

Pathfinding Future PIM Architectures by Demystifying a Commercial PIM Technology

Bongjoon Hyun Taehun Kim Dongjae Lee Minsoo Rhu
KAIST

{bongjoon.hyun, taehun.kim, dongjae.lee, mrhu}@kaist.ac.kr

Abstract—Processing-in-memory (PIM) has been explored for decades by computer architects, yet it has never seen the light of day in real-world products due to its high design overheads and lack of a killer application. With the advent of critical memory-intensive workloads, several commercial PIM technologies have been introduced to the market, ranging from domain-specific PIM architectures to more general-purpose PIM architectures. In this work, we deepdive into UPMEM’s commercial PIM technology, a general-purpose PIM-enabled parallel computing architecture that is highly programmable. Our first key contribution is the development of a flexible simulation framework for PIM. The simulator we developed (aka uPIMulator) enables the compilation of UPMEM-PIM source codes into its compiled machine-level instructions, which are subsequently consumed by our cycle-level performance simulator. Using uPIMulator, we demystify UPMEM’s PIM design through a detailed characterization study. Finally, we identify some key limitations of the current UPMEM-PIM system through our case studies and present some important architectural features that will become critical for future PIM architectures to support.

I. INTRODUCTION

“We’ve investigated applying PIM to our workloads and determined there are several challenges to using these approaches. Perhaps the biggest challenge of PIM is its programmability. It is hard to anticipate future model compression methods, so programmability is required to adapt to these. PIM must also support flexible parallelization since it is hard to predict how much each dimension (of embedding tables) will scale in the future.”

“First-Generation Inference Accelerator Deployment at Facebook”, Facebook, 2021 [1]

Emerging workloads in the areas of scientific computing, graph processing, and machine learning pose unprecedented demand for larger data. However, the well-known memory *bandwidth* wall causes a critical performance bottleneck for these memory-bound workloads, due to the widening performance gap between processor and memory. Processing-in-memory (PIM) architectures have been explored extensively for decades [2], [3], [4], [5], as they help alleviate the memory bandwidth bottleneck by moving compute logic closer to memory. Unfortunately, the computing industry has so far been lukewarm in commercializing PIM architectures, primarily because of their high design overheads (e.g., regression in DRAM density, thermal issues [6]) and their intrusiveness to the software stack (e.g., programming model [7], [8], [9], [10],

managing address space and data coherence [11], [12], [13]), rendering PIM mostly an academic pursuit.

Interestingly, such sentiment towards PIM has changed dramatically over the past couple of years with several commercial PIM systems introduced to the market. These PIM designs can broadly be grouped into two categories: 1) domain-specific PIM and 2) general-purpose PIM. Domain-specific PIM designs have been driven by key memory vendors like Samsung [14], [15], [16], [17] and SK Hynix [18], [19], which focus on specializing their PIM design by supporting key compute primitives for a targeted application domain (e.g., matrix multiplication for accelerating deep neural networks), reigniting people’s interest in PIM designs [14], [17], [18], [20], [21], [22], [23], [24], [25], [26], [27], [28], [29], [30], [31], [32], [33], [34], [35], [36], [37], [38], [39], [40], [41], [42], [43], [44]. At the other end of the spectrum, the PIM solution offered by UPMEM [45] (henceforth referred to as UPMEM-PIM) takes a different approach by providing a *general-purpose* parallel programming language with an LLVM-based compiler stack [46], [47], providing application developers the flexibility to write any parallel program to be executed using PIM. Thanks to its high programmability and flexibility, several recent work studied the applicability of UPMEM-PIM for accelerating a variety of application domains, e.g., graph algorithms, machine learning, bioinformatics, etc. [48], [49], [50], [51], [52]. Similar to how GPUs have transformed themselves into a first-class computing citizen after years of hardware/software refinement, we believe that it is possible for such general-purpose PIM design to similarly evolve into an important computing device (or at a minimum provide valuable insights in designing future general-purpose PIM) as its hardware/software stack matures.

Given this landscape, our key motivation is to demystify industry’s first general-purpose PIM design through a detailed characterization study, understanding the unique properties of UPMEM-PIM and identifying important research domains that computer architects can explore. To this end, we first develop an UPMEM-PIM ISA compatible simulation framework that utilizes UPMEM’s open-source compiler stack to compile *any* UPMEM-PIM program, from its C-level source code down to its machine level instructions. The compiled UPMEM-PIM binary is then consumed by our cycle-level hardware performance simulator, which we carefully cross-validate against a real UPMEM-PIM system (Section III). Simulators are, by design, immensely flexible and customizable, so they

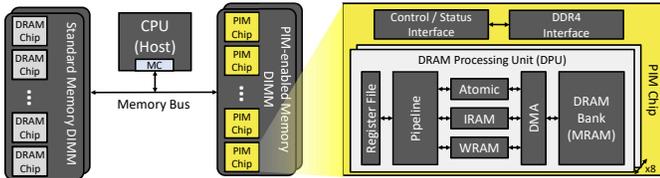


Fig. 1: UPMEM-PIM hardware system overview.

enable us to understand the fine-grained details of the runtime execution of a *real* (UPMEM) PIM program. Using our PIM simulator (henceforth referred to as UPMEM-PIM simulator, aka uPIMulator), we conduct a workload characterization study and provide a number of interesting insights that cannot be easily uncovered using UPMEM-PIM chip’s hardware performance counters or profiling tools (Section IV). Finally, we uncover some critical limitations of the current UPMEM-PIM system through our case studies and propose several key architectural features required for PIM to become more performant, robust, and secure (Section V). These features include the need for vector processing and ILP (instruction-level parallelism) enhancing microarchitectures, architectural support for multi-tenant execution, and the support for on-demand caching rather than solely relying on scratchpads. Overall, we expect our in-depth exploration of UPMEM-PIM using our uPIMulator to open up important research directions for computer system designers¹, paving the way for PIM to evolve into a truly general-purpose computing device.

II. UPMEM-PIM ARCHITECTURE

A. Hardware Architecture

System overview. Figure 1 provides a high-level overview of an UPMEM-PIM based system containing a host-side CPU communicating with a group of standard regular DIMMs and another group of PIM-enabled memory DIMMs (UPMEM-PIM modules). An UPMEM-PIM module is based on a standard DDR4-2400 [53] DIMM form factor containing 8 UPMEM-PIM DRAM chips per each rank. Within each UPMEM-PIM DRAM chip, there are 8 DPUs (DRAM Processing Units), one DPU per each DRAM bank. Each DPU has direct access to a dedicated 64 MB DRAM bank (referred to as *Main RAM*, aka MRAM), a 64 KB SRAM-based scratchpad memory (aka *Working RAM*, WRAM), and 24 KB instruction memory (aka *Instruction RAM*, IRAM). Before an UPMEM-PIM program (i.e., the PIM kernel) is to be executed, the host CPU must explicitly offload both the PIM kernel and the input data from CPU’s conventional memory address space (i.e., regular DIMMs) to DPU’s UPMEM-PIM address space. The real PIM system we study in this work contains 20 double-ranked UPMEM-PIM DIMMs, so a total of $(20 \times 2 \times 8 \times 8) = 2,560$ DPUs concurrently execute as co-processors to the CPU.

DPU architecture. The DPU is designed as an in-order 14-stage pipelined processor with a RISC-based ISA, implementing fine-grained multi-threading. A total of 24 threads

¹uPIMulator is open-sourced at <https://github.com/VIA-Research/uPIMulator>.

```

1 #define VECTOR_SIZE 1024
2 #define NUM_DPUS 64
3 #define DPU_BINARY "../dpu/VA"
4
5 int main() {
6     struct dpu_set_t dpu_set, dpu;
7
8     dpu_alloc(NUM_DPUS, NULL, &dpu_set);
9     dpu_load(dpu_set, DPU_BINARY, NULL);
10
11     int *A = malloc(VECTOR_SIZE * sizeof(int));
12     int *B = malloc(VECTOR_SIZE * sizeof(int));
13     int *C = malloc(VECTOR_SIZE * sizeof(int));
14
15     int size_per_dpu = VECTOR_SIZE / NUM_DPUS;
16     int i;
17
18     DPU_FOREACH(dpu_set, dpu, i) { dpu_prepare_xfer(dpu, &size_per_dpu); }
19     dpu_push_xfer(dpu_set, DPU_XFER_TO_DPU, "size_per_dpu", 0, sizeof(int),
20                 DPU_XFER_DEFAULT);
21
22     DPU_FOREACH(dpu_set, dpu, i) { dpu_prepare_xfer(dpu, A + size_per_dpu * i); }
23     dpu_push_xfer(dpu_set, DPU_XFER_TO_DPU, DPU_MRAM_HEAP_POINTER_NAME, 0,
24                 size_per_dpu * sizeof(int), DPU_XFER_DEFAULT);
25
26     DPU_FOREACH(dpu_set, dpu, i) { dpu_prepare_xfer(dpu, B + size_per_dpu * i); }
27     dpu_push_xfer(dpu_set, DPU_XFER_TO_DPU, DPU_MRAM_HEAP_POINTER_NAME,
28                 size_per_dpu * sizeof(int), size_per_dpu * sizeof(int),
29                 DPU_XFER_DEFAULT);
30
31     dpu_launch(dpu_set, DPU_SYNCHRONOUS);
32
33     DPU_FOREACH(dpu_set, dpu, i) { dpu_prepare_xfer(dpu, C + size_per_dpu * i); }
34     dpu_push_xfer(dpu_set, DPU_XFER_FROM_DPU, DPU_MRAM_HEAP_POINTER_NAME,
35                 2 * size_per_dpu * sizeof(int), size_per_dpu * sizeof(int),
36                 DPU_XFER_DEFAULT);
37
38     return 0;
39 }

```

(a) Host-side code.

```

1 #define NUM_TASKLETS 16
2
3 __host int size_per_dpu;
4
5 BARRIER_INIT(my_barrier, NR_TASKLETS);
6
7 void vector_addition(int *A, int *B, int *C, int size_per_tasklet) {
8     for (int i = 0; i < size_per_tasklet; i++) {
9         C[i] = A[i] + B[i];
10    }
11 }
12
13 int main() {
14     int tasklet_id = me();
15     if (tasklet_id == 0) {
16         mem_reset();
17     }
18     barrier_wait(&my_barrier);
19
20     int size_per_tasklet = size_per_dpu / NUM_TASKLETS;
21
22     int *A_mram = (int *)
23     (DPU_MRAM_HEAP_POINTER +
24      (tasklet_id * size_per_tasklet) * sizeof(int));
25
26     int *B_mram = (int *)
27     (DPU_MRAM_HEAP_POINTER +
28      (size_per_dpu + tasklet_id * size_per_tasklet) * sizeof(int));
29
30     int *C_mram = (int *)
31     (DPU_MRAM_HEAP_POINTER +
32      (2 * size_per_dpu + tasklet_id * size_per_tasklet) * sizeof(int));
33
34     int *A_wram = (int *)mem_alloc(size_per_tasklet * sizeof(int));
35     int *B_wram = (int *)mem_alloc(size_per_tasklet * sizeof(int));
36     int *C_wram = (int *)mem_alloc(size_per_tasklet * sizeof(int));
37
38     mram_read((__mram_ptr void *)A_mram, A_wram, size_per_tasklet * sizeof(int));
39     mram_read((__mram_ptr void *)B_mram, B_wram, size_per_tasklet * sizeof(int));
40
41     vector_addition(A_wram, B_wram, C_wram, size_per_tasklet);
42
43     mram_write(C_wram, (__mram_ptr void *)C_mram, size_per_tasklet * sizeof(int));
44
45     return 0;
46 }

```

(b) DPU-side code.

Fig. 2: An element-wise vector addition program written for UPMEM-PIM: (a) host-side and (b) DPU-side program.

(called *tasklets* by UPMEM) can concurrently execute within a DPU, all of which share the scratchpad (WRAM), instruction memory (IRAM), and per-bank DRAM (MRAM). The UPMEM DPU has a peculiar thread scheduling rule where two consecutive instructions within the same thread must be dispatched 11 cycles apart (UPMEM refers to such

microarchitecture as the *revolver pipeline* [54]). UPMEM states that such scheduling constraint is enforced to obviate the need to implement complicated circuitry for data forwarding and pipeline interlocks [55]. Another unique aspect of the DPU microarchitecture is in its register file (RF) design: the RF is split into an even and odd RF and a thread cannot access multiple even or odd registers at the same cycle (e.g., r0 and r2 cannot be accessed at the same cycle) due to a structural hazard (i.e., RF conflict).

B. Programming Model

UPMEM-PIM follows the single-program multiple-data (SPMD) programming paradigm. A single program written by the programmer gets executed by all the software threads (i.e., tasklets) that are instantiated, but each individual thread can take its own control flow and access different parts of the data using its thread ID (tasklet ID). Since there can be up to 2,560 DPUs and 24 threads per DPU, the programmer must carefully partition the input data, not only across the DPUs (Figure 2(a), line 18-20, 22-24, and 26-29) but also across the threads within each DPU (Figure 2(b), line 22-29). We use Figure 2 as a running example to highlight some of the important programming semantics of UPMEM-PIM.

Host-side programming. Any program that is written in UPMEM’s C-like programming language can be compiled into its machine code by using the LLVM-based compiler toolchain [46] developed by UPMEM [47]. Similar to NVIDIA’s CUDA [56], UPMEM-PIM follows the *co-processor* computing model where the CPU *offloads* the memory-intensive task to the DPU, functioning as an arbiter of the PIM program’s execution. Consequently, the UPMEM compiler generates two binaries, one that runs on the host and the other that runs across all the DPUs. In the host-side code (Figure 2(a)), the programmer must (1) allocate the desired number of DPUs to be used (line 8), (2) offload the program binary to all the DPUs (line 9), (3) partition and send input data to the DPU’s scratchpad (line 18-20) and per-bank DRAM (line 22-24, and 26-29), (4) ask the host to send commands to the DPUs to execute the PIM program (line 31), and (5) once the PIM program terminates, retrieve back the results from DPU memory back to the host CPU’s memory address space (line 33-36).

