

RESEARCH ARTICLE | DECEMBER 15 2025

OptiCAM: An optical content-addressable memory architecture for ultra-fast pattern matching

Special Collection: [Optical Computing Systems](#)Md Abdullah-Al Kaiser   ; Ankur Singh  ; Shay Snyder  ; Hamed Poursiami  ; Maryam Parsa  ; Akhilesh Jaiswal

APL Photonics 10, 126116 (2025)

<https://doi.org/10.1063/5.0288417>

Articles You May Be Interested In

Loose fruit detection for autonomous loose fruit collector

AIP Conf. Proc. (August 2024)

3D analysis of semiconductor devices: A combination of 3D imaging and 3D elemental analysis

J. Appl. Phys. (January 2018)

A predictive analytics and space allocation method using machine learning in a smart parking system based on internet of things-smart space

AIP Conf. Proc. (November 2025)

APL Photonics

Special Topics Open for Submissions

[Learn More](#)

OptiCAM: An optical content-addressable memory architecture for ultra-fast pattern matching

Cite as: APL Photon. 10, 126116 (2025); doi: 10.1063/5.0288417

Submitted: 1 July 2025 • Accepted: 30 November 2025 •

Published Online: 15 December 2025



Md Abdullah-Al Kaiser,^{1,a)}  Ankur Singh,¹  Shay Snyder,²  Hamed Poursiami,²  Maryam Parsa,²  and Akhilesh Jaiswal¹

AFFILIATIONS

¹ University of Wisconsin-Madison, Madison, 500 Lincoln Dr, Madison, Wisconsin 53706, USA

² George Mason University, 4400 University Dr, Fairfax, Virginia 22030, USA

Note: This paper is part of the Special Topic on Optical Computing Systems.

a) Author to whom correspondence should be addressed: mkaiser8@wisc.edu

ABSTRACT

Content-Addressable Memory (CAM) circuits enable inherently parallel and high-speed search operations, making them essential for data-intensive applications, such as packet classification in high-speed networking, pattern recognition, and artificial intelligence. However, as CMOS technology continues to scale, conventional electrical CAMs suffer from increasing interconnect delays, hindering their performance and energy efficiency in large-scale deployments. In contrast, photonic systems leverage light as the information carrier, offering ultra-fast data transmission with minimal signal degradation due to their length-independent impedance characteristics. In this work, we introduce OptiCAM—an optical CAM architecture that employs silicon photonic microring resonators and photodiode-based latch structures to store data and execute matching entirely in the optical domain. Furthermore, we present a novel sectioned array architecture that provides a pathway to build scalable CAM arrays, effectively overcoming the wavelength limitations imposed by the microring resonators' free spectral range. Our OptiCAM design, evaluated on a 16×16 photonic CAM array using the commercial GlobalFoundries' 45SPCLO technology node, achieves a remarkable search latency of 100 ps and an energy consumption of 100.9 fJ/bit per search.

© 2025 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0288417>

I. INTRODUCTION

Conventional memory systems organize data in binary form across rows and columns. To perform read or write operations, they rely on column-parallel peripheral circuits.^{1,2} When data need to be accessed, a specific row is activated based on its address, and the data are transferred through shared column lines at the memory's edge. However, many modern applications—such as high-speed networking for packet classification, pattern recognition, and artificial intelligence—often require searching for specific content within memory, rather than retrieving it by address.^{3,4} In such cases, the traditional approach of activating one row at a time and comparing data externally leads to high latency and energy consumption, particularly with large datasets. To address this limitation, Content-Addressable Memories (CAMs) offer a powerful solution. CAMs

can compare input data against all stored entries simultaneously and return the matching address in a single clock cycle, significantly improving both speed and energy efficiency. Their inherent parallelism makes CAMs especially suited for high-bandwidth, low-latency workloads that are increasingly common in today's computing systems.^{5,6}

Traditionally, Content-Addressable Memories (CAMs) have been widely used in networking hardware, such as routers, switches, and Smart Network Interface Cards (SmartNICs), to enable fast packet classification and forwarding.^{7,8} Over time, their use has expanded into areas, such as security analytics, data compression, caching, neuromorphic computing, and more recently, deep learning accelerators.^{9–15} Most conventional CAMs are built using Static Random-Access Memory (SRAM) technology because of its fast access and reliable performance.^{16–21}

However, with continued CMOS scaling, SRAM-based CAMs face limitations, such as RC-delay bottlenecks, static leakage power, and substantial area overhead due to peripheral logic.²² Dynamic Random-Access Memory (DRAM)-based CAMs offer higher density but suffer from slower access times.^{23,24} To mitigate these issues, researchers have explored CAM architectures based on emerging non-volatile memories (NVMs), including phase-change memory (PCM),^{25–27} spin-transfer torque magnetic RAM (STT-MRAM),^{28–30} flash memory,³¹ Ferroelectric Field Effect Transistors (FeFETs),³² and memristors.^{33–35} These technologies promise reduced standby power and improved density but are constrained by higher write energy, device variability, and limited endurance. Moreover, even with these improvements, electrical CAMs (E-CAMs) remain constrained by the energy and latency limitations imposed by long interconnects and electrical parasitics.

This speed mismatch between electrical CAMs and modern optical interconnects is especially critical in high-performance computing and data center environments, where optical signal rates routinely reach tens or even hundreds of gigabits per second.³⁶ In contrast, current E-CAMs typically operate in the range of a few gigabits per second, necessitating optoelectronic conversion (OEC) and data-rate downconversion at the optical–electrical interface. These two stages not only introduce latency but also significantly increase energy consumption at a cost that is frequently excluded from reported CAM efficiency metrics.^{37,38} The need for such costly conversion stages disrupts high-speed photonic data flows and introduces a critical bottleneck in photonic memory subsystems. In addition, interconnect-dominated energy and delay components increasingly dominate the system-level performance profile of E-CAM arrays. This includes capacitive RC-delay, power-hungry drivers, and layout-induced variability at advanced CMOS nodes. As a result, the performance limits of modern CAMs are often dictated more by interconnect physics than by the underlying memory or logic elements, creating a fundamental architectural mismatch between photonic interconnects and electronic search engines.

To address these challenges, several optical CAM (O-CAM) architectures have been developed. Early efforts employed InP-based semiconductor optical amplifier Mach–Zehnder interferometer (SOA-MZI) configurations to implement optical flip-flops and XOR gates for binary and ternary CAMs.^{39,40} While these demonstrated up to 10 Gbps search speeds nearly 10× faster than the typical E-CAMs, they suffered from large device footprints, significant energy consumption, and scalability limitations.⁴¹ More recent photonic CAM designs have explored the use of silicon photonic microring resonators (MRRs), leveraging their small footprint, high Q-factor, and wavelength selectivity to enable dense, energy-efficient CAM arrays.^{42,43} These MRR-based CAMs use wavelength-division multiplexing (WDM) to encode search and stored data, allowing for parallel content matching within a compact form factor. Despite these advantages, WDM-based photonic CAMs are inherently constrained by the free spectral range (FSR) of the microring resonators, which limits the number of usable wavelengths within a given optical bandwidth. In addition, higher-Q resonators have narrower linewidths and therefore require smaller voltage shifts to switch resonance, but this comes at the expense of modulator bandwidth. Moreover, most photonic CAM demonstrations to date continue to rely on electronic subsystems for memory write operations, which

limit the degree to which they can overcome the inherent limitations of electrical CAM architectures.

