TransPimLib: Efficient Transcendental Functions for Processing-in-Memory Systems

Maurus Item Geraldo F. Oliveira Juan Gómez-Luna Yuxin Guo Mohammad Sadrosadati ETH Zürich

Onur Mutlu

ABSTRACT

Processing-in-memory (PIM) promises to alleviate the data movement bottleneck in modern computing systems. However, current real-world PIM systems have the inherent disadvantage that their hardware is more constrained than in conventional processors (CPU, GPU), due to the difficulty and cost of building processing elements near or inside the memory. As a result, general-purpose PIM architectures support fairly limited instruction sets and struggle to execute complex operations such as transcendental functions and other hard-to-calculate operations (e.g., square root). These operations are particularly important for some modern workloads, e.g., activation functions in machine learning applications.

In order to provide support for transcendental (and other hardto-calculate) functions in general-purpose PIM systems, we present TransPimLib, a library that provides CORDIC-based and LUT-based methods for trigonometric functions, hyperbolic functions, exponentiation, logarithm, square root, etc. We develop an implementation of TransPimLib for the UPMEM PIM architecture and perform a thorough evaluation of TransPimLib's methods in terms of performance and accuracy, using microbenchmarks and three full workloads (Blackscholes, Sigmoid, Softmax). We open-source all our code and datasets at https://github.com/CMU-SAFARI/transpimlib.

KEYWORDS

processing-in-memory, processing-near-memory, transcendental functions, activation functions, machine learning

1 INTRODUCTION

Processor performance increasing more rapidly than memory performance for decades has caused a wide gap between processing units and memory units in terms of latency and energy consumption. As a result, access to data has become a major bottleneck in current computing systems [20, 21, 72, 98, 104, 106]. Processingin-memory (PIM) is a promising solution to this data movement bottleneck. PIM consists of equipping memory with computing capabilities, either with small functional units near the memory arrays or by using the analog operational properties of memory cells themselves [52, 98, 114]. PIM provides access to data at significantly higher bandwidth, lower latency, and lower energy consumption than conventional compute-centric processors (e.g., CPUs, GPUs).

Explored for more than 50 years [69, 125], PIM's materialization into real products was delayed in part by fundamentally different requirements of fabrication of memory units and processing units [27, 30, 108, 120, 138, 144]. Only now we are witnessing the arrival of the first PIM commercial products and prototypes. UP-MEM [135], for example, introduced the first general-purpose commercial PIM architecture [46, 47, 50, 131, 135], which integrates

a small in-order core next to each DRAM bank in a DRAM chip. HBM-PIM [84, 88] and Acceleration DIMM (AxDIMM) [70] are Samsung's real PIM prototypes. HBM-PIM features a SIMD unit that supports multiply-add (MAD) and multiply-accumulate (MAC) operations between every two banks in HBM [65, 86] layers. HBM-PIM is designed to accelerate neural network inference. AxDIMM is a near-rank solution that places an FPGA fabric on a DDRx module to accelerate specific workloads (e.g., recommendation inference). Accelerator-in-Memory (AiM) [87] is a GDDR6-based PIM architecture from SK Hynix with specialized units for multiply-accumulate and activation functions for deep learning. HB-PNM [101] is a 3D-stacked-based PIM architecture from Alibaba, which stacks together a layer of LPDDR4 memory and a logic layer with specialized accelerators for recommendation systems. Among other characteristics, these PIM systems have in common that (1) they place PIM processing elements near memory, and (2) these PIM processing elements are relatively simple and natively support limited instruction sets (e.g., integer arithmetic in UPMEM, 16-bit floating-point MAC/MAD in HBM-PIM or AiM).

Given that instruction sets of current real-world PIM architectures are limited, some complex operations should be emulated by the runtime library (e.g., integer multiplication/division and floating-point arithmetic in UPMEM PIM [50, 134]) or are not even supported. This is usually true for transcendental functions (e.g., trigonometric, hyperbolic, exponentiation, logarithm) and other hard-to-calculate functions (e.g., square root) in PIM architectures. These functions are used in a variety of important applications, such as activation functions in machine learning applications [55, 57], finite element methods [16], ray tracing [44], and option pricing in the stock market [92]. Prior work [126, 128] has identified the presence of transcendental functions that are executed thousands of times in several workloads from various benchmark suites (e.g., SPEC CPU2006 [123], SPLASH-2 [140]). Many of these applications suffer from the data movement bottleneck in conventional processors and, thus, can potentially benefit from PIM. For example, a recent PIM benchmarking study [104] identified seven memorybound applications that use transcendental functions.

To cope with the need for efficient support for transcendental functions and other hard-to-calculate functions in PIM architectures, we analyze possible alternatives to compute these functions. In current computing systems with PIM capabilities [30, 46, 47, 50, 84, 87, 88, 134, 135], there is a host processor (e.g., CPU, GPU) that offloads memory-bound computations to the PIM processing elements. Offloading may require moving data from the standard main memory to the PIM-enabled memory. We identify three possible alternatives to compute transcendental functions in such systems.

First, the PIM hardware can integrate a special hardware unit for transcendental functions. As Figure 1(a) represents, a PIM core can invoke a "transcendental unit", which executes the desired function. Such special units are common in conventional processors, e.g., specialized co-processors in CPUs [63] and special function units in GPUs [102]. To our knowledge, the only real PIM architecture that provides this kind of support is AiM [87], which employs lookup tables (LUTs) and a hardware interpolation unit for different activation functions. Such support may *not* (1) be affordable for all PIM designs (given the area cost of additional hardware units [48, 147]), (2) provide a flexible implementation (i.e., with support for different methods which can better suit different functions), and (3) provide flexible control on the desired precision.



Figure 1: Three options for calculation of transcendental functions in a PIM system.

Second, we can compute transcendental functions on the host processor, as shown in Figure 1(b). This approach has two drawbacks: (1) programmers of computing systems with PIM capabilities need to properly partition the applications, and (2) data needs to move back-and-forth between PIM cores and the host processor, which hampers the potential benefit from PIM.

Third, we can implement fast and self-contained calculation methods for transcendental functions using existing instructions (native or emulated) in the PIM core, as Figure 1(c) shows. While polynomial approximations [28, 67, 124] are frequently used such methods, there is no study of other methods, such as CORDIC [136], or a comprehensive exploration of LUT-based methods [94, 111, 112, 129].

Our **goal** is to explore different methods for calculating transcendental and other hard-to-calculate functions in PIM systems. We develop our methods for the UPMEM PIM architecture [30, 46, 47, 50], the first publicly-available PIM architecture, which consists of general-purpose in-order cores placed near DRAM banks. As a result of our study, we present *TransPimLib*, an open-source library of transcendental and other hard-to-calculate functions for PIM. TransPimLib implements CORDIC-based methods, LUT-based methods, and combinations and variations of them (with and without interpolation). In total, TransPimLib uses eight different methods for trigonometric functions (sine, cosine, tangent), hyperbolic functions (sinh, cosh, tanh), exponentiation, logarithm, square root, and GELU (Gaussian Error Linear Unit) [56].¹

We compare all of TransPimLib's methods in terms of accuracy, execution cycles in the PIM cores, setup time in the host CPU, and memory consumption. TransPimLib's methods provide accuracy of up to 10^{-9} root-mean-square absolute error (RMSE). LUT-based methods, in particular our *LDEXP-based Fuzzy Lookup Table (L-LUT)* method, demonstrate the best tradeoff between performance and accuracy. The non-interpolated L-LUT method requires no multiplication or other complex operations while achieving RMSE= 10^{-7} . The interpolated L-LUT method has a accuracy of RMSE= 10^{-9} at the expense of just one multiplication.

We evaluate TransPimLib's functions for three full workloads: Blackscholes [92], Sigmoid [55], and Softmax [6]. The fastest PIM version of Blackscholes outperforms a 32-thread CPU baseline by 62%. The PIM versions of Sigmoid and Softmax, which are typically used as activation functions in neural networks, provide competitive performance to their 32-thread CPU counterparts and prove that TransPimLib can reduce data movement between PIM cores and the host CPU (as it is needed in Figure 1(b), where transcendental functions run in the host CPU).

Our main contributions are as follows:

- We present TransPimLib, the first library of transcendental and other hard-to-calculate functions for PIM architectures. TransPimLib contains CORDIC-based and LUT-based methods, and combinations and variations of them.
- We propose new LUT-based methods, called *L-LUT*, *D-LUT*, and *DL-LUT* (Section 3.2), which demonstrate good suitability for general-purpose PIM architectures with limited instruction sets, such as the UPMEM PIM system.
- We evaluate TransPimLib's methods in terms of accuracy, execution cycles in the PIM cores, setup time in the host CPU, and memory consumption. We also evaluate TransPimLib's functions for three real workloads (Blackscholes, Sigmoid, Softmax).
- We open-source TransPimLib, as well as all codes and datasets used for evaluation, in our GitHub repository [110].

2 BACKGROUND

This section provides the necessary background on current realworld processing-in-memory systems (Section 2.1), and on transcendental functions and calculation methods (Section 2.2).

2.1 Processing-in-Memory (PIM)

Processing-in-memory (PIM) is a computing paradigm that advocates for memory-centric computing systems, where memory becomes an active system component with computing capabilities. These capabilities can be either (1) small processing elements (general-purpose cores and/or accelerators) placed *near* the memory arrays (e.g., [3–5, 14, 15, 19, 22, 24, 35, 41, 43, 49, 51, 58, 59,

¹We open source TransPimLib to facilitate reproducibility and future research [110].

73, 76, 99, 105, 107, 118, 119, 145, 148], or (2) mechanisms that compute by *using* the analog operational properties of memory components (e.g., by simultaneously activating multiple memory cells [2, 7–12, 17, 23, 25, 26, 33, 37, 39, 40, 52–54, 60, 68, 74, 75, 81–83, 89–91, 103, 114–117, 137, 141–143, 146]). First proposed more than 50 years ago [69, 125], processing-in-memory is a compelling solution to alleviate the *data movement bottleneck* [96–98], which is caused by the need for moving data between memory units and compute units in processor-centric systems. Such bottleneck has only become worse over the years due to the faster development of processor performance over memory performance.

