

DARTH-PUM: A Hybrid Processing-Using-Memory Architecture

Ryan Wong

Univ. of Illinois Urbana-Champaign
Urbana, IL, USA

Ben Feinberg

Sandia National Laboratories
Albuquerque, NM, USA

Saugata Ghose

Univ. of Illinois Urbana-Champaign
Urbana, IL, USA

Abstract

Analog *processing-using-memory* (PUM; a.k.a. *in-memory computing*) makes use of electrical interactions inside memory arrays to perform bulk *matrix–vector multiplication* (MVM) operations. However, many popular matrix-based kernels need to execute non-MVM operations, which analog PUM cannot directly perform. To retain its energy efficiency, analog PUM architectures augment memory arrays with CMOS-based domain-specific fixed-function hardware to provide complete kernel functionality, but the difficulty of integrating such specialized CMOS logic with memory arrays has largely limited analog PUM to being an accelerator for machine learning inference, or for closely related kernels. An opportunity exists to harness analog PUM for general-purpose computation: recent works have shown that memory arrays can also perform Boolean PUM operations, albeit with very different supporting hardware and electrical signals than analog PUM.

We propose *DARTH-PUM*, a general-purpose hybrid PUM architecture that tackles key hardware and software challenges to integrating analog PUM and digital PUM. We propose optimized peripheral circuitry, coordinating hardware to manage and interface between both types of PUM, an easy-to-use programming interface, and low-cost support for flexible data widths. These design elements allow us to build a practical PUM architecture that can execute kernels fully in memory, and can scale easily to cater to domains ranging from embedded applications to large-scale data-driven computing. We show how three popular applications (AES encryption, convolutional neural networks, large-language models) can map to and benefit from *DARTH-PUM*, with speedups of 59.4×, 14.8×, and 40.8× over an analog+CPU baseline.

CCS Concepts: • **Computer systems organization** → **Heterogeneous (hybrid) systems**; • **Hardware** → *Memory and dense storage*.

Keywords: Processing-Using-Memory; In-Memory Computing; Hybrid Architecture

ACM acknowledges that this contribution was authored or co-authored by an employee, contractor, or affiliate of the United States government. As such, the United States government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for government purposes only. Request permissions from owner/author(s).

ASPLOS '26, Pittsburgh, PA, USA

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2359-9/2026/03

<https://doi.org/10.1145/3779212.3790151>

ACM Reference Format:

Ryan Wong, Ben Feinberg, and Saugata Ghose. 2026. DARTH-PUM: A Hybrid Processing-Using-Memory Architecture. In *Proceedings of the 31st ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS '26)*, March 22–26, 2026, Pittsburgh, PA, USA. ACM, New York, NY, USA, 20 pages. <https://doi.org/10.1145/3779212.3790151>

1 Introduction

Modern systems often suffer from the overheads of moving data between the CPU cores and the memory subsystem, with latencies and energy costs that are several orders of magnitude greater than those of computational instructions [12, 21, 101]. These overheads are particularly harmful for many modern and emerging workloads from machine learning [22, 33, 135], to databases [23, 49, 113], to scientific computing [58, 109, 110, 118]. *Processing-using-memory* (PUM; also known as *in-memory computing* or *in situ computing*) offers a potential solution to circumvent data movement overheads, by enabling intelligent memories that leverage the electrical interactions of the memory devices to realize a meaningful computation.

PUM has attracted substantial interest in both academia [38, 82, 102, 137, 166] and industry [59, 60, 90] due to its potential for large performance and energy improvements. PUM can be classified as analog or digital, with both types having been explored in a variety of applications [25, 79, 80, 122, 133, 148, 165, 166] and memory device types [1, 26, 30, 34, 35, 71, 120].

Analog PUM is usually characterized by its ability to store multiple bits per cell and compute fast, though sometimes approximate, *matrix–vector* multiplies by using a combination of Ohm’s Law and Kirchhoff’s Current Laws. Although promising, analog PUM is susceptible to noise and other analog non-idealities; therefore, analog PUM has primarily been applied as a *matrix–multiply* accelerator for algorithms resilient to such errors (e.g., machine learning [2, 9, 19, 20, 127] and scientific computing [25, 126, 152, 171]). While analog PUM has spurred significant research and commercial prototypes [59, 60, 90], it requires expensive peripheral circuits (e.g., analog-to-digital and digital-to-analog converters), and either dedicated specialized function units or constant host–PUM communication to perform non-matrix operations required for machine learning and scientific computing. This highly constrains analog PUM’s applicability and benefits.

In contrast, digital PUM (i.e., Boolean PUM) is often much more error-resilient and can perform general-purpose computations. Notably, digital PUM uses single-bit cells and often

realizes a Boolean operator between memory devices [71, 121, 138]. When executed sequentially, these sequences are able to perform *any* computation, similar to a CPU. Thus, digital PUM has been applied to a variety of data-intensive workloads [38, 102, 137]. Despite having been explored in a variety of memory devices, including SRAM [1, 24, 29, 61, 143, 144], DRAM [31, 32, 38, 51, 81, 98–100, 102, 119–121, 163, 164], ReRAM [71, 79, 137, 138], MRAM [4–6, 35, 158], PCM [14, 34, 41, 82], and NAND flash [15, 30, 62, 104, 124, 139, 140], digital PUM often performs substantially worse compared to its analog counterparts when executing matrix-based applications.

Prior works have shown that analog PUM can be specialized for specific applications (e.g., hyper-dimensional computing [85], transformers [80, 125]) with the addition of digital PUM built to support application-specific functionality. Although these works leverage both domains of PUM computation, they do not harness the generality of digital PUM in a way that can be extended to other applications. Our goal in this work is to demonstrate that with careful design, a true hybrid PUM architecture can integrate both analog and digital PUM such that the same architecture can provide a best-of-both-worlds solution across many domains.

We propose DARTH-PUM (*D*igital-*A*nalog *R*esistive *T*iles for *H*ybrid *P*rocessing-*U*sing-*M*emory), which combines both analog and digital PUM onto the same chip. DARTH-PUM avoids the need for external (i.e., PUM-CPU) data movement or application-specific functional units, by taking advantage of the generalizable properties of digital PUM. We explore the design choices and functionality required to enable a hybrid PUM chip capable of executing a variety of applications without hardware redesign. DARTH-PUM can execute a variety of applications (and datasets) by supporting auxiliary operations needed by analog PUM in nearby digital PUM units. DARTH-PUM provides flexible functionality, with support for previously proposed techniques that optimize both analog (e.g., accuracy [125], noise mitigation [151], bit-slicing [149]) and digital (e.g., vector computing [1], bit-pipelining [137]) PUM independently. Furthermore, DARTH-PUM supports a variety of bit width operands in both digital and analog PUM.

We make a series of key insights about flexible hybrid PUM. (1) It can substantially reduce communication costs, as well as capacity overheads, of digital PUM. (2) It can support a wide variety of bit-width operands without the need for re-designing expensive post-processing circuitry; furthermore, digital PUM can support all auxiliary operations needed by analog-PUM-accelerated kernels. (3) Bottlenecks with digital-to-analog conversions and vice-versa can originate from both analog and digital periphery. The rate at which analog and digital PUM arrays produce and consume data can directly affect the number of required peripheries (e.g., ADCs). We analyze which functionalities are required for hybrid PUM (e.g., memory I/O), overheads (e.g., inter-tile communication) when executing these applications, and how to cleanly rate match the analog-to-digital conversions (and vice-versa).

We demonstrate the performance benefits of DARTH-PUM and highlight its flexible nature by evaluating it on a collection of three different workloads: cryptography, convolutional neural networks, and large-language models. Compared to an analog-only PUM accelerator coupled with a state-of-the-art CPU, DARTH-PUM achieves a 59.4×, 14.8×, and 40.8× performance improvement while reducing energy by 39.6×, 51.2×, and 110.7× for the aforementioned workloads, respectively.

2 Background

In this section, we provide a brief summary of background information related to analog PUM and digital PUM.

2.1 Processing-in-Memory

Processing-in-memory (PIM) broadly encompasses the technique of bringing computation closer to the memory (or storage) substrate. PIM can be broadly divided into two categories: *processing-near-memory* (PNM) and *processing-using-memory* (PUM). PUM [48, 56, 57, 63, 64, 75, 92, 106, 114] adds compute capabilities to the memory by integrating conventional logic close to the memory arrays. In contrast, PUM enables computation in the memory by leveraging the electrical characteristics of the memory devices to realize meaningful computation. This work focuses on the trade-offs associated with an architecture capable of both analog and digital PUM.

2.2 Processing-Using-Memory

Both analog and digital PUM have been explored for a number of application contexts [15, 17, 25, 45, 77, 122, 123, 154], across a wide range of memory technologies [1, 30, 34, 35, 73, 120, 121, 150]. In this work, we use ReRAM because of its relatively mature support for both analog and digital operations.

2.2.1 Analog PUM. Analog PUM is often characterized by its ability to store multiple bits per device and has been explored as both an accelerator for *matrix-vector multiplication* (MVM) and as an analog content-addressable memory (CAM; Section 8). In this work, we focus on using analog PUM for efficient MVM, as DARTH-PUM builds upon this functionality.

Computation in analog PUM leverages the shared bitline within the memory array to execute add operations between all devices along the bitline in parallel. Therefore, the bitline can efficiently realize a multiply-and-accumulate (MAC) operation.¹ Furthermore, by activating the entire analog array in parallel, full vector-matrix (or matrix-vector) multiplies can be realized. Despite its promise as an efficient MVM accelerator, analog computation is susceptible to noise, limiting its applicability. Therefore, analog PUM has primarily been explored as part of application domains such as machine learning [11, 20, 44, 52, 122, 147, 162], cloud analytics [17], scientific computing [25, 126, 152, 171], neuromorphic computing [45], and graph applications [84, 128].

¹Analog PUM can also use single-level devices (i.e., one bit per device), with the computation remaining a MAC.

Figure 1 shows an example of a 2×2 matrix multiplied by a 2×1 input vector. To execute the MVM, the matrix values are first programmed into the crossbar, usually in the form of resistances (or conductance, $G = \frac{1}{R}$). Each element of the input vector is then applied as a voltage to the wordlines (1). Using Ohm's Law ($I = \frac{V}{R}$), an element-wise multiply can be realized between the input and device (2a). The current through each device is proportional to the input voltage times the device conductance (2b). Similarly, by Kirchhoff's Current Law, the current through each device along the bitline is summed, realizing the accumulation. Finally, to read out the value, analog PUM frequently employs the use of analog-to-digital converters (ADCs), which convert the analog current or voltages into the discrete digital value of the MVM output (3).

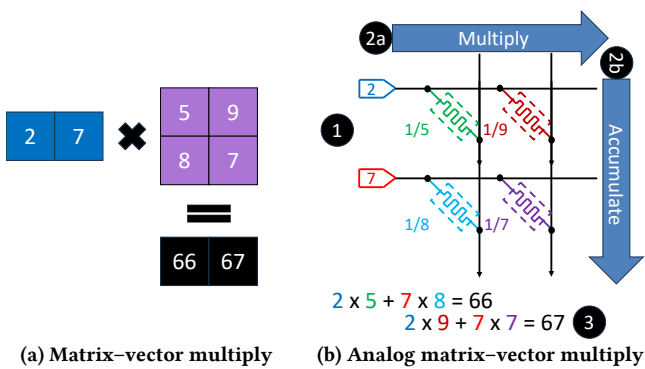


Figure 1. Matrix-vector multiply with analog PUM.

Bit-Slicing. For an analog PUM device, its effective maximum bit precision is approximately 6–12 bits, often constrained by the precision of the ADC used for the write-verify programming schemes. To realize larger-width operands, *bit-slicing* techniques [122, 149] split each stored N -bit value into smaller $\frac{N}{M}$ bit values, where M is the number of bits that can reliably be stored in a single device. Figure 2 shows how a 4-bit value can be bit-sliced into two 2-bit values, which are then separated into different arrays (e.g., Array 1 stores Value[3:2], while Array 0 stores Value[1:0]). During MVM, each array executes the multiply independently, resulting in multiple partial products. To recombine the partial products into a single result, a post-processing step shifts the partial products based on the relative bit positions of each bit slice. In the example, the products from Array 1 must be shifted left by two bit-positions to account for bit-slicing. In contrast, the products from Array 0 are not shifted (i.e., shift by zero). The results are then added together forming the full output.

Bit-slicing can also be applied to input values, alleviating the need for high-range digital-to-analog converters (DACs). An N -bit input can be sequentially applied, which results in N partial products, which are handled during post-processing. Notably, this recombination step is similar to the long-multiplication algorithm. An in-depth discussion on bit-slicing for analog PUM can be found in prior work [149].




Figure 2. Bit-slicing matrix values.

Peripheral Circuitry. Peripheral circuitry remains a substantial challenge for many analog PUM accelerators. First, the array must be accompanied by costly analog-to-digital converters (ADCs), which convert the analog output current (or voltage) back into the discrete digital bit values. Second, digital-to-analog converters (DACs) must be used to convert the digital inputs into analog input voltages. Together, these periphery consume a non-trivial amount of energy and area, degrading the analog PUM efficiency. Furthermore, analog PUM requires additional logic for post-processing operations beyond the shift-and-add operation. For example, additional logic, often in the form of custom CMOS components, must be added to the analog chip to execute auxiliary functions (e.g., activation functions in machine learning workloads), reducing the performance and flexibility of analog PUM.

Analog-to-Digital Converters. For analog PUM, the type of ADC used can affect performance, power, and precision. *Ramp* ADCs [130, 150] operate by comparing the input voltage with a predetermined reference voltage. After each comparison, the reference voltage is increased and compared again. By iterating linearly over N increments, ramp ADCs take N cycles to fully digitize the analog input. While ramp ADCs are often slower than other ADC types, they can share the power-hungry reference generator, increasing efficiency in applications where parallel readout is beneficial. *Successive-approximation register* (SAR) ADCs [70] by contrast use binary search over the ADC range, requiring $\log(N)$ comparisons to find the correct output value. As an iterative binary search cannot be easily parallelized, SAR ADCs used for analog PUM typically use a high-speed analog multiplexer to iterate over all the outputs. In the context of bit-sliced systems, this approach has another advantage, since only a single value is output from the ADC at a time, thus only a single set of shift-and-add logic is required.

Analog Non-Idealities. Analog devices are susceptible to noise. For example, many devices [142, 155] exhibit non-linear resistances (and therefore non-linear I-V curves), causing imprecision during programming. Other non-idealities may affect device state over time (e.g., drift), and random perturbations on the bitline or device may affect the computation result. Thus, analog PUM is often treated as an approximate solution

and has primarily been applied to applications with inherent error tolerance (e.g., ML [26, 116]). These non-idealities have been widely observed in a variety of prior works [151].

