO tuning power (P_{EO})	4 μW/nm [17]
Max. power at GST cell	1 mW
Intra-subarray SOA power	1.4 mW [20]

A. Modeling COMET Architecture

As described in Section III.D, *COMET* has a capacity of $(B \times N_r \times N_c \times b)$ bits, with the array of $N_r \times N_c$ GST memory cells divided into subarrays of size $M_r \times M_c$ memory cells for enabling parallel and energy-efficient access. For bit density b, there are works which show the GST cell should be able to achieve up to 5 bits/cell [11]. However, considering data distribution across cells, it is practical to consider bit densities in the multiples of 2, which allows for an even distribution of data across cells in a row and equal loss and crosstalk considerations to be made across all columns. But as discussed in Section III.C, at an architecture level, to allow high bit density per cell, several considerations must be made. To determine the optimal b value in *COMET*, we analyze how power, latency, bandwidth, and EPB varies for $b = \{1, 2, 4\}$, for an 8 GB memory. We opt to reduce M_c $(= N_c)$ over N_r $(= S_r \times M_r)$ as b increases, as reduced N_c results in a reduction of both WDM-degree requirement and significant intra-subarray SOA power reduction. Modifying N_c also allows *COMET* to retain its cache line capacity and bandwidth across the designs.

The resulting stacked power plots are shown in Fig. 4(a). Based on these analyses, we have chosen $b = 4$ to keep the power overhead relatively low. For further comparison studies

we consider a low EPB configuration of the *COMET* architecture, which was found through conducting a design-space exploration for $(B \times S_r \times M_r \times M_c \times 4)$. This configuration was found to be $(4 \times 4096 \times 256 \times 512 \times 4)$, and is used in the rest of the studies in this paper.

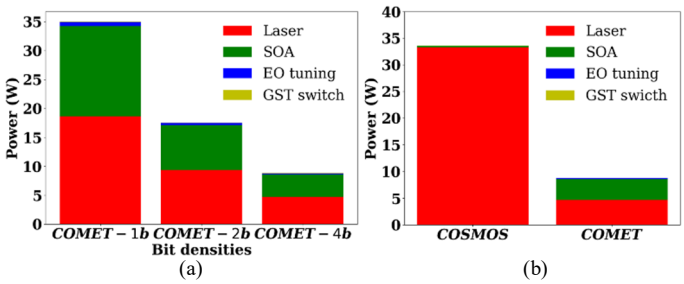

Fig. 4. (a) Power dissipation of COMET with bit density (b) of 1 (*COMET*-1b), 2 (*COMET*-2b), and 4 (*COMET*-4b); (b) Power dissipation of *COSMOS* and *COMET*.

B. Modeling COSMOS Architecture

We model *COSMOS* and update its design assumptions to avoid the data corruption scenario depicted in Fig. 2. But before we address the crosstalk and data corruption issues in *COSMOS*, we must address the energy delivery assumptions. The GST cell design in *COSMOS* is taken from [11], which requires 5 mW laser pulses over 50 ns to 150 ns to deliver 250 pJ to 750 pJ in energy for phase transitions. Unfortunately, *COSMOS* reported cells needing only 0.5 mW laser pulses while retaining the timing parameters from [11].

TABLE II: ARCHITECTURAL DETAILS OF PHOTONIC MEMORY SYSTEMS.

COMET	4 banks, 1 rank/channel, 1 device/rank bus width = 256 bits, burst length = 4 max. write time = 170 ns, erase time = 210 ns, read time = 10 ns, data burst time = 1 ns, electrical interface delay= 105 ns
COSMOS	8 banks, 1 rank/channel, 1 device/rank bus width = 128 bits, burst length = 8 write time = 1.6 μs, erase time = 250 ns, read time = 25 ns, data burst time = 1 ns, electrical interface delay= 105 ns

To ensure that the energy required for phase transitions is delivered to the memory cell, we have remodeled the timing constraints and assume 5 mW laser pulse power to ensure that the required energy is delivered to the GST cells. We have opted to increase the timing parameters as increasing the power value would make the total power consumption entirely too high for an 8 GB main memory. The modified parameters for *COSMOS* (and also *COMET*) are summarized in Table II.