PIM architectures are becoming a reality, with the commercialization of the UPMEM PIM architecture [47, 131, 135], and the announcement of HBM-PIM [84, 88], AxDIMM [70], AiM [87], and HB-PNM [101] (all four prototyped and evaluated in real systems). UPMEM places a small general-purpose in-order core (called DPU) near each memory bank of a DDR4 DRAM chip. HBM-PIM features a SIMD unit (called PCU) between every two banks in the memory layers of an HBM stack [65]. HBM-PIM is designed for machine learning inference. Thus, its SIMD units execute only a reduced set of instructions (i.e., 16-bit floating-point multiplication and addition). AxDIMM [70] is a DIMM-based solution with an FPGA inside the buffer chip of the DIMM. The FPGA can accelerate memory-bound workloads, such as recommendation inference [70] or database operations [85]. AiM [87] is a GDDR6-based PIM architecture with a near-bank processing unit (called PU) that executes multiply-and-accumulate operations and activation functions. Same as HBM-PIM, AiM targets machine learning workloads. HB-PNM [101] is a 3D-stacked based PIM solution for recommendation systems. HB-PNM stacks one layer of LPDDR4 DRAM [66] on one logic layer connected through hybrid bonding (HB) technology [38]. The logic layer embeds two types of specialized engines for matching and ranking, which are the memory-bound steps of the evaluated recommendation system.

These five real-world PIM systems have some important common characteristics, as depicted in Figure 2. First, there is a host processor (CPU or GPU), typically with a deep cache hierarchy, which has access to (1) standard main memory, and (2) PIM-enabled memory (i.e., UPMEM DIMMs, HBM-PIM stacks, AxDIMM DIMMs, AiM GDDR6, HB-PNM LPDDR4). Second, the PIM-enabled memory chip contains multiple PIM processing elements (PIM PEs), which have access to memory (either memory banks or ranks) with much higher bandwidth and lower latency than the host processor. Third, the PIM processing elements (either general-purpose cores, SIMD units, FPGAs, or specialized processors) run at only a few hundred megahertz, and have a small number of registers and relatively small (or no) cache or scratchpad memory. Fourth, PIM PEs may not be able to communicate directly with each other (e.g., UPMEM DPUs, HBM-PIM PCUs or AiM PUs in different chips), and communication between them happens via the host processor. Figure 2 shows a high-level view of such a state-of-the-art processing-in-memory system.

In this work, we use the UPMEM PIM architecture [30, 47, 131, 133–135], the first PIM architecture to be commercialized in real hardware. The UPMEM PIM architecture uses conventional 2D DRAM arrays and combines them with general-purpose processing cores, called *DRAM Processing Units* (*DPUs*), on the same chip. There



Figure 2: Example state-of-the-art processing-in-memory system with general-purpose PIM cores. The host CPU has access to *m* standard memory modules and *n* PIM-enabled memory modules.

are 8 DPUs and 8 DRAM banks per chip, and 16 chips per DIMM (8 chips/rank). DPUs are relatively deeply pipelined and fine-grained multithreaded [121, 122, 130]. DPUs run software threads, called *tasklets*, which are programmed in *Single Program Multiple Data* (*SPMD*) manner.

DPUs have a 32-bit RISC-style general-purpose instruction set [134]. They feature native support for 32-bit integer, but some complex operations (e.g., 32-bit integer multiplication/division) and floating-point operations are emulated [47].

Each DPU has exclusive access to its own (1) 64-MB DRAM bank, called *Main RAM (MRAM)*, (2) 24-KB instruction memory, called *Instruction RAM (IRAM)*, and (3) 64-KB scratchpad memory, called *Working RAM (WRAM)*. The host CPU can access the MRAM banks for copying input data (from main memory to MRAM) and retrieving results (from MRAM to main memory). These CPU-DPU/DPU-CPU transfers can be performed in parallel (i.e., concurrently across multiple MRAM banks), if the size of the buffers transferred from/to all MRAM banks is the same. Otherwise, the data transfers should be performed serially. Since there is no direct communication channel between DPUs, all inter-DPU communication takes place through the host CPU by using DPU-CPU and CPU-DPU data transfers. We refer the reader to our prior work [46, 47, 50] for a comprehensive introduction to and analysis of the UPMEM PIM system.

Throughout this paper, we use generic terminology, since our implementation strategies are applicable to PIM systems like the generic one described in Figure 2, and not exclusive to the UPMEM PIM architecture. Thus, we use the terms *PIM core*, *PIM thread*, *DRAM bank*, *scratchpad*, and *Host-PIM/PIM-Host transfer*, which correspond to DPU, tasklet, MRAM bank, WRAM, and CPU-DPU/DPU-CPU transfer in UPMEM's terminology [134].

2.2 Transcendental Functions

Transcendental functions [34, 80] are functions that do *not* satisfy a polynomial equation. As a result, they cannot be exactly expressed with a finite number of algebraic operations (e.g., addition, sub-traction, multiplication, division, power, root). Commonly-used transcendental functions are trigonometric functions (e.g., sine, co-sine), exponential functions, and logarithm functions. In this work, we also target other functions that are not transcendental, but hard to calculate (e.g., square root).

There are various methods to calculate transcendental functions, such as *Taylor approximation* [67, 124], *minimax polynomials* [67],

CORDIC [79, 95, 136], and *table-based methods* [94, 111, 112, 129]. In this work, we focus on (1) CORDIC [136] and (2) *fuzzy lookup tables* [94, 111, 112, 129], since these methods have low usage of floating-point multiplication (more expensive than addition/subtraction in current general-purpose PIM architectures [50, 133]). We introduce both methods in Sections 2.2.1 and 2.2.2. In Section 2.2.3, we describe how we extend the input range of these methods.

2.2.1 CORDIC. CORDIC [136] is an iterative method that uses only bit-shifts, additions, and table lookups. The maximum error shrinks exponentially with the number of iterations [79]. We use CORDIC in *rotation mode* [136], which computes a function value (e.g., sine, cosine) for a given input (an angle θ). CORDIC starts with an angle $\theta_0 = \theta$ and a vector $v_0 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$. In each iteration (i > 0), we update θ_i and rotate the vector by multiplying it with a 2×2 matrix $M_i: v_{i+1} = k_i \cdot M_i \cdot v_i$. k_i is a stretching factor. M_i represents a rotation by an angle ϕ_i , which decreases in each iteration. The values of ϕ_i can be precalculated and kept in a table. After *n* iterations, the output vector v_n contains function values for θ . Table 1 shows common rotation matrices, angles, stretching factors, and functions they are used for.

Table 1: CORDIC's Rotation Matrices, Angles, and Stretching Factors

Matrix	Mi	ϕ_i	ki	Functions			
Circular	$\begin{bmatrix} 1 & \mp 2^{-i} \\ \pm 2^{-i} & 1 \end{bmatrix}$	$\pm arctan(2^{-i})$	$\sqrt{1+2^{-2i}}$	sin, cos, tan, arctan			
Hyperbolic	$\begin{bmatrix} 1 & \pm 2^{-i} \\ \pm 2^{-i} & 1 \end{bmatrix}$	$\pm atanh(2^{-i})$	$\frac{2^{-i-1}+1}{2^{-i-1}}$	sinh, cosh, tanh, exp, log, sqrt, atanh			
Linear	$\begin{bmatrix} 1 & 0 \\ \pm 2^{-i} & 1 \end{bmatrix}$	$\pm 2^{-i}$	1	multiplication, division			

2.2.2 Fuzzy Lookup Tables. Lookup tables are widely used to generate a variety of functions [94, 111, 112, 129]. Since lookup tables are limited in size, it is not always possible to return exact output values, but easier to return *fuzzy* (or approximate) matches. For a given input *x*, we obtain an output $\tilde{f}(x)$ that approximates the original function f(x). First, we use a function a(x) that returns an address. a(x) typically maps a range of inputs *x* to a single address. For example, a(x) may simply round down, such that all values in the range [0, 1) map to address 0. Second, we access the lookup table with a(x) as the address, and the table returns $l(a(x)) = \tilde{f}(x)$.

In order to generate the lookup table, we need a helper function $a^{-1}()$, which is the pseudo-inverse of a(x) in the sense that $x = a(a^{-1}(x))$, $\forall x$. For the previous example, $a^{-1}(0)$ returns one specific value in the range [0, 1). Thus, address 0 of the lookup table will contain $f(a^{-1}(0))$, i.e., the lookup result for any input in [0, 1). The function $a^{-1}()$ determines the spacing between table entries and, thus, the maximum error. For example, in [0, 1) the maximum error will be different if the table contains f(0) or f(0.5). A good spacing places more entries where the function's slope is steeper, in order to minimize the error. As a result, the spacing is in proportion to the function's first derivative (f'(x)). For this reason, the selection of a good $a^{-1}()$ is not always trivial. However, $a^{-1}()$ is only used during table generation, not during table lookups, which allows us to improve the accuracy (minimize the maximum error) without affecting performance.

We can use *interpolation* to further improve the accuracy. In this case, we query two lookup table entries, e.g., l(a(x)) and l(a(x)+1). Then, we approximate f(x) as $\tilde{f}(x) = l(a(x)) + l(a(x) + 1) - l(a(x)) \cdot \Delta$, where Δ represents where the input resides between the two lookup table entries. For example, $\Delta = 0.5$ means that the input *x* is exactly in between the two entries. In the general case, $\Delta = \frac{x-a^{-1}(a(x))}{a^{-1}(a(x)+1)-a^{-1}(a(x))}$, which needs to be calculated for each input *x*. This calculation can be greatly simplified for our target functions, as we show in Section 3.2.1. Regarding the spacing between entries $(a^{-1}())$, the error grows when a function's slope changes quickly (i.e., when the *rate of change* of the function's first derivative is high). Thus, a desirable spacing should follow the function's second derivative (f''(x)).

2.2.3 Range Extensions. Both CORDIC and lookup tables support limited ranges of inputs. For some transcendental functions, it is possible to extend the range by performing a conversion that depends on the function itself. For example, trigonometric functions (e.g., sine, cosine) take advantage of the fact that the output repeats every 2π . For other functions, such as square root, logarithm, and exponentiation, we can separate exponent and mantissa. For example, to calculate the logarithm of $x = 2^{exponent(x)} \cdot mantissa(x)$, we convert $log(x) = log(2) \cdot exponent(x) + log(mantissa(x))$.