Handling Negative Numbers. Negative number representation can be a challenge for analog arrays as resistance is a strictly positive quantity. Two widely used representations are offset subtraction [7, 45, 46, 91, 122, 159] and differential cell pairs [19, 36, 55, 88, 160]. Offset subtraction (Figure 3a) splits the conductance range into two regions, wherein a conductance below the midpoint is a negative value, while one above the midpoint is positive. To compute the final value, an offset is subtracted from the result of the ADC. Differential cell pairs (Figure 3b) utilize two discrete devices to represent negative and positive values while applying input voltages of opposite polarity. Unlike offset subtraction, differential cell pairs do not require an offset scalar as the resulting current is directly proportional to the result. This work utilizes differential pairs, which has been shown to have increased resilience against parasitic effects [153]. For a more detailed discussion on number representations, we defer to prior work [151].




Figure 3. Number representations for the range [-3, +3].

2.2.2 Digital PUM. Digital (i.e., Boolean) PUM has demonstrated potential to accelerate data-intensive applications [102, 119, 121, 137] due to its ability to implement bulk bitwise operators with minimal errors. Two features that often separate digital and analog PUM are that digital PUM (1) often stores only a single bit per device, and (2) operates by realizing a Boolean logic operator between two input cells, creating discrete outputs of ‘0’ and ‘1’. Digital PUM has been explored in a wide variety of contexts including SRAM [1, 24, 29], DRAM [31, 38, 81, 102, 121], ReRAM [71, 137, 138], MRAM [35, 158], PCM [14, 34, 82], and NAND flash [30, 104]. Each digital PUM implementation chooses one of several *logic families* [37, 71, 138], which correspond to the specific device, Boolean primitives, and corresponding voltages and operating parameters needed to realize the digital PUM operations.

One recently proposed logic family, OSCAR [138], implements the NOR and OR primitives using ReRAM devices. Figure 4 shows an example of the OSCAR NOR primitive, which uses four ReRAM devices as two inputs, an output, and a load resistor. To execute the NOR primitive, the two input bitlines are driven to V_{NOR} , the output bitline is driven to $V_{NOR} + \Delta$, and the load resistor bitline is set to GND (with the resistor itself set between R_{on} and R_{off}). The wordline is then electrically disconnected, leaving it in a floating state; thus, the wordline is perturbed by the interactions between the input cell states, the output cell state, V_{NOR} , and $V_{NOR} + \Delta$. Depending on the cell states, a current flows through the output device towards ground and changes the state of the output device.² To improve performance, entire columns of devices may be activated in parallel (with all wordlines in the array floated), allowing for entire vectors of operands to simultaneously execute the *same* primitive.




Figure 4. Boolean PUM operation with OSCAR NOR.

Beyond OSCAR NOR, prior works have realized several different Boolean operators (e.g., AND [104, 119], OR [119], NOR [71, 138], XOR [37], NOT [1, 121], IMPLY [72, 73]), which provide different trade-offs and expressiveness. Because logic families are typically Boolean complete, the PUM operations can be chained together to realize more complex operations (e.g., add, sub, cosine, square root). Digital PUM often relies on bit-serial computation, wherein each bit of the output is computed sequentially, which significantly increases the latency of a single operation [137]. Despite this overhead, digital PUM can execute *any* arbitrary function, enabling the same flexibility afforded to conventional CPUs.

Bit-Pipelining. To improve throughput, RACER [137] proposes *bit-pipelining* for digital PUM operations. In RACER, an N -bit pipeline is composed of N discrete digital PUM arrays. Accordingly, an N -bit value is *bit-striped* across all arrays (i.e., each array contains a single bit position of the value). Therefore, each array in the pipeline can independently execute a different operation, enabling N times improvement in throughput. Figure 5 shows how data is striped in a 4-bit pipeline. Each *vector register* (VR) contains four *elements* (A–D in our example), which are distributed to different rows of the pipeline. Each element is bit-striped across the arrays (e.g., for all elements in VR2, bit 0 is placed in column 2 of array 0). Because operands are striped across different arrays in the

²By using a fourth resistor, OSCAR can properly balance the voltage division across the four devices, increasing fidelity of the operation.




Figure 5. Data layout for digital PUM pipeline.

same row, different rows may contain different values; therefore, a pipeline with M rows can execute an M -element vector in parallel. Due to design constraints (e.g., limited peripheral circuitry), vector register elements can only interact with other bits along the same row (e.g., VR_0 element B can only execute a Boolean primitive with VR_N element B). Although bit-pipelining can substantially improve the performance for many data-parallel applications, this technique alone leaves substantial performance opportunities compared to analog PUM for matrix-based algorithms.

3 Motivation: The Potential of Hybrid PUM

The limited flexibility of analog PUM has led researchers to primarily focus on application- or domain-specific accelerators, designed and tuned with specific peripherals and specialized function units (SFUs). In contrast, digital PUM has enabled a wide range of general-purpose accelerators, but leaves a large gap in performance for MVM-style kernels.

Hybrid PUM can address these issues by combining the complementary behavior of analog and digital PUM, without the need for SFUs. However, as we show in this section, simply combining two high-performance schemes is insufficient to realize the full potential of hybrid PUM, as combining the two domains requires careful balancing and coordination between them. To illustrate why, we compare three different architectures, as shown in Figure 6: (1) digital PUM (D), a RACER pipelined microarchitecture [137] that performs NOR operations using the OSCAR logic family [138] in ReRAM arrays; (2) analog PUM (A), a ReRAM-based version of Xiao et al. [150], where non-MVM operations are performed by a 4 GHz 8-core Arm CPU with 256-bit vector extensions; and (3) nine configurations of a naively combined hybrid PUM (H), each with different digital PUM and analog PUM array counts, where digital PUM performs non-MVM operations. These comparisons are iso-area to the Arm CPU (except for A, whose analog PUM area is treated as free).




Figure 6. Motivating architectures.

We explore the shortcomings of these architectures using AES (the Advanced Encryption Standard), a widely used block cipher [93, 94].³ To encrypt data, a 16 B input (plaintext) is organized as a 4×4 matrix. The AES algorithm consists of four main steps: (1) SubBytes, which replaces each byte of the plaintext block with a byte from the substitution matrix (S-box); (2) ShiftRows, which cyclically left-shifts each row based on a predefined value; (3) MixColumns, which is a matrix multiply between the plaintext and a predefined matrix; and (4) AddRoundKey, which executes an XOR operation between the generated key and the plaintext. Figure 7 shows the iso-area throughput of AES-128 encryption for our three architectures. We make three observations from the figure.




Figure 7. Throughput of AES-128 encryption with digital (D), analog (A), and naive hybrid (H) PUM with the OSCAR and ideal logic families, normalized to D with OSCAR.

First, analog PUM performs only 18% better than digital PUM, even though AES involves many MVMs. This is because three of the four AES steps cannot be performed using analog PUM, and must be offloaded to the CPU. To analyze this further, we study the CPU performance of AES using the gem5 simulator [87], and find that non-MVM operations consume a majority of the execution time (even when including the data-movement overheads), being bottlenecked by the limited parallelism in the CPU compared to the analog PUM chip.

Second, most hybrid PUM configurations outperform both our pure digital PUM and pure analog PUM configurations. We note that as we start from a fully digital PUM chip and convert some of the area into analog PUM arrays, the performance increases up to a peak point (configuration H-5), as the inclusion of analog PUM significantly accelerates the bottleneck step (MixColumns). Beyond the peak, as more digital PUM area is converted to analog PUM, fewer digital PUM arrays are available, shifting the bottleneck to the inability to perform enough plaintext encryptions in parallel. At its peak, hybrid PUM performs 3.54× better than digital PUM.

Third, the choice of logic family has minimal impact on most hybrid PUM configurations. We introduce an *ideal* logic family, capable of performing any two-input Boolean operation within a single cycle, to observe if a NOR-only architecture is a limiting factor. This builds upon architectures such as FELIX [37], which show benefits for digital PUM by diversifying the number of available Boolean operators. While the

³For a detailed explanation of AES, see Section 5.3.

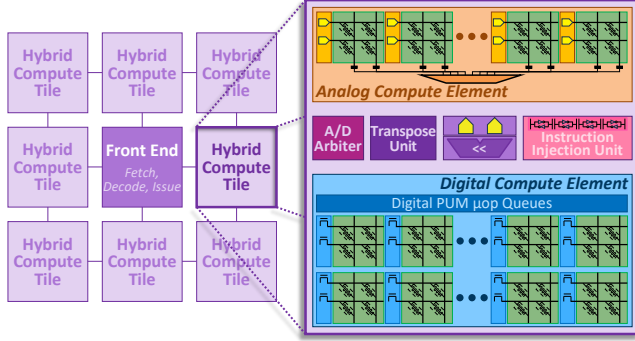


Figure 8. DARTH-PUM architecture.

pure digital PUM configuration sees a $2.1\times$ improvement in throughput, due to significant reductions in executing the $Mi\times Columns$ MVM operations, this still falls short of a hybrid PUM configuration with just a small number of analog arrays and only OSCAR support. In the best case for hybrid PUM, an ideal logic family increases throughput over OSCAR by only 3.2%. This means that we can build a hybrid PUM architecture without needing to support multiple Boolean operators, for which each additional operator increases the peripheral circuit area and complexity [37, 138].

From our analysis, we make two conclusions: (1) hybrid PUM can offer significant performance benefits by offering the best of both digital PUM and analog PUM in an iso-area footprint; and (2) hybrid PUM does not need complex logic families or SFUs to achieve its performance, improving its memory density and significantly simplifying the chip design. However, we note that for our naive hybrid architecture to be successful, we would have to do a similar sweep of configurations for each separate architecture. In Section 4, we show how we can avoid such customization by instead introducing co-designed hardware–software support to design and optimize a practical multi-domain hybrid PUM architecture.

4 DARTH-PUM System Design

Based on our conclusions from Section 3, we propose DARTH-PUM, a fixed architecture capable of providing best-of-both-worlds (i.e., digital PUM and analog PUM) benefits across many application domains. DARTH-PUM consists of two parts, as shown in Figure 8: (1) a front-end controller, which can fetch and decode hybrid PUM instructions (Section 4.4) and issue PUM-specific μ ops; and (2) a back end composed of multiple *hybrid compute tiles* (HCTs).

Hybrid Compute Tile. Each HCT contains (1) an *analog compute element* (ACE), (2) a *digital compute element* (DCE), and (3) auxiliary components that manage the flow of μ ops and optimize interactions between the ACE and DCE. An ACE contains 64 analog arrays and peripheral circuitry required for computation. A DCE contains 64 RACER-based digital PUM pipelines, provisioned to provide storage and sufficient output pipelines for analog MVMs, and RACER control circuitry that

includes a digital issue queue to dispatch μ ops to digital arrays within each pipeline. We connect the ACE and DCE using an I/O data transfer network capable of transferring 8 B per cycle, which is chosen to rate-match analog-to-digital converter (ADC) throughput with DCE write bandwidth.

We introduce four types of auxiliary components in an HCT, as shown in Figure 8, to coordinate kernel execution and address issues with the naive hybrid PUM architecture from Section 3. First, we introduce an analog–digital interface that consists of shift units and ADCs that carefully balance ACE and DCE throughput (Section 4.1). Second, we introduce an analog–digital arbiter and an instruction injection unit to sequence operations and data movement between the ACE and DCE (Section 4.2). Third, we introduce an error compensation scheme specialized for the HCT (Section 4.3). Fourth, we expose a novel software interface for DARTH-PUM (Section 4.4), with per-HCT instruction injection units to reduce the overhead of μ op expansion and mitigate front-end stalls.

Matrix–Vector Multiplication Walkthrough. Figure 9 shows how a 2×1 vector is multiplied with a 2×2 matrix using the analog side of DARTH-PUM, with three-bit inputs and four-bit matrix elements. The matrix resides inside the ACE, and the input vector is bit-sliced and placed at the ACE’s input buffers (1 in Figure 9). At each cycle, a new bit of the input is applied to each wordline (horizontal), performing a one-bit MVM, and generating a vector of partial products along each bitline (2). The resulting partial products are then transferred from the ACE to the DCE via the on-chip network and bit-striped across different arrays within a DCE pipeline (3). Each element of the partial product vector corresponds to a different element in the final output vector; therefore, each element is stored on a separate row within the same DCE vector register (4). As the ACE computes the remainder of the input bit positions, shift and add instructions are interleaved and added to the Digital Issue Queues (5).

As the result of each bit-sliced input is generated by the ACE and placed in the DCE, shift and add operations are executed to reduce the partial products. For clarity, we show the registers without the bit-striped values (Figure 9b), though the data remains in bit-striped format. After each bit slice is processed by the ACE, the partial product is shifted by its corresponding bit position using digital μ ops, e.g., the partial products of bit position 1 (stored in register R1) must be shifted once into R3 (6). Similarly, the result of bit position 2 (R2) must be shifted twice into R4 (7). Between each shift, addition instructions accumulate a partial sum. Once all partial sums have been accumulated, the final output vector is stored in R6 (8).

4.1 Interfacing the ACE With the DCE

To support input bit-slicing, analog PUM architectures typically introduce dedicated shifter and adder SFUs that can quickly recombine the partial products at the cost of significant chip area. Since our unoptimized hybrid PUM lacks

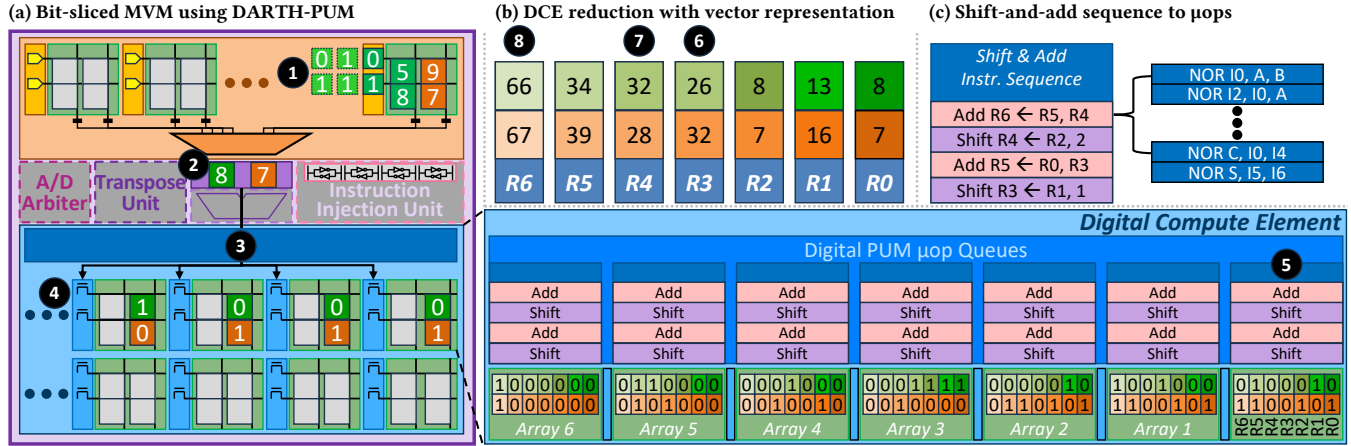


Figure 9. Unoptimized matrix–vector multiply with DARTH-PUM using ACE and DCE (gray components are unused).

explicit shifters and adders, the data must be written, shifted, and added using Boolean PUM operations. This creates a new performance bottleneck, as all three operations cannot interfere with one another. Specifically, partial products cannot be added until they have been shifted into the proper bit position. Simultaneously, the analog arrays are producing new partial products that must be written into the digital PUM arrays.