Data corruption in *COSMOS* is a result of the thermo-optic effect and the losses that the optical signals experience as they traverse the crossbar. We decided to not re-architect *COSMOS* to enable cell isolation as this will depart too far from the design in [13]. We instead opt to change the intermediate level count to allow for higher tolerance to thermo-optic effect. This necessitates dropping *COSMOS* bit density b from 4 to 2. This results in a memory architecture where $(B \times N_r \times N_c \times b)$ is $(16 \times 16384 \times 16384 \times 2)$ and $S_r \times M_r = S_c \times M_c = 512 \times 32$. For this version of *COSMOS*, we make the generous assumption that the MDM losses are negligible.

We analyzed the minimum transmission-level separation required to avoid thermo-optic corruption. Our results showed that 4 asymmetric transmission levels (0.99, 0.90, 0.81, 0.72) separated by 9% transmission can be used to represent the 4 levels necessary to represent 2-bit data, without errors. Keeping the subarray size of 32×32 from *COSMOS*, the worst-case loss of 1.4 dB (from transmission level 0.72) and the 15.2 dB gain for SOAs [20], this also requires 6 SOA arrays (when considering both row and column losses) per subarray. To avoid excessive splitter loss and hence laser power consumption, we also assume a PCM cell-based subarray row access control in *COSMOS* (which we proposed for *COMET* in this work), separating the subarray rows. Fig. 4(b) shows the power stack comparison between *COSMOS* and *COMET*.

C. Performance Evaluation

We compare *COMET* in terms of bandwidth and EPB against 2D_DDR3, 3D_DDR3, 2D_DDR4, 3D_DDR4, the EPCM-MM, and *COSMOS* modified from [13], as discussed in the previous subsection. The absence of refreshes along with higher bandwidth provided by the silicon photonic interconnects allow both *COSMOS* and *COMET* to outperform their electronic counterparts in terms of bandwidth, (Fig. 5(a)). *COMET* achieves 100.3×, 47.2×, 58.7×, 42.1×, 40.6×, and 5.1× higher bandwidth on average than 2D_DDR3, 3D_DDR3, 2D_DDR4, 3D_DDR4, EPCM-MM, and *COSMOS*, respectively.

In terms of EPB, it can be observed from Fig. 5(b) that the 3D and PCM electronic counterparts are able to outperform both fully photonic memory systems. This can be attributed to the fact that the entire power consumption depicted in Fig. 4(b) is utilized for orchestrating reads and writes for photonic memory. This is different from the case in an electronic memory where only a very small portion of the overall power is required to orchestrate a read/write. However, the high read/write bandwidth in comparison, enables *COMET* to outperform the 2D DRAM platforms. *COMET* is able to achieve 4.1×, 2.3×, 12.9× lower EPB than 2D_DDR3, 2D_DDR4, and *COSMOS*.

COMET is able to contribute towards better overall energy efficiency of systems it is part of, owing to the combination of high bandwidth and comparable EPB performance to other main

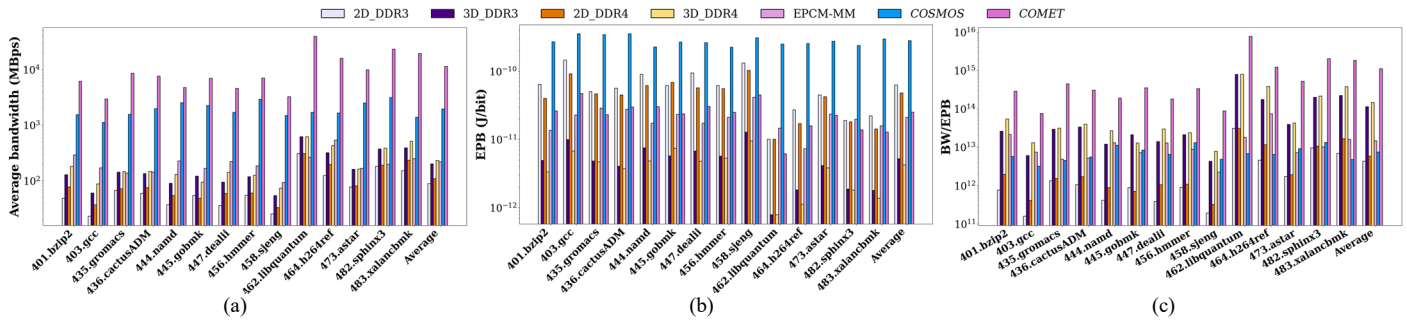

Fig. 5. (a) Average bandwidth (BW); (b) energy-per-bit (EPB); and (c) BW/EPB, of applications across memory architectures.

memory platforms. We showcase this using a BW/EPB metric (see Fig. 5(c)), where *COMET* achieves 6.5× and 65.8× better BW/EPB over 3D_DDR4 (best electronic platform) and *COSMOS*, respectively.