3 TRANSPIMLIB: IMPLEMENTATION

Given their relatively simple hardware, general-purpose PIM cores do not support complex operations (e.g., [50, 84, 88, 133]), such as transcendental (e.g., trigonometric, logarithm, exponentiation) and other hard-to-calculate functions (e.g., square root). To fill this gap, we present *TransPimLib*, a library of transcendental and other hard-to-calculate functions for general-purpose PIM cores. TransPimLib leverages the implementation methods introduced in Section 2.2 (e.g., CORDIC, lookup tables), and variations and combinations of them. Table 2 shows TransPimLib's implementation methods and supported functions.²

Table 2: TransPimLib's Implementation Methods and Supported Functions

	Supported Functions										
Implementation Method	sin	cos	tan	sinh	cosh	tanh	exp	log	sqrt	GELU	
CORDIC	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark		
M-LUT	\checkmark	\checkmark	\checkmark				\checkmark	\checkmark	\checkmark		
M-LUT+Interpolation	\checkmark	\checkmark	\checkmark				\checkmark	\checkmark	\checkmark		
L-LUT	\checkmark	\checkmark	\checkmark				\checkmark	\checkmark	\checkmark		
L-LUT+Interpolation	\checkmark	\checkmark	\checkmark				\checkmark	\checkmark	\checkmark		
D-LUT+Interpolation	\checkmark					\checkmark				~	
DL-LUT+Interpolation	\checkmark					\checkmark				~	
CORDIC+LUT	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark				

TransPimLib provides files [110] to be included, using the #include directive, in the host CPU code (for the necessary setup, e.g., loading lookup tables) and the PIM core code for the different implementation methods. The APIs are simple and intuitive. For example, for the sine function: float sinf (float x);

 $^{^2}$ We provide all implementation methods for sine on the UPMEM PIM architecture [135]. Based on our preliminary analysis of these methods, we also provide the most suitable methods for each of the other supported functions. Future work can extend TransPimLib with new supported functions.

3.1 CORDIC-based Implementations

TransPimLib contains CORDIC implementations of trigonometric (sin, cos, tan) and hyperbolic (sinh, cosh, tanh) functions, exponentiation, logarithm, and square root.

We illustrate their implementation using the sine function as an example. As Figure 3(a) shows, the calculation takes six steps. First, we translate the input value (angle) to the range 0 to 2π . Second, if the input is a floating-point value, we convert it to fixed-point format.³ Our fixed-point format uses 28 bits for the fractional part, 3 bits for the integer part (enough to represent up to 2π), and 1 sign bit. Third, the range is further reduced to the range from 0 to $\frac{\pi}{2}$ (the quadrant, i.e., 0 to $\frac{\pi}{2}$, $\frac{\pi}{2}$ to π , π to $\frac{3\pi}{2}$, $\frac{3\pi}{2}$ to 2π , is also saved to not lose information). Fourth, the CORDIC algorithm iterates until it obtains the sine of the input angle. Fifth, we adjust the sine value based on the quadrant of the input angle. Sixth, we convert the output value back to floating point format.



Figure 3: TransPimLib's CORDIC-based implementations of sine and square root functions.

Other functions may require additional steps. For example, our square root implementation (Figure 3(b)) uses a final range extension step.

TransPimLib does *not* implement multiplication and division (see Table 1), which are not natively supported by the UPMEM PIM architecture, because the emulations of the runtime library [50] have similar complexity as potential CORDIC-based implementations and no range limitations.

3.2 Lookup Table-based Implementations

TransPimLib implements several lookup table-based methods that differ in their address generation functions (a() and $a^{-1}()$). These methods provide different tradeoffs, as we explain next.

3.2.1 Multiplication-based Fuzzy Lookup Table (M-LUT). This method defines regular spacing between table entries [139]. We define the address generation function $a(x) = round((x - p) \cdot k)$, where p and k are constants. k represents the density of the lookup table, which is the inverse of the spacing. p defines what input

value *x* corresponds to address 0 of the lookup table. For example, to map the interval [0, 5] to a 12-entry M-LUT, we can use $k = \frac{12}{5} = 2.4$ and $p = \frac{5}{2\cdot 12} = 0.20834$. Thus, an input x = 3 translates to $a(3) = round((3 - 0.20834) \cdot 2.4) = round(6.7) = 7$. Figure 4(a) depicts the coverage of a 12-entry M-LUT with density 2.4 for the interval [0, 5].



Figure 4: Example lookup table density (y axis) for the input range [0, 5] (x axis) and 12 LUT entries for TransPimLib's four different LUT implementations.

For the M-LUT, the inverse operation is $a^{-1}(a(x)) = \frac{a(x)}{k} + p$. For our previous example, address 7 of the M-LUT stores the exact value for the input x = 7/2.4 + 0.20834 = 3.125.

The M-LUT's address generation needs one subtraction, one multiplication, and one step of rounding or truncation. For interpolated M-LUTs, we use the address generation function $a(x) = floor((x - p) \cdot k)$ to get the next smaller lookup table address. The calculation of Δ simplifies to $\Delta = floor((x - p) \cdot k) - (x - p) \cdot k$. As a result, with respect to the M-LUT, the interpolated M-LUT needs one extra lookup table query, one extra multiplication (to compute $\tilde{f}(x)$), and one extra subtraction (to compute Δ).

3.2.2 LDEXP-based Fuzzy Lookup Table (L-LUT). Multiplication is generally expensive [47, 50], but we can make it cheaper if we multiply by 2^n . We use the function 1dexp(arg, exp), which is common in standard math libraries [1], to perform the operation $arg \cdot 2^{exp}$. This function is not available among UPMEM library functions, but we implemented it in accordance with the C99 standard [36].

For the L-LUT, we define the address generation function as $a(x) = round((x - p) \cdot 2^n)$, which loses some freedom to design lookup tables (the density *k* must be a power-of-two), but avoids costly multiplication [47, 50]. For the example of a 12-entry lookup table for the interval [0, 5], we can no longer have density k = 2.4, but k = 2 (i.e., a power of 2). This results in an L-LUT of lower precision than the M-LUT, but expands the range to the interval [0, 6] (Figure 4(b)).

3.2.3 Direct Float Conversion-based Fuzzy Lookup Table (D-LUT). As we mention in Section 2.2.2, a good spacing depends on the approximated function to minimize the error. This may need non-linear address generation functions, which are generally more computationally expensive than multiplications. We circumvent this issue by exploiting the natural non-linearity of the floating-point format. We propose an address generation function that uses (1) the last *n* bits of the exponent, and (2) *p* bits of the manitissa. This results in a piece-wise linear density with 2^n steps of 2^p addresses each. The density of each step is inversely proportional to the input value, i.e., high density for small inputs and low density for large inputs. For example, for our 12-entry table, we can use n = 2 (thus, exponents 2^0 , 2^1 , 2^2) and p = 2 (thus, 4 entries per exponent). As

³CORDIC can operate with fixed-point arithmetic [79], which is natively supported by the UPMEM PIM architecture [50].

Figure 4(c) shows, the resulting density is 4 in [1, 2), 2 in [2, 4), and 1 in [4, 8).

The limitation of the D-LUT is that there are no LUT entries between the smallest exponent (e.g., 2^0) and 0, which may cause large inaccuracy for LUT queries near 0. To deal with this issue, we propose a combined L-LUT + D-LUT method called *DL-LUT* (Section 3.3.1).

3.3 Combined Implementations

In addition to the previous implementations, TransPimLib combines pairs of implementation methods to leverage the strengths of two different implementation methods.

3.3.1 Direct Float Conversion + LDEXP-based Fuzzy Lookup Tables (DL-LUT). This combination solves the limitation of D-LUT (i.e., no entries between the smallest exponent and 0) by combining a D-LUT with an L-LUT. The DL-LUT uses (1) an L-LUT between 0 and the smallest exponent, and (2) a D-LUT for larger inputs, providing a density pattern as depicted in Figure 4(d).

3.3.2 CORDIC + LDEXP-based Fuzzy Lookup Table (CORDIC+LUT). Prior work [13] proposes to replace the first few iterations of CORDIC with a LUT (while still updating θ_i). This provides a flexible tradeoff between computing cost, table size, and precision, within the bounds of pure CORDIC and pure LUT approaches (TransPim-Lib uses L-LUT for this combined method).

4 EVALUATION

This section presents our evaluation of TransPimLib on a real-world PIM system. Section 4.1 introduces our evaluation methodology. Section 4.2 presents a microbenchmark-based analysis of the different calculation methods used in TransPimLib. Section 4.3 presents the evaluation of TransPimLib for three real workloads on the PIM system and the comparison to their multithreaded CPU implementations of the three real methods.

4.1 Methodology

We evaluate TransPimLib on a real-world system with the UPMEM PIM architecture. The system consists of a host CPU (2-socket Intel Xeon with 32 cores at 2.10 GHz), standard main memory (128 GB), and 20 UPMEM PIM DIMMs (159 GB and 2545 PIM cores at 350 MHz) [132].

4.1.1 Microbenchmarks. In Section 4.2, we evaluate TransPimLib's CORDIC, LUT-based, and combined implementations on a single PIM core (DPU in UPMEM terminology) using microbenchmarks, in order to compare them in terms of (1) performance, (2) accuracy, (3) memory consumption, and (4) setup time. Our microbenchmarks compute all our implemented transcendental functions (i.e., all versions of all functions in Table 2) for the elements of an array (of 2¹⁶ floating-point values with random uniform distribution) that resides in a DRAM bank (MRAM in UPMEM terminology). The PIM core moves chunks of the array into the scratchpad memory (WRAM in UPMEM terminology) and operates on each element.

For performance comparison, we measure total execution cycles using a hardware counter [134].

For accuracy comparison, we compare to the output of the host CPU, computed with the standard math library. We obtain the root-mean-square absolute error (RMSE), which we analyze in Section 4.2. Maximum absolute error and error in terms of units of last place (ULP) [45] show very similar trends to the RMSE.

For memory consumption, we account for all tables and variables that we allocate in the DRAM bank of a PIM core.

For setup time, we include the generation of any tables or variables on the host CPU, and their transfers to the DRAM bank of a PIM core.

4.1.2 Benchmarks. In Section 4.3, we implement three full workloads that use transcendental functions (Blackscholes, Sigmoid, and Softmax) on the UPMEM PIM architecture, and compare them to their CPU-only versions.

Blackscholes [18, 92] calculates the prices for a portfolio of options with a partial differential equation. This benchmark uses several functions that benefit from TransPimLib: exponentiation, logarithm, square root, and cumulative normal distribution function (CNDF). The original benchmark implements CNDF using polynomial approximation. We also implement CNDF using TransPimLib's LUT-based methods. In our experiments, we use an input vector of 10M elements.

Sigmoid [55] is a bounded differentiable function whose derivative is always positive. A general equation of the Sigmoid function with a scalar input *x* is defined as $S(x) = \frac{1}{1+e^{-x}}$. Sigmoid is commonly used in logistic regression [57] to compute the probability of an output event. It is also frequently used as an activation function in neural networks. Our Sigmoid benchmark takes an input vector of 30M elements and computes the Sigmoid output of each input element.

Softmax [6] is a function that turns a vector of *K* real values into a vector of *K* real values that sum to 1 ($\sigma(\mathbf{z})_j = \frac{e^{z_j}}{\sum_{k=1}^{K} e^{z_k}}$). Softmax is frequently used as the last layer of neural networks and in reinforcement learning.⁴ In our experiments, the input vector contains 30M values.