Figure 10a shows the timeline of the unoptimized MVM shown in Figure 9. The inputs are bit-sliced. Thus, the ACE executes three discrete MVM operations (orange operations in Figures 9a and 10), producing three separate outputs that must be written into the DCE. However, the DCE’s pipelines can only write one row of data per cycle. Thus, for a pipeline with N elements per vector register, it would take N cycles to write the partial products to the digital pipeline. Furthermore, it would take another M cycles to shift the data into the proper position (blue operations). This is due to each input bit i needing to be shifted i bit positions to the left. Therefore, for a pipeline with a width of N and a depth of M , it would require N cycles to write followed by M cycles to shift the data before the next set of partial products can be written to the digital arrays. Furthermore, the digital PUM would then need to execute bit-pipelined ADD operations prior to reading the next set of data (red operations in Figure 10). Importantly, these operations (write, shift, add) cannot be interleaved or pipelined without disruptions to the DCE’s bit-pipelining execution model, enforcing serialization.

To alleviate this bottleneck in DARTH-PUM, we implement a series of shift units that can place data into the correct bit positions during each transfer from the ACE to the DCE, and wait until the transfer completes before executing the add operations in the DCE. Because the shift amount can be predetermined based on the number of bits applied (e.g., 1-bit inputs require a single shift), the shift units can write the data into the DCE in a fixed manner without the need for software intervention or reconfigurable interconnects. By shifting during the data transfer, we gain three major benefits.

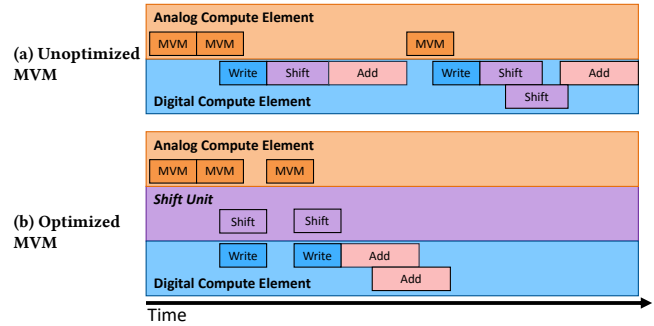


Figure 10. Sequence of operations for MVM.

First, waiting for all partial products can avoid paying the warmup and cooldown costs of a shift operation per partial product. Second, the rate at which ADCs produce data can be matched with the rate at which the DCE can write data to its arrays. Third, all outstanding ADD operations *can* be pipelined, leveraging the benefits of digital pipelined PUM. Figure 10b shows the resulting timeline after optimization.

Balancing Analog and Digital Array Counts. In addition to balancing the rate at which partial products are produced and processed, we need to carefully decide how to allocate the number of analog arrays in the ACE, and the number of digital arrays in the DCE. To optimize the balance between analog and digital arrays, sufficient arrays should first be allocated to hold the matrices (or sub-matrices) needed for the computation. Notably, this can mitigate the need for writing to the multi-bit devices during computation, as storing data in analog devices can incur substantial latency overheads [7, 45]. Therefore, underallocating analog arrays may cause disruptions during the computation.

As a second consideration, the number of analog and digital arrays also influences the amount of peripheral circuitry required for analog computation. Specifically, when only few arrays contain data required at any point in the computation (e.g., a single layer of a neural network), fewer peripheral circuits are required as fewer arrays are active at any one point

in time (i.e., arrays containing later layers may be inactive while processing the earlier layers of the network).

The amount of parallel computation executed during the application may also influence the type of peripheral circuits required (e.g., SAR ADC vs. ramp ADC). In the case of the AES algorithm, the matrix contains many rows, but few columns; therefore, an ideal system would be able to convert many bitline outputs in parallel. While a ramp ADC consumes greater area than an SAR ADC, it requires less power; we evaluate the impact of ADC choice in Section 7.3.

The outputs produced by the analog arrays are effectively wide-width row vectors (where each element is generated as a bitline output). To recombine these into a single value, the vectors must be added using the DCE. Therefore, the DCE can be treated as an SIMD vector unit, where the number of elements is equal to the number of rows. To account for this, sufficient bandwidth as well as capacity must be allocated in the DCE. For example, increasing the number of active arrays will increase the number of digital pipelines required to store and process the output data. Similarly, fewer active digital PUM arrays may enable the architecture to tolerate longer ADC latencies or fewer ADC units, as fewer outputs can be processed per cycle.

4.2 Coordinating ACE and DCE Operations

Whereas analog-only PUM can be easily enabled via a series of library calls (e.g., fixed-function fixed-width MVM), hybrid PUM requires a more sophisticated interface. Our goal with hybrid PUM is to deploy entire applications to the device. Therefore, we propose a full ISA to execute analog and digital operations and an accompanying library (Section 4.4). We observe that digital ISA instructions can be made to only interact with digital arrays; however, analog-type instructions require coordination between *both* analog and digital arrays. Therefore, we develop instructions to coordinate which can be used to prevent analog and digital operations from interfering with one another.

With our proposed ISA, instructions can be dispatched to analog and digital arrays from a single instruction stream. However, this creates two problems, which we describe below, (1) analog and digital instruction interference, and (2) contention for the digital PUM pipelines due to shift-and-add.

Mitigating Instruction Interference. First, analog and digital instructions have different latencies and dependencies. Notably, analog instructions may take hundreds of cycles to execute (due to ADC overheads and I/O from the array); in contrast, digital operations may consume only up to tens of cycles (in the case of a Boolean primitive). Thus, even though the analog instruction may consume hundreds of cycles (e.g., MVM), the entire MVM should appear atomic (i.e., new digital instructions should not interleave with the reduction sequence in Figure 9c), without disruption from digital instructions.

Therefore, to manage the instruction sequence and prevent instructions from colliding, we provide an arbiter in hardware, which limits an array to either analog or digital operations. Importantly, this also provides serialization, ensuring that a younger instruction cannot execute until an older instruction has completed. Therefore, a digital instruction that has dependence on a prior analog instruction (e.g., ReLU after an MVM) will be stalled until after the analog MVM has completed.

Over the course of the analog matrix multiply, numerous partial products are generated and placed in the digital vector registers. This creates a potential for existing data to be overwritten in hybrid PUM. Recall that partial products are shifted and placed into the digital arrays before any products are summed. Thus, these products may require up to N temporary vector registers for an N -bit input. Without a mechanism to determine which vector registers are available or unavailable, the analog arrays may overwrite data inside “live” vector registers. Unfortunately, dynamically determining which vector registers may and may not be used to store intermediate values could add substantial cost. To prevent data corruption, we introduce a *pipeline reserve* instruction for DARTH-PUM, which allocates a pipeline, effectively marking all data stored in the pipeline as “dead.” Notably, this data can be efficiently moved to another pipeline prior to executing the MVM instruction, and this constraint can easily be enforced by either a compiler or runtime system.

Reducing Issue Logic Stalls. In addition to the shift-and-add, the add sequence may consume hundreds of additional micro-ops to combine the partial results. Recall that digital PUM implements Boolean operations; therefore, a single ADD instruction may be decomposed into tens of Boolean operators (Figure 9c). Therefore, relying on the issue logic to recombine the partial products would consume substantial bandwidth and cause the issue and dispatch logic to stall on every MVM until all partial products have been generated and combined. To mitigate these overheads, we introduce a hardware *instruction injection unit* (IIU) in DARTH-PUM. We leverage the observation that the same instruction, ADD, is executed sequentially and repeatedly but with the only change being the vector register arguments. Thus, the IIU can be implemented as a small table and a counter, which is able to issue the Boolean operation sequence directly to the digital pipelines. By adding the IIU, DARTH-PUM can free up its front-end issue logic to issue instructions to other HCTs.

Moving Data Between the ACE and the DCE. Even with optimizations of allocations, applications require shuffling of data between analog and digital arrays; however, the operands in the two discrete systems operate on different axes. Specifically, analog PUM applies a row-wise (wordline) input, but combines in a columnar direction (bitline). In contrast, digital PUM stripes data column-wise, where each piece of data is a different bitline, but computes row-wise (e.g., activating two different bitlines along the same row produces a new bit in

a new bitline). Therefore, data that crosses between analog and digital requires transpositions.

To accelerate this, we introduce a *transposition unit*, which assists during both MVM computation and data storage. First, when input bit-slicing is enabled, the output of the analog array contains a vector of partial products. To combine the subsequent partial products (e.g., $\text{prod}_i + \text{prod}_{i+1}$), the vector must be shifted and subsequently added to the next output vector. Importantly, each *element* of the output vector i must interact with the corresponding element of vector $i+1$. To place the outputs correctly in the digital arrays, the analog *row vector* must be transposed into a *column vector* (i.e., placed into a vector register). Notably, this transposition must occur for each partial product. Second, when data (e.g., a copy of a matrix) is moved from digital to analog or vice-versa, that data must be transposed. When a matrix is stored in the digital arrays, each vector register (column) is allocated to a column of the matrix. When the matrix is programmed into an analog array, each row contains the data stored in the digital vector register.

Supporting Non-Vector Operations. RACER only provides support for vector copy operations, thereby incurring high overheads when only specific elements of a vector register are required. To retrieve a singular element, the digital PUM must first copy a vector, apply an element-wise mask, and execute an AND operation. Furthermore, because the digital PUM is pipelined, the AND must be applied sequentially to every bit before the next operation can proceed. This overhead can be prohibitively expensive in applications that require loading only a single element based on an address specified in the vector register (e.g., AES SubBytes). Notably, this would require the entire vector copy and AND sequence for every element in the vector register (e.g., $64 \times \text{vector copy} + 64 \times \text{AND}$ operations for a 64-element vector register).

To alleviate this, we implement element-wise loads and stores, which sequentially fetches the address specified by each element in the register. On each cycle, a row of digital PUM is read out, which is used as an address. The digital PUM can then read the data from an adjacent pipeline (again in one cycle) and transfer it back to the original pipeline. In the case of AES, other pipelines may be pre-loaded with the S-box (e.g., for SubBytes) or other plaintexts to encrypt. For other applications (e.g., ML), these pipelines may store weight matrices for layers not actively in use. To reduce complexity, we limit the address range of the load and store to only digital pipelines within the same Hybrid Compute Tile (see Section 4).

Expanding to Large-Width Operands. Analog accelerators usually target specific bit-width operands (e.g., 8b), and design their post-processing digital logic with these outputs in mind; however, given the flexible nature of digital PUM, a hybrid system can be dynamically configured to handle both different bit width operands and support analog devices with a different number of bits per cell. For example, an application may elect to use 8b operands, which are distributed across

$2 \times 4b$ arrays. Alternatively, the same application may require higher-precision, thereby limiting the devices to 2 bits per cell, requiring $4 \times 2b$ arrays. Notably, the only change in post-processing is the shift lengths and arguments needed for the ADD operation.

To support variable bit-widths, we introduce the idea of a *virtual analog core* (vACore). A vACore logically combines multiple analog arrays within a single analog compute element together to support wide-width operands. Additionally, allocating a vACore configures the shift units as well as the instruction injection unit to transparently implement the shift-and-add sequence. The vACores can be tracked through firmware and can be changed when the application requests either a different bit width, or the number of bits per cell stored in the arrays. Using the vACore abstraction, applications can dynamically adjust precision (based on the number of bits per cell) and support larger matrix values. Note that to reduce complexity, the HCT can only have vACores of the same bit width at a time.

4.3 Compensating for Errors

As noted in Section 2.2.1, analog MVMs are susceptible to noise. Therefore, for accuracy-critical applications, errors during the computation in the analog domain can lead to incorrect results. For example, during AES encryption, an error during the MixColumns operation can change the ciphertext in a way that prevents decryption. In our analysis, we observe that using the analog PUM arrays in SLC mode is insufficient to ensure an error-free approach. To mitigate these concerns, we develop a parasitic compensation scheme to mitigate against analog non-idealities.

Our parasitic compensation scheme consists of two components: (1) a remapping technique, in which the matrix values are stored using alternative representations; and (2) a post-processing error-correcting code (ECC) compensation factor, which is applied using the DCE. Figure 11 shows the original values, data remappings, and compensation factor.

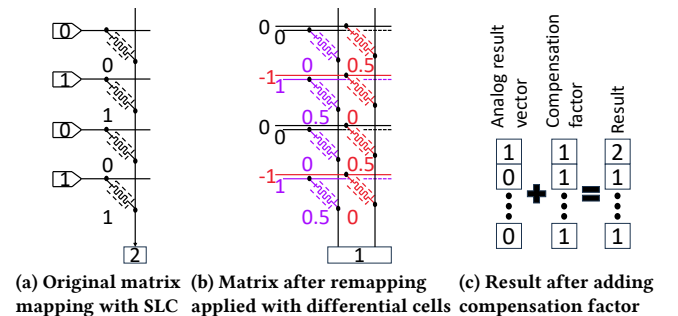


Figure 11. Steps for parasitic compensation scheme.

Our remapping technique leverages the observation that the AES matrix is strictly positive. This means that each device along a bitline contains a bit value of 0 or 1. Recall that we use differential cell pairs (Section 2.2.1); therefore, because

the matrix is entirely positive, the device storing the negative value is always storing a 0. Thus, the current down each column becomes relatively large causing large IR (current/resistance; i.e., Ohmic) drops along the positive bitline [153]. To mitigate these behaviors, we remap the bit values 0/1 to -1/+1, respectively, using differential cell pairs [151]. By remapping values from strictly positive to both positive and negative, the currents down each bitline are reduced, which brings the effect of IR drop below one ADC least significant bit level (i.e., it will no longer cause the ADC to read out an incorrect value).

To further reduce the effects on non-idealities, we leverage the sparsity of the matrix and input vector, as well as the fixed number of inputs, to reduce the magnitude of the error. Specifically, we combine the observation that the input vectors are sparse and contain a fixed number of 1s. For example, AES has a total of four 1s in the input vector. Additionally, we can scale the range of the matrix values from [-1, 1] to [-0.5, +0.5], effectively reducing the computed value by half. By reducing the range of the values, we can simultaneously reduce the magnitude of the noise. To re-scale the output to the correct value, a post-MVM factor can be applied to the result of the MVM. In the case of AES, a scale factor of 2 is applied as an addition during post-processing (4×0.5).⁴ Importantly, this post-processing is applied to the entire output vector (i.e., all output values of the ACE). Note that this operation can be efficiently realized using a vector operation in the nearby DCE, which can then be used to execute subsequent PUM operations. Our parasitic compensation scheme leverages algorithm insights to effectively mitigate noise by leveraging both analog and digital PUM, further highlighting the benefits of hybrid PUM.