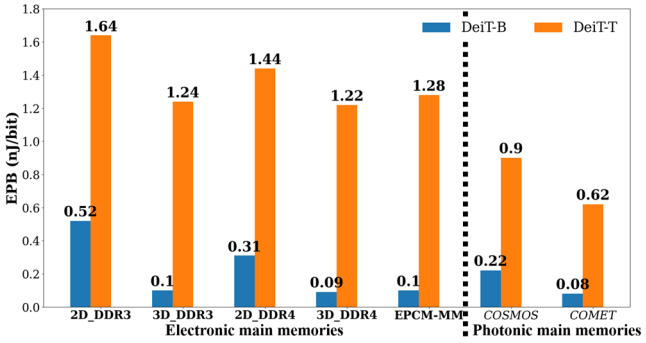

Fig. 6. EPB of DOTA accelerator with different main memories.

D. Photonic AI Accelerator Case Study

Photonic main memory architectures, such as our proposed *COSMOS*, are an excellent fit for emerging optical computing platforms. To quantify the benefits, we analyzed how different electronic and photonic main memories impact the operation of DOTA [31], a photonic tensor engine-based transformer accelerator [32]. For this analysis, we considered the two transformer models DeiT-T and DeiT-B, as used in [31]. Fig. 6 shows the EPB results for the various main memory architectures considered. Photonic memory architectures not only provide higher bandwidth (Fig. 5) than electronic memories, but also have the added benefit of being able to inject data directly into the photonic tensor engine, without the need for energy-hungry electro-photonic conversion stages. *COMET*+DOTA achieves 1.3× and 2.06× lower EPB against 3D_DDR4+DOTA and 2.7× and 1.45× better EPB against *COSMOS*+DOTA. These results highlight the promise of photonic main memory for improving the performance of emerging optical computing platforms.

V. Conclusions

In this work, we presented *COMET*, a low-loss, low-latency, and high throughput OPCM-based main memory architecture that makes use of GST material integrated with silicon photonic waveguides. We described the cross-layer design and optimization of our *COMET* architecture from the material and device level to the architecture level. Our GST cell offered a high transmission contrast between crystalline and amorphous state of the cell (≈96%), enabling 16 distinctive transmission levels which makes *COMET* tolerant to transmission drift. Crosstalk-free, reliable memory operation is enabled with various loss-aware optimizations at the architecture level, owing to which, *COMET* consumes only 26% of the power when compared to the best-known prior work on OPCM-based main memory design. The low power overhead along with high-speed GST cell operations enable *COMET* to offer 5.1× better bandwidth, 12.9× lower EPB, and 65.8× better BW/EPB than the state-of-the-art.

References

[1] W. Zhao, *et al.* "A survey of large language models." arXiv preprint arXiv:2303.18223, 2023.

[2] X. Gao, *et al.*, "An adaptive ensemble machine learning model for intrusion detection," in *IEEE Access*, vol. 7, pp. 82512–82521, 2019.

[3] J. E. Gonzalez, *et al.*, "PowerGraph: Distributed graph-parallel computation on natural graphs," in *OSDI,* 2012.

[4] S. K. Kim and M. Popovici, "Future of dynamic random-access memory as main memory," in *MRS Bulletin*, vol. 43, pp. 334–339, 2018.

[5] H. Ishiwara, "FeFET and ferroelectric random access memories," in *Multifunctional Oxide Heterostructures*, 2012.

[6] Y. Chen, "ReRAM: History, Status, and Future," in *IEEE Transactions on Electron Devices*, vol. 67, pp. 1420-1433, 2020.

[7] A. Shafiee, S. Pasricha, and M. Nikdast, "A Survey on Optical Phase-Change Memory: The Promise and Challenges", in *IEEE Access*, 2023, Jan 31;11:11781-803..

[8] M. Wuttig, *et al.*, "Phase-change materials for non-volatile photonic applications," in *Nat. Photon.*, vol. 11, no. 8, pp. 465–476, 2017.