On the UPMEM PIM architecture, we implement PIM baselines of these three workloads that do *not* use TransPimLib functions. These baseline PIM implementations are instead based on polynomial approximation [67, 124]. For the PIM implementations that use TransPimLib, we use interpolated M-LUT and L-LUT methods for both Blackscholes, Sigmoid, and Softmax. Additionally, we implement a version of Blackscholes that operates on fixed-point values.

4.2 Microbenchmarks Results

We use microbenchmarks to evaluate all implementation methods and supported functions in Table 2. In this section, we analyze the different implementation methods for the sine function, as a representative function. We evaluate floating-point versions of all implementation methods, and also fixed-point versions of L-LUT methods. Our observations and takeaways are applicable to the other functions as well.

We run experiments for all CORDIC, LUT-based, and combined methods where we tune the methods (e.g., number of iterations

⁴At the time of writing, there are no public implementations of neural networks (e.g., convolutional neural networks) or reinforcement learning for the UPMEM PIM architecture where we can test TransPimLib-based Softmax. Such implementations are out of the scope of this work but TransPimLib can be a useful resource for them.

in CORDIC-based methods, LUT size in LUT-based methods) to obtain an accuracy range, i.e., different root-mean-square absolute errors between 10^{-4} and 10^{-9} . The desired accuracy impacts the execution time (on the PIM side), the setup time (on the host CPU side), and the memory consumption (on the PIM side). Hence, we analyze (1) performance (execution cycles on the PIM core), (2) setup time (seconds), and (3) memory consumption (bytes) as a function of the root-mean-square absolute error (RMSE).

4.2.1 Execution Cycles. Figure 5 shows the execution cycles (per input element) on a PIM core for the different TransPimLib methods that implement the sine function as a function of the accuracy provided by the method. LUT-based versions place the LUT in either the PIM core's DRAM bank (MRAM, solid line) or the scratchpad (WRAM, dashed line). We make five observations.



Figure 5: Execution cycles (y axis, linear scale) per input element on one PIM core as a function of root-mean-square absolute error values (x axis, logarithmic scale) for different TransPimLib implementations of the sine function.

First, each LUT-based method consumes the same number of cycles for any RMSE, because the number of LUT accesses and operations is independent of the LUT size. LUT-based methods with the same number of floating-point multiplications take similar number of execution cycles, i.e., the number of floating-point multiplications determines the number of execution cycles. The slowest of the LUT-based methods is the interpolated M-LUT method (yellow line), which executes two floating-point multiplications per input element. Non-interpolated M-LUT (purple) and interpolated L-LUT (green) execute one multiplication, while non-interpolated L-LUT (cyan) needs no multiplications. As a result, L-LUT methods (with our custom LDEXP operation, Section 3.2.2) outperform M-LUT methods for the entire accuracy range. Interpolated L-LUT reduces the cycle count by ~50% over interpolated M-LUT, and non-interpolated L-LUT reduces the count by ~80% over non-interpolated M-LUT. The fixed-point version of the non-interpolated L-LUT method (teal) does not improve the performance over its floating-point counterpart (none of them use multiplications). However, the fixedpoint version of the interpolated L-LUT (orange) doubles the performance of the floating-point version of the interpolated L-LUT (green) on the UPMEM PIM architecture, where floating-point multiplications are significantly more costly than fixed-point multiplications [47, 50].

Second, CORDIC-based methods take more execution cycles to provide higher accuracy, because accuracy increases with each iteration of the CORDIC algorithm. CORDIC+LUT runs faster than pure CORDIC, as it replaces the initial iterations of the CORDIC algorithm with an L-LUT query. Third, L-LUT methods obtain the best tradeoff in terms of performance and accuracy. Note that curves that are closer to the bottomleft corner of the plot (e.g., interpolated floating-point L-LUT, interpolated fixed-point L-LUT) are better. Interpolated floating-point and fixed-point L-LUT methods are only 2-3 times more costly than a single-precision floating-point and an integer multiplication, respectively. As a result, L-LUT methods have an inherent advantage on PIM architectures over other methods, such as Taylor approximation [67, 124], where one floating-point multiplication is needed for each bit of precision (e.g., ~28 multiplications for RMSE= 10^{-9}).

Fourth, there is no significant performance difference between allocating and accessing LUTs in the DRAM bank (MRAM, solid line) or in the scratchpad (WRAM, dashed line), and this observation holds for any number of PIM threads running on the PIM core. However, the scratchpad is significantly smaller than the DRAM bank. Thus, the accuracy is limited by the maximum possible LUT size (most noticeably for non-interpolated methods, e.g., in Figure 5, non-interpolated fixed-point L-LUT results in RMSE= $\sim 10^{-6}$ in WRAM (teal, dashed) but $\sim 10^{-7}$ in MRAM (teal, solid)). As such, placing LUTs in MRAM can be a good choice to save WRAM space for input/output operands.

Fifth, at around RMSE= 10^{-9} , further increasing the LUT size or the number of CORDIC iterations does *not* provide further accuracy increase. In our experiments, we also measure the maximum absolute error to be around 10^{-7} (not shown in Figure 5). Intuitively, this is due to the precision of floating-point values being limited to $4 \cdot 2^{-24}$ ($2.38 \cdot 10^{-7}$) for inputs in the range [4, 8) (inputs values are in the range [0, 2π]). For fixed-point values, we use a 28-bit fractional part, such that the precision of 2^{-28} ($3.7 \cdot 10^{-9}$) is sufficient to match the accuracy provided by floating-point values. This makes fixed-point L-LUT methods a good choice for general-purpose PIM architectures such as UPMEM PIM, where floating-point computation is not natively supported.

Key Takeaway 1. Interpolated L-LUT methods (lookup table with LDEXP operation) offer the best tradeoff in terms of performance and accuracy.

4.2.2 Setup Time in Host CPU. Figure 6 shows the setup time on the host CPU for each implementation method as a function of the accuracy provided by the method. In combination with the execution cycles on the PIM core, the setup time gives us a more complete understanding about when to use CORDIC-based or LUT-based methods. We make two observations.

First, CORDIC-based methods have flat setup times while LUTbased methods have setup times that increase with the table size. We observe that pure CORDIC implementations can provide higher performance (i.e., lower setup time) than LUT-based methods when the total number of transcendental operations in a workload is low. For example, we can compare L-LUT and pure CORDIC at RMSE= 10^{-9} : (1) CORDIC takes 5380 cycles more than L-LUT on the PIM core, and (2) L-LUT's setup time on the host CPU is $\sim 5 \cdot 10^{-4}$ seconds longer than CORDIC's setup time. With a PIM core running at 425 MHz, a PIM kernel would need to execute ~40 sine operations for the L-LUT sine to amortize the setup time with respect to the CORDIC-based sine. Thus, CORDIC appears to be preferable for kernels computing just a few transcendental functions (e.g., less than 40 sine operations in the previous example).



Figure 6: Setup time in seconds (y axis, linear scale) on the host CPU as a function of root-mean-square absolute error values (x axis, logarithmic scale) for different TransPim-Lib implementations of the sine function.

Second, CORDIC+LUT has higher setup times than pure CORDIC due to the use of a LUT for the initial iterations. However, CORDIC+LUT's setup times are mostly flat due to the use of CORDIC for later iterations. Compared to CORDIC and L-LUT, CORDIC+LUT is a good choice for kernels with very few transcendental functions that require high accuracy.

Key Takeaway 2. CORDIC-based methods are preferable when a PIM kernel needs to execute just a few transcendental functions (e.g., less than 40 sine operations in a kernel running on the UPMEM PIM architecture) due to their low setup time in the host CPU.

4.2.3 *Memory Consumption.* Figure 7 shows the memory consumption (in bytes) per PIM core of all LUT- and CORDIC-based implementation methods as a function of root-mean-square absolute error values. We make several observations.



Figure 7: Memory consumption in bytes (y axis, logarithmic scale) per PIM core as a function of root-mean-square absolute error values (x axis, logarithmic scale) for different TransPimLib implementations of the sine function.

First, accuracy of non-interpolated LUT-based methods is limited by the amount of available memory (DRAM bank or scratchpad). Thus, they are only recommended when fast calculation is needed but lower accuracy is acceptable.

Second, CORDIC and CORDIC+LUT methods have the advantage that their memory consumption does *not* grow exponentially. Thus, they are recommended for applications that require high accuracy, where the amount of memory that can be devoted to TransPimLib is limited. For example, this can happen in applications with large datasets, where we need to allocate most of the space in the PIM core's DRAM bank to input/output operand arrays. Third, interpolation is an effective way of increasing accuracy without increasing LUT size. Overall, interpolated L-LUT offers a good tradeoff in terms of accuracy, execution cycles, and memory consumption. For example, at the maximum accuracy that non-interpolated LUT-based methods can provide ($\sim 10^{-7}$), interpolated L-LUT needs less memory than CORDIC+LUT and it is significantly faster than pure CORDIC (see Figure 5).

Key Takeaway 3. Interpolated L-LUT methods offer a good tradeoff in terms of accuracy, execution cycles, and memory consumption. However, CORDIC and CORDIC+LUT methods are recommended for applications that require high accuracy, where the available memory is needed for large datasets (i.e., not available for lookup tables required for the necessary accuracy).

4.2.4 Other Supported Functions. The general trends for other functions supported by TransPimLib are similar to those of the sine function, which we discuss above. Some major differences that are worth highlighting are as follows.

First, methods for tangent calculation take around 2-3 times more cycles than the same methods for sine. This is explained by the fact that tangent needs (1) calculation of sine *and* cosine, and (2) a floating-point division (much costlier than a floating-point multiplication on UPMEM [47, 50]).

Second, some supported functions may require range reduction and/or range extension (Section 2.2.3), e.g., sine/cosine, exponentiation, logarithm, square root. The cost of these operations largely differs between functions, because it depends on the specific operations needed for the conversion (e.g., mathematical identity that applies to each function, Section 2.2.3). Figure 8 shows the execution cycles per input element for range reduction/extension in *sin, exp, log*, and *sqrt*. Note that range reduction/extension is only necessary depending on the range of input values. For example, our experiments with the sine function (Figures 5 to 7) use input values in $[0, 2\pi)$.



Figure 8: Execution cycles per input element for range reduction/extension of TransPimLib implementations of the sine, exponential, logarithm, and square root functions.

Third, functions that do *not* need range reduction/extension are cheaper to calculate. This is the case for activation functions such as *tanh* and GELU (Gaussian Error Linear Unit) [56], which are approximately linear in most parts. D-LUT and DL-LUT methods are particularly well-suited for *tanh* and GELU, unlike sine (Figure 5). They are ~2× faster than, e.g., interpolated L-LUT for sine, while providing similar accuracy.