As noted in Section 2.2.1, differential cell pairs are not exclusive to AES and can be broadly applied to all applications of DARTH-PUM. To generalize, the remapping technique leverages prior work in parasitic compensation [150, 153], which highlights the effects of current cancellation. Similarly, the parasitic compensation factor can be extrapolated by knowing the relative sparsity of the input vector and the relation the stored matrix.

4.4 Software Support

Tracking the internal state of HCTs, including parameters, across the entire chip can provide a burdensome task for programmers. Therefore, to ease programmer complexity when using DARTH-PUM, we provide both a library and a full hybrid ISA. The library is comprised of two types of commands: application-agnostic and application-specific. Using these library API calls, the explicit handling of vACores can be made transparent to the programmer.

Application-Agnostic Commands. Application-agnostic commands provide general-purpose MVM support and transparently manage the ACE–DCE dataflow without making any

⁴Figure 11 uses a factor of 1 (2×0.5).

Table 1. List of DARTH-PUM library calls.

API Call	Description
<i>Application-Agnostic Calls</i>	
<code>allocVACore()</code>	Allocates a vACore on an HCT, based on <i>elementSize</i> and bits per cell. Initializes the HCTs with the shift-and-add μ op sequence.
<code>setMatrix()</code>	Allocate and store a matrix with the required number of HCTs.
<code>execMVM()</code>	Execute MVM between matrix and vector.
<code>updateCol()</code>	Updates a column of a matrix on HCTs.
<code>updateRow()</code>	Updates a row of a matrix on HCTs.
<code>disableAnalogMode()</code>	Disables ACE and copies matrix from analog to digital arrays.
<code>disableDigitalMode()</code>	Disables DCE post-processing.
<i>Application-Specific Calls</i>	
<code>AES_initArrays()</code>	Reserve HCTs. Copies S-box to DCE. Sets the matrix needed for MixColumns in the analog arrays.
<code>AES_<en/de>crypt()</code>	Encrypt/decrypt. Optional argument for input key.
<code>CNN_setModel()</code>	Allocates and stores the model layers to HCTs. Mapping is done as a per-layer distribution.
<code>CNN_runInference()</code>	Executes inference operation.
<code>CNN_changeActivation()</code>	Changes the activation function used between layers during inference.
<code>LLM_build<En/De>coder()</code>	Allocates and stores encoder/decoder on HCTs.
<code>LLM_runInference()</code>	Execute inference operation.
<code>LLM_changeActivation()</code>	Changes the activation function used between layers during inference.

assumptions on the intended workloads. These commands can be used by programmers with minimal a priori knowledge of the underlying hardware (e.g., that it has analog and digital PUM). As seen in Table 1, these commands enable programmers to allocate and interface with the stored matrices in DARTH-PUM. For example, to store a matrix in the HCT, a user would use the `setMatrix()` operation, which takes three arguments: the matrix, *elementSize* (which refers to the bit width of each matrix element), and the level of bit precision required by the programmer. To reduce complexity, we elect to have bit precision be a scale (e.g., 0–2), which we map to 1 bit per device, half of all possible bits per device, and the maximum possible number of bits per device. For an 8b device, this would correspond to 1, 4, and 8 bits per cell, respectively. Using this information, the runtime system can allocate the number of HCTs required to store the matrix *without any more programmer input*.

Application-Specific Commands. For our evaluated workloads, we develop application-specific commands better tailored to high-level APIs. We develop these commands for users with minimal to no knowledge of the underlying hardware. Therefore, these API calls are designed to closely model high-level functions, while including sufficient meta-information to enable the system to map the computation to DARTH-PUM. For example, the `setModel()` command used for CNNs can build the DARTH-PUM mapping by allocating HCTs on a per-layer basis. `setModel()`-specific layer types (e.g., fully

connected, convolution) are mapped to analog arrays within the HCT, while leaving other computations (e.g., bias, activations) in the digital PUM. Additionally, to elide the complexities of analog PUM and precision, the bit-precision parameter is abstracted as an “accuracy” target, which is exposed as a scale to the programmer. We design the accuracy parameter to mirror the performance–accuracy trade-off between storing more bits per cell and fewer HCTs, versus fewer bits per cell and more HCTs, which affects total throughput.

Assisting Expert Programmers. To support new applications and enable greater flexibility for expert programmers, we include a full hybrid ISA. Additionally, we provide an additional set of commands to directly manage the vACores on the HCT, enabling (1) the ability to allocate the vACores on an HCT, (2) the precision at which they are operating, and (3) updates to the data stored inside the HCT. Note that these commands are hardware-centric, enabling the programmer to directly change the configuration of the HCT, whereas the application-agnostic commands are matrix-centric, enabling programmers to work at both the individual array or software levels. The combination of the ISA and application-agnostic commands provides developers with the flexibility to create their own custom mappings from applications to DARTH-PUM. Similarly, these functions enable expert programmers to leverage the wealth of prior work on analog and digital mapping techniques [7, 25, 45, 122, 125, 151] to improve different metrics (e.g., throughput, precision, latency, peripheries). For even greater flexibility, we provide options to enable or disable the ACE or DCE within each HCT, allowing developers to more precisely customize where computation is performed.

5 Applications

We describe how DARTH-PUM can be used to improve the performance of three different applications, convolutional neural networks, LLM encoders, and AES encryption.

5.1 Convolutional Neural Networks

Convolutional neural networks (CNN) have been widely explored as part of the analog PUM literature. Although various analog accelerators exist [19, 122, 150], they are often engineered for a specific set of functions (e.g., activation functions) or neural network topology. Of the main operations needed for CNN execution (convolution, fully connected or FC layers, downsampling, pooling, activation functions, and batch normalizations), only convolution and FC layers can make use of analog MVM. Therefore, each layer requires some degree of extra digital logic (i.e., specialized function units), in some cases [40, 91, 122] requiring periphery for dedicated shift-and-add network, in addition to the expensive analog periphery.

In contrast, with DARTH-PUM, the entire CNN pipeline can be executed on a single chip without the need for SFUs. We place the layers of the neural network (matrices) into the

analog arrays, while reserving digital pipelines for input vectors. Convolution layers leverage a Toeplitz expansion [132] that maximizes the number of rows. To handle layers which exceed the size of the arrays, we decompose the MVM into a series of MVMs, which are issued sequentially. Once the operations have completed, including the shift-and-add, the results are combined in the digital pipelines by using vector COPY instructions from different pipelines followed by ADD operations. We leverage digital PUM to implement the remaining auxiliary functions that analog PUM would have needed external support for (e.g., pooling, activation functions). For increased efficiency, inputs can be batched by storing multiple sets of inputs in the inactive digital pipelines. This presents a cost trade-off, as indefinitely increasing the batch size may affect performance due to the numerous in-flight operations.

5.2 LLM Encoders (Transformers)

LLM encoders are widely used in modern ML applications [22, 134] (e.g., LLMs). Specifically, these encoders rely on attention mechanisms [141], which require softmax, square root, and layer normalization. These operations cannot be accelerated using only analog PUM. Therefore, in place of dedicated digital logic, DARTH-PUM relies on its DCE to realize the non-MVM operations using IBERT [65] algorithms.

Unfortunately, not all of the encoder’s MVM operations can be easily mapped to DARTH-PUM’s ACE. Specifically, the matrices used in the attention mechanism rely on dynamic updates, which can create disruptions, as analog PUM requires more time and energy to program compared to digital PUM, potentially leading to long overheads. Therefore, we execute the computations needed by the attention mechanism in the DCE, while executing the feed-forward network (FFN) using the ACE, without the need for additional SFUs.

5.3 Advanced Encryption Standard

As noted in Section 3, the AES algorithm encrypts a 16 B plaintext using successive rounds of four main functions, SubBytes, ShiftRows, MixColumns, and AddRoundKey, which begin after an AddRoundKey initialization step. In the last (10th) round, only SubBytes, ShiftRows, and AddRoundKey are executed. AES-128 executes 10 rounds (for a 128-bit key), while AES-192 and AES-256 execute 12 and 14 rounds, respectively. Figure 12 shows how each step of AES is mapped to DARTH-PUM. SubBytes (1), ShiftRows (2), and AddRoundKey (4) are all executed using the DCE. The MixColumns step (3) leverages the ACE for efficient execution.

During SubBytes, each byte of the plaintext is replaced with a byte from the S-box. Therefore, we store the S-box in an unused pipeline within each HCT. Because the requested bytes may not be contained in a single vector register, the DCE uses the new element-wise load instruction (Section 4.2) to access the S-box pipeline data.

ShiftRows requires a cyclic left shift depending on the row of the block plaintext (e.g., row 1 is left shifted 1 full byte). We

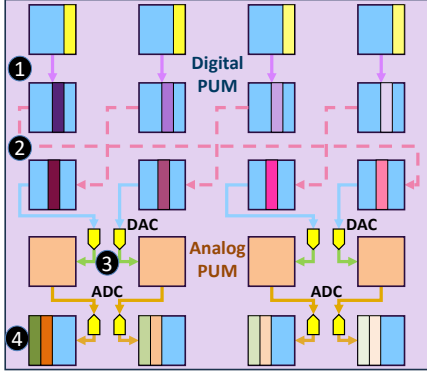


Figure 12. Dataflow for AES encryption.

implement this using pipelined left shifts; however, because there is no left terminal buffer, which could be used to shift the most-significant bit to the least-significant bit, we implement pipeline reversal and right shift operation. As part of this pipeline reversal macro, the entire digital pipeline is drained, followed by shift operations propagating in reverse mode. Note that we must fully drain the pipeline to preserve data integrity (e.g., reversing the pipeline while instructions are still present in the issue logic may cause only a subset of the bits to have the instruction applied).

MixColumns can be implemented as an MVM between the input and a predetermined matrix, followed by a modulus operator over a Galois field [13]. Alternatively, lookup tables may be used, where the indexes are denoted by a vector. Post-lookup (or MVM), an XOR is applied to the resulting bits to compute the final result. We use this insight to perform the MixColumns step in the analog domain by pre-storing the fixed matrix inside the analog arrays using 1-bit cells. Therefore, each bitline corresponds to an output of the MixColumns operation. Our key insight leverages the observation that although we could convert the full N-bit output from the ADC, only 1-bit is needed due to the subsequent XOR. Specifically, this implies that the ADC can be limited to 2-bits (e.g., SAR ADC), or can terminate after iterating through the 2-bits (e.g., ramp ADC), saving both energy and time.

6 Methodology

We model our system using a heavily modified version of MASTODON [8], a cross-stack simulator for digital PUM. We integrate MASTODON with CrossSim [28], a high-fidelity analog crossbar simulator, which faithfully models peripherals (e.g., ADCs, DACs) as well as analog non-idealities. We use the MILO model [89] with CrossSim to capture programming noise and parasitics. We model a DARTH-PUM chip running at 1 GHz, based on prior SPICE-based circuit simulations. Table 2 shows the configuration of each HCT in DARTH-PUM.

We compare DARTH-PUM to a Baseline configuration that contains an Intel Core i7-13700 [50] CPU and a 1.5 GB ReRAM analog PUM accelerator. We perform iso-area comparisons, where DARTH-PUM fits in the area of only the CPU

Table 2. Hybrid compute tile configuration.

1 Digital Compute Element	
Number of Pipelines	64
Pipeline Depth	64 arrays
ReRAM Array Size	64 × 64
1 Analog Compute Element	
Number of Arrays	64
ReRAM Array Size	64 × 64
Number of ADCs	SAR: 2; Ramp: 1
ADC Latency	SAR: 1 cycle; Ramp: 256 cycle

(2.57 cm²), with all hardware parameters scaled to 15 nm to be consistent with prior work [102, 131, 137]. Table 3 shows the area breakdown for a DARTH-PUM HCT. We derive our ADC parameters from prior work on ramp ADCs [150] and SAR ADCs [70], and incorporate a modified front-end module from the MPU [136]. With SAR ADCs, an iso-area DARTH-PUM chip contains 1860 HCTs (total memory capacity: 4.1 GB), while with ramp ADCs it contains 1660 HCTs (3.7 GB).

Table 3. Area and power for HCT hardware.

DCE	Area (μm ²)	ACE	Area (μm ²)
ReRAM Array	240	ReRAM Array	240
Pipeline Control	74000	Input Buffers	27000
IO Ctrl	9600	Row Periphery	13000
Decode & Drive	280	SAR/RAMP ADC	600/3800
Pipeline Select	64	Sample & Hold	62
HCT	Area (μm ²)	HCT	Area (μm ²)
Shift Unit	946	Transpose Unit	1760
A/D Arbiter	0.6	Instr. Injection Unit	42
Front End (shared by 8 HCTs; 63 mW power)		87000	
Component	Power (mW)	Component	Power (mW)
Array (Bool Ops)	8	Row Periphery	0.7
Pipeline Ctrl	1.6	SAR ADC	1.5
S&H (Analog)	2.1 × 10 ⁻⁵	Ramp ADC	1.2

To demonstrate the flexibility of DARTH-PUM, we evaluate it with three applications: (1) AES [94] encrypt, (2) ResNet-20 [39, 112] inference with the CIFAR-10 dataset [69], (3) an LLM encoder [141]. We use PyTorch [28, 105] implementations of ResNet-20 and the LLM encoder, and use open-source implementations for AES on CPUs [103] and GPUs [76].

We compare DARTH-PUM to four architectures in our evaluation: (1) *Baseline*, our Intel CPU with an analog PUM accelerator; (2) *DigitalPUM*, an iso-area 5.3 GB RACER [137, 138] chip with one front end per 8 clusters (running at 1 GHz using the OSCAR logic family [138], with two pipelines active per cluster to stay within thermal limits); (3) *AppAccel*, which compares against different application-specific accelerators; and (4) *GPU*, an NVIDIA GeForce RTX 4090 [97]. CPU and GPU runs are performed on real hardware, with data collected using performance counters.

AppAccel represents a different design for each application. For AES, AppAccel uses Intel’s AES-NI [115] accelerator. For ResNet-20, AppAccel uses a ReRAM version of a state-of-the-art CNN accelerator [150], which uses a ramp ADC and current integrators for the shift-and-add and peripheral ALUs. For LLM encoding, AppAccel is modeled after an ISAAC-style accelerator [122], with SAR ADCs and an SFU from recent work [125] that implements all supplemental functionality (e.g., shift, add, sqrt, ReLU, layernorm).

7 Evaluation

For most of the evaluation, we compare *DARTH-PUM* to *Baseline*, *DigitalPUM*, and *AppAccel*, iso-area to the Baseline CPU. We use SAR ADCs for all four configurations as it leads to overall improved performance; we show the effect of using ramp ADCs in Section 7.3. We conduct iso-area comparisons with *GPU* in Section 7.4.