In this work, we introduce two designs of OptiCAM bitcells—an optical content-addressable memory built using silicon photonic microring resonators (MRRs) and photodiodes (PDs). In the first design, binary data are encoded using two different wavelengths, one for “1” and another for “0.” In the second design, a single wavelength is used to encode data. The data are stored in an electro-optic latch, formed by the cross-coupling between MRRs and a balanced photodiode structure. In addition, we propose a sectioned architecture for the OptiCAM array that includes intermediate optoelectronic conversion stages to overcome wavelength scaling limitations imposed by the free spectral range (FSR) of MRRs, enabling scalability to larger arrays without degrading performance. Our design, built as a 16×16 array using the GlobalFoundries’ 45SPCLO process, achieves a search speed of 10 GHz and energy consumption of 100.9 fJ/bit per search.

II. OptiCAM CIRCUIT AND ARCHITECTURE

In this section, we begin by introducing our photonic storage element, followed by a detailed explanation of the two proposed OptiCAM designs, using a 1×4 array as the core macro unit. We then describe how this macro can be scaled to larger CAM arrays through our novel sectioned array architecture, which enables efficient scaling without being limited by the number of available wavelengths within a single FSR.

A. Photonic-SRAM

Figure 1 shows the schematic of the proposed photonic SRAM (pSRAM) bitcell, where data are stored using a cross-coupled




FIG. 1. Differential cross-coupled photonic SRAM (pSRAM) bitcell used as the storage element in our OptiCAM architecture. The data and their complement are stored at the optically modulated electric nodes Y and YB.

configuration of microring resonators (MRRs) and photodiodes (PDs).^{44,45} The through and drop ports of MRR M1 (M2) are connected to the waveguides of PDs P1 (P3) and P2 (P4), respectively. An optical laser at wavelength λ_{IN} is coupled into the circuit through an input splitter (PS1) at port IN, distributing light evenly to M1 and M2. The input wavelength λ_{IN} is chosen such that applying VDD to both MRRs aligns their resonances with the input laser. Electrical drivers D1 and D2, connected to nodes Y and YB, drive M1 and M2, respectively, forming a cross-coupled loop that latches binary data at Y and YB as long as optical and electrical biases are maintained.

For instance, when $Y = \text{VDD}$ and $YB = \text{GND}$, M1 is tuned into resonance with the input light, directing optical power toward P2. The photocurrent generated at P2 keeps YB at GND, detuning M2 and allowing light to pass to P3, which in turn maintains Y at VDD. Conversely, when $YB = \text{VDD}$ and $Y = \text{GND}$, M2 resonates with the input light, steering power to P4 and reinforcing $Y = \text{GND}$, while M1 directs light to P1 to sustain YB at VDD. In this way, the cross-coupled MRR-PD configuration provides robust bistable storage.

To perform a write operation, a differential optical pulse is applied to the write bitline (WBL) and write bitline bar (WBLB). Assume that the cell initially holds $Y = 0$ and $YB = 1$, and the goal is to flip the state to $Y = 1$ and $YB = 0$. A high-intensity optical pulse is launched into WBL, while WBLB remains inactive. The splitter PS2 distributes the pulse to PDs P3 and P2, lowering the resistance paths to VDD for Y and to GND for YB, thereby forcing the latch to flip. Similarly, to write the complementary state, an optical pulse is applied to WBLB, while WBL is left dark. The splitter PS3 delivers

the light to P1 and P4, enabling the complementary transition ($Y = 0$ and $YB = 1$) and ensuring the latch updates to the correct state.

This photonic-SRAM bitcell has been fully designed, fabricated, and experimentally validated on the GlobalFoundries' 45 nm CMOS-silicon photonics platform.⁴⁶ This structure supports stable data retention and fast optical switching, making it well-suited for high-speed photonic memory applications. Furthermore, a voltage-based automatic calibration technique has been developed to ensure robust and drift-tolerant microring resonance alignment.⁴⁷ This pSRAM holds the stored data, and dedicated read/search optical ports are embedded alongside it to enable CAM operations by comparing the stored data with input data from the periphery, as will be detailed in the later sections.

B. OptiCAM bitcell and array structure

In this subsection, we present two designs of our proposed OptiCAM bitcell, each offering distinct trade-offs between circuit area and optical signaling complexity. In the first design, binary data are encoded using two separate wavelengths, where one represents a logical "1" and the other a "0." This approach requires only a single microring resonator per bit to perform the search operation, resulting in a more compact footprint. By contrast, the second design employs a single wavelength to encode data, simplifying the optical signaling but requiring two compute MRRs per bit, which leads to a comparatively larger footprint. In both cases, data are stored within the pSRAM structure described earlier in Subsection II A, and




FIG. 2. Schematic and operation of the dual-wavelength OptiCAM design. (a) Schematic and functional overview of the OptiCAM-DW bitcell. Here, P_0 denotes the dark state and P_X represents the optical power of the input data, with P_X indicating logic "1" and P_0 indicating logic "0" in the optical domain. In the electrical domain, VDD and GND correspond to logic "1" and "0," respectively. The pSRAM serving as the cross-coupled photonic memory element is detailed in Fig. 1. (b) Transmission spectra of the microring resonators (MRRs) as their length varies, showing four distinct wavelength channels. (c) Schematic of a 1×4 OptiCAM-DW array structure.

the compute rings are augmented on top of the pSRAM to enable in-memory optical search operations.

1. OptiCAM-DW: OptiCAM dual-wavelength

Figure 2(a) shows the schematic of the dual-wavelength OptiCAM (OptiCAM-DW) bitcell design. This structure employs a cross-coupled pSRAM (depicted as a square box in the figure) to store data and uses a single compute microring resonator (MRR), labeled M3, to compare the stored and input data. The integrated PN-junction phase shifter of M3 is electrically connected to the storage node (Y) of the pSRAM (as shown in Fig. 1). In this design, two distinct optical wavelengths encode the input logic states: λ_1 represents logic “1” and λ_{1B} represents logic “0.” The microring M3 is tuned to resonate with the incoming light of wavelength λ_1 via X when its phase shifter is biased at VDD, and to resonate with λ_{1B} when the bias is at GND. The wavelength encoded input data are applied via the X port. This dual-resonance capability enables wavelength-selective transmission depending on the stored value in the pSRAM, facilitating high-speed optical comparison.

When the input and stored data match, the output remains dark, signaling a successful match. For example, if both are logic “0” (input wavelength λ_{1B} at X, stored voltage Y = GND), the compute MRR M3 resonates with λ_{1B} , diverting the light away so that no power reaches the output Z. Likewise, if both are logic “1” (input wavelength λ_1 at X, stored voltage Y = VDD), M3 resonates with

λ_1 and again blocks light from reaching Z. In contrast, during a mismatch, light passes through to the output node. For instance, if the input is “1” (λ_1) and the stored data are “0” (Y = GND), M3 is out of resonance with λ_1 , allowing the light to continue through to the output and producing high optical power at Z. A similar situation occurs when the input is “0” (λ_{1B}) and the stored data are “1” (Y = VDD), with M3 out of resonance and letting light reach Z.

Using the OptiCAM-DW bitcell, we can build a 1×4 array, where the outputs from each bitcell are combined through multimode interferometers (MMIs) as shown in Fig. 2(c). To enable non-coherent optical addition, this design leverages wavelength division multiplexing (WDM). Each compute microring [M3, shown in Fig. 2(a)] is tuned to resonate at a specific wavelength, with spacing carefully chosen within a single free spectral range (FSR) to reduce crosstalk between adjacent channels. Figure 2(b) illustrates the transmission spectra of compute MRRs with different geometries, adjusted by modifying the microring length (dL). As depicted in Fig. 2(c), the red, green, blue, and orange rings are tuned to resonate at λ_1 , λ_2 , λ_3 , and λ_4 , respectively, when biased at VDD, and shift to λ_{1B} , λ_{2B} , λ_{3B} , and λ_{4B} when the bias is set to GND.