DPU-side programming. A unique aspect of UPMEM-PIM’s programming model is that all the PIM kernel’s working set *must* be staged through DPU’s scratchpad using DMA instructions. Consider the code snippet in Figure 2(b). Any thread executing within the DPU can only load (store) data from (to) the scratchpad (WRAM) address space but it is not able to address data in the per-bank DRAM (MRAM) address space directly (line 7, 42). In effect, DPUs operate over *two* distinct memory address spaces, the slower but larger per-bank DRAM region and the faster yet smaller scratchpad region. Only when the programmer explicitly requests data movements from the per-bank DRAM region to the scratchpad region (using DMA instructions via `mram_read()`, line 39-40) can the DPU threads access the necessary data from

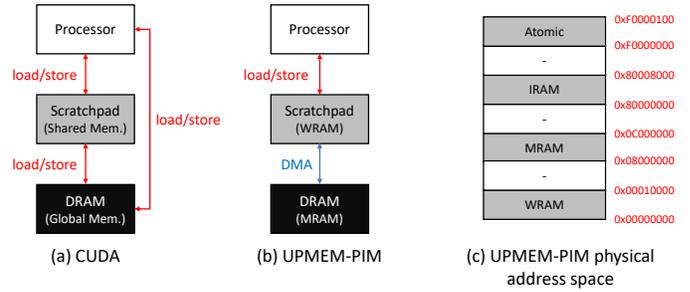


Fig. 3: Memory model of (a) CUDA and (b) UPMEM-PIM. (c) The (physical) address map of UPMEM-PIM.

the scratchpad using load/store instructions (line 9, notice the pointers to the arrays A, B, C are dynamically allocated at the scratchpad WRAM via `mem_alloc()` calls in line 35-37). This is similar to NVIDIA’s CUDA programming model [56] where the programmer must explicitly orchestrate data movements across the CPU memory and the GPU memory using `cudaMemcpy()` (unless the programmer employs Unified Memory [57]). CUDA, however, does allow threads to directly load (store) from (to) *both* its scratchpad and its DRAM, unlike UPMEM’s memory model which only allows load/store semantics over the scratchpad (Figure 3(a,b)). In the remainder of this paper, we refer to such a model as UPMEM’s *scratchpad-centric* programming model.

Data sharing and synchronization. Threads executing within the same DPU can share data over the DPU scratchpad or its local DRAM bank (MRAM). They can also synchronize with each other by using mutexes, barriers, or semaphores allocated in UPMEM-PIM’s atomic memory region (Figure 3(c)), all of which are supported by UPMEM’s SDK [58].

However, threads executing in different DPUs cannot share data or synchronize with each other directly. This is because 1) all the DPUs’ per-bank DRAM across the UPMEM-PIM DIMM are not virtualized within a single, *shared* memory address space (further discussed in Section II-C) 2) nor is there a direct communication datapath among them. If data sharing or synchronization across different DPUs is in need, the programmer must first explicitly copy back the shared data from the producer DPU’s memory to the CPU memory after kernel terminates. The CPU can then copy back this shared data from its CPU memory region to the consumer DPU’s memory region during the next PIM kernel execution.

C. System Software for Memory Management

UPMEM-PIM does not have a memory management unit (MMU) to virtualize its physical memory, so the DPU uses *physical* addresses when accessing WRAM, IRAM, and MRAM, as illustrated in Figure 3(c). In other words, when moving data across UPMEM-PIM’s memory hierarchy using (1) load/store instructions (for scratchpad \leftrightarrow RF) or (2) DMA instructions (for DRAM \leftrightarrow scratchpad), the memory addresses generated by executing an instruction are used *as-is*, without any address translation process involved (Figure 3(b)). Consider the example in Figure 2. When the input array B is being copied from the CPU to DPU’s per-bank DRAM

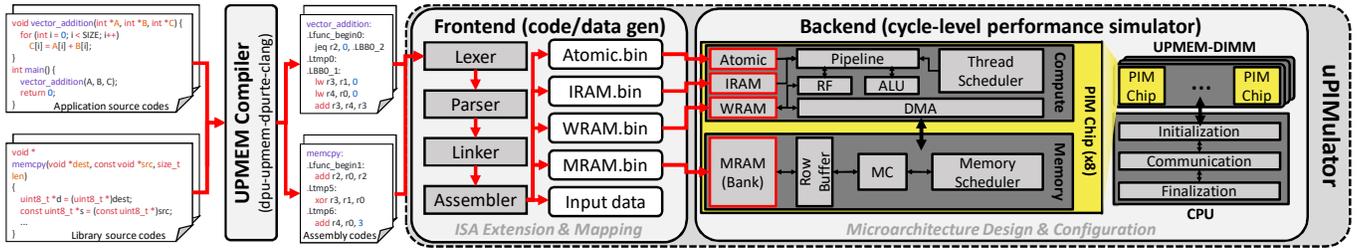


Fig. 4: uPIMulator simulation framework overview.

(Figure 2(a), line 26-29) and then from DPU’s DRAM to DPU’s scratchpad (Figure 2(b), line 40), the programmer must carefully program the pointer value to use as the destination (for CPU→DPU’s DRAM) and source address (for DPU’s DRAM→DPU’s scratchpad) within per-bank DRAM (MRAM) by using `DPU_MRAM_HEAP_POINTER_NAME` (Figure 2(a), line 27-29) or `DPU_MRAM_HEAP_POINTER` (Figure 2(b), line 26-29) as the base physical address.

Overall, the lack of a virtual memory support leaves the programmer with the burden of reasoning about where the input (output) data should be copied over to (from) within DPU’s DRAM, hurting user productivity. Section V-C further discusses the architectural implication of an MMU-less PIM.

III. UPIMULATOR SIMULATION FRAMEWORK

Figure 4 provides an overview of uPIMulator, which consists of two key components: (1) a compiler toolchain that supports execution-driven simulation of UPMEM ISA-compatible, machine-level instructions, and (2) a hardware performance simulator cross-validated against a real UPMEM-PIM. Together, these dual components reduce the effort required to model UPMEM’s general-purpose PIM architecture with high accuracy, enabling architectural exploration of any PIM program written with UPMEM’s programming model.

A. Simulator Development

Software compilation toolchain. The open-source UPMEM SDK [58] provides an LLVM [46]-based compiler toolchain [47] (`dpu-upmem-dpurte-clang`) that takes in (1) the programmer-written source codes and (2) `glibc`-style, UPMEM-PIM compatible C library (e.g., `mem_alloc()` for `malloc` in DPU scratchpad, `memcpy()`, `printf()`) to preprocess, compile, and assemble into binary objects, finally linking them into an UPMEM-PIM binary executable. uPIMulator utilizes UPMEM SDK’s preprocessor and compiler *as-is* to first lower UPMEM program source files into multiple assembly-level codes. These assembly codes are then fed into our custom-designed linker (which is based on the open-source ANTLR’s lexer and parser [59], [60], [61]) to go through the lexical and syntax analyses to resolve the def-use relationships of all the functions, code labels, etc. for linking. Finally, our custom-designed assembler generates the final binary files to upload into UPMEM-PIM’s atomic (i.e., mutex), IRAM (i.e., the UPMEM-PIM program), WRAM, and MRAM (i.e., the input data) address spaces (Figure 4).

The reason why uPIMulator employs a custom-designed linker and assembler is as follows. We observe that the current version of UPMEM linker is specifically tied to UPMEM-PIM’s microarchitecture, preventing us from exploring alternative PIM architectures. For instance, UPMEM’s linker generates a linking error when the compiled program’s instruction memory or scratchpad usage exceeds the physical IRAM or WRAM capacity. As detailed later in Section V-D, this paper presents a case study to evaluate the trade-offs of employing an on-demand cache for UPMEM-PIM, as opposed to UPMEM’s current scratchpad-centric design. Under UPMEM’s programming model, this requires us to write the UPMEM-PIM program that has a working set allocated in the scratchpad (WRAM) space exceeding its 64 KB size, which is subsequently re-mapped to the per-bank DRAM region in our cycle-level hardware performance simulator. This allows us to treat a DPU thread’s load/store transactions to scratchpad as if they are to DRAM, so plugging in a cache simulator in between the DPU and scratchpad (which is emulated as DRAM) enables us to study the performance of caches vs. scratchpads (Section V-D details our methodology for this study). None of these features are available with UPMEM’s current linker design, motivating us to implement our own linker and assembler for a flexible simulator development and design space exploration.

Overall, uPIMulator enjoys LLVM’s mature compiler stage optimizations (e.g., common subexpression elimination [62]) by leveraging UPMEM’s existing preprocessor/compiler *as-is* while also enabling diverse architectural explorations through our custom-designed linker/assembler.

Hardware performance simulator. We implement a cycle-level performance simulator of UPMEM DPU by referring to both UPMEM’s user manual and publicly available information and discussion about the DPU’s microarchitecture [47], [54], [55], [58], [63], [64]. As summarized in Table I, the baseline DPU architecture is modeled as a 14-stage in-order pipelined processor, faithfully modeling its revolver pipeline scheduling algorithm and the structural hazard enforced at the odd/even RF accesses (Section II-A). uPIMulator functionally executes the instructions to update its architectural state, allowing us to verify the correctness of PIM program’s execution.

As for modeling the DRAM subsystem, rather than employing a highly accurate cycle-level DRAM simulator [65], [66], [67], we base our implementation on GPGPU-Sim’s cycle-level DRAM simulator for fast simulation time [68] (our simulator runs $2.5\times$ slower when interfaced with Ramulator [66], which

TABLE I: uPIMulator default configuration.

DPU processor architecture	
Operating frequency	350 MHz
Number of pipeline stages	14
Revolver scheduling cycles	11
WRAM / IRAM size	64 KB / 24 KB
WRAM / IRAM access latency	1 cycle
WRAM / IRAM access granularity	4 / 6 B per clock
WRAM / IRAM access bandwidth	1,400 / 2,100 MB/sec
Atomic memory size	256 Bits
DRAM system	
MRAM size	64 MB
DDR specification	DDR4-2400 [53]
Memory scheduling policy	FR-FCFS
Row buffer size	1 KB
tRCD, tRAS, tRP, tCL, tBL	16, 39, 16, 16, 4 cycles
Communication	
CPU→DPU bandwidth (per rank)	0.296 GB/s per DPU
CPU←DPU bandwidth (per rank)	0.063 GB/s per DPU
Software architecture	
Number of general-purpose registers	24
Maximum number of threads	24
Stack size (per thread)	2 KB
Heap size	4 KB

is known to be the fastest among popular DRAM simulators). Because the details of UPMEM-PIM’s memory scheduling policy is not publicly available, we employ a first-row, first-come-first-serve (FR-FCFS [69]) algorithm for scheduling memory transactions. The communication latency of transferring data over the CPU↔DPU channel is simulated by employing a fixed bandwidth model as its communication channel (i.e., communication latency = transfer size/communication bandwidth), whose value is carefully tuned by profiling a real UPMEM-PIM system (Table I).

Because UPMEM-PIM implements the CPU↔DPU communication using Intel AVX read (CPU←DPU) and write (CPU→DPU) instructions [70], we observe asymmetric CPU↔DPU communication bandwidth (i.e., the synchronous AVX reads have lower throughput than the asynchronous AVX writes), a phenomenon also reported in [48].

B. Simulator Availability and Extensibility

uPIMulator is designed to cleanly decouple the SPMD-based frontend code/data generation from the backend performance model with its modular design (Figure 4). Such design philosophy is inspired by GPGPU-Sim [68] which similarly utilizes NVIDIA’s CUDA compiler and PTX assembler as its frontend to generate CUDA code/data, which is subsequently consumed by its backend cycle-level GPU microarchitecture simulator. As such, uPIMulator can easily be extended to model and evaluate architecture designs with alternative software/hardware architectures (we later demonstrate uPIMulator’s extensibility via our case study in Section V). For instance, one can modify uPIMulator’s frontend code/data generation stage to flexibly map the code/data binaries at arbitrary locations in the memory address space, a feature we utilize to generate the proper instructions/data in our “cache vs. scratchpad” case study discussed later in Section V-D. Similarly, uPIMulator’s backend performance model can also

TABLE II: PrIM benchmarks configurations used for the characterization and case studies conducted in this work.

Benchmark	Dataset (single DPU)	Dataset (multiple DPUs)
BFS	2K vertices, 15K edges	16K vertices, 120K edges
BS	32K elem., 4K queries	128K elem., 16K queries
GEMV	2K x 64, 64 x 1 elem.	8K x 64, 64 x 1 elem.
HST-L	128K elem., 256 bins	512K elem., 256 bins
HST-S	128K elem., 256 bins	512K elem., 256 bins
MLP	3 layers, 256 neurons	3 layers, 1K neurons
NW	256 gene sequence	512 gene sequence
RED	512K elem.	2M elem.
SCAN-RSS	256K elem.	1M elem.
SCAN-SSA	256K elem.	1M elem.
SEL	512K elem.	2M elem.
SpMV	12K x 12K, 80519 elem.	14K x 14K, 316740 elem.
TRNS	128K elem.	256K elem.
TS	2K elem., 64 queries	64K elem., 64 queries
UNI	512K elem.	2M elem.
VA	1M elem.	4M elem.

be extended to execute UPMEM-PIM’s SPMD code over alternative hardware architectures. For instance, one can maintain the same UPMEM-PIM’s code to execute over an NVIDIA GPU style SIMD processor architecture by modifying the backend performance model to be implemented as a SIMT (single-instruction-multiple-thread) [56] vector processor microarchitecture model, a case study we conduct in Section V-A.