7.1 Performance

Figure 13 shows the iso-area throughput of each configuration across the three workloads that we examine. We make four observations from the figure. First, we observe that *DARTH-PUM* provides an average throughput improvement of 31.4 \times over Baseline, with sizable benefits for all three applications over both Baseline and *DigitalPUM*. This illustrates *DARTH-PUM*'s ability to adapt across a variety of applications.




Figure 13. Throughput normalized to Baseline. AppAccel-AES uses an AES-NI accelerator (yellow), while the other AppAccel configurations use analog PUM accelerators (red).

Second, we observe that for AES, *DARTH-PUM* outperforms *AppAccel* by 36.9 \times . To investigate the source of these benefits, Figure 14 breaks down the execution time spent by each configuration on the individual kernels of the AES algorithm. We find that the latency of a single encryption in *DARTH-PUM* improves by 53.7% over Baseline. These savings predominantly come from *DARTH-PUM*'s ability to (1) reduce inter-kernel data movement; and (2) employ analog PUM to significantly reduce the *MixColumns* kernel latency, for which it provides an 11.5 \times improvement over *DigitalPUM*. Combined with Figure 13, we conclude that *DARTH-PUM*'s throughput benefits over Baseline come from a combination of reduced execution time and increased parallelism.

Third, we observe that *DARTH-PUM* approaches *AppAccel*'s throughput for ResNet-20, coming within 26.2%. As we can see from Figure 15, similar to AES, *DARTH-PUM*'s improvement over Baseline throughput is not just a result of improved inference latency (which decreases by 40.0%). In the case of ResNet-20, *AppAccel* sacrifices a significant amount of its chip area for the SFUs needed to compute each of the model layers, reducing the amount of parallel inference calculations it can perform in an iso-area comparison. *DARTH-PUM*'s HCT design is significantly more area-efficient than *AppAccel*'s analog PUM plus SFU design, allowing it to make up the throughput gap.




Figure 14. Kernel latency breakdown for AES, normalized to Baseline.

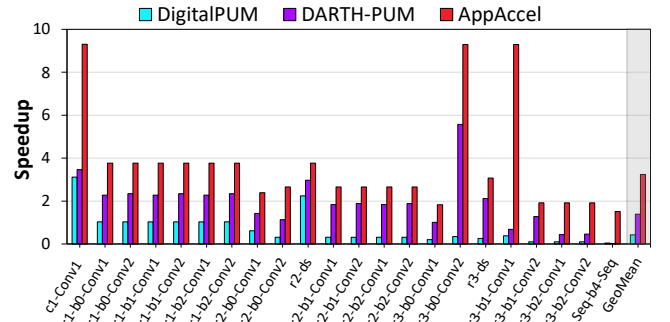


Figure 15. Per-layer speedup for ResNet-20 during inference, normalized to Baseline.

Fourth, we observe that while *DARTH-PUM* still falls short of *AppAccel* for LLMEnc, it significantly bridges the gap compared to other prior architectures. The main reason that *DARTH-PUM* falls short is because 71% of the total execution time is spent on non-MVM operations (e.g., softmax, ReLU, layernorm), which *AppAccel* performs with multiple highly specialized SFUs. Compared to Baseline, *DARTH-PUM* achieves a 45.6 \times benefit.

We conclude that *DARTH-PUM* is effective at providing large performance benefits over other multi-domain architectures, and captures many of the benefits of application-specific architectures without the need for any SFUs.

7.2 Energy

Processing data locally using *DARTH-PUM* can provide substantial energy savings compared to Baseline, which relies on the CPU to do non-matrix operations. Figure 16 shows the normalized energy required to execute each application. From the figure, we observe that *DARTH-PUM* achieves an average savings of 66.8 \times compared to Baseline. When we compare *DARTH-PUM* to *DigitalPUM*, we observe energy savings of 2.0 \times . Notably, this reduction in energy comes from the substantial reduction in digital Boolean operations required to execute the MVM and matrix multiply operations. Across the three applications, we observe that *DARTH-PUM* achieves competitive energy savings with *AppAccel*, though its biggest shortfall is for ResNet-20, due to its MVM-heavy nature.




Figure 16. Energy savings, normalized to Baseline (y-axis in log scale). AppAccel-AES uses an AES-NI accelerator (yellow), while the other AppAccel configurations use state-of-the-art analog PUM accelerators (red).

Overall, we conclude that DARTH-PUM delivers its throughput improvements with significant energy savings across our mix of applications.

7.3 Impact of Analog-to-Digital Converter Choice

In this section, we explore whether using ramp ADCs instead of SAR ADCs produces any benefit for those architectures that include analog PUM (see Section 4.1). Figures 17a and 17b show the performance and energy savings, respectively, with both SAR and ramp ADCs. We make two observations from the figure. First, for DARTH-PUM, we find that SAR ADCs outperform ramp ADCs by 1.5×, while achieving 99% of the energy savings. This is because across our applications, Boolean PUM operations are responsible for over 88.1% of the total energy consumption (9.4% is consumed by the front-end unit), minimizing the ADC’s impact on energy. The application that benefits most from using SAR ADCs is ResNet-20, due to its extensive use of analog PUM. Second, the only application for which ramp ADCs outperform SAR ADCs is AES, due to an additional optimization opportunity in the MixColumns kernel. This is due to two reasons: (1) we can terminate the ramp ADCs early because we need to check only four bit states (bringing its latency down from 256 cycles to four); and (2) while the ramp ADCs can convert all 64 bitlines concurrently, the SAR ADCs are multiplexed across multiple bitlines, reducing their throughput.




Figure 17. Comparison of SAR ADCs vs. ramp ADCs, normalized to Baseline: SAR.

7.4 Iso-Area Comparison With GPU

Figure 18 shows an iso-area comparison between DARTH-PUM and the NVIDIA GeForce RTX 4090 [97], a state-of-the-art GPU. From the figure we make three observations. First, on average, DARTH-PUM achieves a 11.8× and 7.5× throughput and energy improvement, respectively, over the GPU. Second, DARTH-PUM’s improvements for AES are significantly lower than those over Baseline, because the AES lookup tables are small enough to be cache-resident in the GPU, enabling it to achieve high throughput. Third, LLM encoding observes a very large energy savings over the GPU with both DARTH-PUM and DigitalPUM, because it can accelerate the full encoder operations using PUM. We conclude that DARTH-PUM offers promising benefits over GPUs across our range of applications.




Figure 18. Iso-area comparisons to state-of-the-art GPU.

7.5 Impact of Device Non-Idealities & Noise

Because of its use of many conductance levels per cell, analog computation in memory is highly susceptible to inherent noise and non-idealities that occur in many memory devices. Key sources of these errors are programming noise, various device parasitics (e.g., IR drop), read noise (i.e., error sources during read/MVM operations), cell drift (i.e., the device conductance can change over time, and for a few device technologies, at random), “stuck-at” faults (i.e., devices may become stuck at a fixed conductance state and cannot be overwritten correctly), and manufacturing process variation.⁵ While CrossSim models the first two error sources in detail, all five can contribute to errors in DARTH-PUM.

Various mitigation techniques, including previously proposed ones for analog PUM and our parasitic compensation scheme (Section 4.3), can address these errors sufficiently. For example, recent works have demonstrated how to increase ReRAM programming precision [129], how to partially mitigate read noise using input bit-slicing [151] (which we incorporate in DARTH-PUM), and how to use digital-style error-correcting codes (ECC) to mitigate “stuck-at” faults [86]. Even without these techniques, for ResNet-20 using the CIFAR-10 dataset [69], we observe an end-to-end accuracy of 75.4%,

⁵For an in-depth discussion on noise and mitigation techniques, we refer the reader to prior work [149, 151].

matching the accuracy of Baseline and AppAccel. Similar to prior work [53, 74, 111], for use cases that need very high accuracy, select MVM operations can be migrated to digital PUM. To faithfully report a comprehensive reliability for DARTH-PUM requires detailed chip-level variation/noise models and extensive datasets for device fabrication metrology, which we leave for future work.

8 Related Works

Hybrid Analog–Digital PUM. Prior works have proposed hybrid analog–digital PUM as application-specific accelerators. Jin et al. [54] propose a hybrid system for training CNNs. Similarly, Liu et al. [85] use hybrid PUM for hyperdimensional computing. Even though some digital operations are realized in PUM, many of these proposals still use additional peripherals (e.g., ALUs, SFUs). Works such as SPRINT [161] focus on accelerating the attention mechanism inside the encoder, but leave opportunities to improve the feed-forward network. DARTH-PUM uses high-performance digital PUM pipelines to realize non-MVM operations and balances data transfer and processing between analog and digital PUM.

Other works [53, 111] attempt to increase reliability by decomposing applications or operands into analog or digital PUM based on their importance in the overall computation. This remapping can be done at the application level [18, 125] wherein low-precision analog PUM is used for computations less critical to end-to-end accuracy, and high-precision digital PUM is used for computations more critical to the computation. DARTH-PUM is amenable to such optimizations by carefully mapping the data into the ACE and DCE.

Recent works mix multiple memory device technologies [10, 68, 145, 156, 167, 168] for accelerating specific applications. These approaches introduce integration complexities that can significantly raise fabrication prices and hurt yield. DARTH-PUM uses a single memory device technology to improve manufacturability, while benefiting multiple applications with the same chip.

Analog Accuracy. Non-idealities present a reliability problem for systems leveraging analog MVM. Therefore, many prior works have examined improving the reliability of analog MVM for different applications, from ML [3, 26, 42, 43, 74, 125, 146, 151] to scientific computing [25, 27, 126]. DARTH-PUM is amenable to these prior works, as it can support different data layouts or additional circuitry for reliability. Furthermore, this work develops a novel parasitic compensation scheme that leverages both analog and digital PUM to achieve high performance with low error counts.

Analog CAMs. Beyond MVM, prior work [78, 107, 108, 169, 170] has proposed to use analog devices for content-addressable memories (CAMs). At their core, these devices operate similar to SRAM-based CAMs, though analog CAMs are

capable of multi-bit matching, increasing both their throughput and densities. Therefore, these devices have been applied to workloads including tree-based machine learning [107, 108], CNNs and DNNs [169], transformers [170], and network routing [78]. Although promising, these devices are susceptible to noise, limiting their generalizability to noise-susceptible applications (e.g., AES), and are unable to handle control flow, both of which are achieved by DARTH-PUM.

ADC-Less Analog. The overheads associated with analog periphery can reduce the overall benefit of analog MVM accelerators. Therefore, prior work [95, 96] has attempted to mitigate these overheads by co-designing the hardware and algorithm to use lower-precision ADCs [47, 67], or in some cases removing them altogether [16, 66, 83, 95, 96, 117, 157]. However, these proposals often rely on specific functions known a priori, or assumptions on the specific application of interest and datasets (e.g., approximators for ReLU in DNNs). In contrast, to maintain flexibility, DARTH-PUM continues to use conventional ADCs, enabling it to be broadly used by a wide variety of applications.

9 Conclusion

Both analog and digital processing-using-memory continues to be a promising solution towards alleviating costs of data movement. DARTH-PUM provides a complete hybrid analog–digital PUM system with ISA, which can be used to accelerate a variety of applications from machine learning to cryptography. We demonstrate that DARTH-PUM can provide speedups of 59.4×, 14.8×, and 40.8× over an analog+CPU baseline for AES, ResNet-20, and LLM encoders, respectively.

Acknowledgments

We thank T. Patrick Xiao and Aniket Das for their feedback, suggestions, and support throughout this project. This work was supported in part by the Sandia National Laboratories Academic Alliance Program, the Univ. of Illinois Center for Advanced Semiconductor Chips with Accelerated Performance (ASAP; an NSF IUCRC), the Univ. of Illinois DREMES HYBRID Center, and NSF grant CCF-2329096.

This article has been authored by an employee of National Technology & Engineering Solutions of Sandia, LLC under Contract No. DE-NA0003525 with the U.S. Department of Energy (DOE). The employee owns all right, title and interest in and to the article and is solely responsible for its contents. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this article or allow others to do so, for United States Government purposes. The DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<https://www.energy.gov/downloads/doe-public-access-plan>).

