

***SparseP*: Towards Efficient Sparse Matrix Vector Multiplication on Real Processing-In-Memory Systems**

CHRISTINA GIANNOULA, ETH Zürich, Switzerland and National Technical University of Athens, Greece

IVAN FERNANDEZ, ETH Zürich, Switzerland and University of Malaga, Spain

JUAN GÓMEZ-LUNA, ETH Zürich, Switzerland

NECTARIOS KOZIRIS, National Technical University of Athens, Greece

GEORGIOS GOUMAS, National Technical University of Athens, Greece

ONUR MUTLU, ETH Zürich, Switzerland

Several manufacturers have already started to commercialize near-bank Processing-In-Memory (PIM) architectures, after decades of research efforts. Near-bank PIM architectures place simple cores close to DRAM banks. Recent research demonstrates that they can yield significant performance and energy improvements in parallel applications by alleviating data access costs. Real PIM systems can provide high levels of parallelism, large aggregate memory bandwidth and low memory access latency, thereby being a good fit to accelerate the Sparse Matrix Vector Multiplication (SpMV) kernel. SpMV has been characterized as one of the most significant and thoroughly studied scientific computation kernels. It is primarily a memory-bound kernel with intensive memory accesses due its algorithmic nature, the compressed matrix format used, and the sparsity patterns of the input matrices given.

This paper provides the first comprehensive analysis of SpMV on a real-world PIM architecture, and presents *SparseP*, the first SpMV library for real PIM architectures. We make three key contributions. First, we implement a wide variety of software strategies on SpMV for a multithreaded PIM core, including (1) various compressed matrix formats, (2) load balancing schemes across parallel threads and (3) synchronization approaches, and characterize the computational limits of a single multithreaded PIM core. Second, we design various load balancing schemes across multiple PIM cores, and two types of data partitioning techniques to execute SpMV on thousands of PIM cores: (1) 1D-partitioned kernels to perform the complete SpMV computation only using PIM cores, and (2) 2D-partitioned kernels to strive a balance between computation and data transfer costs to PIM-enabled memory. Third, we compare SpMV execution on a real-world PIM system with 2528 PIM cores to an Intel Xeon CPU and an NVIDIA Tesla V100 GPU to study the performance and energy efficiency of various devices, i.e., both memory-centric PIM systems and conventional processor-centric CPU/GPU systems, for the SpMV kernel. *SparseP* software package provides 25 SpMV kernels for real PIM systems supporting the four most widely used compressed matrix formats, i.e., CSR, COO, BCSR and BCOO, and a wide range of data types. *SparseP* is publicly and freely available at <https://github.com/CMU-SAFARI/SparseP>. Our extensive evaluation using 26 matrices with various sparsity patterns provides new insights and recommendations for software designers and hardware architects to efficiently accelerate the SpMV kernel on real PIM systems.

Key Words: high-performance computing, HPC, sparse matrix-vector multiplication, SpMV, SpMV library, multicore, processing-in-memory, near-data processing, memory systems, data movement bottleneck, DRAM, benchmarking, real-system characterization, workload characterization

1 Introduction

Sparse Matrix Vector Multiplication (SpMV) is a fundamental linear algebra kernel for important applications from the scientific computing, machine learning, and graph analytics domains. In commodity systems, it has been repeatedly reported to achieve only a small fraction of the peak performance [59, 60, 63, 83, 84, 115, 123, 161, 245, 247, 248, 261] due to its algorithmic nature, the employed compressed matrix storage format, and the sparsity pattern of the input matrix. SpMV performs indirect memory references as a result of storing the matrix in a compressed format, and irregular memory accesses to the input vector due to sparsity. The matrices involved are very sparse, i.e., the vast majority of elements are zeros [60, 63, 83, 104, 121, 152, 201, 226, 250]. For

example, the matrices that represent Facebook’s and YouTube’s network connectivity contain 0.0003% [121, 152] and 2.31% [121, 226] non-zero elements, respectively. Therefore, in processor-centric systems, SpMV is a memory-bandwidth-bound kernel for the majority of real sparse matrices, and is bottlenecked by data movement between memory and processors [56, 59, 60, 62, 63, 82–84, 94, 115, 123, 137, 138, 161, 197, 245, 247, 248, 259].