C. Simulator Validation

We validate our uPIMulator using PrIM [71], an open-source UPMEM-PIM benchmark suite (Table II). PrIM consists of 16 data-intensive workloads from various application domains such as linear algebra, graph processing, neural networks, etc. We verify uPIMulator’s functional correctness as well as its performance correlation to real UPMEM-PIM hardware by running each individual PrIM benchmark with 1/2/4/8/16/24 threads under various input data sizes, cross-validating both uPIMulator and real UPMEM-PIM’s final output data as well as its execution time. Among the 16 PrIM benchmarks, uPIMulator was able to compile and simulate 13 workloads as-is. However, the remaining 3 workloads (BFS, SpMV, NW) had minor bugs or utilized undisclosed functions within the UPMEM SDK, preventing its simulation and debugging on uPIMulator, so we modified these three benchmarks to provide the same functionality of the original implementation while being executable on top of uPIMulator. As discussed in Section III-A, the CPU↔DPU transfer (used for inter-DPU communication) is modeled as a fixed bandwidth model, so the frequency of inter-DPU communications can affect the accuracy of uPIMulator’s simulated execution time. To separately analyze the fidelity of uPIMulator’s DPU architecture model and the effect CPU↔DPU communication model has on system-level simulations, we separately report the validation results of uPIMulator when running the PrIM benchmark suite (1) with just a single DPU executing without any inter-DPU communication and (2) with multiple DPUs with inter-DPU communication. For the single DPU validation, we used 710 data points whose execution times are within the range of 500 ms, showing 98.4% correlation against UPMEM-PIM with a

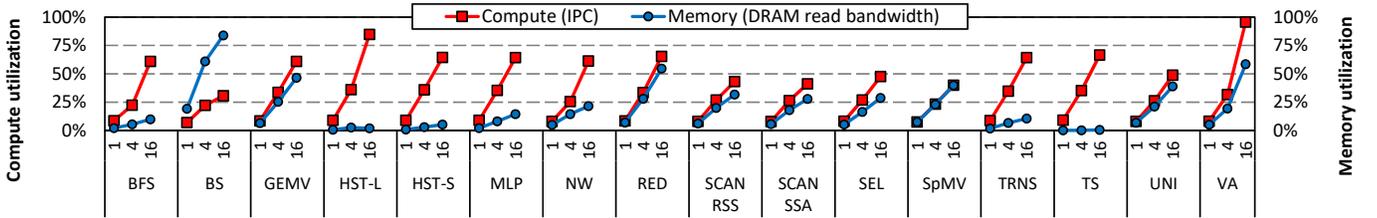


Fig. 5: PrIM’s compute utilization (left axis) and memory read bandwidth utilization (right axis) when executing with 1/4/16 threads. While a DPU’s theoretical maximum DRAM bandwidth is 700 MB/sec, prior work [48] observed that the maximum bandwidth is around 600 MB/sec in real UPMEM-PIM system. We therefore configured uPIMulator’s DRAM bandwidth accordingly. A single DPU’s max compute throughput is set as 1 IPC and compute utilization is the percentage of this max IPC achieved.

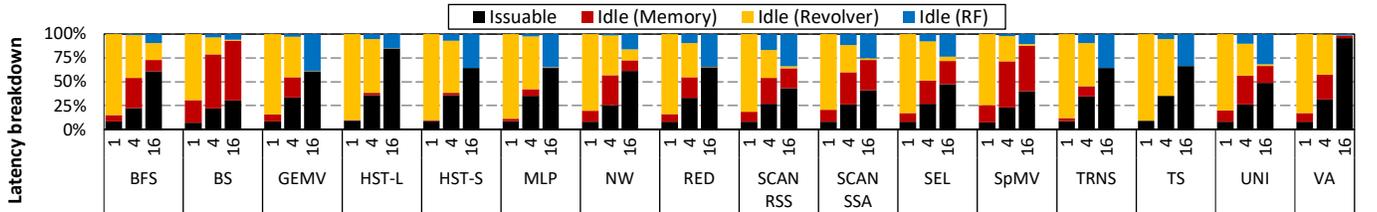


Fig. 6: Breakdown of DPU’s runtime into active (black) and idle (red, yellow, blue) cycles. When all the threads are idle, we categorize each thread’s status based on the reason for its idleness, i.e., memory (red), revolver pipeline scheduling constraint (yellow), and the structural hazard at the RF (blue).

mean absolute error (MAE) of 12.0%. Under the multi-DPU validation, uPIMulator shows 83.6% correlation with MAE of 26.9% under 387 data points, with relatively larger absolute errors observed when the inter-DPU communication time is more pronounced.

D. Simulation Rate

Developing a detailed execution-driven simulator often comes at the expense of increased simulation time. While uPIMulator is not multi-threaded, it achieves an average simulation rate of 3 KIPS (kilo-instructions-per-second), which is on par with other popular execution driven simulators like GPGPU-Sim [68]. Because of UPMEM’s current programming model and how its communication/synchronization primitives work (Section II-B), DPUs mostly operate independently as a standalone processor, so we expect parallelizing uPIMulator with multi-threading will lead to significant simulation rate improvements. We leave the support of multi-threaded uPIMulator implementation as future work.

IV. DEMYSTIFYING UPMEM-PIM WITH UPIMULATOR

This section utilize uPIMulator and the PrIM benchmark suite [71] to demystify the internal runtime characteristics of UPMEM-PIM, showcasing the applicability of uPIMulator for architectural exploration. We first focus on simulating PrIM under a single DPU setting in Section IV-A, identifying its bottleneck in Section IV-B, and finally discussing multi-DPU execution with strong scaling in Section IV-C. Table II summarizes the PrIM benchmarks and its input data sizes we explore in this paper. Due to space constraints, when sweeping the number of threads that execute a given PrIM benchmark (collected over 1/2/4/8/16/24 threads), we only show the results with 1/4/16 threads for brevity.

A. Analyzing Runtime Performance

Figure 5 shows the compute and memory bandwidth utilization as a function of the number of concurrent threads instantiated (1/4/16 threads). With the exception of BS and SpMV, PrIM benchmarks generally exhibit a compute-bound behavior, having a relatively higher compute utilization than DRAM bandwidth utilization. PrIM targets data-intensive workloads that are traditionally categorized as memory-bound under von-Neumann CPU/GPU architectures. As such, the results in Figure 5 highlight the unique value proposition of PIM vs. CPUs/GPUs, i.e., the performance bottleneck is now shifted from the memory-bound regime to the compute-bound territory. We observe similar performance results over real UPMEM-PIM systems (prior work in [48] reports similar observations), demonstrating the fidelity of our uPIMulator.

B. Identifying Bottlenecks

While the workloads in PrIM generally exhibit a compute-bound behavior, the results in Figure 5 imply that there are still some performance left on the table. Using uPIMulator, we now root-cause the key bottlenecks in UPMEM-PIM’s microarchitecture that cause such performance loss.

Latency breakdown. In Figure 6, we breakdown DPU’s execution time into two categories: (1) the time when the thread scheduler has one or more threads to *issue* into the pipeline (black bar), and (2) when the scheduler is left idle with *zero* threads to issue (all non-black bars) because all the threads are either (2-a) waiting for a memory operation to be serviced, (2-b) stalled due to the UPMEM’s revolver pipeline scheduling constraint, or (2-c) stalled due to the structural hazard at the odd/even register file (see Section II-A for revolver pipeline & RF hazard). As the number of threads increases, the DPU scheduler is provided with more thread-level parallelism to populate its 14-stage pipeline, leading to larger fraction of the

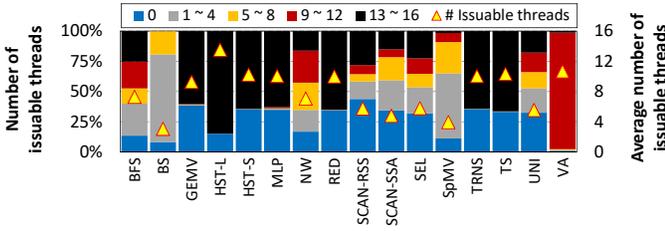


Fig. 7: Number of issuable threads by DPU scheduler each cycle, binned per each category (left axis) and average number of issuable threads (right axis) when executing with 16 threads.

runtime executing instructions. Nonetheless, several workloads still suffer from non-negligible portion of its execution time with idle cycles due to memory-side bottlenecks (BS, SpMV), the structural hazards caused by the revolver pipeline and odd/even RF (GEMV, HST-S, MLP, RED, TRNS, TS), or both (BFS, NW, SCAN-RSS, SCAN-SSA, SEL, UNI). While pipeline stalls due to memory operations are a fundamental one that cannot be resolved easily through processor-side optimizations, idle cycles originating from the revolver pipeline scheduling constraint or odd/even RF hazard is an artifact of UPMEM-PIM’s specific processor microarchitecture.

Thread-level parallelism (in space and time). To analyze UPMEM-PIM’s performance bottleneck from a different dimension, we measure the amount of thread-level parallelism (TLP) available to the DPU scheduler in space (Figure 7) and in time (Figure 8). In Figure 7, we categorize the number of issuable threads available to the DPU scheduler to issue into the pipeline by categorizing which bin it falls under. As depicted, workloads suffering from sub-optimal performance generally exhibit a higher fraction of its runtime with less TLP (i.e., ‘0’ issuable threads in the left axis of Figure 7), rendering the DPU to lose compute throughput (Figure 5). While insightful, the analysis in Figure 7 cannot capture the temporal variation in TLP or any phase behavior at runtime, which can add another level of insights for architectural exploration. uPIMulator enables the analysis of how TLP fluctuates as execution progresses, as shown in Figure 8. Although some workloads consistently exhibit low (BS) or high (GEMV) TLP, others exhibit a mix of high-and-low TLP behavior (SCAN-SSA), providing valuable insights to understand the runtime dynamics of a workload.

Instruction mix. Finally, Figure 9 shows the instruction mix of PrIM when executed with a 1/4/16 threaded single DPU. uPIMulator uncovers a couple of interesting observations as follows. First, with the exception of BFS, the number of load/store instructions to the scratchpad memory (red) generally outweighs the number of DMA instructions to the per-bank DRAM (yellow). This is an artifact of the scratchpad-centric programming model of UPMEM-PIM, i.e., the register data operands can only be loaded from the scratchpad and the programmer must manually initiate DRAM→scratchpad copies to move the working set closer to the processor. Consequently, to make sure the scratchpad accesses do not cause a performance bottleneck, the DPU microarchitecture is designed to guarantee a short, single cycle

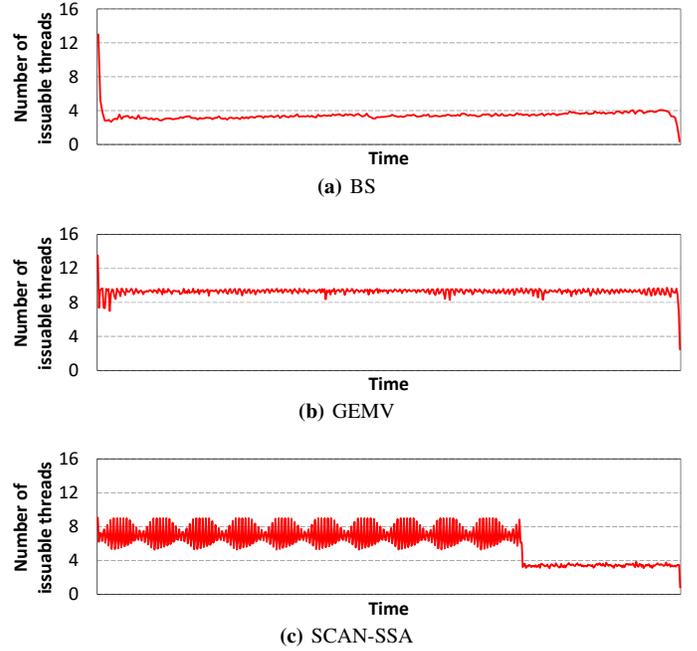


Fig. 8: Changes in the number of issueable threads (y-axis) in time (x-axis) during the course of (a) BS, (b) GEMV, and (c) SCAN-SSA’s execution. For clear visualization, the y-axis shows the number of issuable threads averaged over 10,000 consecutive cycles (i.e., cycles with zero issuable threads are not shown clearly as they are smoothed out while averaging).

latency in handling load/store instructions. Second, although the compute utilization of certain workloads like HST-L and TRNS seemingly look decent (Figure 5), a significant portion of its runtime is effectively wasted as it is busy waiting to acquire locks (e.g., `mutex_lock`). This is illustrated by the large fraction of the instructions executed in HST-L and TRNS dedicated to synchronization instructions (e.g., `acquire`, `release` in UPMEM ISA). Future UPMEM ISA extensions that enable busy waiting threads to transition into sleep mode and only resume execution when they are woken up can potentially reduce such inefficiency.

C. Strong Scaling with Multi-DPUs

Figure 10 shows the latency breakdown when each PrIM benchmark is parallelized across 1, 16, and 64 DPUs using strong-scaling, i.e., benchmark’s working set remains identical, so perfect strong-scaling would reduce latency proportional to the number of DPUs. In general, the majority of PrIM’s performance scales well when parallelized across multiple DPUs because the communication size per DPU gets proportionally reduced as a function of the DPUs concurrently executing. BFS, BS, and NW, however, exhibit sub-linear scaling because the communication size gets larger as the number of DPUs is increased. It is also worth noting that for some benchmarks like SCAN-RSS, SCAN-SSA, SEL, UNI, and VA, transferring input (CPU→DPU) and output (DPU→CPU) data dominates the total execution time. For these benchmarks, the latencies to copy the input/output data over CPU↔DPU channel are not being effectively hidden by overlapping it with DPU’s

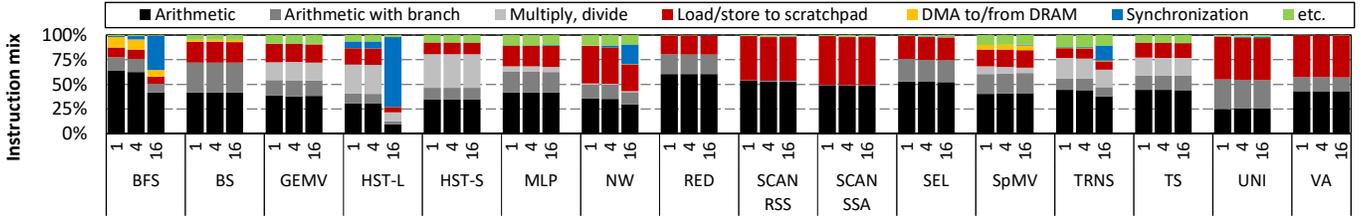


Fig. 9: Instruction mix when executing with a single DPU.