References

- [1] S. Aga, S. Jeloka, A. Subramaniyan, S. Narayanasamy, D. Blaauw, and R. Das. 2017. Compute Caches. In *HPCA*.
- [2] V. Agrawal, T. P. Xiao, C. H. Bennett, B. Feinberg, S. Shetty, K. Ramkumar, H. Medu, K. Thekkekara, R. Chettuvetty, S. Leshner, Z. Luzada, L. Hinh, T. Phan, M. J. Marinella, and S. Agarwal. 2022. Subthreshold Operation of SONOS Analog Memory to Enable Accurate Low-Power Neural Network Inference. In *IEDM*.
- [3] T. Andrulis, J. S. Emer, and V. Sze. 2023. RAELLA: Reforming the Arithmetic for Efficient, Low-Resolution, and Low-Loss Analog PIM: No Retraining Required!. In *ISCA*.
- [4] S. Angizi, Z. He, and D. Fan. 2018. PIMA-Logic: A Novel Processing-in-Memory Architecture for Highly Flexible and Energy-Efficient Logic Computation. In *DAC*.
- [5] S. Angizi, Z. He, A. S. Rakin, and D. Fan. 2018. CMP-PIM: An Energy-Efficient Comparator-Based Processing-in-Memory Neural Network Accelerator. In *DAC*.
- [6] S. Angizi, J. Sun, W. Zhang, and D. Fan. 2019. AlignS: A Processing-in-Memory Accelerator for DNA Short Read Alignment Leveraging SOT-MRAM. In *DAC*.
- [7] A. Ankit, I. E. Hajj, S. R. Chalamalasetti, G. Ndu, M. Foltin, R. S. Williams, P. Faraboschi, W. m. Hwu, J. P. Strachan, K. Roy, and D. S. Milojicic. 2019. PUMA: A Programmable Ultra-Efficient Memristor-Based Accelerator for Machine Learning Inference. In *ASPLOS*.
- [8] ARCANA Research Group at Univ. of Illinois Urbana-Champaign. 2024. MASTODON — GitHub Repository. <https://github.com/ARCANA-Research/MASTODON/>.
- [9] C. H. Bennett, T. P. Xiao, R. Dellana, B. Feinberg, S. Agarwal, M. J. Marinella, V. Agrawal, V. Prabhakar, K. Ramkumar, L. Hinh, S. Saha, V. Raghavan, and R. Chettuvetty. 2020. Device-Aware Inference Operations in SONOS Non-Volatile Memory Arrays. In *IRPS*.
- [10] A. Bhattacharjee, A. Moitra, and P. Panda. 2023. HyDe: A Hybrid PCM/FeFET/SRAM Device-Search for Optimizing Area and Energy-Efficiencies in Analog IMC Platforms. *JETCAS* (Oct. 2023).
- [11] M. N. Bojnordi and E. Ipek. 2016. Memristive Boltzmann Machine: A Hardware Accelerator for Combinatorial Optimization and Deep Learning. In *HPCA*.
- [12] A. Boroumand, S. Ghose, Y. Kim, R. Ausavarungnirun, E. Shiu, R. Thakur, D. Kim, A. Kuusela, A. Knies, P. Ranganathan, and O. Mutlu. 2018. Google Workloads for Consumer Devices: Mitigating Data Movement Bottlenecks. In *ASPLOS*.
- [13] L. Carlitz. 1932. The Arithmetic of Polynomials in a Galois Field. *Am J. Math.* (Jan. 1932).
- [14] M. Cassinerio, N. Ciochini, and D. Ielmini. 2013. Logic Computation in Phase Change Materials by Threshold and Memory Switching. *Advanced Materials* (Aug. 2013).
- [15] J. Chen, C. Gao, Y. Lu, Y. Zhang, and J. Shu. 2024. Ares-Flash: Efficient Parallel Integer Arithmetic Operations Using NAND Flash Memory. In *MICRO*.
- [16] P. Chen, M. Wu, Y. Ma, L. Ye, and R. Huang. 2023. RIMAC: An Array-level ADC/DAC-Free ReRAM-Based In-Memory DNN Processor With Analog Cache and Computation. In *ASPDAC*.
- [17] X.-J. Chen, H.-P. Chen, and C.-L. Yang. 2024. PointCIM: A Computing-in-Memory Architecture for Accelerating Deep Point Cloud Analytics. In *MICRO*.
- [18] Y.-C. Chen, S. Ando, D. Fujiki, S. Takamaeda-Yamazaki, and K. Yoshioka. 2024. OSA-HCIM: On-the-Fly Saliency-Aware Hybrid SRAM CIM With Dynamic Precision Configuration. In *ASPDAC*.
- [19] P. Chi, S. Li, C. Xu, T. Zhang, J. Zhao, Y. Liu, Y. Wang, and Y. Xie. 2016. PRIME: A Novel Processing-in-Memory Architecture for Neural Network Computation in ReRAM-Based Main Memory. In *ISCA*.
- [20] T. Chou, W. Tang, J. Botimer, and Z. Zhang. 2019. CASCADE: Connecting RRAMs to Extend Analog Dataflow in an End-to-End In-Memory Processing Paradigm. In *MICRO*.
- [21] B. Dally. 2015. Challenges for Future Computing Systems. Keynote talk at HiPEAC.
- [22] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. 2019. BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding. arXiv:1810.04805 [cs.CL].
- [23] Domo, Inc. 2023. Data Never Sleeps 11.0. <https://www.domo.com/learn/infographic/data-never-sleeps-11>.
- [24] C. Eckert, X. Wang, J. Wang, A. Subramaniyan, R. Iyer, D. M. Sylvester, D. T. Blaauw, and R. Das. 2018. Neural Cache: Bit-Serial In-Cache Acceleration of Deep Neural Networks. In *ISCA*.
- [25] B. Feinberg, U. K. R. Vengalam, N. Whitehair, S. Wang, and E. Ipek. 2018. Enabling Scientific Computing on Memristive Accelerators. In *ISCA*.
- [26] B. Feinberg, S. Wang, and E. Ipek. 2018. Making Memristive Neural Network Accelerators Reliable. In *HPCA*.
- [27] B. Feinberg, R. Wong, T. P. Xiao, C. H. Bennett, J. N. Rohan, E. G. Boman, M. J. Marinella, S. Agarwal, and E. Ipek. 2021. An Analog Preconditioner for Solving Linear Systems. In *HPCA*.
- [28] B. Feinberg, T. P. Xiao, C. J. Brinker, C. H. Bennett, M. J. Marinella, and S. Agarwal. 2025. CrossSim: Accuracy Simulation of Analog In-Memory Computing. <https://github.com/sandialabs/cross-sim/>
- [29] D. Fujiki, S. Mahlke, and R. Das. 2019. Duality Cache for Data Parallel Acceleration. In *ISCA*.
- [30] C. Gao, X. Xin, Y. Lu, Y. Zhang, J. Yang, and J. Shu. 2021. ParaBit: Processing Parallel Bitwise Operations in NAND Flash Memory Based SSDs. In *MICRO*.
- [31] F. Gao, G. Tziantzioulis, and D. Wentzlaff. 2019. ComputeDRAM: In-Memory Compute Using Off-the-Shelf DRAMs. In *MICRO*.
- [32] F. Gao, G. Tziantzioulis, and D. Wentzlaff. 2022. FracDRAM: Fractional Values in Off-the-Shelf DRAM. In *MICRO*.
- [33] Gemini Team at Google. 2023. Gemini: A Family of Highly Capable Multimodal Models. arXiv:2312.11805 [cs.CL].
- [34] Q. Guo, X. Guo, Y. Bai, and E. Ipek. 2011. A Resistive TCAM Accelerator for Data-Intensive Computing. In *MICRO*.
- [35] Q. Guo, X. Guo, R. Patel, E. Ipek, and E. G. Friedman. 2013. AC-DIMM: Associative Computing With STT-MRAM. In *ISCA*.
- [36] X. Guo, F. Merrikh Bayat, M. Bavandpour, M. Klachko, M. R. Mahmoodi, M. Prezioso, K. K. Likharev, and D. B. Strukov. 2017. Fast, Energy-Efficient, Robust, and Reproducible Mixed-Signal Neuromorphic Classifier Based on Embedded NOR Flash Memory Technology. In *IEDM*.
- [37] S. Gupta, M. Imani, and T. Rosing. 2018. FELIX: Fast and Energy-Efficient Logic in Memory. In *ICCAD*.
- [38] N. Hajinazar, G. F. Oliveira, S. Gregorio, J. D. Ferreira, N. M. Ghiasi, M. Patel, M. Alser, S. Ghose, J. Gomez-Luna, and O. Mutlu. 2021. SIMDRAM: A Framework for Bit-Serial SIMD Processing Using DRAM. In *ASPLOS*.
- [39] K. He, X. Zhang, S. Ren, and J. Sun. 2016. Deep Residual Learning for Image Recognition. In *CVPR*.
- [40] M. He, C. Song, I. Kim, C. Jeong, S. Kim, I. Park, M. Thottethodi, and T. N. Vijaykumar. 2020. Newton: A DRAM-Maker's Accelerator-in-Memory (AiM) Architecture for Machine Learning. In *MICRO*.
- [41] B. Hoffer, N. Wainstein, C. M. Neumann, E. Pop, E. Yalon, and S. Kvatinisky. 2022. Stateful Logic Using Phase Change Memory. *JXDC* (Nov. 2022).
- [42] Y. Hou, Z. Liu, G. Gagnon, H. Tsai, K. El Maghraoui, G. W. Burr, and L. Liu. 2025. SAGE: Saliency-Aware Grouping for Efficient Mapping of LLMs on Analog Compute-in-Memory. In *ICCAD*.
- [43] Y. Hou, H. Tsai, K. El Maghraoui, T. Gokmen, G. W. Burr, and L. Liu. 2025. NORA: Noise-Optimized Rescaling of LLMs on Analog Compute-in-Memory Accelerators. In *DATE*.
- [44] H.-W. Hu, W.-C. Wang, Y.-H. Chang, Y.-C. Lee, B.-R. Lin, H.-M. Wang, Y.-P. Lin, Y.-M. Huang, C.-Y. Lee, T.-H. Su, C.-C. Hsieh, C.-M. Hu, Y.-T. Lai, C.-K. Chen, H.-S. Chen, H.-P. Li, T.-W. Kuo, M.-F. Chang, K.-C. Wang, C.-H. Hung, and C.-Y. Lu. 2022. ICE: An Intelligent Cognition Engine With 3D NAND-Based In-Memory Computing for Vector

- Similarity Search Acceleration. In *MICRO*.
- [45] M. Hu, J. P. Strachan, Z. Li, E. M. Grafals, N. Davila, C. Graves, S. Lam, N. Ge, J. J. Yang, and R. S. Williams. 2016. Dot-Product Engine for Neuromorphic Computing: Programming 1T1M Crossbar to Accelerate Matrix-Vector Multiplication. In *DAC*.
- [46] M. Hu, J. P. Strachan, Z. Li, and R. S. Williams. 2016. Dot-Product Engine as Computing Memory to Accelerate Machine Learning Algorithms. In *ISQED*.
- [47] S. Huang, A. Ankit, P. Silveira, R. Antunes, S. R. Chalamalasetti, I. El Hajj, D. E. Kim, G. Aguiar, P. Bruel, S. Serebryakov, C. Xu, C. Li, P. Faraboschi, J. P. Strachan, D. Chen, K. Roy, W.-m. Hwu, and D. Milojkic. 2021. Mixed Precision Quantization for ReRAM-Based DNN Inference Accelerators. In *ASPDAC*.
- [48] B. Hyun, T. Kim, D. Lee, and M. Rhu. 2024. Pathfinding Future PIM Architectures by Demystifying a Commercial PIM Technology. In *HPCA*.
- [49] J.-F. Im, K. Gopalakrishna, S. Subramaniam, M. Shrivastava, A. Tumbde, X. Jiang, J. Dai, S. Lee, N. Pawar, J. Li, and R. Aringunram. 2018. Pinot: Realtime OLAP for 530 Million Users. In *SIGMOD*.
- [50] Intel Corp. 2023. Intel® Core™ i7-13700 Processor. <https://www.intel.com/content/www/us/en/products/sku/230490/intel-core-i7-13700-processor-30m-cache-up-to-5-20-ghz/specifications.html>
- [51] Z. Jahshan and L. Yavits. 2024. MajorK: Majority Based kmer Matching in Commodity DRAM. *CAL* (Apr. 2024).
- [52] Y. Ji, Y. Zhang, X. Xie, S. Li, P. Wang, X. Hu, Y. Zhang, and Y. Xie. 2019. FPSA: A Full System Stack Solution for Reconfigurable ReRAM-Based NN Accelerator Architecture. In *ASPLOS*.
- [53] H. Jiang, S. Huang, X. Peng, and S. Yu. 2020. MINT: Mixed-Precision RRAM-Based IN-Memory Training Architecture. In *ISCAS*.
- [54] H. Jin, C. Liu, H. Liu, R. Luo, J. Xu, F. Mao, and X. Liao. 2022. ReHy: A ReRAM-Based Digital/Analog Hybrid PIM Architecture for Accelerating CNN Training. *TPDS* (Nov. 2022).
- [55] V. Joshi, M. Le Gallo, S. Haefeli, I. Boybat, S. R. Nandakumar, C. Piveteau, M. Dazzi, B. Rajendran, A. Sebastian, and E. Eleftheriou. 2020. Accurate Deep Neural Network Inference Using Computational Phase-Change Memory. *Nat. Commun.* 11 (May 2020).
- [56] L. Ke, U. Gupta, B. Y. Cho, D. Brooks, V. Chandra, U. Diril, A. Firoozshahian, K. Hazelwood, B. Jia, H.-H. S. Lee, M. Li, B. Maher, D. Mudigere, M. Naumov, M. Schatz, M. Smelyanskiy, X. Wang, B. Reagen, C.-J. Wu, M. Hempstead, and X. Zhang. 2020. RecNMP: Accelerating Personalized Recommendation With Near-Memory Processing. In *ISCA*.
- [57] L. Ke, X. Zhang, J. So, J.-G. Lee, S.-H. Kang, S. Lee, S. Han, Y. Cho, J. H. Kim, Y. Kwon, K. Kim, J. Jung, I. Yun, S. J. Park, H. Park, J. Song, J. Cho, K. Sohn, N. S. Kim, and H.-H. S. Lee. 2022. Near-Memory Processing in Action: Accelerating Personalized Recommendation With AxDIMM. *IEEE Micro* (Jul. 2022).
- [58] G. Kestor, R. Gioiosa, D. J. Kerbyson, and A. Hoisie. 2013. Quantifying the Energy Cost of Data Movement in Scientific Applications. In *IISWC*.
- [59] R. Khaddam-Aljameh, M. Stanisavljevic, J. Fornit Mas, G. Karunaratne, M. Braendli, F. Liu, A. Singh, S. M. Müller, U. Egger, A. Petropoulos, T. Antonakopoulos, K. Brew, S. Choi, I. Ok, F. L. Lie, N. Saulnier, V. Chan, I. Ahsan, V. Narayanan, S. R. Nandakumar, M. Le Gallo, P. A. Francese, A. Sebastian, and E. Eleftheriou. 2021. HERMES Core—a 14nm CMOS and PCM-Based In-Memory Compute Core Using an Array of 300ps/LSB Linearized CCO-Based ADCs and Local Digital Processing. In *VLSI Symposia*.
- [60] R. Khaddam-Aljameh, M. Stanisavljevic, J. F. Mas, G. Karunaratne, M. Brändli, F. Liu, A. Singh, S. M. Müller, U. Egger, A. Petropoulos, T. Antonakopoulos, K. Brew, S. Choi, I. Ok, F. L. Lie, N. Saulnier, V. Chan, I. Ahsan, V. Narayanan, S. R. Nandakumar, M. Le Gallo, P. A. Francese, A. Sebastian, and E. Eleftheriou. 2022. HERMES-Core—A 1.59-TOPS/mm² PCM on 14-nm CMOS In-Memory Compute Core Using 300-ps/LSB Linearized CCO-Based ADCs. *JSSC* (Apr. 2022).
- [61] A. Khadem, D. Fujiki, H. Chen, Y. Gu, N. Talati, S. Mahlke, and R. Das. 2025. Multi-Dimensional Vector ISA Extension for Mobile In-Cache Computing. In *HPCA*.
- [62] H. Kim, S. Song, S. Choi, J. Choe, S. Han, J. Park, J. Lee, and J.-J. Kim. 2025. CrossBit: Bitwise Computing in NAND Flash Memory With Inter-Bitline Data Communication. In *MICRO*.
- [63] J. H. Kim, S.-H. Kang, S. Lee, H. Kim, Y. Ro, S. Lee, D. Wang, J. Choi, J. So, Y. Cho, J. Song, J. Cho, K. Sohn, and N. S. Kim. 2022. Aquabolt-XL HBM2-PIM, LPDDR5-PIM With In-Memory Processing, and AXDIMM With Acceleration Buffer. *IEEE Micro* (May 2022).
- [64] J. H. Kim, S.-H. Kang, S. Lee, H. Kim, W. Song, Y. Ro, S. Lee, D. Wang, H. Shin, B. Phuah, J. Choi, J. So, Y. Cho, J. Song, J. Choi, J. Cho, K. Sohn, Y. Sohn, K. Park, and N. S. Kim. 2021. Aquabolt-XL: Samsung HBM2-PIM With In-Memory Processing for ML Accelerators and Beyond. In *HCS*.
- [65] S. Kim, A. Gholami, Z. Yao, M. W. Mahoney, and K. Keutzer. 2021. I-BERT: Integer-only BERT Quantization. In *PMLR*.
- [66] S. Kim, S. Kim, S. Um, S. Kim, K. Kim, and H.-J. Yoo. 2023. Neuro-CIM: ADC-Less Neuromorphic Computing-in-Memory Processor With Operation Gating/Stopping and Digital-Analog Networks. *JSSC* (May 2023).
- [67] Y. Kim, H. Kim, and J.-J. Kim. 2022. Extreme Partial-Sum Quantization for Analog Computing-In-Memory Neural Network Accelerators. *JETC* (Oct. 2022).
- [68] G. Krishnan, Z. Wang, I. Yeo, L. Yang, J. Meng, M. Liehr, R. V. Joshi, N. C. Cady, D. Fan, J.-S. Seo, and Y. Cao. 2022. Hybrid RRAM/SRAM in-Memory Computing for Robust DNN Acceleration. *IEEE TCAD* (Aug. 2022).
- [69] A. Krizhevsky. 2009. *Learning Multiple Layers of Features From Tiny Images*. Technical Report. Univ. of Toronto.
- [70] L. Kull, T. Toifl, M. Schmatz, P. A. Francesc, C. Menolfi, M. Brändli, M. Kossel, T. Morf, T. M. Andersen, and Y. Leblebici. 2013. A 3.1 mW 8b 1.2 GS/s Single-Channel Asynchronous SAR ADC With Alternate Comparators for Enhanced Speed in 32 nm Digital SOI CMOS. *JSSC* (Sep. 2013).
- [71] S. Kvatinisky, D. Belousov, S. Liman, G. Satat, N. Wald, E. G. Friedman, A. Kolodny, and U. C. Weiser. 2014. MAGIC: Memristor-Aided Logic. *TCAS II* (Sep. 2014).
- [72] S. Kvatinisky, A. Kolodny, U. C. Weiser, and E. G. Friedman. 2011. Memristor-Based IMPLY Logic design Procedure. In *ICCD*.
- [73] S. Kvatinisky, G. Satat, N. Wald, E. G. Friedman, A. Kolodny, and U. C. Weiser. 2014. Memristor-Based Material Implication (IMPLY) Logic: Design Principles and Methodologies. *TVLSI* (2014).
- [74] C. Lammie, Y. Wang, F. Ponzina, J. Klein, H. Benmeziane, M. Zapater, I. Boybat, A. Sebastian, G. Ansaloni, and D. Atienza. 2025. LionHeart: A Layer-Based Mapping Framework for Heterogeneous Systems With Analog In-Memory Computing Tiles. *IEEE Trans. Emerg. Top. Comput.* (Mar. 2025).
- [75] D. Lee, B. Hyun, T. Kim, and M. Rhu. 2024. Analysis of Data Transfer Bottlenecks in Commercial PIM Systems: A Study With UPMEM-PIM. *CAL* (Apr. 2024).
- [76] Wai-Kong Lee, Hwa Jeong Seo, Seog Chung Seo, and Seong Oun Hwang. 2022. AES-GPU. <https://github.com/benlwk/AES-GPU/>
- [77] B. Li, L. Xia, P. Gu, Y. Wang, and H. Yang. 2015. Merging the Interface: Power, Area and Accuracy Co-Optimization for RRAM Crossbar-Based Mixed-Signal Computing System. In *DAC*.
- [78] C. Li, C. E. Graves, X. Sheng, D. Miller, M. Foltin, G. Pedretti, and J. P. Strachan. 2020. Analog Content-Addressable Memories With Memristors. *Nat. Commun.* 11 (Apr. 2020).
- [79] H. Li, H. Jin, L. Zheng, and X. Liao. 2020. ReSQM: Accelerating Database Operations Using ReRAM-Based Content Addressable Memory. *IEEE TCAD* (Nov. 2020).
- [80] H. Li, Z. Li, Z. Bai, and T. Mitra. 2024. ASADI: Accelerating Sparse Attention Using Diagonal-Based In-Situ Computing. In *HPCA*.