One promising way to alleviate the data movement bottleneck is the Processing-In-Memory (PIM) paradigm [1–4, 6–9, 11–13, 16, 23, 29, 30, 33, 39, 41, 43, 46, 47, 52, 54, 55, 58, 66–69, 73–79, 81, 82, 90, 94, 96–98, 101, 102, 109, 112, 122, 124, 126, 127, 130–132, 141–143, 145, 151, 153, 156, 157, 166, 167, 176–179, 181, 191, 192, 194, 195, 209, 217, 218, 220–225, 227, 229, 233, 249, 253, 255, 258, 260, 266–268, 271, 278]. PIM moves computation close to application data by equipping memory chips with processing capabilities [176, 178]. Prior works [3, 30–32, 47, 52, 57, 67, 68, 76–78, 81, 110, 111, 133, 150, 166, 167, 167, 180, 181, 206, 269, 271, 281] propose PIM architectures wherein a processor logic layer is tightly integrated with DRAM memory layers using 2.5D/3D-stacking technologies [95, 120, 149]. Nonetheless, the 2.5D/3D integration itself might not always be able to provide significantly higher memory bandwidth for processors than standard DRAM [16, 151]. To provide even higher bandwidth for the in-memory processors, *near-bank* PIM designs have been explored [1, 9, 16, 44–46, 55, 82, 89, 90, 94, 140, 145, 151, 156, 179, 199, 213, 222, 241, 266]. *Near-bank* PIM designs tightly couple a PIM core with each DRAM bank, exploiting bank-level parallelism to expose high on-chip memory bandwidth of standard DRAM to processors. Moreover, manufacturers of near-bank PIM architectures avoid disturbing the key components (i.e., subarray and bank) of commodity DRAM to provide a cost-efficient and practical way for silicon materialization. Two *real* near-bank PIM architectures are Samsung’s FIMDRAM [145, 151] and the UPMEM PIM system [55, 82, 94, 240].

Most near-bank PIM architectures [16, 44, 45, 55, 82, 89, 94, 140, 145, 151, 179, 199, 241] support several PIM-enabled memory chips connected to a host CPU via memory channels. Each memory chip comprises multiple PIM cores, which are low-area and low-power cores with relatively low computation capability [82, 94], and each of them is located close to a DRAM bank [16, 44, 45, 55, 82, 89, 94, 140, 145, 151, 179, 199, 241]. Each PIM core can access data located on their local DRAM banks, and typically there is no direct communication channel among PIM cores. Overall, near-bank PIM architectures provide high levels of parallelism and very large memory bandwidth, thereby being a very promising computing platform to accelerate memory-bound kernels. Recent works leverage near-bank PIM architectures to provide high performance and energy benefits on bioinformatics [82, 94, 147, 148], skyline computation [282], compression [185] and neural network [44, 82, 89, 94, 151] kernels. A recent study [82, 94] provides PrIM benchmarks [87], which are a collection of 16 kernels for evaluating near-bank PIM architectures, like the UPMEM PIM system. However, there is *no* prior work to thoroughly study the widely used, memory-bound SpMV kernel on a real PIM system.

Our work is the first to efficiently map the SpMV execution kernel on near-bank PIM systems, and understand its performance implications on a real PIM system. Specifically, our **goal** in this work is twofold: (i) design efficient SpMV algorithms to accelerate this kernel in current and future PIM systems, while covering a wide variety of sparse matrices with diverse sparsity patterns, and (ii) provide an extensive characterization analysis of the widely used SpMV kernel on a real PIM architecture. To this end, we provide a wide variety of SpMV implementations for real PIM architectures, and conduct a rigorous experimental analysis of SpMV kernels in the UPMEM PIM system, the first publicly-available real-world PIM architecture.

We present the *SparseP* library [88] that includes 25 SpMV kernels for real PIM systems, supporting various (1) data types, (2) data partitioning techniques of the sparse matrix to PIM-enabled memory, (3) compressed matrix formats, (4) load balancing schemes across PIM cores, (5) load

balancing schemes across threads of a multithreaded PIM core, and (6) synchronization approaches among threads within PIM core. We support a wide range of data types, i.e., 8-bit integer, 16-bit integer, 32-bit integer, 64-bit integer, 32-bit float and 64-bit float data types to cover a wide variety of real-world applications that employ SpMV as their underlying kernel. We design two types of well-crafted data partitioning techniques: (i) the 1D partitioning technique to perform the complete SpMV computation only using PIM cores, and (ii) the 2D partitioning technique to strive a balance between computation and data transfer costs to PIM-enabled memory. In the 1D partitioning technique, the matrix is horizontally partitioned across PIM cores, and the *whole* input vector is copied into the DRAM bank of *each* PIM core, while PIM cores directly compute the elements of the final output vector. In the 2D partitioning technique, the matrix is split in 2D tiles, the number of which is equal to the number of PIM cores, and a *subset* of the elements of the input vector is copied into the DRAM bank of each PIM core. However, in the 2D partitioning technique, PIM cores create a large number of partial results for the elements of the output vector which are gathered and merged by the host CPU cores to assemble the final output vector. We support the most popular compressed matrix formats, i.e., CSR [25, 205], COO [205, 216], BCSR [114], BCOO [205], and for each compressed format we implement various load balancing schemes across PIM cores to provide efficient SpMV execution for a wide variety of sparse matrices with diverse sparsity patterns. Finally, we design several load balancing schemes and synchronization approaches among parallel threads within a PIM core to cover a variety of real PIM systems that provide multithreaded PIM cores.