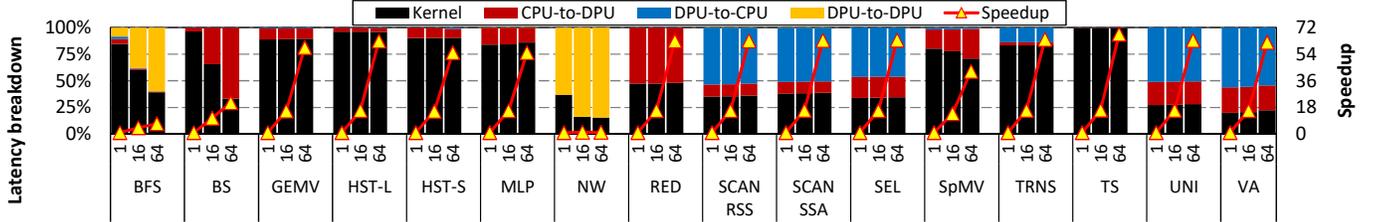


Fig. 10: Multi-DPU’s latency breakdown (left axis) and achieved speedups (right axis) when strong-scaling PRIM across 1/16/64 DPUs. All non-black bars represent communication latency. The DPU-to-DPU bar (yellow) in 1 DPU shows latency incurred in copying input/output data in/out of the DPU across kernel execution boundaries.

kernel execution time. Future versions of UPMEM SDK which provide programming semantics that facilitate flexible kernel partitioning and task scheduling (e.g., CUDA stream, CUDA dynamic parallelism [56]) will likely enable further performance improvements.

V. PATHFINDING FUTURE PIM ARCHITECTURES

In this section, we uncover some key limitations of the current UPMEM-PIM system through a series of case studies and demonstrate how uPIMulator can be utilized to explore architectural support required for *future* PIM architectures to become more performant, robust, and secure.

A. Case Study #1: UPMEM-PIM with SIMT Processing

The baseline UPMEM-PIM employs a scalar processor leveraging thread-level parallelism to maximize performance. Recent domain-specific PIMs [14], [19], on the other hand, leverage data-level parallelism by employing vector processing to boost their performance for key machine learning primitives (e.g., matrix multiplication). We observe that UPMEM’s SPMD programming model renders its hardware architecture to similarly reap out performance benefits of data-parallel execution by employing a SIMT (single-instruction-multiple-thread) vector processor [56]. In this subsection, we augment the baseline UPMEM-PIM as follows to analyze the performance benefits of employing SIMT vector processing. First, the processor pipeline is augmented with a vector register file which an N -way vector unit reads (write) vector operands from (to). Similar to the notion of “warps” in CUDA, we group N consecutive tasklets as the (grouped) thread scheduling granularity to the N -way vector unit which executes N scalar instructions in lockstep for vector processing. Similar to SIMT GPUs, a *memory address coalescing* operation [56] is applied among the grouped N scalar load/store instructions which helps maximize memory bandwidth utilization by minimizing the effect of SIMT memory divergence [72], [73]. SIMT

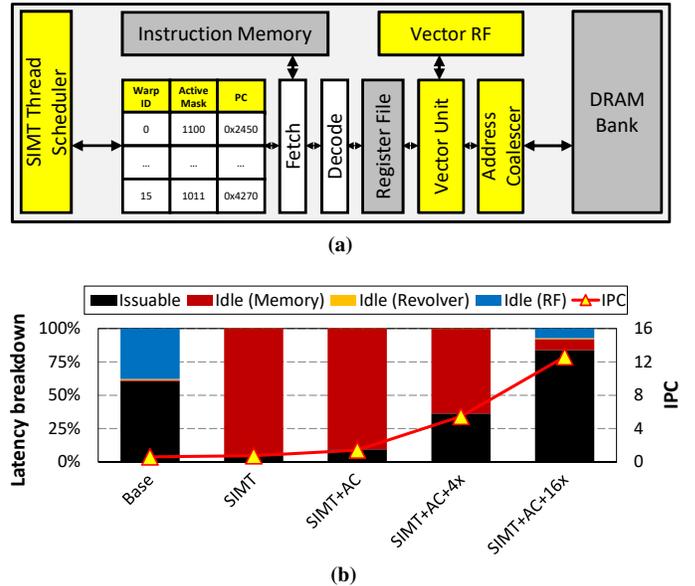


Fig. 11: (a) SIMT based DPU architecture modeled using uPIMulator, (b) performance (right axis) achieved for GEMV. The max IPC of Base and all SIMT designs are 1 and 16, respectively.

control divergence [73], [74], [75], [76], [77], [78], [79] is handled dynamically at runtime using each individual thread’s program counter values to only execute scalar threads executing the same instruction over the vector lanes, masking out inactive threads from execution as appropriate, similar to how recent NVIDIA GPUs (post Volta GPU) handle SIMT control divergence [80].

Figure 11 shows the performance achieved for GEMV, a key primitive in machine learning which recent domain-specific PIMs are optimized for. The figure first shows baseline UPMEM-PIM (Base), which is additively augmented with 1) 16-way SIMT vector unit *without* memory address coalescing (SIMT) and 2) *with* address coalescing (SIMT+AC). Both

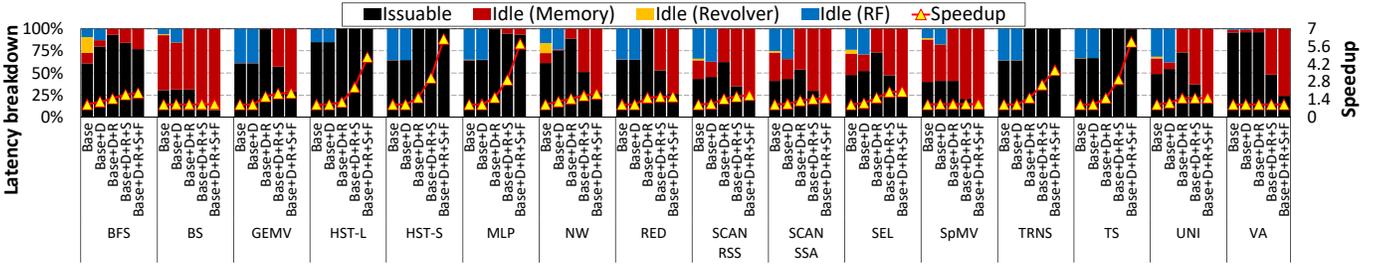


Fig. 12: Ablation study to explore UPMEM-PIM’s possible performance improvements when baseline DPU with 16 threads is additively enhanced with data forwarding logic (D), unified RF with $2\times$ read bandwidth to remove hazards at RF (R), 2-way superscalar in-order pipeline (S), and doubling the operating frequency to 700 MHz (F).

of these SIMT design points have MRAM read bandwidth identical to $Base$. Finally, another design point that scales up MRAM read bandwidth by increasing DRAM operating frequency by $4\times/16\times$ ($SIMT+AC+4\times/16\times$) is explored to evaluate the upperbound performance with SIMT. As depicted, augmenting UPMEM-PIM with a 16-way vector unit ($SIMT$) provides a mere $2.6\times$ speedup as performance is heavily bottlenecked by MRAM read bandwidth. Adding the memory coalescer ($SIMT+AC$) helps better utilize memory bandwidth and provides an additional $1.9\times$ speedup vs. $SIMT$ ($4.6\times$ vs. $Base$). Even with memory address coalescing (AC), however, the memory-boundedness of $SIMT$ execution persists with $SIMT+AC$, leaving performance left on the table, one which is only alleviated by the more aggressive design which scales up MRAM bandwidth further with $SIMT+AC+4\times/16\times$.

Key takeaways: UPMEM-PIM’s SPMD programming model makes its hardware architecture amenable to data-parallel processing via SIMT vector execution. UPMEM-PIM’s baseline memory system, however, is not sufficiently provisioned to sustain the higher DRAM read bandwidth requirements of vector execution, resulting in limited speedup with a naively implemented SIMT PIM design. Properly optimizing the PIM memory system to maximize bandwidth utilization (e.g., memory address coalescing, higher DRAM read bandwidth) will thus be crucial for future SIMT vector based PIM designs to fully unlock the potential of SIMT.

B. Case Study #2: ILP-enhanced PIM Architectures

Today’s commercial PIM processors employ a simple, in-order pipeline without any sophisticated microarchitectures to extract ILP for higher performance (e.g., superscalar, super-pipelining, ...) [14], [15], [18], [55]. As discussed in Figure 6, such a wimpy PIM processor design point leaves significant performance left on the table, as conventionally memory-bound workloads now fall under the compute-bound regime with PIM (Section IV-A). We believe such design decision was inevitable because current generation of PIM processors are fabricated on a density-optimized technology node (e.g., ≥ 20 nm DRAM technology for HBM-PIM and UPMEM-PIM [15], [48]) posing several design constraints that prevent advanced microarchitecture designs. That being said, future PIM architectures with more flexibility in area, power, and thermal budget can certainly consider relatively more

aggressive, performance-oriented design points with higher operating frequency and ILP-enhancing microarchitectures.

In this case study, we use uPIMulator to see how much performance can be unlocked in PrIM’s “compute-bound” workload by augmenting UPMEM-PIM’s DPU with ILP enhancing optimizations. Figure 12 summarizes our ablation study on how much the baseline UPMEM-PIM’s performance (denoted “Base”) can be improved by adding the following features in an additive manner: (D) addressing the scheduling constraint enforced with baseline revolver pipeline by enabling data forwarding across instructions without data dependencies within the same thread to execute, (R) merging the odd/even RF into a single one but doubling the read bandwidth to eliminate baseline RF’s structural hazard, (S) 2-way superscalar in-order pipeline to better leverage parallelism, and finally (F) doubling DPU’s operating frequency to 700 MHz. As depicted, the addition of these microarchitecture techniques substantially improve the performance of PrIM’s compute-bound workloads (avg $2.7\times$, max $6.2\times$ speedup) as it successfully addresses the performance bottlenecks discussed in Figure 6. Interestingly, with the addition of (D+R+S) features to baseline UPMEM-PIM, several workloads become relatively more memory-bound (i.e., larger fraction of Idle(Memory)) so the benefits of higher operating frequency (F) are less pronounced for these workloads (e.g., GEMV, VA). A fundamental reason why baseline UPMEM-PIM cannot fully reap out the potential of these ILP optimizations is because of the large performance gap between WRAM bandwidth (2,800 MB/sec) vs. MRAM-to-WRAM bandwidth (600-700 MB/sec). More concretely, for those workloads exhibiting low data locality, the performance becomes relatively MRAM access bound and renders any optimizations that resolve the compute-boundness of a workload ineffective (e.g., all data points exhibiting high fraction of Idle(Memory) in Figure 12). Note that the existing 600-700 MB/sec of MRAM-to-WRAM bandwidth is not a fundamental constraint because the maximum memory bandwidth that can be reaped out at the MRAM (DRAM) “bank” level is much higher (up to several GB/sec of bandwidth), i.e., the limited 600-700 MB/sec of MRAM bandwidth is simply a design point pursued by UPMEM-PIM architects for this particular PIM design. Using uPIMulator, we conduct a sensitivity study that scales up the MRAM-to-WRAM read bandwidth and analyze its performance implication for memory-bound workloads. As

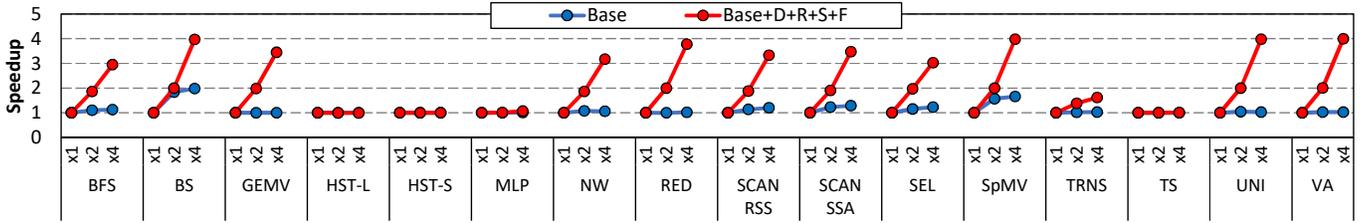


Fig. 13: Speedup achieved when scaling up the MRAM-to-WRAM bandwidth by four times ($\times 1$ to $\times 4$). The evaluated design points are 1) baseline UPMEM-as-is (Base) and 2) UPMEM with all the ILP optimizations (Base+(D+R+S+F)) discussed in Figure 12.

shown in Figure 13, the speedup is more pronounced with the ILP-enhanced UPMEM-PIM designs (red lines) because they exhibit more memory-boundedness as shown in Figure 12. Contrarily, benefit of MRAM bandwidth scaling is minimal for workloads still exhibiting compute-boundedness even under ILP optimizations (HST-L, HST-S, MLP, TRNS, TS). Same principle holds for the baseline UPMEM as-is (blue lines) where the only noticeable speedup with MRAM scaling is observed only for BS and SpMV which are already heavily memory-bound even without ILP optimizations, the other remaining compute-bound workloads achieving little speedup.