- [81] S. Li, D. Niu, K. T. Malladi, H. Zheng, B. Brennan, and Y. Xie. 2017. DRISA: A DRAM-Based Reconfigurable In-Situ Accelerator. In *MICRO*.
- [82] S. Li, C. Xu, Q. Zou, J. Zhao, Y. Lu, and Y. Xie. 2016. Pinatubo: A Processing-in-Memory Architecture for Bulk Bitwise Operations in Emerging Non-volatile Memories. In *DAC*.
- [83] Y. Li, J. Chen, L. Wang, W. Zhang, Z. Guo, J. Wang, Y. Han, Z. Li, F. Wang, C. Dou, X. Xu, J. Yang, Z. Wang, and D. Shang. 2023. An ADC-Less RRAM-Based Computing-in-Memory Macro With Binary CNN for Efficient Edge AI. *TCAS-II* (Jan. 2023).
- [84] C. Liu, H. Liu, H. Jin, X. Liao, Y. Zhang, Z. Duan, J. Xu, and H. Li. 2022. ReGNN: A ReRAM-Based Heterogeneous Architecture for General Graph Neural Networks. In *DAC*.
- [85] C. Liu, K. Wu, H. Liu, H. Jin, X. Liao, Z. Duan, J. Xu, H. Li, Y. Zhang, and J. Yang. 2025. A ReRAM-Based Processing-In-Memory Architecture for Hyperdimensional Computing. *IEEE TCAD* (Feb. 2025).
- [86] M. Liu, L. Xia, Y. Wang, and K. Chakrabarty. 2018. Fault Tolerance for RRAM-Based Matrix Operations. In *ITC*.
- [87] J. Lowe-Power, A. M. Ahmad, A. Akram, M. Alian, R. Amslinger, M. Andreozzi, A. Armejach, N. Asmussen, S. Bharadwaj, G. Black, G. Bloom, B. R. Bruce, D. R. Carvalho, J. Castrillón, L. Chen, N. Derumigny, S. Diestelhorst, W. Elsasser, M. Fariborz, A. F. Farahani, P. Fotouhi, R. Gambord, J. Gandhi, D. Gope, T. Grass, B. Hanindhito, A. Hansson, S. Haria, A. Harris, T. Hayes, A. Herrera, M. Horsnell, S. A. R. Jafri, R. Jagtap, H. Jang, R. Jeyapaul, T. M. Jones, M. Jung, S. Kannoth, H. Khaleghzadeh, Y. Kodama, T. Krishna, T. Marinelli, C. Menard, A. Mondelli, T. Mück, O. Naji, K. Nathella, H. Nguyen, N. Nikoleris, L. E. Olson, M. S. Orr, B. Pham, P. Prieto, T. Reddy, A. Roelke, M. Samani, A. Sandberg, J. Setoain, B. Shingarov, M. D. Sinclair, T. Ta, R. Thakur, G. Travaglini, M. Upton, N. Vaish, I. Vougioukas, Z. Wang, N. Wehn, C. Weis, D. A. Wood, H. Yoon, and E. F. Zulian. 2020. The gem5 Simulator: Version 20.0+. <https://arxiv.org/abs/2007.03152>.
- [88] M. J. Marinella, S. Agarwal, A. Hsia, I. Richter, R. Jacobs-Gedrim, J. Niroula, S. J. Plimpton, E. Ipek, and C. D. James. 2018. Multiscale Co-Design Analysis of Energy, Latency, Area, and Accuracy of a ReRAM Analog Neural Training Accelerator. *JETCAS* 8 (Jan. 2018).
- [89] V. Milo, F. Anzalone, C. Zambelli, E. Pérez, M. K. Mahadevaiah, Ó. G. Ossorio, P. Olivo, C. Wenger, and D. Ielmini. 2021. Optimized Programming Algorithms for Multilevel RRAM in Hardware Neural Networks. In *IRPS*.
- [90] Mythic, Inc. [n. d.]. M1076 Analog Matrix Processor. <https://mythic.ai/products/m1076-analog-matrix-processor/>.
- [91] A. Nag, R. Balasubramonian, V. Srikumar, R. Walker, A. Shafiee, J. P. Strachan, and N. Muralimanohar. 2018. Newton: Gravitating Towards the Physical Limits of Crossbar Acceleration. *IEEE Micro* (Sep. 2018).
- [92] R. Nair, S. F. Antao, C. Bertolli, P. Bose, J. R. Brunheroto, T. Chen, C.-Y. Cher, C. H. A. Costa, J. Doi, C. Evangelinos, B. M. Fleischer, T. W. Fox, D. S. Gallo, L. Grinberg, J. A. Gunnels, A. C. Jacob, P. Jacob, H. M. Jacobson, T. Karkhanis, C. Kim, J. H. Moreno, J. K. O'Brien, M. Ohmacht, Y. Park, D. A. Prener, B. S. Rosenburg, K. D. Ryu, O. Sallenave, M. J. Serrano, P. D. M. Siegl, K. Sugavanam, and Z. Sura. 2015. Active Memory Cube: A Processing-in-Memory Architecture for Exascale Systems. *IBM JRD* (Mar.–May 2015).
- [93] National Institute of Standards and Technology. 2001. Advanced Encryption Standard (AES). <https://nvlpubs.nist.gov/nistpubs/fips/nist.fips.197.pdf>
- [94] J. Nechvatal, E. Barker, L. Bassham, W. Burr, M. Dworkin, J. Foti, and E. Roback. 2001. Report on the Development of the Advanced Encryption Standard (AES). *NIST* (Nov. 2001).
- [95] S. Negi, U. Saxena, D. Sharma, and K. Roy. 2025. HciM: ADC-Less Hybrid Analog-Digital Compute in Memory Accelerator for Deep Learning Workloads. In *ASPAC*.
- [96] S. Negi, U. Saxena, D. Sharma, J. Victor, I. Ahmed, S. K. Gupta, and K. Roy. 2024. Algorithm Hardware Co-Design for ADC-Less Compute In-Memory Accelerator. *TCASAI* (Dec 2024).
- [97] NVIDIA Corp. [n. d.]. GeForce RTX 4090. <https://www.nvidia.com/en-us/geforce/graphics-cards/40-series/rtx-4090/>
- [98] A. Olgun, J. Gómez-Luna, K. Kanellopoulos, B. Salami, H. Hassan, O. Ergin, and O. Mutlu. 2022. PiDRAM: A Holistic End-to-End FPGA-Based Framework for Processing-in-DRAM. *TACO* (Nov. 2022).
- [99] A. Olgun, H. Hassan, A. G. Yağlıkcı, Y. C. Tuğrul, L. Orosa, H. Luo, M. Patel, O. Ergin, and O. Mutlu. 2023. DRAM Bender: An Extensible and Versatile FPGA-Based Infrastructure to Easily Test State-of-the-Art DRAM Chips. *IEEE TCAD* (Dec. 2023).
- [100] A. Olgun, M. Patel, A. G. Yağlıkcı, H. Luo, J. S. Kim, F. N. Bostancı, N. Vijaykumar, O. Ergin, and O. Mutlu. 2021. QUAC-TRNG: High-Throughput True Random Number Generation Using Quadruple Row Activation in Commodity DRAM Chips. In *ISCA*.
- [101] G. F. Oliveira, J. Gómez-Luna, L. Orosa, S. Ghose, N. Vijaykumar, I. Fernandez, M. Sadrosadati, and O. Mutlu. 2021. DAMOV: A New Methodology and Benchmark Suite for Evaluating Data Movement Bottlenecks. *IEEE Access* (2021).
- [102] G. F. Oliveira, A. Olgun, A. G. Yağlıkcı, F. N. Bostancı, J. Gómez-Luna, S. Ghose, and O. Mutlu. 2024. MIMDRAM: An End-to-End Processing-Using-DRAM System for High-Throughput, Energy-Efficient and Programmer-Transparent Multiple-Instruction Multiple-Data Computing. In *HPCA*.
- [103] OpenSSL Software Foundation Inc. [n. d.]. OpenSSL. <https://github.com/openssl/openssl>
- [104] J. Park, R. Azizi, G. F. Oliveira, M. Sadrosadati, R. Nadig, D. Novo, J. Gómez-Luna, M. Kim, and O. Mutlu. 2022. Flash-Cosmos: In-Flash Bulk Bitwise Operations Using Inherent Computation Capability of NAND Flash Memory. In *MICRO*.
- [105] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *NeurIPS*.
- [106] J. T. Pawlowski. 2011. Hybrid Memory Cube (HMC). In *HCS*.
- [107] G. Pedretti, C. E. Graves, S. Serebryakov, R. Mao, X. Sheng, M. Foltin, C. Li, and J. P. Strachan. 2021. Tree-Based Machine Learning Performed In-Memory With Memristive Analog CAM. *Nat. Commun.* (Oct. 2021).
- [108] G. Pedretti, S. Serebryakov, J. P. Strachan, and C. E. Graves. 2022. A General Tree-Based Machine Learning Accelerator With Memristive Analog CAM. In *ISCAS*.
- [109] R. A. Pielke Sr. 2013. *Mesoscale Meteorological Modeling* (2nd ed.). Academic Press.
- [110] F. Pop. 2014. High Performance Numerical Computing for High Energy Physics: A New Challenge for Big Data Science. *Adv. High Energy Phys.* (Feb. 2014).
- [111] M. R. H. Rashed, S. K. Jha, and R. Ewetz. 2021. Hybrid Analog-Digital In-Memory Computing. In *ICCAD*.
- [112] V. J. Reddi, C. Cheng, D. Kanter, P. Mattson, G. Schmuelling, C.-J. Wu, B. Anderson, M. Breughe, M. Charlebois, W. Chou, R. Chukka, C. Coleman, S. Davis, P. Deng, G. Diamos, J. Duke, D. Fick, J. S. Gardner, I. Hubara, S. Idgunji, T. B. Jablin, J. Jiao, T. St. John, P. Kanwar, D. Lee, J. Liao, A. Lohmotov, F. Massa, P. Meng, P. Micikevicius, C. Osborne, G. Pekhimenko, A. T. R. Rajan, D. Sequeira, A. Sirasao, F. Sun, H. Tang, M. Thomson, F. Wei, E. Wu, L. Xu, K. Yamada, B. Yu, G. Yuan, A. Zhong, P. Zhang, and Y. Zhou. 2020. MLPerf Inference Benchmark. In *ISCA*.
- [113] D. Reinsel, J. Gantz, and J. Rydning. 2018. *Data Age 2025: The Digitization of the World From Edge to Core*. Technical Report. IDC.
- [114] S. Rhyner, H. Luo, J. Gómez-Luna, M. Sadrosadati, J. Jiang, A. Olgun, H. Gupta, C. Zhang, and O. Mutlu. 2024. PIM-Opt: Demystifying Distributed Optimization Algorithms on a Real-World Processing-in-Memory System. In *PACT*.
- [115] J. K. Rott. 2012. *Intel® Advanced Encryption Standard Instructions (AES-NI)*. Technical Report. © Intel Corporation. <https://www.intel.com/content/www/us/en/developer/articles/>