We conduct an extensive characterization analysis of *SparseP* kernels on the UPMEM PIM system [55, 82, 94, 241] analyzing the SpMV execution using (1) one single multithreaded PIM core, (2) thousands of PIM cores, and (3) comparing it with that achieved on conventional processor-centric CPU and GPU systems. First, we characterize the limits of a single multithreaded PIM core, and show that (i) high operation imbalance across threads of a PIM core can impose high overhead in the core pipeline, and (ii) fine-grained synchronization approaches to increase parallelism cannot outperform a coarse-grained approach, if PIM hardware serializes accesses to the local DRAM bank. Second, we analyze the end-to-end SpMV execution of 1D and 2D partitioning techniques using thousands of PIM cores. Our study indicates that the performance (i) of the 1D partitioning technique is limited by data transfer costs to *broadcast* the whole input vector into *each* DRAM bank of PIM cores, and (ii) of the 2D partitioning technique is limited by data transfer costs to *gather* partial results for the elements of the output vector from PIM-enabled memory to the host CPU. Such data transfers incur high overheads, because they take place via the narrow memory bus. In addition, our detailed study across a wide variety of compressed matrix formats and sparse matrices with diverse sparsity patterns demonstrates that (i) the compressed matrix format determines the data partitioning strategy across DRAM banks of PIM-enabled memory, thereby affecting the computation balance across PIM cores with corresponding performance implications, and (ii) there is *no one-size-fits-all* solution. The load balancing scheme across PIM cores (and across threads within a PIM core) and data partitioning technique that provides the best-performing SpMV execution depends on the characteristics of the input matrix and the underlying PIM hardware. Finally, we compare the SpMV execution on a state-of-the-art UPMEM PIM system with 2528 PIM cores to state-of-the-art CPU and GPU systems, and observe that SpMV on the UPMEM PIM system achieves a much higher fraction of the machine’s peak performance compared to that on the state-of-the-art CPU and GPU systems. Our extensive evaluation provides programming recommendations for software designers, and suggestions and hints for hardware and system designers of future PIM systems.

Our most significant recommendations for PIM software designers are:

- (1) Design algorithms that provide high load balance across threads of PIM core in terms of computations, loop control iterations, synchronization points and memory accesses.

- (2) Design compressed data structures that can be effectively partitioned across DRAM banks, with the goal of providing high computation balance across PIM cores.
- (3) Design *adaptive* algorithms that trade off computation balance across PIM cores for lower data transfer costs to PIM-enabled memory, and adapt their configuration to the particular patterns of each input given, as well as the characteristics of the PIM hardware.

Our most significant suggestions for PIM hardware and system designers are:

- (1) Provide low-cost synchronization support and hardware support to enable concurrent memory accesses by multiple threads to the local DRAM bank to increase parallelism in a multithreaded PIM core.
- (2) Optimize the broadcast collective operation in data transfers from main memory to PIM-enabled memory to minimize overheads of copying the input data into all DRAM banks in the PIM system.
- (3) Optimize the gather collective operation *at DRAM bank granularity* for data transfers from PIM-enabled memory to the host CPU to minimize overheads of retrieving the output results.
- (4) Design high-speed communication channels and optimized libraries for data transfers to/from thousands of DRAM banks of PIM-enabled memory.

Our *SparseP* software package is freely and publicly available [88] to enable further research on SpMV in current and future PIM systems. The main contributions of this work are as follows:

- We present *SparseP*, the first open-source SpMV software package for real PIM architectures. *SparseP* includes 25 SpMV kernels, supporting the four most widely used compressed matrix formats and a wide range of data types. *SparseP* is publicly available at [88], and can be useful for researchers to improve multiple aspects of future PIM hardware and software.
- We perform the first comprehensive study of the widely used SpMV kernel on the UPMEM PIM architecture, the first real commercial PIM architecture. We analyze performance implications of SpMV PIM execution using a wide variety of (1) compressed matrix formats, (2) data types, (3) data partitioning and load balancing techniques, and (4) 26 sparse matrices with diverse sparsity patterns.
- We compare the performance and energy of SpMV on the state-of-the-art UPMEM PIM system with 2528 PIM cores to state-of-the-art CPU and GPU systems. SpMV execution achieves less than 1% of the peak performance on processor-centric CPU and GPU systems, while it achieves on average 51.7% of the peak performance on the UPMEM PIM system, thus better leveraging the computation capabilities of underlying hardware. The UPMEM PIM system also provides high energy efficiency on the SpMV kernel.