When the input data X{1}, X{2}, X{3}, and X{4} exactly match the stored data Y{1}, Y{2}, Y{3}, and Y{4}, the output ports Z{1}, Z{2}, Z{3}, and Z{4} receive no light, so the combined output from the MMIs stays in the dark state. These outputs connect to a balanced photodiode structure, where the top photodiode is a dual detector linked to the outputs of two MMIs, while the bottom




FIG. 3. Schematic and operation of the single-wavelength OptiCAM design. (a) Schematic and functional overview of the OptiCAM-SW bitcell. Here, P_0 denotes the dark state and P_X represents the optical power of the input data, with P_X indicating logic “1” and P_0 indicating logic “0” in the optical domain. In the electrical domain, VDD and GND correspond to logic “1” and “0,” respectively. The pSRAM serving as the cross-coupled photonic memory element is detailed in Fig. 1. (b) Transmission spectra of the microring resonators (MRRs) as their length varies, showing four distinct wavelength channels. (c) Schematic of a 1×4 OptiCAM-SW array structure.

photodiode receives a threshold power P_{TH} to enable binary thresholding. When all bits match, the optical power at the top photodiode remains below P_{TH} , causing the transimpedance amplifier (TIA) and the variable gain amplifier (VGA) to output VDD. The TIA, implemented using an inverter with a feedback resistor, amplifies the small current from the balanced photodiodes into a voltage, which is then boosted by the cascaded amplifier to a full rail-to-rail binary level.⁴⁵

This VGA output drives another MRR M5, which resonates with the input light P_C when its integrated pn junction is at VDD, resulting in no light at the match line ML, indicating a match. However, if there is any bit mismatch, one or more Z outputs carry light, raising the power above P_{TH} at the balanced photodiode. This causes the amplifier output to drop to GND, tuning the MRR M5 out of resonance. As a result, the match line ML shows high optical power, indicating a mismatch in the 1×4 array. Note that the 1×4 array can be easily scaled to any arbitrary vector size by adding more pSRAM cells and compute rings [such as M3 shown in Fig. 2(a)], along with additional MMIs and multiple photodiodes connected in parallel at the top side of the balanced photodiode structure. This allows binary CAM operations to be performed on larger vectors.

2. OptiCAM-SW: OptiCAM single-wavelength

Figure 3(a) illustrates the OptiCAM structure utilizing single wavelength for data encoding. It consists of the basic pSRAM latch structure to hold the stored data and two additional MRRs [denoted as M3 and M4 in Fig. 3(a)]. The compute rings M3 and M4 are driven by the storage nodes Y and YB, respectively. The OptiCAM-SW can perform the CAM operation between the input data (X) and the stored data (Y). The input data and their complement are applied through the read/search waveguides connected to the X and XB ports.

When the input and stored data match, the output (Z) stays dark, indicating a match. For instance, if both are logic “0” ($X = P_0$, $XB = P_X$, $Y = GND$, and $YB = VDD$), MRR M4 resonates with the input light coming through XB and diverts it to be absorbed, so no light reaches the output Z. Here, P_0 denotes the dark-state optical input (no light), while P_X represents the optical power supplied per data bit. A similar effect occurs when both are logic “1” ($X = P_X$, $XB = P_0$, $Y = VDD$, and $YB = GND$), where MRR M3 resonates and blocks the light coming from the X port, preventing it from reaching Z. In contrast, during a mismatch, light passes through to the output Z. For example, if the input is “1” ($X = P_X$, $XB = P_0$) and the stored data are “0” ($Y = GND$ and $YB = VDD$), MRR M3 is out of resonance, letting light continue to the coupler C1 and creating high optical power at Z. Likewise, on the contrary, where the input is “0” ($X = P_0$ and $XB = P_X$) and the stored data are “1” ($Y = VDD$ and $YB = GND$), MRR M4 does not resonate with the incoming light from XB, allowing it to pass through the coupler and produce high optical power at Z.

Integrating the OptiCAM-SW bitcell in an array configuration [shown in Fig. 3(c)], we can construct an array of vector size 1×4 utilizing four different wavelength channels as shown in Fig. 3(b) by tuning the length of the compute microrings, as discussed in Subsection II B 1. When the input data matches the stored data, the match line ML will be in a dark state because it generates power lower than the threshold power P_{TH} , causing the ring M5 connected to the match line to resonate and make the ML dark. In contrast, if

there is any mismatch, the balanced photodiode structure outputs a high signal, which drives the TIA and VGA chains to GND, pushing the ring connected to the ML out of resonance and yielding a high optical output power. Similar to the OptiCAM-DW structure, the OptiCAM-SW can be scaled to support arbitrary vector sizes.

In summary, both the OptiCAM-DW and OptiCAM-SW bitcells output high optical power at Z when the input and stored data do not match, and remain dark when they do—effectively performing the CAM operation at high speed. The OptiCAM-DW uses two wavelengths to execute CAM searches, which may require additional tuning of the source laser, but relies on a single compute microring, making it more compact. In contrast, the OptiCAM-SW operates with a single input wavelength but employs two MRRs and an MMI coupler, resulting in a larger footprint than OptiCAM-DW.

C. Two-dimension (2D) OptiCAM array architecture

Figure 4 illustrates a two-dimensional (2D) 4×4 CAM array constructed using our proposed OptiCAM-SW structure; however, a similar 2D configuration can also be implemented with the OptiCAM-DW design. To scale up from the one-dimensional (1D) arrays described in Subsections II B 1 and II B 2 into a 2D architecture, multiple 1D arrays of vector size n (here, $n = 4$) are arranged along several rows (m), each with its own dedicated match line (i.e., ML(1), ML(2), ML(3), and ML(4)) to independently perform the CAM operation on the corresponding row-wise stored data. On the input side, a $1 \times m$ splitter is used to divide the input data into m separate rows, routing the signals to the corresponding rows of the 2D array for each column. For example, in the figure, m is chosen as 4 to simultaneously deliver $X(1)$, $X(2)$, $X(3)$, and $X(4)$, each with $4 \times$ the optical power per data line. This design enables parallel CAM operations across all rows and supports scalability to larger vector sizes.

D. Sectioned analog OptiCAM array architecture

Figure 5 shows a sectioned OptiCAM array architecture built using our proposed OptiCAM-SW bitcell; however, the same design can also be realized with the OptiCAM-DW bitcell. In this setup, we employ the 1×4 macro block described in Subsection II B 2 to compute the CAM operation outputs independently through wavelength-division-multiplexed compute rings connected to the match line (ML). Here, each section (for example, a 1×4 array) carries out the CAM operation as detailed earlier. Depending on the match condition per section, the MRRs (M5, M6, M7, and M8 depicted as red, green, blue, and orange for section-1, 2, 3, and 4, respectively) linked to the ML either enter resonance—absorbing specific wavelengths of light—or remain out of resonance, allowing those wavelengths to pass through the ML.

The sectioned architecture is inherently scalable and can be integrated in a hierarchical manner to expand the array size as an integer multiple of the macro block. In this figure, we use a 1×4 array as the basic macro, selected based on the number of wavelengths that can fit within a single FSR to minimize adjacent channel crosstalk and interference. Each macro or section performs the CAM operation independently of the others. Since this configuration employs separate compute rings (M5, M6, M7, and M8 in Fig. 5), the resulting optical power on the ML supports analog matching with a resolution determined by the macro or section size.



It is worth noting that the sectioned analog OptiCAM array architecture is inherently scalable. By leveraging the 16-bit analog CAM block shown in Fig. 5, the array size can be expanded to




FIG. 5. Sectioned analog OptiCAM array architecture based on the OptiCAM-SW bitcell, employing multiple 1×4 macro blocks with wavelength-division-multiplexed compute rings to perform independent CAM operations per section.