- [technical/advanced-encryption-standard-instructions-aes-ni.html](#)
- [116] K. Roy, A. Jaiswal, and P. Panda. 2019. Towards Spike-Based Machine Intelligence With Neuromorphic Computing. *Nature* (2019).
- [117] U. Saxena and K. Roy. 2023. Partial-Sum Quantization for Near ADC-Less Compute-In-Memory Accelerators. In *ISLPED*.
- [118] B. Schölkopf, K. Tsuda, and J.-P. Vert. 2004. *Kernel Methods in Computational Biology*. MIT press.
- [119] V. Seshadri, K. Hsieh, A. Boroum, D. Lee, M. A. Kozuch, O. Mutlu, P. B. Gibbons, and T. C. Mowry. 2015. Fast Bulk Bitwise AND and OR in DRAM. *CAL* (Jul. 2015).
- [120] V. Seshadri, Y. Kim, C. Fallin, D. Lee, R. Ausavarungnirun, G. Pekhimenko, Y. Luo, O. Mutlu, P. B. Gibbons, M. A. Kozuch, and T. C. Mowry. 2013. RowClone: Fast and Energy-Efficient In-DRAM Bulk Data Copy and Initialization. In *MICRO*.
- [121] V. Seshadri, D. Lee, T. Mullins, H. Hassan, A. Boroumand, J. Kim, M. A. Kozuch, O. Mutlu, P. B. Gibbons, and T. C. Mowry. 2017. Ambit: In-Memory Accelerator for Bulk Bitwise Operations Using Commodity DRAM Technology. In *MICRO*.
- [122] A. Shafiee, A. Nag, N. Muralimanohar, R. Balasubramonian, J. P. Strachan, M. Hu, R. S. Williams, and V. Srikumar. 2016. ISAAC: A Convolutional Neural Network Accelerator With In-Situ Analog Arithmetic in Crossbars. In *ISCA*.
- [123] M. Sharad, D. Fan, and K. Roy. 2013. Ultra Low Power Associative Computing With Spin Neurons and Resistive Crossbar Memory. In *DAC*.
- [124] W. Shim and S. Yu. 2022. GP3D: 3D NAND Based In-Memory Graph Processing Accelerator. *JETCAS* (Jun. 2022).
- [125] C. E. Song, P. Bhatnagar, Z. Xia, N. S. Kim, T. S. Rosing, and M. Kang. 2025. Hybrid SLC-MLC RRAM Mixed-Signal Processing-in-Memory Architecture for Transformer Acceleration via Gradient Redistribution. In *ISCA*.
- [126] L. Song, F. Chen, H. Li, and Y. Chen. 2023. ReFloat: Low-Cost Floating-Point Processing in ReRAM for Accelerating Iterative Linear Solvers. In *SC*.
- [127] L. Song, X. Qian, H. Li, and Y. Chen. 2017. PipeLayer: A Pipelined ReRAM-Based Accelerator for Deep Learning. In *HPCA*.
- [128] L. Song, Y. Zhuo, X. Qian, H. Li, and Y. Chen. 2018. GraphR: Accelerating Graph Processing Using ReRAM. In *HPCA*.
- [129] W. Song, M. Rao, Y. Li, C. Li, Y. Zhou, F. Cai, M. Wu, W. Yin, Z. Li, Q. Wei, S. Lee, H. Zhu, L. Gong, M. Barnell, Q. Wu, P. A. Beerel, M. S.-W. Chen, N. Ge, M. Hu, Q. Xia, and J. J. Yang. 2024. Programming Memristor Arrays With Arbitrarily High Precision for Analog Computing. *Science* (Feb. 2024).
- [130] M. Spear, J. E. Kim, C. H. Bennett, S. Agarwal, M. J. Marinella, and T. P. Xiao. 2023. The Impact of Analog-to-Digital Converter Architecture and Variability on Analog Neural Network Accuracy. *JXCDC* (2023).
- [131] A. Stillmaker and B. Baas. 2017. Scaling Equations for the Accurate Prediction of CMOS Device Performance From 180nm to 7nm. *Integration* (2017).
- [132] G. Strang. 1986. A Proposal for Toeplitz Matrix Calculations. *Stud. Appl. Math.* (Apr. 1986).
- [133] Y. Sun, Y. Wang, and H. Yang. 2017. Energy-Efficient SQL Query Exploiting RRAM-Based Process-in-Memory Structure. In *NVMSA*.
- [134] M. O. Topal, A. Bas, and I. van Heerden. 2021. Exploring Transformers in Natural Language Generation: GPT, BERT, and XLNet. [arXiv:2102.08036 \[cs.CL\]](#).
- [135] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, A. Rodriguez, A. Joulin, E. Grave, and G. Lample. 2023. LLaMA: Open and Efficient Foundation Language Models. [arXiv preprint arXiv:2302.13971](#) (2023).
- [136] M.S.Q. Truong, Y. Sun, D. Xiong, A. Shah, A. Glass, A. Farrell, J. A. Bain, L. R. Carley, and S. Ghose. 2026. The Memory Processing Unit: A Generalized Interface for End-to-End In-Memory Execution. In *HPCA*.
- [137] M. S. Q. Truong, E. Chen, D. Su, L. Shen, A. Glass, L. R. Carley, J. A. Bain, and S. Ghose. 2021. RACER: Bit-Pipelined Processing Using Resistive Memory. In *MICRO*.
- [138] M. S. Q. Truong, L. Shen, A. Glass, A. Hoffmann, L. R. Carley, J. A. Bain, and S. Ghose. 2022. Adapting the RACER Architecture to Integrate Improved In-ReRAM Logic Primitives. *JETCAS* (2022).
- [139] P.-H. Tseng, F.-M. Lee, Y.-H. Lin, L.-Y. Chen, Y.-C. Li, H.-W. Hu, Y.-Y. Wang, C.-C. Hsieh, M.-H. Lee, H.-L. Lung, K.-Y. Hsieh, K.-C. Wang, and C.-Y. Lu. 2020. In-Memory-Searching Architecture Based on 3D-NAND Technology With Ultra-High Parallelism. In *IEDM*.
- [140] P.-H. Tseng, F.-M. Lee, Y.-H. Lin, Y.-Y. Wang, M.-H. Lee, K.-Y. Hsieh, K.-C. Wang, and C.-Y. Lu. 2021. A Hybrid In-Memory-Searching and In-Memory-Computing Architecture for NVM Based AI Accelerator. In *VLSI Symposia*.
- [141] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. 2017. Attention Is All You Need. *NIPS* (2017).
- [142] P. O. Vontobel, W. Robinett, P. J. Kuekes, D. R. Stewart, J. Straznicky, and R. S. Williams. 2009. Writing to and Reading From a Nano-Scale Crossbar Memory Based on Memristors. *Nanotechnology* (Sep. 2009).
- [143] J. Wang, X. Wang, C. Eckert, A. Subramaniyan, R. Das, D. Blaauw, and D. Sylvester. 2019. A Compute SRAM With Bit-Serial Integer/Floating-Point Operations for Programmable In-Memory Vector Acceleration. In *ISSCC*.
- [144] J. Wang, X. Wang, C. Eckert, A. Subramaniyan, R. Das, D. Blaauw, and D. Sylvester. 2020. A 28-nm Compute SRAM With Bit-Serial Logic/Arithmetic Operations for Programmable In-Memory Vector Computing. *JSSC* (Jan. 2020).
- [145] J. Wang, S. Yan, X. Fu, Z. Qian, Z. Li, Z. Guo, Z. Dai, Z. Cong, C. Dou, F. Zhang, J. Yue, and D. Shang. 2025. A High-Density Energy-Efficient CNM Macro Using Hybrid RRAM and SRAM for Memory-Bound Applications. *TVLSI* (Jun. 2025).
- [146] S. Wang, Z. Liu, C. Ding, C. Zhang, T. Wu, J. Zhou, and G. Wong. 2025. NoiseZO: RRAM Noise-Driven Zeroth-Order Optimization for Efficient Forward-Only Training. In *DAC*.
- [147] Y. Wang, Z. Zhu, F. Chen, M. Ma, G. Dai, Y. Wang, H. Li, and Y. Chen. 2021. RREC: In-ReRAM Acceleration With Access-Aware Mapping for Personalized Recommendation. In *ICCAD*.
- [148] R. Wong, N. Kim, A. Das, K. Higgs, E. Ipek, S. Agarwal, S. Ghose, and B. Feinberg. 2025. ANVIL: An In-Storage Accelerator for Name-Value Data Stores. In *ISCA*.
- [149] T. P. Xiao, C. H. Bennett, B. Feinberg, S. Agarwal, and M. J. Marinella. 2020. Analog Architectures for Neural Network Acceleration Based on Non-Volatile Memory. *APR* (Jul. 2020).
- [150] T. P. Xiao, B. Feinberg, C. H. Bennett, V. Agrawal, P. Saxena, V. Prabhakar, K. Ramkumar, H. Medu, V. Raghavan, R. Chettuvetty, S. Agarwal, and M. J. Marinella. 2022. An Accurate, Error-Tolerant, and Energy-Efficient Neural Network Inference Engine Based on SONOS Analog Memory. *TCAS-I* (Jan. 2022).
- [151] T. P. Xiao, B. Feinberg, C. H. Bennett, V. Prabhakar, P. Saxena, V. Agrawal, S. Agarwal, and M. J. Marinella. 2023. On the Accuracy of Analog Neural Network Inference Accelerators. *CAS* (Jan. 2023).
- [152] T. P. Xiao, B. Feinberg, D. K. Richardson, M. Cannon, H. Medu, V. Agrawal, M. J. Marinella, S. Agarwal, and C. H. Bennett. 2024. Analog Fast Fourier Transforms for Scalable and Efficient Signal Processing. [arXiv:2409.19071 \[cs.ET\]](#).
- [153] T. P. Xiao, B. Feinberg, J. N. Rohan, C. H. Bennett, S. Agarwal, and M. J. Marinella. 2021. Analysis and Mitigation of Parasitic Resistance Effects for Analog In-Memory Neural Network Acceleration. *SST* 36 (Oct. 2021).
- [154] M. Xie, S. Li, A. O. Glova, J. Hu, Y. Wang, and Y. Xie. 2018. AIM: Fast and Energy-Efficient AES In-Memory Implementation for Emerging Non-volatile Main Memory. In *DATE*.
- [155] C. Xu, D. Niu, N. Muralimanohar, R. Balasubramonian, T. Zhang, S. Yu, and Y. Xie. 2015. Overcoming the Challenges of Crossbar Resistive Memory Architectures. In *HPCA*.

- [156] L. Xu, S. Yuan, D. Wang, Y. Chen, X. Li, and Y. Sun. 2024. HEIRS: Hybrid Three-Dimension RRAM- and SRAM-CIM Architecture for Multi-task Transformer Acceleration. In *DAC*.
- [157] B. Yan, J.-L. Hsu, P.-C. Yu, C.-C. Lee, Y. Zhang, W. Yue, G. Mei, Y. Yang, Y. Yang, H. Li, Y. Chen, and R. Huang. 2022. A 1.041-Mb/mm² 27.38-TOPS/W Signed-INT8 Dynamic-Logic-Based ADC-less SRAM Compute-in-Memory Macro in 28nm With Reconfigurable Bitwise Operation for AI and Embedded Applications. In *ISSCC*.
- [158] B. Yan, Z. Li, Y. Chen, and H. Li. 2016. RAM and TCAM Designs by Using STT-MRAM. In *NVMTS*.
- [159] T.-H. Yang, H.-Y. Cheng, C.-L. Yang, I.-C. Tseng, H.-W. Hu, H.-S. Chang, and H.-P. Li. 2019. Sparse ReRAM Engine: Joint Exploration of Activation and Weight Sparsity in Compressed Neural Networks. In *ISCA*.
- [160] P. Yao, H. Wu, B. Gao, J. Tang, Q. Zhang, W. Zhang, J. J. Yang, and H. Qian. 2020. Fully Hardware-Implemented Memristor Convolutional Neural Network. *Nature* 577 (Jan. 2020).
- [161] A. Yazdanbakhsh, A. Moradifrouzabadi, Z. Li, and M. Kang. 2022. Sparse Attention Acceleration With Synergistic In-Memory Pruning and On-Chip Recomputation. In *MICRO*.
- [162] G. Yuan, P. Behnam, Z. Li, A. Shafiee, S. Lin, X. Ma, H. Liu, X. Qian, M. N. Bojnordi, Y. Wang, and C. Ding. 2021. FORMS: Fine-Grained Polarized ReRAM-Based In-Situ Computation for Mixed-Signal DNN Accelerator. In *ISCA*.
- [163] İ. E. Yüksel, Y. C. Tuğrul, F. N. Bostancı, G. F. Oliveira, A. G. Yağlıkçı, A. Olgun, M. Soysal, H. Luo, J. Gómez-Luna, M. Sadrosadati, and O. Mutlu. 2024. Simultaneous Many-Row Activation in Off-the-Shelf DRAM Chips: Experimental Characterization and Analysis. In *DSN*.
- [164] İ. E. Yüksel, Y. C. Tuğrul, A. Olgun, F. N. Bostancı, A. G. Yağlıkçı, G. F. Oliveira, H. Luo, J. Gómez-Luna, M. Sadrosadati, and O. Mutlu. 2024. Functionally-Complete Boolean Logic in Real DRAM Chips: Experimental Characterization and Analysis. In *HPCA*.
- [165] Y. Zha and J. Li. 2018. Liquid Silicon: A Data-Centric Reconfigurable Architecture Enabled by RRAM Technology. In *FPGA*.
- [166] Y. Zha and J. Li. 2018. Liquid Silicon-Monona: A Reconfigurable Memory-Oriented Computing Fabric With Scalable Multi-Context Support. In *ASPLOS*.
- [167] F. Zhang, A. Sridharan, W. Tsai, Y. Chen, S. X. Wang, and D. Fan. 2024. Efficient Memory Integration: MRAM-SRAM Hybrid Accelerator for Sparse On-Device Learning. In *DAC*.
- [168] F. Zhang, L. Yang, and D. Fan. 2024. Hyb-Learn: A Framework for On-Device Self-Supervised Continual Learning With Hybrid RRAM/SRAM Memory. In *DAC*.
- [169] L. Zhao, L. Buonanno, A. Gajjar, J. Moon, A. Natarajan, S. Serebryakov, R. M. Roth, X. Sheng, Y. Zhang, P. Faraboschi, J. Ignowski, and G. Pedretti. 2025. NL-DPE: An Analog In-Memory Non-Linear Dot Product Engine for Efficient CNN and LLM Inference. arXiv:2511.13950 [cs.AR].
- [170] L. Zhao, L. Buonanno, R. M. Roth, S. Serebryakov, A. Gajjar, J. Moon, J. Ignowski, and G. Pedretti. 2025. RACE-IT: A Reconfigurable Analog CAM-Crossbar Engine for In-Memory Transformer Acceleration. In *ICCD*.
- [171] P. Zuo, Q. Wang, Y. Luo, R. Xie, S. Wang, Z. Cheng, L. Bao, Z. Wang, Y. Cai, R. Huang, and Z. Sun. 2025. Precise and Scalable Analogue Matrix Equation Solving Using Resistive Random-Access Memory Chips. *Nat. Electron.* (Oct. 2025).