Key Takeaway 4. D-LUT and DL-LUT methods are well-suited for activation functions, such as tanh and GELU, which (1) do not require range extension, and (2) are approximately linear in most parts. They are faster than interpolated L-LUT, while providing similar accuracy.

4.3 Real-World Benchmark Results

We implement for the UPMEM PIM architecture three full workloads (Section 4.1.2) that make use of several functions that are supported by TransPimLib. We compare them to their single-thread and 32-thread CPU baselines as well as a PIM baseline that uses polynomial approximation.

Figure 9 shows the execution time of PIM implementations of Blackscholes, Sigmoid, and Softmax (on 2545 PIM cores running 16 PIM threads each), and the CPU baselines (on 1 and 32 CPU cores). We test PIM versions that use (1) polynomial approximation [67, 124] (for comparison to TransPimLib's methods), (2) interpolated M-LUT, and (3) interpolated L-LUT. For Blackscholes, we also test a version with interpolated fixed-point L-LUTs. We include in the measurements all range reduction/extension costs needed for some functions (Figure 8). We make the following observations from Figure 9.



Figure 9: Execution time (s) of Blackscholes (a), Sigmoid (b), and Softmax (c) implementations on PIM (on 2545 PIM cores running 16 PIM threads each), 1 CPU core, and 32 CPU cores.

First, the PIM versions of Blackscholes that use TransPimLib reduce the execution time by $5-10\times$ with respect to the baseline PIM version with polynomial approximation. The M-LUT and L-LUT versions are, respectively, within 75% and 82% the performance of the 32-thread CPU baseline. The fixed-point L-LUT version (on 2545 PIM cores running 16 PIM threads each) is 62% faster than the 32-thread CPU baseline.

Second, both Sigmoid and Softmax show similar qualitative behavior. TransPimLib's methods outperform the PIM version with polynomial approximation by 50-75%. The 32-thread CPU baselines are around 2× faster than TransPimLib's PIM version. Nonetheless, these functions are typically used as part of neural networks and machine learning algorithms, which may run on the PIM cores. Thus, TransPimLib's methods can reduce data movement from PIM cores to the CPU (Figure 1(b)) for applications running on the PIM cores. As a result of saving such PIM-Host and Host-PIM transfers, the execution of transcendental functions in the PIM cores (Figure 1(c)) could be $6 - 8 \times$ faster than the execution in the host CPU (as inferred from Figure 1(b)).

5 RELATED WORK

To our knowledge, TransPimLib is the first library of transcendental (and other hard-to-calculate) functions for general-purpose PIM systems.

5.1 Real Processing-in-Memory Systems

PIM has become a reality in the last few years. UPMEM [135] was first to release their PIM architecture [50, 134]. Since it is the first publicly-available general-purpose PIM architecture, we have implemented TransPimLib for it.

There is a good amount of recent works that analyze the UPMEM PIM architecture and implement important applications for it. An experimental characterization of the UPMEM PIM architecture and a benchmark suite is presented in [46, 50]. SpMV, an important memory-bound kernel, is extensively explored in [42]. The *wavefront algorithm (WFA)* [93], which is currently the state-of-the-art gap-affine pairwise alignment algorithm, a key step in genome analysis, is implemented on the UPMEM PIM architecture in [31, 32].

Besides UPMEM, there have been several prototypes of real PIM chips developed by major vendors in industry, including Samsung, SK Hynix, and Alibaba, in 2021-2022. Samsung introduced HBM-PIM, also known as FIMDRAM, [84, 88], an architecture that embeds one floating-point SIMD unit with a reduced instruction set, called Programmable Compute Unit (PCU), next to two DRAM banks in HBM2 layers. This architecture is targeted to accelerate machine learning inference. The second prototype from Samsung is AxDIMM [70, 85]. AxDIMM is a DIMM-based solution which places an FPGA fabric in the buffer chip of the DIMM. It has been tested for DLRM recommendation inference [70, 100] and in-memory databases [85]. Another major DRAM vendor, SK Hynix, introduced Accelerator-in-Memory [87], a GDDR6-based PIM architecture with specialized units for multiply-and-accumulate and lookup-tablebased activation functions for deep learning applications. AiM [87] uses LUTs and interpolation hardware for activation functions. A key difference with our work is that AiM requires dedicated hardware for address generation and interpolation. Alibaba introduced HB-PNM [101], a PNM system with specialized engines for recommendation systems, which is composed of a DRAM die and a logic die vertically integrated via hybrid bonding [38]. TransPim-Lib can be realized for any PIM architecture that supports addition, subtraction, multiplication, and division. As such, future work can implement new versions of TransPimLib's methods for other current and future PIM architectures.

Though we expect that real PIM systems will continue improving their computing capabilities, integrating processing elements in DRAM technology is challenging and constrains design decisions heavily [30, 71, 78]. For example, DRAM has a lower number of metal layers than CMOS and slower transistors [27, 30, 108, 120, 138, 144]. As a result, manufacturing complex execution units (as those needed for transcendental and other hard-to-calculate functions) using DRAM technology would require the addition of extra (costly) metal layers while resulting in low frequency logic units [71, 78], which will hardly be affordable in PIM systems. Other future PIM systems, e.g., 3D-stacked memories with processing elements in a logic layer [61, 62, 65], can make the integration of complex execution units easier. However, their area and thermal budget will still be constrained. For example, in a forward-looking HBM-based PIM system [77], PIM logic can occupy only ~28% of the logic layer due to the need for peripheral/control logic. Thus, the availability of libraries for complex operations, such as TransPimLib, will likely continue to be necessary.

5.2 Acceleration of Transcendental Functions

Several works [29, 113] aim to improve LUT-based implementation methods (e.g., improving accuracy for a given memory consumption). These works are not specific to PIM systems. Compared to them, TransPimLib is simpler and requires less calculations before and after lookup table queries.

Several other works [109, 127, 128] study the use of memoization to accelerate expensive transcendental function calls in CPUs. Compared to TransPimLib, these approaches are not self-sufficient and, as such, they need additional mechanisms for the cases where a value that has not been memoized is needed. As a result, if implemented for PIM systems, these approaches could lead to excessive data movement, as shown in Figure 1(b).

6 CONCLUSION

Processing-in-memory (PIM) is a promising trend to alleviate the data movement bottleneck in current computing systems. PIM is becoming a reality with the advent of real-world PIM architectures, which place simple processing elements near the memory arrays. These architectures support only limited instruction sets, which makes the execution of complex operations challenging. This is the case of transcendental functions and other hard-to-calculate operations (e.g., square root).

In this work, we present TransPimLib, the first library for PIM systems that provides CORDIC-based and LUT-based methods for trigonometric functions, hyperbolic functions, exponentiation, logarithm, square root, etc. We develop an implementation of TransPimLib for the UPMEM PIM architecture and perform a thorough evaluation of TransPimLib's methods in terms of performance and accuracy, using microbenchmarks and three full workloads (Blackscholes, Softmax, Sigmoid).

We believe that TransPimLib methods can be suitable for other current and future PIM architectures lacking native support for these complex functions. The implementation of these methods for other current and future PIM architectures is subject of future work.

ACKNOWLEDGMENTS

We thank the anonymous reviewers of ISPASS 2023 for feedback. We acknowledge the generous gifts provided by our industrial partners, including ASML, Facebook, Google, Huawei, Intel, Microsoft, and VMware. We acknowledge support from the Semiconductor Research Corporation and the ETH Future Computing Laboratory. A longer version of this paper is available in arXiv [64].