64 bits by hierarchically connecting multiple such blocks through additional opto-electronic conversion stages. In particular, the ML in the figure can be linked to a balanced photodiode with a reference power to binarize the optical output, which then drives

another MRR (similar to M5) connected to a global match line for a 64-bit vector. Likewise, three similar sections can be combined to drive three distinct wavelength-specific MRRs (akin to M6, M7, and M8, but connected to the global ML) to achieve CAM operations

on a 64-bit vector. While this approach maintains analog matching, it does so at a coarser granularity than the original 16-bit sections. In essence, the sectioned design overcomes the limitation on the number of wavelengths per FSR by incorporating an opto-electronic binarization stage, enabling array scaling beyond the constraints imposed by the FSR. This strategy is also broadly applicable to scaling other optical array computations, such as in photonic tensor core architectures.

III. RESULTS AND DISCUSSIONS

In this section, we present the verification results of our proposed OptiCAM structures—both the dual-wavelength and single-wavelength designs—and the sectioned analog OptiCAM architecture, implemented using the commercial GlobalFoundries' 45CLO process design kit (PDK), which enables monolithic integration of silicon photonics with high-performance transistors. These simulations use industry-standard PDK models, calibrated against actual fabricated devices to enable accurate performance prediction via electronic design automation (EDA) tools. The GF45CLO platform integrates photonic devices with RF CMOS circuitry on a 300 mm SOI wafer and includes an eight-metal-layer back-end stack.⁴⁸ We designed the OptiCAM structures using GF45CLO PDK components (including four-port microring resonators, photodiodes, splitters, and transistors) in Cadence Virtuoso Studio and performed Spectre simulations to evaluate and verify their performance.

Figure 6 shows a representative diagram of the Cadence Virtuoso-based schematic setup and simulation environment used in this study, illustrating the complete workflow used to validate the OptiCAM circuits. All photonic and electronic components—including the virtual optical source (irsource), microring resonators, waveguides, MMIs, photodiodes, transistors, and circuit blocks—were instantiated directly from the GF45CLO PDK within the Virtuoso Studio Schematic Editor. The simulations were configured and executed in Analog Design Environment (ADE) explorer, where DC analysis was used for microring power vs wavelength characterization and transient analysis was employed for functional verification, timing extraction, and energy evaluation. ADE explorer relies on encrypted, foundry-verified model libraries containing measured fabrication data and compact models, enabling accurate co-simulation of optical and electrical domains.

Simulation results were visualized in Virtuoso Visualization and Analysis (VIVA) XL, where we monitored electrical node voltages and optical powers using virtual instruments, such as the irmeter.

A. Verification of OptiCAM array structure

Figure 7 shows the verification results for the proposed 1×4 OptiCAM-DW array. To test various matching scenarios, we set the input data to 1111 and sequentially varied the stored data to represent cases ranging from a full match to 1-bit, 2-bit, 3-bit, and complete 4-bit mismatches. As shown in the figure, during the time interval from 500 to 600 ps, the stored data $Y(1:4)$ equals 1111, indicating a match with the input. Consequently, the bottom subplot for the ML shows low optical power output. Note that there is a latency between the input change and the resulting output due to the TIA and VGA stage; however, this does not impact the maximum speed at which the circuit can distinguish match and mismatch conditions. In this simulation, we used a 100 ps clock pulse to sequentially change the input bits. Between 600 and 700 ps, a 1-bit mismatch is introduced ($Y(4) = 0$). In subsequent clock cycles, we simulate 2-, 3-, and 4-bit mismatch scenarios. The bottom subplot clearly shows that under these conditions, the ML remains at a high optical power state, indicating mismatches.

Figure 8 presents the verification results for the 1×4 OptiCAM-SW array, evaluated under the same sequence of matching and mismatching scenarios previously demonstrated for the OptiCAM-DW design. In this simulation, the input data are fixed at 1111, while the stored data are systematically varied from a full match to cases with 1-, 2-, 3-, and 4-bit mismatches. As shown in the figure, during the time interval from 500 to 600 ps—when the stored data fully matches the input—the match line (ML) exhibits low optical power, indicating a successful match. As mismatches are introduced in subsequent clock cycles, the optical power on the ML correspondingly rises, clearly distinguishing mismatches from matches. It is important to note that for both sets of verifications (OptiCAM-DW and OptiCAM-SW), an optical input power of $100 \mu\text{W}$ per channel was used, yielding a 100 ps search latency. This latency can be further improved by increasing the optical input power, demonstrating the design's flexibility in trading off power for speed.




FIG. 6. Representative Cadence Virtuoso simulation setup and EDA workflow used for evaluating the proposed 1×4 OptiCAM-SW circuit.




FIG. 7. Verification of the proposed 1×4 OptiCAM-DW array under different matching conditions considering input data of 1111.




FIG. 8. Verification of the proposed 1×4 OptiCAM-SW array under different matching conditions considering input data of 1111.

B. Verification of sectioned analog OptiCAM array structure

Figure 9 presents the verification of the 1×16 sectioned analog OptiCAM array architecture shown earlier in Fig. 5 utilizing both OptiCAM-DW and OptiCAM-SW bitcell, with the input dataset to all “1”s across the 16 bits. Each section comprises a 1×4 macro array along with an independent opto-electronic stage, including balanced photodiodes, a TIA, and a VGA chain. The output of each section’s VGA drives a compute MRR tuned to $\lambda_1, \lambda_2, \lambda_3$, and λ_4 for sections 1 through 4, respectively. The through ports of these MRRs connect to a shared match line (ML), enabling optical power to appear under mismatch conditions, while remaining in a dark state when their




FIG. 9. Verification of the proposed 1×16 sectioned analog OptiCAM array under different matching conditions considering input data of 0xFFFF.

respective sections detect a match by resonating at the designated wavelengths.

To evaluate different matching scenarios under a worst-case pattern, the stored data were sequentially modified to create cases ranging from a complete match to a 1-bit mismatch introduced in section 1, then in section 2, section 3, and finally section 4. In particular, bits $Y\langle 4 \rangle$, $Y\langle 8 \rangle$, $Y\langle 12 \rangle$, and $Y\langle 16 \rangle$ were switched from “1” to “0” to introduce a 1-bit mismatch in each section one after another, while the remaining bits were held at “1.” When all stored bits matched the input, the match line (ML) exhibited the lowest optical power. As mismatches were introduced section by section, the optical power at the ML gradually increased in analog manner depending on the percentage of mismatches, reaching a maximum when all four sections contained mismatches.

Note that the analog optical power output on the ML reflects the percentage of mismatches across the sections. For example, if only section 1 (out of four sections) contains between 1 and 4 bit mismatches while sections 2, 3, and 4 remain fully matched, this corresponds to roughly 25% mismatches in the overall search, yielding an optical power of approximately one unit (P_C , as illustrated in Fig. 5). Similarly, mismatches in two, three, or all four sections result in optical power levels of $2P_C$, $3P_C$, and the maximum level, respectively, as more MRRs connected to the ML move out of resonance. This approach provides an analog behavior with a granularity defined by the section size. Note that since the GF45CLO models do not account for noise, the reported speed metrics represent an upper bound; in practice, noise may slightly degrade the maximum achievable speed.