It is worth pointing out that, while the two case studies discussed so far have quantified the performance merits of both SIMT and superscalar execution in a PIM architecture, the available power and area budget can limit how aggressively SIMT or superscalar can be employed within PIM. Standalone PIM systems like SK Hynix’s AiM [19], [81], which are integrated as co-processors on top of a PCIe bus, have much larger power and area budget than a DIMM-based PIM solutions like UPMEM-PIM. Therefore, these standalone, domain-specific PIM solutions which have more design flexibility will more likely be prime candidates to embrace SIMT or superscalar based PIM designs that require higher design overheads.

Key takeaways: Many data-intensive workloads exhibit a compute-bound behavior when executed with PIM. Enhancing PIM’s computational throughput will therefore become much more important in future PIM designs. Using uPIMulator, we demonstrate the efficacy of various ILP-enhancing microarchitectural techniques for future PIM architectures, improving the performance of several compute-bound PIM workloads.

C. Case Study #3: Multi-tenant Execution in PIM

Multi-tenancy is one of the most important features for processors to provide for cloud deployment as it helps better saturate the processor’s compute and memory resources, reducing its total cost of ownership. As such, current CPUs/GPUs come with a variety of hardware/software features that support multi-tenancy with performance isolation and security guarantees (e.g., CPU cache partitioning [82], [83], NVIDIA’s multi-instance GPU [84], etc. [85], [86]). Given UPMEM-PIM’s abundant compute and memory throughput (e.g., an aggregate compute and memory throughput of 0.896 TOPS and 2.5 TB/sec of memory bandwidth in a 40 ranked UPMEM-PIM system), having a proper multi-tenancy support will be

vital for future PIM architectures, especially when seeking for industrial adoption by cloud vendors.

Unfortunately, our case study reveals that current commercial PIM chips (whether it be UPMEM-PIM or domain-specific PIMs [14], [15], [18], [19]) are not able to meet the requirements of multi-tenancy, not just from a hardware/software perspective, but also from its programming model’s perspective. Due to space limitations, let us focus our attention on two important conditions to be met for multi-tenancy. First, co-located workloads should securely execute without interfering with each other (i.e., “security” guarantees). Second, co-located workloads must not be aware of the fact that they are concurrently executing (i.e., “transparency” to co-located applications). We discuss each of these challenges below.

Security. One of the fundamental architectural supports that is needed for secure execution is isolating the memory address space of co-located applications using MMU’s *address translation* capability. Practically all commercial PIM systems [14], [15], [18], [55] are designed *without* an MMU, a point we emphasized in Section II-C with UPMEM-PIM’s *physical addressing* based memory model. Note that the decision regarding which granularity multi-tenancy should be employed (e.g., coarse-grained per-DPU vs. fine-grained intra-DPU multi-tenancy) presents interesting tradeoffs in terms of DPU resource contention, virtualization overhead, etc. Such design decision, however, does not obviate the need for the MMU to isolate different tenants by translating virtual addresses. Consider a design point where per-DPU multi-tenancy is implemented, e.g., two different PIM programs (two tenants) execute over non-overlapping groups of DPUs within the same DIMM. If one of the tenants is a malicious attacker, the malicious host-side CPU program can freely access the other victim tenant’s DPU physical memory because current PIM architectures employ (MMU-less) physical addressing when accessing their DRAM banks. Therefore, co-locating multiple workloads with address space isolation is fundamentally impossible in MMU-less PIM architectures.

Aside from such security benefits, PIM chips with an MMU can greatly enhance programmer productivity by *virtualizing* the memory address space, i.e., they can separate the physical allocation of data in PIM DRAM against its logical allocation within the virtual address space. As discussed in Section II-B, copying data from CPU to UPMEM-PIM’s DRAM (MRAM) requires the programmer to painstakingly derive the physical

location in DPU’s DRAM because UPMEM ISA is currently based on physical addressing. Having a proper MMU support will enable more flexible allocation of data across the physical address space and can also provide “pointer-is-a-pointer” semantics to enhance programmability [57], [87], [88], [89].

In this case study, we add a hardware MMU to UPMEM-PIM, using our uPIMulator, and study its performance implications. Our MMU model employs a single page-table walker (page size of 4 KB) backed with a single-level, 16-entry fully-associative TLB. The page-tables are assumed to be stored in DPU’s local DRAM bank and the access latency to the TLB is assumed as a single DPU clock cycle. Aside from how a PIM core (the DPU) handles address translation exceptions, the interaction between a DPU and its MMU largely follows that of a conventional CPU and its MMU. That is, in the event that a DPU accesses memory, the MMU translates all DRAM (MRAM)’s virtual address to its physical address by leveraging the TLB or the page-table. For memory requests that the MMU is not able to handle, however, an assistance from the host CPU is required. This can occur, for example, when a page fault occurs and an update to the page-table is in need. Under such circumstances, the MMU writes the fault information into a fault buffer, which can either be recognized by the host CPU via a polling approach or an interrupt-based approach. Under a polling approach, the host periodically polls each DPU’s fault buffer to fulfill DPU’s service needs. If the interrupt-based approach is to be employed, the MMU can raise an interrupt-like signal via DDR4’s ALERT_N standard protocol to interrupt and notify the host [90]. The host can then recognize the existence of a page fault within the DPU and handle it appropriately through a fault handler, updating the DPU’s page-table before sending a resume signal. We utilize such mechanism to translate all DRAM (MRAM) accesses from its virtual address to its physical address and measure its performance regression. Overall, PrIM experiences an average performance loss of 0.8% (max 14.1%) by adding address translations to DPU’s memory accesses. Such low performance overhead is mainly attributed to UPMEM’s scratchpad-centric memory model where data transfers across DRAM↔scratchpad are orchestrated in coarse-grained chunks (several KBs) over DMA instructions, exhibiting high spatial locality and thus achieving high TLB hit rates. Furthermore, DPU cores are clocked at 350 MHz frequency, rendering their memory access latency to be in the range of several tens of DPU clock cycles (unlike CPUs/GPUs which operate in the GHz range and exhibit hundreds of CPU/GPU processor cycles of memory latency), experiencing much less TLB miss penalty than CPUs/GPUs. Overall, our case study demonstrates the practicality of adding a functional MMU architecture to future PIM technologies.

Transparency. We believe that multi-tenant execution under the *current* UPMEM programming model is not practical because of its scratchpad-centric programming model. Consider a scenario where we seek to co-locate two PrIM benchmarks, a memory-bound BS and a compute-bound TS, which exhibit complementary resource requirements

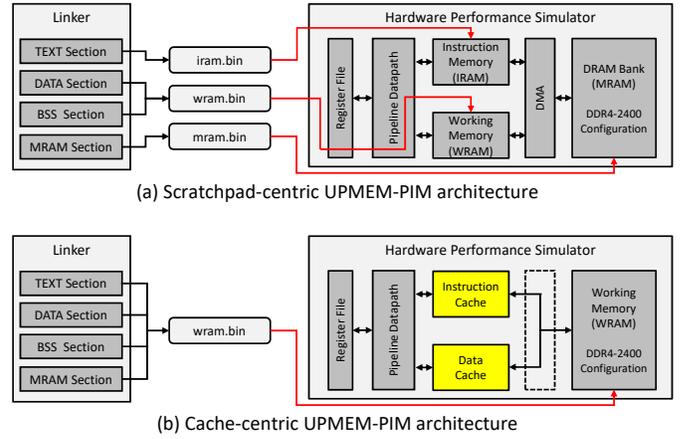


Fig. 14: Modeling a (a) scratchpad-centric and (b) cache-centric UPMEM-PIM architecture in uPIMulator.

(as quantified in Figure 5) and are perhaps the perfect candidates for multi-tenant execution. Unfortunately, the BS and TS each require using the same scratchpad (WRAM)’s heap via a memory allocation API call (`mem_alloc()` in UPMEM SDK, e.g., line 35-37 in Figure 2(b)), which leads to exceeding the total size of scratchpad (WRAM)’s heap size. Consequently, co-locating both of these workloads requires a non-trivial amount of changes to *both* co-located PrIM programs, arguably an unacceptable requirement to enforce on end-user applications. More crucially, it directly violates the *transparency* requirement we previously discussed, rendering a scratchpad-centric PIM programming model ill-suited for multi-tenant execution.

Consequently, our analysis reveals that future PIM should also employ *on-demand caches*, rather than singlehandedly relying on scratchpads, to reap data locality benefits. PIM programs running on top of an on-demand cache will be able to leverage data locality in an opportunistic manner while not having to change the program itself. In the next subsection, our next case study details the feasibility of supporting on-demand cache architectures for future PIM designs.

Key takeaways: Supporting multi-tenancy in PIM requires security and transparency guarantees for the co-located workloads. To enhance security in PIM architectures, we augment uPIMulator with an MMU to quantify the performance overheads of address translations and observe an average 0.8% (max 14.1%) latency increase, demonstrating the practicality of an MMU-enabled PIM design. Guaranteeing transparency to co-located PIM workloads under UPMEM’s current, scratchpad-centric programming model is a different story, however, as it requires significant changes to the co-located programs, a non-option to begin with for transparent multi-tenant execution. Having an on-demand cache architecture supported in PIM can bridge this gap, opening the door for multi-tenant PIM architectures.

D. Case Study #4: On-demand Caches vs. Scratchpads

As discussed in Section II-B, UPMEM’s scratchpad-centric programming model requires the programmer to explicitly

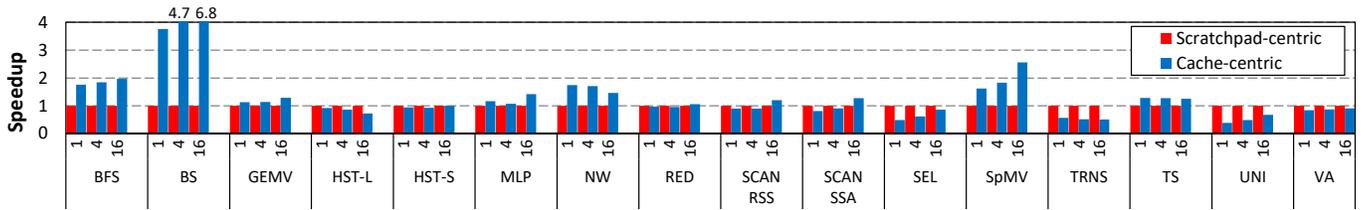


Fig. 15: Performance of scratchpad-centric vs. cache-centric UPMEM-PIM (normalized to scratchpad-centric design). The cache-centric UPMEM-PIM employs a cache line size of 64 bytes with load coalescing features enabled.

orchestrate the data movements across two *distinct* address spaces, the DRAM space and scratchpad space (i.e., MRAM \leftrightarrow WRAM). This is because the DPU threads can only load (store) data from (to) the scratchpad but cannot directly address data mapped in the DRAM space. Using uPIMulator’s custom-designed linker, this subsection conducts the cache vs. scratchpad case study based on the following methodology.

- 1) The open-source UPMEM compiler does not limit the data size the programmer can allocate and copy into WRAM (scratchpad) space. Concretely, compiling an UPMEM-PIM program to an assembly-level code whose scratchpad allocation size exceeds the physical WRAM capacity in itself does not cause any compilation errors. During the linking process, however, if the WRAM data allocation size exceeds the *physical* WRAM capacity, the UPMEM linker generates a linking error as the hardware UPMEM-PIM chip cannot execute the compiled codes properly (see Section III-A for discussion on UPMEM linker’s key properties).
- 2) Because uPIMulator’s linker is designed to flexibly relocate and map a given address region to anywhere in the physical address space, we take the following measures to *emulate* an alternative, cache-centric UPMEM-PIM (a) whose DPU threads can directly address data allocated in DRAM without having to move data to the scratchpad (i.e., there is no notion of scratchpad under this model), while (b) also leveraging data locality by storing recently accessed data within the cache.
- 3) uPIMulator emulates cache-centric UPMEM-PIM as follows. First, the input data is allocated directly in the WRAM (scratchpad) address space, unlike the baseline UPMEM model whose input data must be copied from MRAM (per-bank DRAM) to WRAM using DMA instructions. The WRAM-allocated input data, which is directly addressable by the DPU threads using load/store instructions (as compiled by the original UPMEM compiler), is then relocated by uPIMulator’s linker to be mapped into a physical address region which is backed by our cycle-level hardware performance simulator, modeling it as a DDR4-2400 [53] compatible DRAM device (Figure 14(b)). By adding a cycle-level cache simulator in between the DPU processor and the (DRAM-emulated) WRAM address space, the data referenced by the load/store instructions will be stored on-demand to this cache simulator, allowing us to explore the cache vs. scratchpad design space.

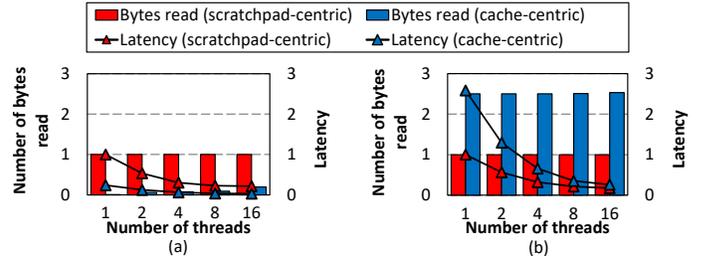


Fig. 16: Bytes read from DRAM (left axis, normalized) and end-to-end execution time (right axis, normalized) for (a) BS and (b) UNI.