[9] A. Cabrini, *et al.*, "Voltage-driven multilevel programming in phase change memories," in *IEEE MTDT*, 2009.

[10] F. Bedeschi, *et al.*, "A multi-level-cell bipolar selected phase-change memory," in *IEEE ISSCC*, 2008.

[11] X. Li, *et al.*, "Fast and reliable storage using a 5 bit, nonvolatile photonic memory cell," in *Optica*, vol. 6, no.1, pp. 1–6, 2019.

[12] F. Sunny, E. Taheri, M. Nikdsat and S. Pasricha, "A survey on silicon photonics for deep learning," in *ACM JETC*, vol. 17, no. 4, pp. 1-57, 2021.

[13] A. Narayan, *et al.*, "Architecting optically controlled phase change memory," in *ACM TACO*, vol. 19, no. 4, pp. 1–26, 2022.

[14] C. Rios, *et al.*, "Integrated all-photonic non-volatile multi-level memory," in *Nat. Photon.*, vol. 9, no. 11, 2015.

[15] J. Feldmann *et al.*, "Integrated 256 cell photonic phase-change memory with 512-bit capacity," in *IEEE JSTQE*, vol. 26, no. 2, pp. 1–7, 2019.

[16] A. Shafiee, B. Charbonnier, S. Pasricha, and M. Nikdast, "Design Space Exploration for PCM-based Photonic Memory," in *ACM GLSVLSI*, 2023.

[17] A. Stefan, *et al.*, "A hybrid barium titanate–silicon photonics platform for ultraefficient electro-optic tuning," in *IEEE JLT*, vol. 34, no. 8, 2016.

[18] Ansys Lumerical. [Online]. Available: https://www.lumerical.com/

[19] P. Guo, *et al.*, "StarLight: a photonic neural network accelerator featuring a hybrid mode-wavelength division multiplexing and photonic nonvolatile memory," in *Opt. Exp.*, vol. 30, pp. 37051-37065, 2022.

[20] T. Lin, *et al.*, "Performance and scalability of a single-stage SOA switch for 10/spl times/10 gb/s wavelength striped packet routing," in *IEEE Photonics Technology Letters*, vol. 18, no. 5, pp. 691-693, March 1, 2006

[21] M. Poremba, *et al.*, "NVMain 2.0: A user-friendly memory simulator to model (non-) volatile memory systems," in *IEEE CAL*, 14:2, 2015.

[22] Y. Choi, *et al.*, "A 20nm 1.8V 8Gb PRAM with 40MB/s program bandwidth," in *IEEE ISSCC*, 2012.

[23] SPEC. [Online]. Available: https://www.spec.org/benchmarks.html

[24] C. Batten, *et al.*, "Building many-core processor-to- DRAM networks with monolithic CMOS silicon photonics," in *IEEE HOTI*, 2008.

[25] M. R. Yahya, *et al.*, "A low insertion loss 5 × 5 optical router for mesh photonic network-onchip topology," in *IEEE CSUDET*, 2019.

[26] S. Pasricha and S. Bahirat, "OPAL: A multi-layer hybrid photonic NoC for 3D ICs," in *IEEE ASPDAC*, 2011.

[27] A.W. Poon, *et al.*, "Cascaded microresonator-based matrix switch for silicon on-chip optical interconnection," in *Proceedings of the IEEE*, vol. 97, no.7, pp.1216-1238, 2009.

[28] L. Zhang, *et al.*, "New-generation silicon photonics beyond the singlemode regime," arXiv:2104.04239, 2021.

[29] M. Behadori, *et al.*, "Universal design of waveguide bends in silicon-on-insulator photonics platform," in *IEEE JLT*, vol. 37, no. 10,Jul. 2019.

[30] E. Taheri, S. Pasricha, and M. Nikdast, "ReSiPI: A Reconfigurable silicon-photonic 2.5d chiplet network with PCMs for energy-efficient interposer communication,"in *IEEE/ACM ICCAD*, 2022.

[31] H. Zhu, *et al.*, "DOTA: A Dynamically-Operated Photonic Tensor Core for Energy-Efficient Transformer Accelerator," arXiv:2305.19533, 2023.

[32] A. Vaswani, *et al.*, "Attention Is All You Need," arXiv:1706.03762 [cs.CL], 2023