C. Performance metrics estimation and comparison

We use a vector size of 16 to estimate the search energy and delay. Based on the simulation results in Figs. 7 and 8, it is evident that both proposed OptiCAM architectures achieve search operations with a latency of 100 ps. To evaluate the total power

consumption, we calculated both the optical and electrical power through simulations. The total required optical power is determined using the following equation:

$$P_{OPT} = n \times P_X + s \times (P_C + P_{TH}) + n \times P_{pSRAM}. \quad (1)$$

Here, P_X , P_C , and P_{TH} denote the optical input power per data channel, the power for the section-wise compute MRRs connected to the match line (ML), and the threshold power needed to differentiate between match and mismatch in each section, respectively. The terms n and s indicate the total number of bits involved in the CAM operation and the number of sections used in the calculation. P_{pSRAM} represents the static optical power required to maintain the stored data in the pSRAM for each bit. We also consider the wall-plug efficiency (WPE) of 0.23,⁴⁹ to estimate the final optical power per search (P_{OPT}/WPE). Assuming $P_X = P_C = 100 \mu\text{W}$ and a threshold power P_{TH} of $5.75 \mu\text{W}$, the total optical power—including wall-plug efficiency (WPE) and static pSRAM dissipation as reported in Ref. 44—is estimated to be 9.49 mW for a 16-bit vector. The nominal input optical power of $100 \mu\text{W}$ was chosen to avoid self-heating while still providing sufficient optical power for memory and CAM operations. In CAM operation, the lowest optical output occurs across the balanced photodiodes when all stored bits match the input, as all corresponding microring resonators are on resonance and divert the incoming light away from the output. In contrast, the worst-case detectable mismatch arises when only 1-bit mismatches, allowing single microring to remain out of resonance and transmit optical power to the output. To reliably separate the “all-match” dark state from the one-bit mismatch condition, we define the threshold power P_{TH} as the midpoint between these two output powers. Considering the microring resonator insertion loss and the two MMI splitters along the optical path (each contributing $\sim 3 \text{ dB}$), from our simulations, this midpoint is $\sim 5.75 \mu\text{W}$, ensuring correct detection of both matches and mismatches.

The electrical power consumption from the TIA and VGA is $\sim 6.65 \text{ mW}$. The reported power consumption value is obtained from circuit-level simulations using the GF45CLO monolithic PDK. In particular, the power was calculated by simulating the average current drawn from the supply (VDD) and multiplying it with the supply voltage. In our design, the TIA is implemented as a shunt-feedback inverter-based structure, while the VGA is realized using cascaded inverters to achieve rail-to-rail swing with a trade-off between design complexity, speed, and power. This results in a total power per search of 16.14 mW for the 16-bit data. Given the 100 ps search latency, the energy per bit search comes out to roughly 100.9 fJ/bit.

Note, the primary speed limitation in our current OptiCAM design arises from the electro-optic conversion stage, specifically the TIA and VGA circuits. The photonic components themselves (MRRs and PDs) are capable of operating at much higher frequencies (over 30 GHz). In this work, we intentionally adopted simple inverter-based TIA and VGA architectures, which provide sufficient bandwidth for 10 Gb/s operation while maintaining low design complexity and power consumption. In principle, higher data rates are feasible if more advanced high-bandwidth TIA and VGA designs, which would allow operation beyond 10 Gb/s.^{50,51} However, such improvements come with trade-offs in terms of higher power consumption, larger area, and increased circuit complexity.

TABLE I. Performance comparison of various CAM architectures.

References	Configuration	Search speed (ps)	Energy/bit (fJ)
TVLSI 2025 ⁵²	10T SRAM	1610.3	1.05
TNANO 2024 ⁵³	RRAM	853	0.71
LPT 2018 ⁵⁴	SOA-MZI optical	100	1000 ^a
Nanophotonics 2023 ⁵⁵	WDM optical	100	243
ACS Photonics 2025 ⁵⁶	PCM optical	100 000 ^b	1855
J. Phys. Photonics 2019 ⁵⁷	SOA optical	100	3×10^5
OFC 2025 ^{58 c}	WDM optical	20	38
This work	Optical	100	100.9

^aEnergy/bit estimated from system power (20 mW) and 100 ps search window, based on Refs. 39 and 59.

^bThe search speed is constrained by the photodetector bandwidth (10 MHz).

^cStorage is controlled by external DC supplies.

Table I summarizes the performance metrics of our proposed OptiCAM alongside prior electrical and optical CAM designs. While electrical CAMs based on SRAM or emerging devices typically achieve lower energy consumption than optical CAMs, they suffer from higher search latency, which can become a bottleneck in high-speed applications. In contrast, photonic CAMs operate at very high speeds (around 10 GHz or more), enabling rapid search operations but generally at the cost of increased energy use. Notably, our proposed OptiCAM achieves lower energy consumption than previously reported optical CAMs while maintaining comparable search latency.

In large-scale arrays, maintaining uniform resonance wavelengths across microrings can pose a significant challenge. To address this, post-fabrication variability can be compensated during an initial calibration step, which is a one-time operation and does not incur continuous power cost, as demonstrated with Ge implant trimming, SiN-film-based permanent post-fabrication trimming robust to process variations and UV-sensitive polymer claddings.^{60–62} For dynamic variations, integrated thermal phase shifters can be employed to actively compensate for wavelength drift and maintain stable operation with additional power consumption. Approximately 7.2 mW of DC power per ring is required to fine-tune a resonator over half of its FSR using the integrated thermal heater, as estimated through simulations conducted with the GF45CLO PDK. In addition, thermal crosstalk can be minimized, in practice, through careful physical design, where microring resonators are placed at optimized distances to limit heater-induced interference. Together, these techniques indicate that resonance non-uniformity and environmental drift can be effectively managed in large-scale OptiCAM arrays.

D. Applications and comparison

OptiCAM, such as content-addressable memory (CAM) architectures more broadly, has widespread use in systems that demand rapid, parallelizable memory access, such as low-latency packet processing,^{7,8} data compression,¹⁰ and caching.⁹ These architectures rely on a powerful principle, where memory is retrieved not by locality, but through information content. This principle also appears

fundamental to associative memory in the mammalian brain. Experimental neuroscience suggests that the hippocampus functions as a biological content-addressable memory.⁶³ Neurons in this region form distributed, high-dimensional codes that aggregate the “what,” “when,” and “where” of experiences into integrated spatiotemporal memories. This enables the hippocampus to efficiently encode, preserve, and infer the sequential order of experiences.

This section explores this application in detail, outlining how neural principles are integrated with OptiCAM to enable real-time, associative event-context recall. In this section, we focus on a specific application of OptiCAM as a biologically inspired memory system grounded in high-dimensional neural codes. Instead of traditional binary keys or headers, OptiCAM stores and retrieves 16-bit vectors that emulate the distributed representations observed in neural circuits.⁶⁴ These vectors encode both event and contextual information, enabling content-driven retrieval of temporal memories. We outline how neural principles are integrated with OptiCAM to enable real-time, associative event-context recall.

The efficacy of OptiCAM for hippocampus-inspired memory mechanisms is evaluated through two complementary task evaluations, each targeting a distinct memory function: associative retrieval and temporal indexing. These tasks leverage high-dimensional neural codes to simulate the binding and recall of episodic memory traces.

Task evaluation 1: Associative event-context recall. In the first task evaluation, OptiCAM’s memory is preloaded with a fixed set of event-context neural codes. Each stored item binds the visual

identity of a shape (e.g., circle) with the contextual information in which it appeared (e.g., time step 0). As a result, every memory row encodes two key elements: the event and its context, regardless of spatial or sequential locality in memory. During recall, OptiCAM is presented with candidate event-context pairs and must classify each as either “in sequence” (i.e., present in memory) or “out of sequence” (i.e., not previously stored). This task highlights OptiCAM’s ability to perform low-latency, parallel associative searches, a key feature of content-addressable memory systems.

Task evaluation 2: Content-based temporal indexing. In the second task evaluation, OptiCAM’s memory is again preloaded with a fixed set of neural codes, but this time, each memory line stores only the event. The temporal context is implicitly encoded by the row index in which the event is stored, corresponding to its occurrence in a sequence (e.g., index 0, 1, 2, ...). During recall, OptiCAM receives an event query and must infer the associated time of occurrence by retrieving the memory index in which the event was stored. This task evaluates OptiCAM’s capacity to infer “when” from “what,” mirroring the hippocampal functionality for content-driven temporal recall.