Figure 15 compares the performance of scratchpad vs. cache in UPMEM-PIM for PrIM. The cache-centric UPMEM-PIM employs an instruction cache and a data cache, each configured as an 8-way set-associative cache with LRU replacement policy and 24 KB and 64 KB capacity, respectively, identical to the instruction memory (IRAM) and scratchpad (WRAM) space provisioned under the baseline UPMEM-PIM. For certain workloads, scratchpad performs better than caches (e.g., UNI) while the opposite also holds true for others (e.g., BS). To better understand the reasons behind such results, Figure 16 shows the number of bytes read from DRAM during the course of BS and UNI’s execution. In general, we can observe that the execution time is highly correlated with the number of bytes read. For example, under the memory-bound BS, the scratchpad based execution with 16 threads incurs $5.1\times$ higher memory read traffic than using caches. For BS, it is challenging to statically estimate the right amount of data to upload into the scratchpad, which results in a severe *overfetching* of useless data and causing a performance bottleneck to this memory-bound workload. Under such scenario, a purely on-demand caching strategy performs much more favorably in terms of fetching (relatively) the right amount of data within the cache and maximizes data locality. In contrast, workloads like UNI performs much better with scratchpads where carefully orchestrating data movements perform better than the opportunistic cache design. Determining which design point is more favorable for PIM architectures is not the purpose of this case study. Rather we seek to demonstrate the practical benefits and feasibility of a cache-centric PIM architecture, motivating future work in this research space.

Key takeaways: Similar to conventional CPUs/GPUs, an on-demand cache design can do a better job in leveraging locality for PIM when its memory access pattern cannot be optimally determined at compile time, a scenario where scratchpad based design points can perform poorly.

E. Other Promising Research Directions

Aside from the case studies we discussed previously, we believe that PIM with better inter-DPU synchronization primitives, high-performance inter-DPU communications, robust system software support for better programmability, and a unified virtual memory across all DPUs are critical components that require attention from PIM architects. We plan on exploring these studies as future work.

VI. SIMULATOR LIMITATIONS AND FUTURE WORK

A. Power and Area Modeling for PIM

Similar to the early efforts on modeling cycle-level performance of programmable GPUs [68], the current version of uPIMulator primarily focuses on simulating the performance aspects of UPMEM-PIM. There exists a rich set of prior work focusing on estimating the power and area of CPUs [91], [92] and GPUs [93], [94], [95] and integrating them with cycle-level CPU/GPU performance model simulators. An important future direction of uPIMulator is to develop a power and area modeling framework targeting PIM and integrate them with our UPMEM-PIM performance model. We leave it as future work as it deserves a detailed exploration on its own.

B. Improving the Fidelity of Inter-DPU Communication

As discussed in Section III-C, using a simple bandwidth model for CPU \leftrightarrow DPU communications renders uPIMulator to exhibit relatively lower correlation with real UPMEM-PIM system when the inter-DPU communication time is more pronounced. A real UPMEM-PIM system implements such communication operator by having the host CPU execute AVX instructions, so improving the fidelity of uPIMulator’s inter-DPU communication requires our simulation framework to be tightly integrated with a detailed cycle-level CPU performance model [96], [97], [98], [99], [100], [101]. Extending uPIMulator to be integrated with high fidelity CPU simulators is left as future work.

VII. RELATED WORK

The initial concept of PIM dates back to the 1970s [102] with numerous follow-on works [2], [3], [4], [5]. With the proliferation of today’s domain-specific architectures, there exists a variety of PIM or near-memory processing studies [14], [16], [18], [20], [21], [22], [23], [24], [25], [26], [27], [28], [29], [30], [31], [32], [33], [34], [35], [36], [37], [38], [39], [40], [41], [42], [103], [104], [105], [106], [107], [108], [109], [110], [111], [112], [113], [114]. There are also several prior works on PIM exploring compiler issues [8], [10], [115], data coherency [11], [12], [13], [116], synchronization [117], QoS aware runtime and scheduling for PIM [118], among many others [119], [120], [121], [122], [123], [124], [125]. This paper focuses on characterizing the first real-world general-purpose PIM via our uPIMulator, pathfinding important research directions for future PIMs. Below we summarize other relevant works on characterizing real-world PIM and PIM simulators.

Analysis on real-world PIM devices. There have been several recent works that characterize commercial PIM technologies [14], [48], [49], [52], [71], [126], [127], [128], [129], [130], [131]. Gómez-Luna et al. [48], [71] developed the PrIM benchmark suite and conducted a workload characterization on UPMEM-PIM. There are also several works exploring the applicability of UPMEM-PIM for accelerating dense/sparse linear algebra, databases, data analytics, graph processing, bioinformatics, image processing, compression, simulation, encryption, and etc [48], [49], [52], [71], [126], [127], [130], [131], with more recent work exploring UPMEM-PIM’s applicability for accelerating machine learning [129]. Lee et al. [14] discusses the hardware/software architecture for Samsung’s HBM-PIM architecture. There is also a recent work by Liu et al. [16] which explores the applicability of Samsung’s near-memory processor AxDIMM for accelerating recommendation models.

Simulators for PIM. PIMSim [132] supports a configurable PIM logic modeling, memory organization, interconnection, and provides co-simulation with other simulation frameworks. Ramulator-PIM [66], [133], [134] integrates ZSim [97] and Ramulator [66] to simulate PIM-enabled memory. MPU-Sim [135] models a near-bank processing architecture which supports NVIDIA CUDA’s SIMT programming model [56]. MultiPIM [136] enables the simulation of PIM systems based on 3D stacked memory with features to explore multi-stack interconnects with virtual memory support. Compared to these existing PIM simulators, the key novelties of uPIMulator are as follows. First, the frontend of our software compilation toolchain employs a custom-designed linker targeting industry’s first general-purpose PIM ISA, which opens up a wide range of hardware/software architectural explorations. Existing PIM simulators primarily rely on conventional software frontends (e.g., x86 in ZSim+Ramulator), making it challenging to change the way the program and data binaries are mapped over the memory address space, a feature critical in some of the case studies we conducted in Section V. Second, uPIMulator’s backend simulator models a real-world per-bank PIM architecture, widely employed and commercialized in both domain-specific [14], [19] and general-purpose PIM designs, unlike popular PIM simulators like MultiPIM or ZSim+Ramulator [66], [133], [134], [136] which assume the PIM cores are placed in the logic layer of a 3D stacked memory (e.g., HMC). Table III summarizes key differences between uPIMulator and other PIM simulators.

VIII. CONCLUSION

In this work, we present a novel simulation framework named uPIMulator which targets UPMEM’s commercial general-purpose PIM architecture. Using uPIMulator, we present our detailed characterization on wide range of real PIM programs and showcase uPIMulator’s applicability for computer architecture research. Furthermore, we identify some important shortcomings of the current UPMEM-PIM system through our case studies and propose some critical research areas that require further investigation from computer

TABLE III: Comparison of uPIMulator vs. other PIM simulators.

	PIMSim [132]	Ramulator -PIM [133]	MultiPIM [136]	MPU-Sim [135]	uPIMulator
ISA	x86, ARM, SPARC	x86	x86	PTX	UPMEM
Implementation	In-house	Zsim + Ramulator	Zsim + Ramulator + BookSim	In-house	In-house
Frontend (Trace vs. Execution)	Trace	Trace + Execution	Trace + Execution	Execution	Execution
ISA & Linker Customization	X	X	X	X	O
Validation Against Real PIM Hardware	X	X	X	X	O
Multi-threaded Simulation	X	O	O	X	X
Lines of Code (LoC)	30 K	35 K	92 K	12 K	52 K
Simulation Rate (KIPS)	N/A	N/A	N/A	N/A	3

architects (e.g., architectural support for vector processing, ILP-enhancing microarchitectures, multi-tenancy, and on-demand caching), which we believe will be vital for future PIM architectures to evolve into first class computing citizens.

ACKNOWLEDGMENT

This research is supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government(MSIT) (NRF-2021R1A2C2091753), Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No. 2022-0-01037, Development of High Performance Processing-in-Memory Technology based on DRAM), and by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2019-0-00075, Artificial Intelligence Graduate School Program(KAIST)). We also appreciate the support from Samsung Electronics (Samsung-KAIST Center for Memory-Centric System Architecture) and Samsung Electronics Co., Ltd (IO201210-07974-01).

REFERENCES

- [1] M. Anderson, B. Chen, S. Chen, S. Deng, J. Fix, M. Gschwind, A. Kalaiah, C. Kim, J. Lee, J. Liang, H. Liu, Y. Lu, J. Montgomery, A. Moorthy, S. Nadathur, S. Naghshineh, A. Nayak, J. Park, C. Petersen, M. Schatz, N. Sundaram, B. Tang, P. Tang, A. Yang, J. Yu, H. Yuen, Y. Zhang, A. Anbudurai, V. Balan, H. Bojja, J. Boyd, M. Breitbart, C. Caldato, A. Calvo, G. Catron, S. Chandwani, P. Christeas, B. Cotel, B. Coutinho, A. Dalli, A. Dhanotia, O. Duncan, R. Dzhbarov, S. Elmir, C. Fu, W. Fu, M. Fulthorp, A. Gangidi, N. Gibson, S. Gordon, B. P. Hernandez, D. Ho, Y.-C. Huang, O. Johansson, S. Juluri, S. Kanaujia, M. Kesarkar, J. Killinger, B. Kim, R. Kulkarni, M. Lele, H. Li, H. Li, Y. Li, C. Liu, J. Liu, B. Maher, C. Mallpedi, S. Mangla, K. K. Matam, J. Mehta, S. Mehta, C. Mitchell, B. Muthiah, N. Nagarkatte, A. Narasimha, B. Nguyen, T. Ortiz, S. Padmanabha, D. Pan, A. Poojary, Y. Qi, O. Raginel, D. Rajagopal, T. Rice, C. Ross, N. Rotem, S. Russ, K. Shah, B. Shan, H. Shen, P. Shetty, K. Skandakumaran, K. Srinivasan, R. Sumbaly, M. Tauberg, M. Tzur, S. Verma, H. Wang, M. Wang, B. Wei, A. Xia, C. Xu, M. Yang, K. Zhang, R. Zhang, M. Zhao, W. Zhao, R. Zhu, A. Mathews, L. Qiao, M. Smelyanskiy, B. Jia, and V. Rao, "First-Generation Inference Accelerator Deployment at Facebook," in *arxiv.org*, 2021.
- [2] D. Elliott, M. Stumm, W. Snelgrove, C. Cojocar, and R. Mckenzie, "Computational RAM: Implementing Processors in Memory," *IEEE Design & Test of Computers*, vol. 16, no. 1, pp. 32–41, 1999.
- [3] D. Patterson, T. Anderson, N. Cardwell, R. Fromm, K. Keeton, C. Kozyrakis, R. Thomas, and K. Yelick, "A Case for Intelligent RAM," *IEEE Micro*, vol. 17, no. 2, pp. 34–44, 1997.
- [4] M. Oskin, F. Chong, and T. Sherwood, "Active Pages: A Computation Model for Intelligent Memory," in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 1998.

- [5] J. Draper, J. Chame, M. Hall, C. Steele, T. Barrett, J. LaCoss, J. Granacki, J. Shin, C. Chen, C. W. Kang, I. Kim, and G. Daglikoca, "The Architecture of the DIVA Processing-In-Memory Chip," in *Proceedings of the 16th International Conference on Supercomputing (ICS)*, 2002.
- [6] N. S. Kim, "Practical Challenges in Supporting Function in Memory," in *IEEE Asian Solid-State Circuits Conference (A-SSCC)*, 2018.
- [7] O. Mutlu, S. Ghose, J. Gómez-Luna, and R. Ausavarungnirun, "A Modern Primer on Processing in Memory," in *arxiv.org*, 2020.
- [8] K. Hsieh, E. Ebrahim, G. Kim, N. Chatterjee, M. O'Connor, N. Vijaykumar, O. Mutlu, and S. W. Keckler, "Transparent Offloading and Mapping (TOM): Enabling Programmer-Transparent Near-Data Processing in GPU Systems," in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2016.
- [9] A. Pattnaik, X. Tang, A. Jog, O. Kayiran, A. K. Mishra, M. T. Kandemir, O. Mutlu, and C. R. Das, "Scheduling Techniques for GPU Architectures with Processing-In-Memory Capabilities," in *Proceedings of the International Conference on Parallel Architectures and Compilation Techniques (PACT)*, 2016.
- [10] A. Devic, S. B. Rai, A. Sivasubramaniam, A. Akel, S. Eilert, and J. Eno, "To PIM or Not for Emerging General Purpose Processing in DDR Memory Systems," in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2022.
- [11] J. Ahn, S. Yoo, O. Mutlu, and K. Choi, "PIM-Enabled Instructions: A Low-Overhead, Locality-Aware Processing-in-Memory Architecture," in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2015.
- [12] A. Boroumand, S. Ghose, M. Patel, H. Hassan, B. Lucia, R. Ausavarungnirun, K. Hsieh, N. Hajinazar, K. T. Malladi, H. Zheng, and O. Mutlu, "CoNDA: Efficient Cache Coherence Support for Near-data Accelerators," in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2019.
- [13] A. Boroumand, S. Ghose, M. Patel, H. Hassan, B. Lucia, K. Hsieh, K. T. Malladi, H. Zheng, and O. Mutlu, "LazyPIM: An Efficient Cache Coherence Mechanism for Processing-in-Memory," *IEEE Computer Architecture Letters*, vol. 16, no. 1, pp. 46–50, 2017.
- [14] S. Lee, S.-h. Kang, J. Lee, H. Kim, E. Lee, S. Seo, H. Yoon, S. Lee, K. Lim, H. Shin, J. Kim, O. Seongil, A. Iyer, D. Wang, K. Sohn, and N. S. Kim, "Hardware Architecture and Software Stack for PIM Based on Commercial DRAM Technology : Industrial Product," in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2021.
- [15] Y.-C. Kwon, S. H. Lee, J. Lee, S.-H. Kwon, J. M. Ryu, J.-P. Son, O. Seongil, H.-S. Yu, H. Lee, S. Y. Kim, Y. Cho, J. G. Kim, J. Choi, H.-S. Shin, J. Kim, B. Phuah, H. Kim, M. J. Song, A. Choi, D. Kim, S. Kim, E.-B. Kim, D. Wang, S. Kang, Y. Ro, S. Seo, J. Song, J. Youn, K. Sohn, and N. S. Kim, "25.4 A 20nm 6GB Function-In-Memory DRAM, Based on HBM2 with a 1.2TFLOPS Programmable Computing Unit Using Bank-Level Parallelism, for Machine Learning Applications," in *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 2021.
- [16] 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, "Near-Memory Processing in Action: Accelerating Personalized Recommendation with AxDIMM," *IEEE Micro*, vol. 42, no. 1, pp. 116–127, 2022.
- [17] 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, "Aquadbolt-XL HBM2-PIM, LPDDR5-PIM with In-Memory Processing, and AXDIMM with Acceleration Buffer," *IEEE Micro*, vol. 42, no. 3, pp. 20–30, 2022.
- [18] M. He, C. Song, I. Kim, C. Jeong, S. Kim, I. Park, M. Thottethodi, and T. N. Vijaykumar, "Newton: A DRAM-maker's Accelerator-in-Memory (AiM) Architecture for Machine Learning," in *Proceedings of the International Symposium on Microarchitecture (MICRO)*, 2020.
- [19] S. Lee, K. Kim, S. Oh, J. Park, G. Hong, D. Ka, K. Hwang, J. Park, K. Kang, J. Kim, J. Jeon, N. Kim, Y. Kwon, K. Vladimir, W. Shin, J. Won, M. Lee, H. Joo, H. Choi, J. Lee, D. Ko, Y. Jun, K. Cho, I. Kim, C. Song, C. Jeong, D. Kwon, J. Jang, I. Park, J. Chun, and J. Cho, "A 1ynm 1.25V 8Gb, 16Gb/s/pin GDDR6-based Accelerator-in-Memory Supporting 1TFLOPS MAC Operation and Various Activation Functions for Deep-Learning Applications," in *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 2022.