REFERENCES

- Idexp, Idexpf, Idexpl. https://en.cppreference.com/w/c/numeric/math/ldexp. Accessed on 2022-11-21.
- [2] Shaizeen Aga, Supreet Jeloka, Arun Subramaniyan, Satish Narayanasamy, David Blaauw, and Reetuparna Das. Compute Caches. In HPCA, 2017.
- [3] Junwhan Ahn, Sungpack Hong, Sungjoo Yoo, Onur Mutlu, and Kiyoung Choi. A Scalable Processing-in-Memory Accelerator for Parallel Graph Processing. In ISCA, 2015.
- [4] Junwhan Ahn, Sungjoo Yoo, Onur Mutlu, and Kiyoung Choi. PIM-Enabled Instructions: A Low-Overhead, Locality-Aware Processing-in-Memory Architecture. In ISCA, 2015.
- [5] Berkin Akin, Franz Franchetti, and James C. Hoe. Data Reorganization in Memory Using 3D-Stacked DRAM. In ISCA, 2015.
- [6] Ethem Alpaydin. Introduction to Machine Learning. 2020.
- [7] Joao Ambrosi, Aayush Ankit, Rodrigo Antunes, Sai Rahul Chalamalasetti, Soumitra Chatterjee, Izzat El Hajj, Guilherme Fachini, Paolo Faraboschi, Martin Foltin, Sitao Huang, et al. Hardware-software Co-design for an Analog-digital Accelerator for Machine Learning. In *ICRC*, 2018.
- [8] S. Angizi, Z. He, and D. Fan. PIMA-Logic: A Novel Processing-in-Memory Architecture for Highly Flexible and Energy-efficient Logic Computation. In DAC, 2018.
- [9] S. Angizi, A. S. Rakin, and D. Fan. CMP-PIM: An Energy-efficient Comparatorbased Processing-in-Memory Neural Network Accelerator. In DAC, 2018.
- [10] S. Angizi, J. Sun, W. Zhang, and D. Fan. AlignS: A Processing-in-Memory Accelerator for DNA Short Read Alignment Leveraging SOT-MRAM. In DAC, 2019.
- [11] Aayush Ankit, Izzat El Hajj, Sai Rahul Chalamalasetti, Sapan Agarwal, Matthew Marinella, Martin Foltin, John Paul Strachan, Dejan Milojicic, Wen-Mei Hwu, and Kaushik Roy. PANTHER: A Programmable Architecture for Neural Network Training Harnessing Energy-efficient ReRAM. IEEE TC, 2020.
- [12] Aayush Ankit, Izzat El Hajj, Sai Rahul Chalamalasetti, Geoffrey Ndu, Martin Foltin, R. Stanley Williams, Paolo Faraboschi, Wen-Mei W Hwu, John Paul Strachan, Kaushik Roy, and Dejan S. Milojicic. PUMA: A Programmable Ultra-Efficient Memristor-Based Accelerator for Machine Learning Inference. In *ASPLOS*, 2019.
- [13] Jason Todd Arbaugh. Table Look-up CORDIC: Effective Rotations Through Angle Partitioning, 2004.
- [14] Hadi Asghari-Moghaddam, Young Hoon Son, Jung Ho Ahn, and Nam Sung Kim. Chameleon: Versatile and Practical Near-DRAM Acceleration Architecture for Large Memory Systems. In *MICRO*, 2016.
- [15] Oreoluwatomiwa O. Babarinsa and Stratos Idreos. JAFAR: Near-Data Processing for Databases. In SIGMOD, 2015.
- [16] Klaus-Jürgen Bathe. Finite Element Method. Wiley Encyclopedia of Computer Science and Engineering, 2007.
- [17] D. Bhattacharjee, R. Devadoss, and A. Chattopadhyay. ReVAMP: ReRAM based VLIW Architecture for In-memory Computing. In DATE, 2017.
- [18] Christian Bienia, Sanjeev Kumar, Jaswinder Pal Singh, and Kai Li. The PARSEC Benchmark Suite: Characterization and Architectural Implications. In PACT, 2008.
- [19] A. Boroumand, S. Ghose, Y. Kim, R. Ausavarungnirun, E. Shiu, R. Thakur, D. Kim, A. Kuusela, A. Knies, P. Ranganathan, and O. Mutlu. Google Workloads for Consumer Devices: Mitigating Data Movement Bottlenecks. In ASPLOS, 2018.
- [20] Amirali Boroumand, Saugata Ghose, Berkin Akin, Ravi Narayanaswami, Geraldo F. Oliveira, Xiaoyu Ma, Eric Shiu, and Onur Mutlu. Google Neural Network Models for Edge Devices: Analyzing and Mitigating Machine Learning Inference Bottlenecks. In PACT, 2021.
- [21] Amirali Boroumand, Saugata Ghose, Youngsok Kim, Rachata Ausavarungnirun, Eric Shiu, Rahul Thakur, Daehyun Kim, Aki Kuusela, Allan Knies, Parthasarathy Ranganathan, et al. Google workloads for consumer devices: mitigating data movement bottlenecks. In ASPLOS, 2018.
- [22] Amirali Boroumand, Saugata Ghose, Minesh Patel, Hasan Hassan, Brandon Lucia, Rachata Ausavarungnirun, Kevin Hsieh, Nastaran Hajinazar, Krishna T. Malladi, Hongzhong Zheng, and Onur Mutlu. CoNDA: Efficient Cache Coherence Support for near-Data Accelerators. In ISCA, 2019.
- [23] Pedro Bruel, Sai Rahul Chalamalasetti, Chris Dalton, Izzat El Hajj, Alfredo Goldman, Catherine Graves, Wen-Mei Hwu, Phil Laplante, Dejan Milojicic, Geoffrey Ndu, et al. Generalize or Die: Operating Systems Support for Memristorbased Accelerators. In *ICRC*, 2017.
- [24] Damla Senol Cali, Gurpreet S Kalsi, Zülal Bingöl, Can Firtina, Lavanya Subramanian, Jeremie S Kim, Rachata Ausavarungnirun, Mohammed Alser, Juan Gomez-Luna, Amirali Boroumand, et al. GenASM: A High-Performance, Low-Power Approximate String Matching Acceleration Framework for Genome

Sequence Analysis. In MICRO, 2020.

- [25] Kevin K. Chang, Prashant J. Nair, Donghyuk Lee, Saugata Ghose, Moinuddin K. Qureshi, and Onur Mutlu. Low-Cost Inter-Linked Subarrays (LISA): Enabling Fast Inter-Subarray Data Movement in DRAM. In *HPCA*, 2016.
- [26] P. Chi, S. Li, C. Xu, T. Zhang, J. Zhao, Y. Liu, Y. Wang, and Y. Xie. PRIME: A Novel Processing-In-Memory Architecture for Neural Network Computation In ReRAM-Based Main Memory. In ISCA, 2016.
- [27] Robert Christy, Stuart Riches, Sujil Kottekkat, Prasanth Gopinath, Ketan Sawant, Anitha Kona, and Rob Harrison. 8.3 A 3GHz ARM Neoverse N1 CPU in 7nm FinFET for Infrastructure Applications. In ISSCC, 2020.
- [28] C. W. Clenshaw. Polynomial Approximations to Elementary Functions. Mathematics of Computation, 1954.
- [29] Hugues de Lassus Saint-Geniès, David Defour, and Guillaume Revy. Exact Lookup Tables for the Evaluation of Trigonometric and Hyperbolic Functions. IEEE TC, 2017.
- [30] F. Devaux. The True Processing In Memory Accelerator. In Hot Chips, 2019.
- [31] Safaa Diab, Amir Nassereldine, Mohammed Alser, Juan Gómez Luna, Onur Mutlu, and Izzat El Hajj. A Framework for High-throughput Sequence Alignment using Real Processing-in-Memory Systems. *Bioinformatics*, 2023.
- [32] Safaa Diab, Amir Nassereldine, Mohammed Alser, Juan Gómez Luna, Onur Mutlu, and Izzat El Hajj. High-throughput Pairwise Alignment with the Wavefront Algorithm using Processing-in-Memory. arXiv preprint arXiv:2204.02085, 2022.
- [33] Charles Eckert, Xiaowei Wang, Jingcheng Wang, Arun Subramaniyan, Ravi Iyer, Dennis Sylvester, David Blaaauw, and Reetuparna Das. Neural Cache: Bit-serial In-cache Acceleration of Deep Neural Networks. In ISCA, 2018.
- [34] European Mathematical Society. Transcendental Function. Encyclopedia of Mathematics. Accessed on 2022-11-21.
- [35] Ivan Fernandez, Ricardo Quislant, Christina Giannoula, Mohammed Alser, Juan Gomez-Luna, Eladio Gutierrez, Oscar Plata, and Onur Mutlu. NATSA: A Near-Data Processing Accelerator for Time Series Analysis. In *ICCD*, 2020.
- [36] International Organization for Standardization (ISO). ISO/IEC 9899: 1999 Programming Languages-C, 1999.
- [37] Daichi Fujiki, Scott Mahlke, and Reetuparna Das. Duality Cache for Data Parallel Acceleration. In ISCA, 2019.
- [38] Bai Fujun, Jiang Xiping, Wang Song, Yu Bing, Tan Jie, Zuo Fengguo, Wang Chunjuan, Wang Fan, Long Xiaodong, Yu Guoqing, et al. A Stacked Embedded DRAM Array for LPDDR4/4X using Hybrid Bonding 3D Integration with 34GB/s/1Gb 0.88 pJ/b Logic-to-Memory Interface. In *IEDM*, 2020.
- [39] P.-E. Gaillardon, L. Amaru, A. Siemon, and et al. The Programmable Logic-in-Memory (PLiM) Computer. In DATE, 2016.
- [40] Fei Gao, Georgios Tziantzioulis, and David Wentzlaff. ComputeDRAM: In-Memory Compute Using Off-the-Shelf DRAMs. In MICRO, 2019.
- [41] Mingyu Gao and Christos Kozyrakis. HRL: Efficient and Flexible Reconfigurable Logic for Near-Data Processing. In HPCA, 2016.
- [42] Christina Giannoula, Ivan Fernandez, Juan Gómez-Luna, Nectarios Koziris, Georgios Goumas, and Onur Mutlu. Towards Efficient Sparse Matrix Vector Multiplication on Real Processing-in-Memory Architectures. In SIGMETRICS, 2022.
- [43] Christina Giannoula, Nandita Vijaykumar, Nikela Papadopoulou, Vasileios Karakostas, Ivan Fernandez, Juan Gómez-Luna, Lois Orosa, Nectarios Koziris, Georgios Goumas, and Onur Mutlu. SynCron: Efficient Synchronization Support for Near-Data-Processing Architectures. In HPCA, 2021.
- [44] Andrew S Glassner. An Introduction to Ray Tracing. 1989.
- [45] David Goldberg. What Every Computer Scientist Should Know About Floating-Point Arithmetic. ACM Comput. Surv., 1991.
- [46] Juan Gómez-Luna, Izzat El Hajj, Ivan Fernandez, Christina Giannoula, Geraldo F. Oliveira, and Onur Mutlu. Benchmarking Memory-Centric Computing Systems: Analysis of Real Processing-In-Memory Hardware. In *IGSC*, 2021.
- [47] Juan Gómez-Luna, Izzat El Hajj, Ivan Fernández, Christina Giannoula, Geraldo F. Oliveira, and Onur Mutlu. Benchmarking a New Paradigm: An Experimental Analysis of a Real Processing-in-Memory Architecture. arXiv:2105.03814 [cs.AR], 2021.
- [48] M6800 Technical Poblications Group. MC68881 Floating Point Coprocessor User Manual. FREESCALE, 6051 William Cannon Drive, Austin Texas, 2 edition, dec 1993.
- [49] Q. Guo, N. Alachiotis, B. Akin, F. Sadi, G. Xu, T. M. Low, L. Pileggi, J. C. Hoe, and F. Franchetti. 3D-Stacked Memory-Side Acceleration: Accelerator and System Design. In *WoNDP*, 2014.
- [50] Juan Gómez-Luna, Izzat El Hajj, Ivan Fernandez, Christina Giannoula, Geraldo F. Oliveira, and Onur Mutlu. Benchmarking a New Paradigm: Experimental Analysis and Characterization of a Real Processing-in-Memory System. *IEEE* Access, 2022.
- [51] Ramyad Hadidi, Lifeng Nai, Hyojong Kim, and Hyesoon Kim. CAIRO: A Compiler-assisted Technique for Enabling Instruction-level Offloading of Processing-in-Memory. ACM TACO, 14, 2017.
- [52] Nastaran Hajinazar, Geraldo F Oliveira, Sven Gregorio, João Dinis Ferreira, Nika Mansouri Ghiasi, Minesh Patel, Mohammed Alser, Saugata Ghose, Juan

Gómez-Luna, and Onur Mutlu. SIMDRAM: A Framework for Bit-Serial SIMD Processing Using DRAM. In *ASPLOS*, 2021.