The mathematical formalism behind neural codes for event-context encoding in this work is grounded in Hyperdimensional Computing (HDC), a brain-inspired computational paradigm that represents information using high-dimensional binary vectors.⁶⁵ Rooted in cognitive science and Vector Symbolic Architectures (VSAs), HDC is motivated by the observation that the brain encodes information through distributed activation across large populations




FIG. 10. Example application of OptiCAM for hippocampus-inspired memory mechanisms. (a) Events and their associated contexts encoded as distinct shapes and recognition times, respectively. (b) Events and contexts transformed into high-dimensional binary hypervectors. (c) Event-context pairs bound via the XOR operation and compared using Hamming distance. (d) Task evaluation 1: event-context associations uncorrelated with memory locality. (e) Task evaluation 2: event-context associations correlated with memory locality.

of neurons. This process results in representations that are inherently robust to noise, generalize well, and support massively parallel operations.⁶⁴

In HDC, discrete symbols, such as the events and contexts relevant to hippocampus-inspired learning, are represented by randomly generated high-dimensional vectors that are quasi-orthogonal, enabling distinct and unambiguous representations. Complex structures, such as event-context associations, are formed through binding operations. For instance, a temporal memory trace can be created by binding an event vector $\mathcal{E}$ with a time or context vector $\mathcal{C}$ using a bitwise XOR operation. Critically, because HDC representations preserve distributed similarity, memory recall is performed by comparing a query vector to stored vectors using a similarity metric, such as Hamming distance. A low Hamming distance between a query and a stored memory trace suggests a likely match. This property aligns directly with the hardware-level operations of OptiCAM, which perform massively parallel bitwise comparisons between stored and queried vectors to support fast and efficient associative recall.

As illustrated in Fig. 10, both task evaluations are scaled to support a full space of 16 distinct visual events (shapes) and 16 unique temporal contexts, resulting in 256 possible event-context combinations. Our task evaluations leverage the aforementioned HDC encoding methods, where each unique event and context is stored in a randomly generated binary vector. Event and context associations are created by binding the individual event and context vectors with an element-wise XOR operation. In task evaluation 1, OptiCAM accurately identifies whether a given event-context pair was previously stored in memory, regardless of its storage location, demonstrating robust associative recall without regard for memory locality. In task evaluation 2, OptiCAM successfully retrieves the correct temporal index for each queried event, confirming its ability to perform content-based temporal inference where memory locality encodes additional information. We created a Python-based software simulation, where OptiCAM achieves 100% accuracy in both associative event-context recall and event-based contextual indexing. These results, along with the latency and energy estimations from Table I, underscore OptiCAM's capacity to encode, store, and retrieve structured episodic traces in a biologically plausible and hardware-efficient manner with an order of magnitude reduction in search latency compared to the existing electronic CAM implementations.

IV. CONCLUSION

In summary, we propose two optical content-addressable memory (OptiCAM) architectures capable of performing massively parallel search operations between stored and input data. In one approach, we use dual wavelengths to encode the input data with a single compute microring resonator, while in the other, a single wavelength is employed alongside two compute microring resonators to execute the CAM operation. Our CAM designs are inherently scalable and can be configured for arbitrary vector sizes by aggregating bitcell outputs through the electrical domain via photodiodes to produce binary decisions. In addition, we introduce a sectioned analog CAM architecture—a novel approach that not only enhances scalability for CAM implementations but also serves as a general solution to circumvent the wavelength constraints of a single

FSR, making it valuable for broader photonic tensor core applications. Our CAM achieves search operations within 100 ps and consumes ~ 100.9 fJ per bit. Overall, this work offers a promising path toward enabling high-speed, energy-efficient associative search and compute in emerging photonic systems.

ACKNOWLEDGMENTS

This work was supported by the Defense Advanced Research Projects Agency (DARPA) under Grant No. N660012424003 and, in part, by the National Science Foundation under Award No. CCF-2319617.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Ankur Singh contributed equally to this work.

Md Abdullah-Al Kaiser: Conceptualization (lead); Investigation (equal); Methodology (equal); Validation (equal); Writing – original draft (equal); Writing – review & editing (equal). **Ankur Singh:** Conceptualization (equal); Investigation (equal); Validation (equal); Writing – original draft (equal); Writing – review & editing (equal). **Shay Snyder:** Validation (supporting); Writing – original draft (supporting). **Hamed Poursiami:** Validation (supporting). **Maryam Parsa:** Supervision (equal). **Akhilesh Jaiswal:** Supervision (equal); Writing – review & editing (equal).

DATA AVAILABILITY

The data that support the findings of this study are available within the article.

REFERENCES

- ¹ M. Kang, S. K. Gonugondla, A. Patil, and N. R. Shanbhag, “A multi-functional in-memory inference processor using a standard 6T SRAM array,” *IEEE J. Solid-State Circuits* **53**, 642–655 (2018).
- ² J. Zhang, Z. Wang, and N. Verma, “In-memory computation of a machine-learning classifier in a standard 6T SRAM array,” *IEEE J. Solid-State Circuits* **52**, 915–924 (2017).
- ³ A. Azab, M. Khasawneh, S. Alrabaei, K.-K. R. Choo, and M. Sarsour, “Network traffic classification: Techniques, datasets, and challenges,” *Digital Commun. Networks* **10**, 676–692 (2024).
- ⁴ A. Kazemi, M. M. Sharifi, A. F. Laguna, F. Müller, R. Rajaei, R. Olivo, T. Kämpfe, M. Niemier, and X. S. Hu, “In-memory nearest neighbor search with FeFET multi-bit content-addressable memories,” in *2021 Design, Automation and Test in Europe Conference and Exhibition (DATE)* (IEEE, 2021), pp. 1084–1089.
- ⁵ D. E. Taylor, “Survey and taxonomy of packet classification techniques,” *ACM Comput. Surv.* **37**, 238–275 (2005).
- ⁶ W. Li, X. Li, and H. Li, “MEET-IP: Memory and energy efficient TCAM-based IP lookup,” in *2017 26th International Conference on Computer Communication and Networks (ICCCN)* (IEEE, 2017), pp. 1–8.