- [20] M. Zhang, Y. Zhuo, C. Wang, M. Gao, Y. Wu, K. Chen, C. Kozyrakis, and X. Qian, "GraphP: Reducing Communication for PIM-based Graph Processing with Efficient Data Partition," in *Proceedings of the International Symposium on High-Performance Computer Architecture (HPCA)*, 2018.
- [21] G. Dai, T. Huang, Y. Chi, J. Zhao, G. Sun, Y. Liu, Y. Wang, Y. Xie, and H. Yang, "GraphH: A Processing-in-Memory Architecture for Large-Scale Graph Processing," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 38, no. 4, pp. 640–653, 2019.
- [22] Y. Zhuo, C. Wang, M. Zhang, R. Wang, D. Niu, Y. Wang, and X. Qian, "GraphQ: Scalable PIM-Based Graph Processing," in *Proceedings of the International Symposium on Microarchitecture (MICRO)*, 2019.
- [23] P. Yao, L. Zheng, Y. Huang, Q. Wang, C. Gui, Z. Zeng, X. Liao, H. Jin, and J. Xue, "ScalaGraph: A Scalable Accelerator for Massively Parallel Graph Processing," in *Proceedings of the International Symposium on High-Performance Computer Architecture (HPCA)*, 2022.
- [24] J. Ahn, S. Hong, S. Yoo, O. Mutlu, and K. Choi, "A Scalable Processing-in-Memory Accelerator for Parallel Graph Processing," in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2015.
- [25] M. Gao, J. Pu, X. Yang, M. Horowitz, and C. Kozyrakis, "TETRIS: Scalable and Efficient Neural Network Acceleration with 3D Memory," in *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2017.
- [26] D. Kim, J. Kung, S. Chai, S. Yalamanchili, and S. Mukhopadhyay, "Neurocube: A Programmable Digital Neuromorphic Architecture with High-Density 3D Memory," in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2016.
- [27] Y. Kwon, Y. Lee, and M. Rhu, "TensorDIMM: A Practical Near-Memory Processing Architecture for Embeddings and Tensor Operations in Deep Learning," in *Proceedings of the International Symposium on Microarchitecture (MICRO)*, 2019.
- [28] 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, "RecNMP: Accelerating Personalized Recommendation with Near-Memory Processing," in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2020.
- [29] J. Park, B. Kim, S. Yun, E. Lee, M. Rhu, and J. H. Ahn, "TRiM: Enhancing Processor-Memory Interfaces with Scalable Tensor Reduction in Memory," in *Proceedings of the International Symposium on Microarchitecture (MICRO)*, 2021.
- [30] B. Asgari, R. Hadidi, J. Cao, D. E. Shim, S.-K. Lim, and H. Kim, "FAFNIR: Accelerating Sparse Gathering by Using Efficient Near-Memory Intelligent Reduction," in *Proceedings of the International Symposium on High-Performance Computer Architecture (HPCA)*, 2021.
- [31] 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 *Proceedings of the International Conference on Computer Design (ICCD)*, 2016.
- [32] A. Boroumand, S. Ghose, G. F. Oliveira, and O. Mutlu, "Polynesia: Enabling High-Performance and Energy-Efficient Hybrid Transactional/Analytical Databases with Hardware/Software Co-Design," in *Proceedings of the International Conference on Data Engineering (ICDE)*, 2022.
- [33] Q. Deng, L. Jiang, Y. Zhang, M. Zhang, and J. Yang, "DrAcc: A DRAM Based Accelerator for Accurate CNN Inference," in *Design Automation Conference (DAC)*, 2018.
- [34] A. Shafiee, A. Nag, N. Muralimanohar, R. Balasubramonian, J. P. Strachan, M. Hu, R. S. Williams, and V. Srikumar, "ISAAC: A Convolutional Neural Network Accelerator with In-Situ Analog Arithmetic in Crossbars," in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2016.
- [35] G. Dai, Z. Zhu, T. Fu, C. Wei, B. Wang, X. Li, Y. Xie, H. Yang, and Y. Wang, "DIMMining: Pruning-Efficient and Parallel Graph Mining on Near-Memory-Computing," in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2022.
- [36] X. Xie, Z. Liang, P. Gu, A. Basak, L. Deng, L. Liang, X. Hu, and Y. Xie, "SpaceA: Sparse Matrix Vector Multiplication on Processing-in-Memory Accelerator," in *Proceedings of the International Symposium on High-Performance Computer Architecture (HPCA)*, 2021.
- [37] M. Lenjani, A. Ahmed, M. Stan, and K. Skadron, "Gearbox: A Case for Supporting Accumulation Dispatching and Hybrid Partitioning in PIM-based Accelerators," in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2022.
- [38] I. Fernandez, R. Quisilant, E. Gutiérrez, O. Plata, C. Giannoula, M. Alser, J. Gómez-Luna, and O. Mutlu, "NATSA: A Near-Data Processing Accelerator for Time Series Analysis," in *Proceedings of the International Conference on Computer Design (ICCD)*, 2020.
- [39] S. Li, D. Niu, K. T. Malladi, H. Zheng, B. Brennan, and Y. Xie, "DRISA: A DRAM-based Reconfigurable In-Situ Accelerator," in *Proceedings of the International Symposium on Microarchitecture (MICRO)*, 2017.
- [40] H. Shin, D. Kim, E. Park, S. Park, Y. Park, and S. Yoo, "McDRAM: Low Latency and Energy-Efficient Matrix Computations in DRAM," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 37, no. 11, pp. 2613–2622, 2018.
- [41] S. Cho, H. Choi, E. Park, H. Shin, and S. Yoo, "McDRAM v2: In-Dynamic Random Access Memory Systolic Array Accelerator to Address the Large Model," *IEEE Access*, vol. 8, pp. 135223–135243, 2020.
- [42] W. Huangfu, X. Li, S. Li, X. Hu, P. Gu, and Y. Xie, "MEDAL: Scalable DIMM Based Near Data Processing Accelerator for DNA Seeding Algorithm," in *Proceedings of the International Symposium on Microarchitecture (MICRO)*, 2019.
- [43] Y. Kwon, Y. Lee, and M. Rhu, "Tensor Casting: Co-Designing Algorithm-Architecture for Personalized Recommendation Training," in *Proceedings of the International Symposium on High-Performance Computer Architecture (HPCA)*, 2021.
- [44] Y. Lee, J. Chung, and M. Rhu, "SmartSAGE: Training Large-scale Graph Neural Networks using In-Storage Processing Architectures," in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2022.
- [45] UPMEM, 2022.
- [46] C. Lattner and V. Adve, "LLVM: A Compilation Framework for Lifelong Program Analysis and Transformation," in *International Symposium on Code Generation and Optimization (CGO)*, 2004.
- [47] UPMEM, "LLVM-based Compiler Stack Developed by UPMEM," 2021.
- [48] J. Gómez-Luna, I. E. Hajj, I. Fernandez, C. Giannoula, G. F. Oliveira, and O. Mutlu, "Benchmarking a New Paradigm: An Experimental Analysis of a Real Processing-in-Memory Architecture," in *arxiv.org*, 2021.
- [49] J. Nider, C. Mustard, A. Zoltan, J. Ramsden, L. Liu, J. Grossbard, M. Dashti, R. Jodin, A. Ghiti, J. Chauzi, and A. Fedorova, "A Case Study of Processing-in-Memory in Off-the-Shelf Systems," in *USENIX Annual Technical Conference (ATC)*, 2021.
- [50] C. Lim, S. Lee, J. Choi, J. Lee, S. Park, H. Kim, J. Lee, and Y. Kim, "Design and Analysis of a Processing-in-DIMM Join Algorithm: A Case Study with UPMEM DIMMs," in *Proceedings of the International Conference on Management of Data (SIGMOD)*, 2023.
- [51] L.-C. Chen, S.-Q. Yu, C.-C. Ho, Y.-H. Chang, D.-W. Chang, W.-C. Wang, and Y.-M. Chang, "RNA-seq Quantification on Processing in memory Architecture: Observation and Characterization NVM," in *Proceedings of the IEEE Non-Volatile Memory Systems and Applications Symposium (NVMSA)*, 2022.
- [52] J. Nider, J. Dagger, N. Gharavi, D. Ng, and A. Fedorova, "Bulk JPEG Decoding on In-Memory Processors," in *Proceedings of the 15th ACM International Conference on Systems and Storage (SYSTOR)*, 2022.
- [53] Samsung, "8Gb C-die DDR4 SDRAM x16," 2017.
- [54] UPMEM, "Instruction Set Architecture," 2021.
- [55] F. Devaux, "The True Processing In Memory Accelerator," in *Hot Chips: A Symposium on High Performance Chips*, 2019.
- [56] NVIDIA, "CUDA, release: 10.2.89," 2020.
- [57] N. Sakharnykh, "Everything You Need to Know about Unified Memory," in *NVIDIA GPU Technology Conference (GTC)*, 2018.
- [58] UPMEM, "UPMEM SDK," 2021.
- [59] T. Parr, *The Definitive ANTLR 4 Reference*. Pragmatic Bookshelf, 2012.
- [60] T. Parr, "ANTLR," 2022.
- [61] T. Parr, "ANTLR 4," 2022.
- [62] A. V. Aho, M. S. Lam, R. Sethi, and J. D. Ullman, *Compilers: Principles, Techniques, and Tools (2nd Edition)*. Addison-Wesley Longman Publishing Co., Inc., 2006.
- [63] UPMEM, "UPMEM DPU ABI," 2021.