- [53] S. Hamdioui, S. Kvatinsky, and et al. G. Cauwenberghs. Memristor for Computing: Myth or Reality? In DATE, 2017.
- [54] S. Hamdioui, L. Xie, H. A. D. Nguyen, and et al. Memristor Based Computationin-Memory Architecture for Data-intensive Applications. In DATE, 2015.
- [55] Jun Han and Claudio Moraga. The Influence of the Sigmoid Function Parameters on the Speed of Backpropagation Learning. In *IWANN*, 1995.
- [56] Dan Hendrycks and Kevin Gimpel. Gaussian Error Linear Units (GELUs). arXiv preprint arXiv:1606.08415, 2016.
- [57] David W Hosmer Jr, Stanley Lemeshow, and Rodney X Sturdivant. Applied Logistic Regression. 2013.
- [58] K. Hsieh, S. Khan, N. Vijaykumar, K. K. Chang, A. Boroumand, S. Ghose, and O. Mutlu. Accelerating Pointer Chasing in 3D-Stacked Memory: Challenges, Mechanisms, Evaluation. In *ICCD*, 2016.
- [59] Kevin Hsieh, Eiman Ebrahimi, Gwangsun Kim, Niladrish Chatterjee, Mike O'Conner, Nandita Vijaykumar, Onur Mutlu, and Stephen Keckler. Transparent Offloading and Mapping (TOM): Enabling Programmer-Transparent Near-Data Processing in GPU Systems. In *ISCA*, 2016.
- [60] Sitao Huang, Aayush Ankit, Plinio Silveira, Rodrigo Antunes, Sai Rahul Chalamalasetti, Izzat El Haji, Dong Eun Kim, Glaucimar Aguiar, Pedro Bruel, Sergey Serebryakov, et al. Mixed Precision Quantization for ReRAM-based DNN Inference Accelerators. In ASP-DAC, 2021.
- [61] Hybrid Memory Cube Consortium. HMC Specification 1.1, 2013.
- [62] Hybrid Memory Cube Consortium. HMC Specification 2.0, 2014.
- [63] Intel. Intel 64 and IA-32 Architectures Software Developer's Manual Volume 3B, 2007.
- [64] Maurus Item, Juan Gómez-Luna, Yuxin Guo, Geraldo F. Oliveira, Mohammad Sadrosadati, and Onur Mutlu. TransPimLib: A Library for Efficient Transcendental Functions on Processing-in-Memory Systems. arXiv:2304.01951 [cs.MS], 2023.
- [65] JEDEC. High Bandwidth Memory (HBM) DRAM. Standard No. JESD235, 2013.
- [66] JEDEC. JESD209-4D LPDDR4 SDRAM standard, 2021.
- [67] Hao Jiangwei, Xu Jinchen, Guo Shaozhong, Xia Yuanyuan, and Liu Dan. Design and implementation of variable precision algorithm for transcendental functions. *Journal of Physics: Conference Series*, 1325:012119, 10 2019.
- [68] Mingu Kang, Min-Sun Keel, Naresh R Shanbhag, Sean Eilert, and Ken Curewitz. An Energy-Efficient VLSI Architecture for Pattern Recognition via Deep Embedding of Computation in SRAM. In *ICASSP*, 2014.
- [69] W. H. Kautz. Cellular Logic-in-Memory Arrays. IEEE TC, 1969.
- [70] Liu Ke, Xuan Zhang, Jinin So, Jong-Geon Lee, Shin-Haeng Kang, Sukhan Lee, Songyi Han, Yeongon Cho, Jin Hyun Kim, Yongsuk Kwon, et al. Near-Memory Processing in Action: Accelerating Personalized Recommendation with AxDIMM. *IEEE Micro*, 2021.
- [71] Doris Keitel-Schulz and Norbert Wehn. Embedded DRAM Development: Technology, Physical Design, and Application Issues. *IEEE Design & Test of Computers*, 2001.
- [72] G. Kestor, R. Gioiosa, D. J. Kerbyson, and A. Hoisie. Quantifying the Energy Cost of Data Movement in Scientific Applications. In *IISWC*, 2013.
- [73] G. Kim, N. Chatterjee, M. O'Connor, and K. Hsieh. Toward Standardized Near-Data Processing with Unrestricted Data Placement for GPUs. In SC, 2017.
- [74] J. Kim, M. Patel, H. Hassan, and O. Mutlu. The DRAM Latency PUF: Quickly Evaluating Physical Unclonable Functions by Exploiting the Latency–Reliability Tradeoff in Modern DRAM Devices. In HPCA, 2018.
- [75] J. Kim, M. Patel, H. Hassan, L. Orosa, and O. Mutlu. D-RaNGe: Using Commodity DRAM Devices to Generate True Random Numbers with Low Latency and High Throughput. In *HPCA*, 2019.
- [76] J. S. Kim, D. Senol, H. Xin, D. Lee, S. Ghose, M. Alser, H. Hassan, O. Ergin, C. Alkan, and O. Mutlu. GRIM-Filter: Fast Seed Location Filtering in DNA Read Mapping Using Processing-in-Memory Technologies. *BMC Genomics*, 2018.
- [77] Seongguk Kim, Subin Kim, Kyungjun Cho, Taein Shin, Hyunwook Park, Daehwan Lho, Shinyoung Park, Kyungjune Son, Gapyeol Park, Seungtaek Jeong, et al. Signal Integrity and Computing Performance Analysis of a Processingin-Memory of High Bandwidth Memory (PIM-HBM) Scheme. *IEEE T-CPMT*, 2021.
- [78] Ytong-Bin Kim and Tom W Chen. Assessing Merged DRAM/Logic Technology. Integration, 1999.
- [79] K. Kota and J.R. Cavallaro. Numerical Accuracy and Hardware Tradeoffs for CORDIC Arithmetic for Special-purpose Processors. IEEE TC, 1993.
 - [80] Adolf Kratzer and Walter Franz. Transzendente Funktionen. 1960.
 - [81] S. Kvatinsky, D. Belousov, S. Liman, G. Satat, N. Wald, E. G. Friedman, A. Kolodny, and U. C. Weiser. MAGIC–Memristor-Aided Logic. *IEEE TCAS II: Express Briefs*, 2014.
 - [82] S. Kvatinsky, A. Kolodny, U. C. Weiser, and E. G. Friedman. Memristor-Based IMPLY Logic Design Procedure. In *ICCD*, 2011.
 - [83] S. Kvatinsky, G. Satat, N. Wald, E. G. Friedman, A. Kolodny, and U. C. Weiser. Memristor-Based Material Implication (IMPLY) Logic: Design Principles and Methodologies. *TVLSI*, 2014.

- [84] Young-Cheon Kwon, Suk Han Lee, Jaehoon Lee, Sang-Hyuk Kwon, Je Min Ryu, Jong-Pil Son, O Seongil, Hak-Soo Yu, Haesuk Lee, Soo Young Kim, et al. 25.4 A 20nm 6GB Function-In-Memory DRAM, Based on HBM2 with a 1.2 TFLOPS Programmable Computing Unit Using Bank-Level Parallelism, for Machine Learning Applications. In ISSCC, 2021.
- [85] Donghun Lee, Jinin So, MINSEON AHN, Jong-Geon Lee, Jungmin Kim, Jeonghyeon Cho, Rebholz Oliver, Vishnu Charan Thummala, Ravi shankar JV, Sachin Suresh Upadhya, et al. Improving In-Memory Database Operations with Acceleration DIMM (AxDIMM). In *DaMoN*, 2022.
- [86] Donghyuk Lee, Saugata Ghose, Gennady Pekhimenko, Samira Khan, and Onur Mutlu. Simultaneous Multi-Layer Access: Improving 3D-Stacked Memory Bandwidth at Low Cost. TACO, 2016.
- [87] Seongju Lee, Kyuyoung Kim, Sanghoon Oh, Joonhong Park, Gimoon Hong, Dongyoon Ka, Kyudong Hwang, Jeongje Park, Kyeongpil Kang, Jungyeon Kim, Junyeol Jeon, Nahsung Kim, Yongkee Kwon, Kornijcuk Vladimir, Woojae Shin, Jongsoon Won, Minkyu Lee, Hyunha Joo, Haerang Choi, Jaewook Lee, Donguc Ko, Younggun Jun, Keewon Cho, Ilwoong Kim, Choungki Song, Chunseok Jeong, Daehan Kwon, Jieun Jang, Il Park, Junhyun Chun, and Joohwan Cho. A Iynm 1.25V 8Gb, 16Gb/s/pin GDDR6-based Accelerator-in-Memory supporting 1TFLOPS MAC Operation and Various Activation Functions for Deep-Learning Applications. In ISSCC, 2022.
- [88] Sukhan Lee, Shin-haeng Kang, Jaehoon Lee, Hyeonsu Kim, Eojin Lee, Seungwoo Seo, Hosang Yoon, Seungwon Lee, Kyounghwan Lim, Hyunsung Shin, et al. Hardware Architecture and Software Stack for PIM Based on Commercial DRAM Technology: Industrial Product. In ISCA, 2021.
- [89] Yifat Levy, Jehoshua Bruck, Yuval Cassuto, Eby G. Friedman, Avinoam Kolodny, Eitan Yaakobi, and Shahar Kvatinsky. Logic Operations in Memory Using a Memristive Akers Array. *Microelectronics Journal*, 2014.
- [90] S. Li, D. Niu, K. T. Malladi, H. Zheng, B. Brennan, and Y. Xie. DRISA: A DRAM-Based Reconfigurable In-Situ Accelerator. In *MICRO*, 2017.
- [91] Shuangchen Li, Cong Xu, Qiaosha Zou, Jishen Zhao, Yu Lu, and Yuan Xie. Pinatubo: A Processing-in-Memory Architecture for Bulk Bitwise Operations in Emerging Non-Volatile Memories. In DAC, 2016.
- [92] James D MacBeth and Larry J Merville. An Empirical examination of the Black-Scholes Call Option Pricing Model. *The Journal of Finance*, 1979.
- [93] Santiago Marco-Sola, Juan Carlos Moure, Miquel Moreto, and Antonio Espinosa. Fast Gap-affine Pairwise Alignment using the Wavefront Algorithm. *Bioinfor-matics*, 2021.
- [94] Vasilios Mavroudis. Computing Small Discrete Logarithms using Optimized Lookup Tables. http://koclab.cs.ucsb.edu/teaching/ecc/project/2015Projects/ Mavroudis.pdf, 2015.
- [95] Pramod K. Meher, Javier Valls, Tso-Bing Juang, K. Sridharan, and Koushik Maharatna. 50 Years of CORDIC: Algorithms, Architectures, and Applications. IEEE TCAS-I: Regular Papers, 2009.
- [96] O. Mutlu et al. Processing Data Where It Makes Sense: Enabling In-Memory Computation. MicPro, 2019.
- [97] Onur Mutlu. Intelligent Architectures for Intelligent Computing Systems. In DATE, 2021.
- [98] Onur Mutlu et al. A Modern Primer on Processing in Memory. Emerging Computing: From Devices to Systems - Looking Beyond Moore and Von Neumann, 2023. https://arxiv.org/pdf/2012.03112.pdf.
- [99] Lifeng Nai, Ramyad Hadidi, Jaewoong Sim, Hyojong Kim, Pranith Kumar, and Hyesoon Kim. GraphPIM: Enabling Instruction-Level PIM Offloading in Graph Computing Frameworks. In *HPCA*, 2017.
- [100] Maxim Naumov, Dheevatsa Mudigere, Hao-Jun Michael Shi, Jianyu Huang, Narayanan Sundaraman, Jongsoo Park, Xiaodong Wang, Udit Gupta, Carole-Jean Wu, Alisson G Azzolini, et al. Deep Learning Recommendation Model for Personalization and Recommendation Systems. arXiv preprint arXiv:1906.00091, 2019.
- [101] Dimin Niu, Shuangchen Li, Yuhao Wang, Wei Han, Zhe Zhang, Yijin Guan, Tianchan Guan, Fei Sun, Fei Xue, Lide Duan, et al. 184QPS/W 64Mb/mm2 3D Logic-to-DRAM Hybrid Bonding with Process-Near-Memory Engine for Recommendation System. In *ISSCC*, 2022.
- [102] NVIDIA. CUDA C++ programming guide, 2022.
- [103] Ataberk Olgun, Minesh Patel, Abdullah Giray Yağlıkçı, Haocong Luo, Jeremie S. Kim, F. Nisa Bostancı, Nandita Vijaykumar, Oğuz Ergin, and Onur Mutlu. QUAC-TRNG: High-Throughput True Random Number Generation Using Quadruple Row Activation in Commodity DRAMs. In ISCA, 2021.
- [104] Geraldo F. Oliveira, Juan Gómez-Luna, Lois Orosa, Saugata Ghose, Nandita Vijaykumar, Ivan Fernandez, Mohammad Sadrosadati, and Onur Mutlu. DAMOV: A New Methodology and Benchmark Suite for Evaluating Data Movement Bottlenecks. IEEE Access, 2021.
- [105] Geraldo F. Oliveira, Juan Gómez-Luna, Lois Orosa, Saugata Ghose, Nandita Vijaykumar, Ivan Fernandez, Mohammad Sadrosadati, and Onur Mutlu. DAMOV: A New Methodology and Benchmark Suite for Evaluating Data Movement Bottlenecks. arXiv:2105.03725 [cs.AR], 2021.
- [106] Dhinakaran Pandiyan and Carole-Jean Wu. Quantifying the energy cost of data movement for emerging smart phone workloads on mobile platforms. In IISWC,