- ⁷A. J. McAuley and P. Francis, "Fast routing table lookup using CAMs," in *IEEE INFOCOM'93 The Conference on Computer Communications, Proceedings* (IEEE, 1993), pp. 1382–1391.
- ⁸C. Xu, S. Chen, J. Su, S.-M. Yiu, and L. C. K. Hui, "A survey on regular expression matching for deep packet inspection: Applications, algorithms, and hardware platforms," *IEEE Commun. Surv. Tutorials* **18**, 2991–3029 (2016).
- ⁹S. Moradi, N. Qiao, F. Stefanini, and G. Indiveri, "A scalable multicore architecture with heterogeneous memory structures for dynamic neuromorphic asynchronous processors (DYNAPs)," *IEEE Trans. Biomed. Circuits Syst.* **12**, 106–122 (2017).
- ¹⁰K.-J. Lin and C.-W. Wu, "A low-power CAM design for LZ data compression," *IEEE Trans. Comput.* **49**, 1139–1145 (2000).
- ¹¹J. P. Wade and C. G. Sodini, "A ternary content addressable search engine," *IEEE J. Solid-State Circuits* **24**, 1003–1013 (1989).
- ¹²R. Karam, R. Puri, S. Ghosh, and S. Bhunia, "Emerging trends in design and applications of memory-based computing and content-addressable memories," *Proc. IEEE* **103**, 1311–1330 (2015).
- ¹³G. Pedretti, C. E. Graves, S. Serebryakov, R. Mao, X. Sheng, M. Foltin, C. Li, and J. P. Strachan, "Tree-based machine learning performed in-memory with memristive analog CAM," *Nat. Commun.* **12**, 5806 (2021).
- ¹⁴K. Ni, X. Yin, A. F. Laguna, S. Joshi, S. Dünkler, M. Trentzsch, J. Müller, S. Beyer, M. Niemier, X. S. Hu, and S. Datta, "Ferroelectric ternary content-addressable memory for one-shot learning," *Nat. Electron.* **2**, 521–529 (2019).
- ¹⁵R. Mao, B. Wen, A. Kazemi, Y. Zhao, A. F. Laguna, R. Lin, N. Wong, M. Niemier, X. S. Hu, X. Sheng *et al.*, "Experimentally validated memristive memory augmented neural network with efficient hashing and similarity search," *Nat. Commun.* **13**, 6284 (2022).
- ¹⁶K. Pagiamtzis and A. Shekholeslami, "Content-addressable memory (CAM) circuits and architectures: A tutorial and survey," *IEEE J. Solid-State Circuits* **41**, 712–727 (2006).
- ¹⁷Y.-J. Chang and Y.-H. Liao, "Hybrid-type CAM design for both power and performance efficiency," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* **16**, 965–974 (2008).
- ¹⁸C.-S. Lin, J.-C. Chang, and B.-D. Liu, "A low-power precomputation-based fully parallel content-addressable memory," *IEEE J. Solid-State Circuits* **38**, 654–662 (2003).
- ¹⁹N. Onizawa, S. Matsunaga, V. C. Gaudet, and T. Hanyu, "High-throughput low-energy content-addressable memory based on self-timed overlapped search mechanism," in *2012 IEEE 18th International Symposium on Asynchronous Circuits and Systems* (IEEE, 2012), pp. 41–48.
- ²⁰N. Onizawa, S. Matsunaga, V. C. Gaudet, W. J. Gross, and T. Hanyu, "High-throughput CAM based on a synchronous overlapped search scheme," *IEEE Electron. Express* **10**, 20130148 (2013).
- ²¹N. Onizawa, S. Matsunaga, V. C. Gaudet, W. J. Gross, and T. Hanyu, "High-throughput low-energy self-timed CAM based on reordered overlapped search mechanism," *IEEE Trans. Circuits Syst. I* **61**, 865–876 (2013).
- ²²K. Cho, H. Choi, I. J. Jung, J. Oh, T. W. Oh, K. Kim, G. Kim, T. Choi, C. Sim, T. Song, and S.-O. Jung, "SRAM write- and performance-assist cells for reducing interconnect resistance effects increased with technology scaling," *IEEE J. Solid-State Circuits* **57**, 1039–1048 (2022).
- ²³V. Lines, A. Ahmed, P. Ma, S. Ma, R. McKenzie, H.-S. Kim, and C. Mar, "66 MHz 2.3 M ternary dynamic content addressable memory," in *Records of the IEEE International Workshop on Memory Technology, Design and Testing* (IEEE, 2000), pp. 101–105.
- ²⁴H. Noda, K. Inoue, M. Kuroiwa, F. Igaue, K. Yamamoto, H. J. Mattausch, T. Koide, A. Amo, A. Hachisuka, S. Soeda *et al.*, "A cost-efficient high-performance dynamic TCAM with pipelined hierarchical searching and shift redundancy architecture," *IEEE J. Solid-State Circuits* **40**, 245–253 (2005).
- ²⁵B. Rajendran, R. W. Cheek, L. A. Lastras, M. M. Franceschini, M. J. Breitwisch, A. G. Schrott, J. Li, R. K. Montoye, L. Chang, and C. Lam, "Demonstration of CAM and TCAM using phase change devices," in *2011 3rd IEEE International Memory Workshop (IMW)* (IEEE, 2011), pp. 1–4.
- ²⁶Q. Guo, X. Guo, Y. Bai, and E. İpek, "A resistive TCAM accelerator for data-intensive computing," in *Proceedings of the 44th Annual IEEE/ACM International Symposium on Microarchitecture* (IEEE, 2011), pp. 339–350.
- ²⁷P. Junsangsri, J. Han, and F. Lombardi, "Design and comparative evaluation of a PCM-based CAM (content addressable memory) cell," *IEEE Trans. Nanotechnol.* **16**, 359–363 (2017).
- ²⁸S. Matsunaga, K. Hiyama, A. Matsumoto, S. Ikeda, H. Hasegawa, K. Miura, J. Hayakawa, T. Endoh, H. Ohno, and T. Hanyu, "Standby-power-free compact ternary content-addressable memory cell chip using magnetic tunnel junction devices," *Appl. Phys. Express* **2**, 023004 (2009).
- ²⁹W. Xu, T. Zhang, and Y. Chen, "Design of spin-torque transfer magnetoresistive ram and CAM/TCAM with high sensing and search speed," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* **18**, 66–74 (2009).
- ³⁰S. Matsunaga, A. Katsumata, M. Natsui, S. Fukami, T. Endoh, H. Ohno, and T. Hanyu, "Fully parallel 6T-2MTJ nonvolatile TCAM with single-transistor-based self match-line discharge control," in *2011 Symposium on VLSI Circuits-Digest of Technical Papers* (IEEE, 2011), pp. 298–299.
- ³¹A. Kazemi, S. Sahay, A. Saxena, M. M. Sharifi, M. Niemier, and X. S. Hu, "A flash-based multi-bit content-addressable memory with Euclidean squared distance," in *2021 IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED)* (IEEE, 2021), pp. 1–6.
- ³²S. George, N. Jao, A. K. Ramanathan, X. Li, S. K. Gupta, J. Sampson, and V. Narayanan, "Integrated CAM-RAM functionality using ferroelectric FETs," in *2020 21st International Symposium on Quality Electronic Design (ISQED)* (IEEE, 2020), pp. 81–86.
- ³³K. Eshraghian, K.-R. Cho, O. Kavehei, S.-K. Kang, D. Abbott, and S.-M. S. Kang, "Memristor MOS content addressable memory (MCAM): Hybrid architecture for future high performance search engines," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* **19**, 1407–1417 (2010).
- ³⁴M. Imani, S. Patil, and T. Š. Rosing, "Approximate computing using multiple-access single-charge associative memory," *IEEE Trans. Emerging Top. Comput.* **6**, 305–316 (2016).
- ³⁵C. Li, C. E. Graves, X. Sheng, D. Miller, M. Foltin, G. Pedretti, and J. P. Strachan, "Analog content-addressable memories with memristors," *Nat. Commun.* **11**, 1638 (2020).
- ³⁶C. E. Graves, C. Li, X. Sheng, W. Ma, S. R. Chalamalasetti, D. Miller, J. S. Ignowski, B. Buchanan, L. Zheng *et al.*, "Memristor TCAMs accelerate regular expression matching for network intrusion detection," *IEEE Trans. Nanotechnol.* **18**, 963–970 (2019).
- ³⁷Y. London, T. Van Vaerenbergh, A. J. Rizzo, P. Sun, J. Hulme, G. Kurczveil, A. Seyed, B. Wang, X. Zeng, Z. Huang *et al.*, "Energy efficiency analysis of comb source carrier-injection ring-based silicon photonic link," *IEEE J. Sel. Top. Quantum Electron.* **26**, 1–13 (2019).
- ³⁸Y. London, T. Van Vaerenbergh, L. Ramini, A. J. Rizzo, P. Sun, G. Kurczveil, A. Seyed, J. Rhim, M. Fiorentino, and K. Bergman, "Performance requirements for terabit-class silicon photonic links based on cascaded microring resonators," *J. Lightwave Technol.* **38**, 3469–3477 (2020).
- ³⁹S. Pitris, C. Vagionas, P. Maniotis, G. T. Kanellos, and N. Pleros, "An optical content addressable memory cell for address look-up at 10 Gb/s," *IEEE Photonics Technol. Lett.* **28**, 1790–1793 (2016).
- ⁴⁰G. Mourgias-Alexandris, C. Vagionas, A. Tsakyridis, P. Maniotis, and N. Pleros, "All-optical 10 Gb/s ternary-CAM cell for routing look-up table applications," *Opt. Express* **26**, 7555–7562 (2018).
- ⁴¹M. A. Z. Quashef and M. K. Alam, "Ultracompact photonic integrated content addressable memory using phase change materials," *Opt. Quantum Electron.* **54**, 182 (2022).
- ⁴²Y. Alkabani, M. Miscuglio, V. J. Sorger, and T. El-Ghazawi, "OE-CAM: A hybrid opto-electronic content addressable memory," *IEEE Photonics J.* **12**, 1–14 (2020).
- ⁴³Y. London, T. Van Vaerenbergh, L. Ramini, C. Li, C. E. Graves, M. Fiorentino, and R. G. Beausoleil, "WDM ternary content-addressable memory for optical links," in *2023 IEEE Silicon Photonics Conference (SiPhotonics)* (IEEE, 2023), pp. 1–2.
- ⁴⁴M. A.-A. Kaiser *et al.*, "Design of energy-efficient cross-coupled differential photonic-SRAM (pSRAM) bitcell for high-speed on-chip photonic memory and compute systems," *arXiv:2503.19544* (2025).
- ⁴⁵M. A.-A. Kaiser, S. Sunder, A. P. Jacob, and A. R. Jaiswal, "A mixed-signal photonic SRAM-based high-speed energy-efficient photonic tensor core with novel electro-optic ADC," in *Proceedings of the 62nd ACM/IEEE Design Automation Conference* (IEEE, 2025), pp. 1–6.