- [64] UPMEM, “UPMEM Processing-In-Memory (PIM) Ultra-Efficient Acceleration for Data-Intensive Applications (White Paper),” 2022.
- [65] P. Rosenfeld, E. Cooper-Balis, and B. Jacob, “DRAMSim2: A Cycle Accurate Memory System Simulator,” *IEEE Computer Architecture Letters*, vol. 10, no. 1, pp. 16–19, 2011.
- [66] Y. Kim, W. Yang, and O. Mutlu, “Ramulator: A Fast and Extensible DRAM Simulator,” *IEEE Computer Architecture Letters*, vol. 15, no. 1, pp. 45–49, 2015.
- [67] “USIMM: The Utah Simulated Memory Module,” 2012.
- [68] A. Bakhoda, G. L. Yuan, W. W. L. Fung, H. Wong, and T. M. Aamodt, “Analyzing CUDA Workloads Using a Detailed GPU Simulator,” in *Proceedings of the International Symposium on Performance Analysis of Systems Software (ISPASS)*, 2009.
- [69] S. Rixner, W. Dally, U. Kapasi, P. Mattson, and J. Owens, “Memory Access Scheduling,” in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2000.
- [70] Intel, “Intel® 64 and IA-32 Architectures Software Developer’s Manual,” 2023.
- [71] “PrIM (Processing-In-Memory Benchmarks),” 2021.
- [72] N. Chatterjee, M. O’Connor, G. H. Loh, N. Jayasena, and R. Balasubramonia, “Managing DRAM Latency Divergence in Irregular GPGPU Applications,” in *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis (SC)*, 2014.
- [73] J. Meng, D. Tarjan, and K. Skadron, “Dynamic Warp Subdivision for Integrated Branch and Memory Divergence Tolerance,” in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2010.
- [74] W. W. Fung, I. Sham, G. Yuan, and T. M. Aamodt, “Dynamic Warp Formation and Scheduling for Efficient GPU Control Flow,” in *Proceedings of the International Symposium on Microarchitecture (MICRO)*, 2007.
- [75] W. W. Fung and T. M. Aamodt, “Thread Block Compaction for Efficient SIMT Control Flow,” in *Proceedings of the International Symposium on High-Performance Computer Architecture (HPCA)*, 2011.
- [76] M. Rhu and M. Erez, “The Dual-Path Execution Model for Efficient GPU Control Flow,” in *Proceedings of the International Symposium on High-Performance Computer Architecture (HPCA)*, 2013.
- [77] M. Rhu and M. Erez, “CAPRI: Prediction of Compaction-Adequacy for Handling Control-Divergence in GPGPU Architectures,” in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2012.
- [78] M. Rhu and M. Erez, “Maximizing SIMD Resource Utilization in GPGPUs with SIMD Lane Permutation,” in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2013.
- [79] A. ElTantawy, J. W. Ma, M. O’Connor, and T. M. Aamodt, “A Scalable Multi-Path Microarchitecture for Efficient GPU Control Flow,” in *Proceedings of the International Symposium on High-Performance Computer Architecture (HPCA)*, 2014.
- [80] NVIDIA, “NVIDIA Tesla V100 GPU Architecture (White Paper): “Independent Thread Scheduling,”” 2017.
- [81] Y. Kwon, G. Kim, N. Kim, W. Shin, J. Won, H. Joo, H. Choi, B. An, G. Shin, D. Yun, J. Kim, C. Kim, I. Kim, J. Park, C. Park, Y. Song, B. Yang, H. Lee, S. Park, W. Lee, S. Lee, K. Kim, D. Kwon, C. Jeong, J. Kim, E. Lim, and J. Chun, “Memory-Centric Computing with SK Hynix’s Domain-Specific Memory,” in *Hot Chips: A Symposium on High Performance Chips*, 2023.
- [82] ARM, “Arm DynamIQ Shared Unit-AE,” 2018.
- [83] Intel, “Introduction to Memory Bandwidth Allocation,” 2019.
- [84] NVIDIA, “Multi-Instance GPU User Guide,” 2020.
- [85] Linux, “Cgroups,” 2004.
- [86] K. Andi, “A NUMA API for Linux,” 2004.
- [87] J. Power, M. D. Hill, and D. A. Wood, “Supporting x86-64 Address Translation for 100s of GPU Lanes,” in *Proceedings of the International Symposium on High-Performance Computer Architecture (HPCA)*, 2014.
- [88] B. Pichai, L. Hsu, and A. Bhattacharjee, “Architectural Support for Address Translation on GPUs: Designing Memory Management Units for CPU/GPUs with Unified Address Spaces,” in *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2014.
- [89] B. Hyun, Y. Kwon, Y. Choi, J. Kim, and M. Rhu, “NeuMMU: Architectural Support for Efficient Address Translations in Neural Processing Units,” in *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2020.
- [90] Micron, “8Gb: x4, x8, x16 DDR4 SDRAM,” 2021.
- [91] D. Brooks, V. Tiwari, and M. Martonosi, “Wattch: A Framework for Architectural-Level Power Analysis and Optimizations,” in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2000.
- [92] S. Li, J. H. Ahn, R. D. Strong, J. B. Brockman, D. M. Tullsen, and N. P. Jouppi, “McPAT: An Integrated Power, Area, and Timing Modeling Framework for Multicore and Manycore Architectures,” in *Proceedings of the International Symposium on Microarchitecture (MICRO)*, 2009.
- [93] J. Leng, T. Hetherington, A. ElTantawy, S. Gilani, N. S. Kim, T. M. Aamodt, and V. J. Reddi, “GPUWattch: Enabling Energy Optimizations in GPGPUs,” in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2013.
- [94] V. Kandiah, S. Peverelle, M. Khairy, J. Pan, A. Manjunath, T. G. Rogers, T. M. Aamodt, and N. Hardavellas, “AccelWattch: A Power Modeling Framework for Modern GPUs,” in *Proceedings of the International Symposium on Microarchitecture (MICRO)*, 2021.
- [95] S. Hong and H. Kim, “An Integrated GPU Power and Performance Model,” in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2010.
- [96] 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 É. F. Zulian, “The Gem5 Simulator: Version 20.0+,” in *arxiv.org*, 2021.
- [97] D. Sanchez and C. Kozyrakis, “ZSim: Fast and Accurate Microarchitectural Simulation of Thousand-Core Systems,” in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2013.
- [98] R. Ubal, B. Jang, P. Mistry, D. Schaa, and D. Kaeli, “Multi2Sim: A Simulation Framework for CPU-GPU Computing,” in *Proceedings of the International Conference on Parallel Architectures and Compilation Techniques (PACT)*, 2012.
- [99] T. E. Carlson, W. Heirman, and L. Eeckhout, “Sniper: Exploring the Level of Abstraction for Scalable and Accurate Parallel Multi-Core Simulation,” in *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis (SC)*, 2011.
- [100] J. E. Miller, H. Kasture, G. Kurian, C. Gruenwald, N. Beckmann, C. Celio, J. Eastep, and A. Agarwal, “Graphite: A Distributed Parallel Simulator for Multicores,” in *Proceedings of the International Symposium on High-Performance Computer Architecture (HPCA)*, 2010.
- [101] M. T. Yourst, “PTLsim: A Cycle Accurate Full System x86-64 Microarchitectural Simulator,” in *Proceedings of the International Symposium on Performance Analysis of Systems Software (ISPASS)*, 2007.
- [102] H. S. Stone, “A Logic-in-Memory Computer,” *IEEE Transactions on Computers*, vol. C-19, no. 1, pp. 73–78, 1970.
- [103] R. Hwang, T. Kim, Y. Kwon, and M. Rhu, “Centaur: A Chiplet-Based, Hybrid Sparse-Dense Accelerator for Personalized Recommendations,” in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2020.
- [104] B. Kim, J. Park, E. Lee, M. Rhu, and J. H. Ahn, “TRIM: Tensor Reduction in Memory,” *IEEE Computer Architecture Letters*, vol. 20, no. 1, pp. 5–8, 2021.
- [105] H. Kim, H. Park, T. Kim, K. Cho, E. Lee, S. Ryu, H.-J. Lee, K. Choi, and J. Lee, “GradPIM: A Practical Processing-in-DRAM Architecture for Gradient Descent,” in *Proceedings of the International Symposium on High-Performance Computer Architecture (HPCA)*, 2021.
- [106] A. Augusta and S. Idreos, “JAFAR: Near-Data Processing for Databases,” in *Proceedings of the International Conference on Management of Data (SIGMOD)*, 2015.
- [107] S. Angizi, Z. He, A. S. Rakin, and D. Fan, “CMP-PIM: An Energy-Efficient Comparator-based Processing-In-Memory Neural Network Accelerator,” in *Design Automation Conference (DAC)*, 2018.

- [108] B. Y. Cho, J. Jung, and M. Erez, "Accelerating Bandwidth-Bound Deep Learning Inference with Main-Memory Accelerators," in *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis (SC)*, 2021.
- [109] S. Kang, S. Lee, B. Kim, H. Kim, K. Sohn, N. S. Kim, and E. Lee, "An FPGA-Based RNN-T Inference Accelerator with PIM-HBM," in *Proceedings of the International Symposium on Field-Programmable Gate Arrays (FPGA)*, 2022.
- [110] H. Asghari-Moghaddam, Y. H. Son, J. H. Ahn, and N. S. Kim, "Chameleon: Versatile and Practical Near-DRAM Acceleration Architecture for Large Memory Systems," in *Proceedings of the International Symposium on Microarchitecture (MICRO)*, 2016.
- [111] H. Asghari-Moghaddam, A. Farmahini-Farahani, K. Morrow, J. H. Ahn, and N. S. Kim, "Near-DRAM Acceleration with Single-ISA Heterogeneous Processing in Standard Memory Modules," *IEEE Micro*, vol. 36, no. 1, pp. 24–34, 2016.
- [112] A. Farmahini-Farahani, J. H. Ahn, K. Morrow, and N. S. Kim, "NDA: Near-DRAM Acceleration Architecture Leveraging Commodity DRAM Devices and Standard Memory Modules," in *Proceedings of the International Symposium on High-Performance Computer Architecture (HPCA)*, 2015.
- [113] M. Alian, S. W. Min, H. Asgharimoghaddam, A. Dhar, D. K. Wang, T. Roewer, A. McPadden, O. O'Halloran, D. Chen, J. Xiong, D. Kim, W.-m. Hwu, and N. S. Kim, "Application-Transparent Near-Memory Processing Architecture with Memory Channel Network," in *Proceedings of the International Symposium on Microarchitecture (MICRO)*, 2018.
- [114] M. Alian and N. S. Kim, "NetDIMM: Low-Latency Near-Memory Network Interface Architecture," in *Proceedings of the International Symposium on Microarchitecture (MICRO)*, 2019.
- [115] R. Hadidi, L. Nai, H. Kim, and H. Kim, "CAIRO: A Compiler-Assisted Technique for Enabling Instruction-Level Offloading of Processing-In-Memory," *ACM Transactions on Architecture and Code Optimization (TACO)*, vol. 14, no. 4, pp. 1–25, 2017.
- [116] B. Y. Cho, Y. Kwon, S. Lym, and M. Erez, "Near Data Acceleration with Concurrent Host Access," in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2020.
- [117] C. Giannoula, N. Vijaykumar, N. Papadopoulou, V. Karakostas, I. Fernandez, J. Gómez-Luna, L. Orosa, N. Koziris, G. Goumas, and O. Mutlu, "SynCron: Efficient Synchronization Support for Near-Data-Processing Architectures," in *Proceedings of the International Symposium on High-Performance Computer Architecture (HPCA)*, 2021.
- [118] S. Chen, Y. Jiang, C. Delimitrou, and J. F. Martínez, "PIMcloud: QoS-Aware Resource Management of Latency-Critical Applications in Clouds with Processing-in-Memory," in *Proceedings of the International Symposium on High-Performance Computer Architecture (HPCA)*, 2022.
- [119] 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 *Proceedings of the International Symposium on Microarchitecture (MICRO)*, 2017.
- [120] J. S. 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 *Proceedings of the International Symposium on High-Performance Computer Architecture (HPCA)*, 2019.
- [121] A. Olgun, M. Patel, A. G. Yağlıkçı, H. Luo, J. S. Kim, F. Nisa Bostancı, N. Vijaykumar, O. Ergin, and O. Mutlu, "QUAC-TRNG: High-Throughput True Random Number Generation Using Quadruple Row Activation in Commodity DRAM Chips," in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, 2021.
- [122] N. Hajinazar, G. F. Oliveira, S. Gregorio, J. a. D. Ferreira, N. M. Ghiasi, M. Patel, M. Alser, S. Ghose, J. Gómez-Luna, and O. Mutlu, "SIMDRAM: A Framework for Bit-Serial SIMD Processing Using DRAM," in *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2021.
- [123] X. Xin, Y. Zhang, and J. Yang, "ELP2IM: Efficient and Low Power Bitwise Operation Processing in DRAM," in *Proceedings of the International Symposium on High-Performance Computer Architecture (HPCA)*, 2020.
- [124] V. Seshadri, Y. Kim, C. Fallin, D. Lee, R. Ausavarungrinun, G. Pekhimenko, Y. Luo, O. Mutlu, P. B. Gibbons, M. A. Kozuch, and T. C. Mowry, "RowClone: Fast and Energy-Efficient In-DRAM Bulk Data Copy and Initialization," in *Proceedings of the International Symposium on Microarchitecture (MICRO)*, 2013.
- [125] K. K. Chang, P. J. Nair, D. Lee, S. Ghose, M. K. Qureshi, and O. Mutlu, "Low-Cost Inter-Linked Subarrays (LISA): Enabling Fast Inter-Subarray Data Movement in DRAM," in *Proceedings of the International Symposium on High-Performance Computer Architecture (HPCA)*, 2016.
- [126] J. Gómez-Luna, I. El Hajj, I. Fernandez, C. Giannoula, G. F. Oliveira, and O. Mutlu, "Benchmarking Memory-Centric Computing Systems: Analysis of Real Processing-in-Memory Hardware," in *International Green and Sustainable Computing Conference (IGSC)*, 2021.
- [127] C. Giannoula, I. Fernandez, J. Gómez-Luna, N. Koziris, G. Goumas, and O. Mutlu, "SparseP: Towards Efficient Sparse Matrix Vector Multiplication on Real Processing-In-Memory Architectures," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 6, no. 1, pp. 1–49, 2022.
- [128] G. F. Oliveira, J. Gómez-Luna, S. Ghose, A. Boroumand, and O. Mutlu, "Accelerating Neural Network Inference with Processing-in-DRAM: From the Edge to the Cloud," *IEEE Micro*, vol. 42, no. 6, pp. 25–38, 2022.
- [129] J. Gómez-Luna, Y. Guo, S. Brocard, J. Legriel, R. Cimadomo, G. F. Oliveira, G. Singh, and O. Mutlu, "An Experimental Evaluation of Machine Learning Training on a Real Processing-in-Memory System," in *arxiv.org*, 2022.
- [130] D. Lavenier, R. Cimadomo, and R. Jodin, "Variant Calling Parallelization on Processor-in-Memory Architecture," in *IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 2020.
- [131] D. Lavenier, J.-F. Roy, and D. Furodet, "DNA Mapping Using Processor-in-Memory Architecture," in *IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 2016.
- [132] S. Xu, X. Chen, Y. Wang, Y. Han, X. Qian, and X. Li, "PIMSim: A Flexible and Detailed Processing-in-Memory Simulator," *IEEE Computer Architecture Letters*, vol. 18, no. 1, pp. 6–9, 2018.
- [133] "ZSim+Ramulator - A Processing-in-Memory Simulation Framework," 2019.
- [134] G. Singh, J. Gómez-Luna, G. Mariani, G. F. Oliveira, S. Corda, S. Stuijk, O. Mutlu, and H. Corporaal, "NAPEL: Near-Memory Computing Application Performance Prediction via Ensemble Learning," in *Design Automation Conference (DAC)*, 2019.
- [135] X. Xie, P. Gu, J. Huang, Y. Ding, and Y. Xie, "MPU-Sim: A Simulator for In-DRAM Near-Bank Processing Architectures," *IEEE Computer Architecture Letters*, vol. 21, no. 1, pp. 1–4, 2021.
- [136] C. Yu, S. Liu, and S. Khan, "MultiPIM: A Detailed and Configurable Multi-Stack Processing-In-Memory Simulator," *IEEE Computer Architecture Letters*, vol. 20, no. 1, pp. 54–57, 2021.