2014.

- [107] Ashutosh Pattnaik, Xulong Tang, Adwait Jog, Onur Kayiran, Asit K. Mishra, Mahmut T. Kandemir, Onur Mutlu, and Chita R. Das. Scheduling Techniques for GPU Architectures with Processing-in-Memory Capabilities. In PACT, 2016.
- [108] Yarui Peng, Bon Woong Ku, Younsik Park, Kwang-Il Park, Seong-Jin Jang, Joo Sun Choi, and Sung Kyu Lim. Design, Packaging, and Architectural Policy Co-optimization for DC Power Integrity in 3D DRAM. In DAC, 2015.
- [109] Pedro Pinto and João M. P. Cardoso. A Methodology and Framework for Software Memoization of Functions. In CF, 2021.
- [110] SAFARI Research Group. TransPimLib: A Library of Transcendental Functions for PIM Systems. https://github.com/CMU-SAFARI/transpimlib.
- [111] Antonio Salazar, G. Bahubalindruno, and Govinda Locharla. A Study on Lookup Table Based Sine Wave Generation. In Proceedings of the Regional Echomail Coordinator, 2011.
- [112] Shadrokh Samavi and Mohammad Reza Jahangir. Reduction of Look Up Tables for Computation of Reciprocal of Square Roots. arXiv preprint arXiv:1710.04688, 2017.
- [113] M.J. Schulte and J.E. Stine. Accurate Function Approximations by Symmetric Table Lookup and Addition. In ASAP, 1997.
- [114] V. Seshadri, D. Lee, T. Mullins, H. Hassan, A. Boroumand, J. Kim, M. A. Kozuch, O. Mutlu, P. B. Gibbons, and T. C. Mowry. Ambit: In-Memory Accelerator for Bulk Bitwise Operations Using Commodity DRAM Technology. In *MICRO*, 2017.
- [115] Vivek Seshadri, Yoongu Kim, Chris Fallin, Donghyuk Lee, Rachata Ausavarungnirun, Gennady Pekhimenko, Yixin Luo, Onur Mutlu, Michael A Kozuch, Phillip B Gibbons, and Todd C. Mowry. RowClone: Fast and Energy-Efficient In-DRAM Bulk Data Copy and Initialization. In *MICRO*, 2013.
- [116] Vivek Seshadri and Onur Mutlu. In-DRAM Bulk Bitwise Execution Engine. arXiv:1905.09822 [cs.AR], 2020.
- [117] A. Shafiee, A. Nag, N. Muralimanohar, and et al. ISAAC: A Convolutional Neural Network Accelerator with In-situ Analog Arithmetic in Crossbars. In ISCA, 2016.
- [118] Gagandeep Singh, Dionysios Diamantopoulos, Christoph Hagleitner, Juan Gomez-Luna, Sander Stuijk, Onur Mutlu, and Henk Corporaal. NERO: A Near High-Bandwidth Memory Stencil Accelerator for Weather Prediction Modeling. In FPL, 2020.
- [119] Gagandeep Singh, Juan Gomez-Luna, Giovanni Mariani, Geraldo F. Oliveira, Stefano Corda, Sander Stujik, Onur Mutlu, and Henk Corporaal. NAPEL: Nearmemory Computing Application Performance Prediction via Ensemble Learning. In DAC, 2019.
- [120] Teja Singh, Sundar Rangarajan, Deepesh John, Carson Henrion, Shane Southard, Hugh McIntyre, Amy Novak, Stephen Kosonocky, Ravi Jotwani, Alex Schaefer, Edward Chang, Joshua Bell, and Michael Co. 3.2 Zen: A Next-generation Highperformance x86 Core. In *ISSCC*, 2017.
- [121] Burton J Smith. A Pipelined, Shared Resource MIMD Computer. In *ICPP*, 1978.[122] Burton J Smith. Architecture and Applications of the HEP Multiprocessor
- Computer System. In *SPIE, Real-Time signal processing IV*, 1981. [123] Standard Performance Evaluation Corp. SPEC CPU2006 Benchmarks. http:
- //www.spec.org/cpu2006/.
 [124] Marvin E. Stick. Maclaurin and taylor series for transcendental functions: A graphing-calculator view of convergence. *The Mathematics Teacher*, 92(9):833– 837, 1999.
- [125] Harold S Stone. A Logic-in-Memory Computer. IEEE TC, 1970.
- [126] Arjun Suresh. Intercepting Functions for Memoization. PhD thesis, Université Rennes 1, 2016.
- [127] Arjun Suresh, Erven Rohou, and André Seznec. Compile-Time Function Memoization. In CC, 2017.
- [128] Arjun Suresh, Bharath Narasimha Swamy, Erven Rohou, and André Seznec. Intercepting Functions for Memoization: A Case Study Using Transcendental Functions. ACM TACO, 2015.
- [129] Ping Tang. Table-lookup Algorithms for Elementary Functions and their Error Analysis. In ARITH, 1991.
- [130] J. E. Thornton. CDC 6600: Design of a Computer. 1970.
- [131] UPMEM. Introduction to UPMEM PIM. Processing-in-memory (PIM) on DRAM Accelerator (White Paper), 2018.
- [132] UPMEM. Intel Server Configuration with 20 UPMEM PIM DIMM Modules. https://www.upmem.com/technology/, 2023.
- [133] UPMEM. UPMEM Software Development Kit (SDK). https://sdk.upmem.com, 2023.
- [134] UPMEM. UPMEM User Manual. Version 2023.1.0, 2023.
- [135] UPMEM. UPMEM Website. https://www.upmem.com, 2023
- [136] Jack E. Volder. The CORDIC Trigonometric Computing Technique. IRE Transactions on Electronic Computers, 1959.
- [137] Yaohua Wang, Lois Orosa, Xiangjun Peng, Yang Guo, Saugata Ghose, Minesh Patel, Jeremie S Kim, Juan Gómez Luna, Mohammad Sadrosadati, Nika Mansouri Ghiasi, et al. FIGARO: Improving System Performance via Fine-Grained In-DRAM Data Relocation and Caching. In MICRO, 2020.

- [138] Detlef Weber, Andreas Thies, U Kahler, M Lepper, and R Schutz. Current and Future Challenges of DRAM Metallization. In *IITC*, 2005.
- [139] Chris Wilcox. A Methodology for Automated Lookup Table Optimization of Scientific Applications. PhD thesis, 2012.
- [140] Steven Cameron Woo, Moriyoshi Ohara, Evan Torrie, Jaswinder Pal Singh, and Anoop Gupta. The SPLASH-2 Programs: Characterization and Methodological Considerations. In ISCA, 1995.
- [141] Yue Xi, Bin Gao, Jianshi Tang, An Chen, Meng-Fan Chang, Xiaobo Sharon Hu, Jan Van Der Spiegel, He Qian, and Huaqiang Wu. In-Memory Learning With Analog Resistive Switching Memory: A Review and Perspective. *Proceedings of the IEEE*, 2020.
- [142] L. Xie, H. A. D. Nguyen, M. Taouil, and et al. Fast Boolean Logic Papped on Memristor Crossbar. In *ICCD*, 2015.
- [143] J. Yu, H. A. D. Nguyen, L. Xie, and et al. Memristive Devices for Computationin-Memory. In DATE, 2018.

- [144] Marcelo Yuffe, Ernest Knoll, Moty Mehalel, Joseph Shor, and Tsvika Kurts. A Fully Integrated Multi-CPU, GPU and Memory Controller 32nm processor. In ISSCC, 2011.
- [145] D. P. Zhang, N. Jayasena, A. Lyashevsky, J. L. Greathouse, L. Xu, and M. Ignatowski. TOP-PIM: Throughput-Oriented Programmable Processing in Memory. In *HPDC*, 2014.
- [146] Le Zheng, Sangho Shin, Scott Lloyd, Maya Gokhale, Kyungmin Kim, and Sung-Mo Kang. RRAM-based TCAMs for pattern search. In *ISCAS*, 2016.
- [147] Jie Zhou, Yong Dou, Yuanwu Lei, Jinbo Xu, and Yazhuo Dong. Double Precision Hybrid-Mode Floating-Point FPGA CORDIC Co-processor. In HPCC, 2008.
- [148] Qiuling Zhu, Tobias Graf, H Ekin Sumbul, Larry Pileggi, and Franz Franchetti. Accelerating Sparse Matrix-Matrix Multiplication with 3D-Stacked Logic-in-Memory Hardware. In HPEC, 2013.