- ⁴⁶M. A. Kaiser, S. Sunder, A. R. Jaiswal, and A. P. Jacob, "Demonstration of a cross-coupled differential regenerative photonic latch for integrated ultra-fast on-chip memory," in *Proceedings of the 2025 IEEE International Electron Devices Meeting (IEDM)* (IEEE, 2025).
- ⁴⁷S. Chakraborty, M. A.-A. Kaiser, and A. R. Jaiswal, "Design of an energy-efficient automatic calibration circuit for high speed photonic cross-coupled memory," *IEEE Photonics Technol. Lett.* **38**, 164 (2025).
- ⁴⁸M. Rakowski, C. Meagher, K. Nummy, A. Aboketaf, J. Ayala, Y. Bian, B. Harris, K. Mclean, K. McStay, A. Sahin *et al.*, "45 nm CMOS—Silicon photonics monolithic technology (45CLO) for next-generation, low power and high speed optical interconnects," in *Optical Fiber Communication Conference* (Optica Publishing Group, 2020), p. T3H-3.
- ⁴⁹S. A. Blokhin *et al.*, "High power single mode 1300-nm superlattice based VCSEL: Impact of the buried tunnel junction diameter on performance," *IEEE J. Quantum Electron.* **58**, 1–15 (2022).
- ⁵⁰L. Aschei, N. Cordioli, P. Rossi, D. Montanari, R. Castello, and D. Manstretta, "A 42-GHz TIA in 28-nm CMOS with less than 1.8% THD for optical coherent receivers," *IEEE Solid-State Circuits Lett.* **3**, 238–241 (2020).
- ⁵¹D.-W. Park, D. R. Utomo, B. H. Lam, S.-G. Lee, and J.-P. Hong, "A 230–260-GHz wideband and high-gain amplifier in 65-nm CMOS based on dual-peak $G_{\text{max-core}}$," *IEEE J. Solid-State Circuits* **54**, 1613–1623 (2019).
- ⁵²Z. Zhang, Z. Chen, J. Wang, G. Xie, and G. Liu, "Reconfigurable 10T SRAM for energy-efficient CAM operation and in-memory computing," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* **33**, 1065 (2025).
- ⁵³M. Kumar, M.-H. Wu, T.-H. Hou, and M. Suri, "CMOS-RRAM based non-volatile ternary content addressable memory (nvTCAM)," *IEEE Trans. Nanotechnol.* **23**, 203–207 (2024).
- ⁵⁴G. Mourgas-Alexandris, C. Vagionas, A. Tsakyridis, P. Maniotis, and N. Pleros, "Optical content addressable memory matchline for 2-bit address look-up at 10 Gb/s," *IEEE Photonics Technol. Lett.* **30**, 809–812 (2018).
- ⁵⁵Y. London, T. Van Vaerenbergh, L. Ramini, A. Descos, L. Buonanno, J. Youn, C. Li, C. E. Graves, M. Fiorentino, and R. G. Beausoleil, "Multiplexing in photonics as a resource for optical ternary content-addressable memory functionality," *Nanophotonics* **12**, 4137–4155 (2023).
- ⁵⁶S. Cheng, X. Li, Y. Li, Y. Feng, W. Jin, Y. Zhang, Z. Liu, Y. Qin, X. Yang, and L. Yuan, "All-optical content-addressable memory (CAM) based on all-fiber platform," *ACS Photonics* **12**, 2997 (2025).
- ⁵⁷C. Vagionas, A. Tsakyridis, G. Mourgas-Alexandris, S. Pitris, P. Maniotis, N. Pleros, and T. Alexoudi, "Optical memory architectures for fast routing address look-up (AL) table operation," *J. Phys.: Photonics* **1**, 044005 (2019).
- ⁵⁸T. Moschos, C. Pappas, S. Kovaos, I. Roumpos, A. Prapas, A. Tsakyridis, M. Moralis-Pegios, C. Vagionas, Y. London, T. Van Vaerenbergh *et al.*, "A 50 Gb/s WDM silicon photonic ternary content addressable memory cell," in *Optical Fiber Communication Conference* (Optica Publishing Group, 2025), p. Th1F-3.
- ⁵⁹P. Maniotis and N. Pleros, "All-optical ternary-content addressable memory (T-CAM) cell and row architectures for address lookup at 20 Gb/s," *Opt. Quantum Electron.* **49**, 348 (2017).
- ⁶⁰H. Jayatilaka, H. Frish, R. Kumar, J. Heck, C. Ma, M. Sakib, D. Huang, and H. Rong, "Post-fabrication trimming of silicon photonic ring resonators at wafer-scale," *J. Lightwave Technol.* **39**, 5083–5088 (2021).
- ⁶¹A. H. Atabaki, A. A. Eftekhari, M. Askari, and A. Adibi, "Accurate post-fabrication trimming of ultra-compact resonators on silicon," *Opt. Express* **21**, 14139–14145 (2013).
- ⁶²L. Zhou, K. Okamoto, and S. J. B. Yoo, "Athermalizing and trimming of slotted silicon microring resonators with UV-sensitive PMMA upper-cladding," *IEEE Photonics Technol. Lett.* **21**, 1175–1177 (2009).
- ⁶³N. J. Fortin, K. L. Agster, and H. B. Eichenbaum, "Critical role of the hippocampus in memory for sequences of events," *Nat. Neurosci.* **5**, 458–462 (2002).
- ⁶⁴C. Eliasmith, *How to Build a Brain: A Neural Architecture for Biological Cognition* (Oxford University Press, 2013).
- ⁶⁵D. Kleyko, A. Rahimi, D. A. Rachkovskij, E. Osipov, and J. M. Rabaey, "Classification and recall with binary hyperdimensional computing: Tradeoffs in choice of density and mapping characteristics," *IEEE Trans. Neural Networks Learn. Syst.* **29**, 5880–5898 (2018).