2 Background and Motivation

2.1 Sparse Matrix Vector Multiplication (SpMV)

The SpMV kernel multiplies a sparse matrix of size $M \times N$ with a dense input vector of size $1 \times N$ to compute an output vector of size $M \times 1$. The SpMV kernel is widely used in a variety of applications including graph processing [22, 35, 80, 121], neural networks [99, 100, 160, 279], machine learning [92, 93, 159, 184, 272, 275], and high performance computing [26, 45, 56, 59, 64, 65, 105]. These applications involve matrices with very high sparsity [60, 63, 83, 104, 121, 152, 201, 226, 250], i.e., a large fraction of zero elements. Thus, using a compression scheme is a straightforward approach to avoid unnecessarily storing zero elements and performing computations on them. For general sparse matrices, the most widely used storage format is the Compressed Sparse Row (CSR) format [25, 205]. Figure 1 presents an example of a compressed matrix using the CSR format (left), and the CSR-based SpMV execution (right), assuming an input vector x and an output vector y .




Fig. 1. (a) CSR representation of a sparse matrix. (b) CSR-based SpMV implementation.

2.1.1 Compressed Matrix Storage Formats

Several prior works [17, 21, 25, 36, 106, 114, 135, 137–139, 146, 163, 165, 169, 170, 174, 204, 205, 211, 214, 216, 248, 261, 263] propose compressed storage formats for sparse matrices, which are typically of two types [121]. The first approach is to design general purpose compressed formats, such as CSR [25, 205], CSR5 [163], COO [205, 216], BCSR [114], and BCOO [205]. Such encodings are general in applicability and are highly-efficient in storage. The second approach is to leverage a certain known structure in a given type of sparse matrix. For example, the DIA format [17] is effective in matrices where the non-zero elements are concentrated along the diagonals of the matrix. Such encodings aim to improve performance of sparse matrix computations by specializing to particular matrix patterns, but they sacrifice generality. In this work, we explore with the four most widely used *general* compressed formats (Figure 2), which we describe in more detail next.




Fig. 2. (a) SpMV with a dense matrix representation, and (b) CSR, (c) COO, (d) BCSR, (e) BCOO formats.

Compressed Sparse Row (CSR) [25, 205]. The CSR format (Figure 2b) sequentially stores values in a row-wise order. A column index array (`colind[]`) and a value array (`values[]`) store the column index and value of each non-zero element, respectively. An array, named `rowptr[]`, stores the location of the first non-zero element of each row within the `values[]` array. The values of an adjacent pair of the `rowptr[]` array, i.e., `rowptr[i, i+1]`, represent a slice of the `colind[]` and `values[]` arrays. The corresponding slice of the `colind[]` and `values[]` arrays stores the column indices and the values of the non-zero elements, respectively, for the i -th row of the original matrix.

Coordinate Format (COO) [205, 216]. The COO format (Figure 2c) stores the non-zero elements as a series of tuples (`tuples[]` array). Each tuple includes the row index, column index, and value of the non-zero element.

Block Compressed Sparse Row (BCSR) [114]. The BCSR format (Figure 2d) is a block representation of CSR. Instead of storing and indexing single non-zero elements, BCSR stores and indexes $r \times c$ sub-blocks with at least one non-zero element. The original matrix is split into $r \times c$ sub-blocks. Figure 2d shows an example of BCSR assuming 4×4 sub-blocks. The original matrix of Figure 2a is split into four sub-blocks, and two of them (highlighted with red color) contain at least one non-zero element. The `bvalues[]` array stores the values of all the *non-zero sub-blocks* of the original matrix. Each non-zero sub-block is stored in the `bvalues[]` array with a dense representation, i.e., padding with zero values when needed. The `bcolind[]` array stores the block-column index of

each non-zero sub-block. The `browptr[]` array stores the location of the first non-zero sub-block of each block row within the `bcolind[]` array, assuming a block row represents r consecutive rows of the original matrix, where r is the vertical dimension of the sub-block.

Block Coordinate Format (BCOO) [205]. The BCOO format is the block counterpart of COO. The `browind[]`, `bcolind[]` and `bvalues[]` arrays store the row indices, column indices and values of the non-zero sub-blocks, respectively. Figure 2e shows an example of BCOO, assuming 4×4 sub-blocks.

2.1.2 SpMV in Processor-Centric Systems

Many prior works [59–61, 63, 84, 115, 123, 245, 247, 248, 250, 261] generally show that SpMV performs poorly on commodity CPU and GPU systems, and achieves a small fraction of the peak performance (e.g., 10% of the peak performance [247]) due to its algorithmic nature, the employed compressed matrix storage format and the sparsity pattern of the matrix.

The SpMV kernel is highly bottlenecked by the memory subsystem in processor-centric CPU and GPU systems due to three reasons. First, due to its algorithmic nature there is *no* temporal locality in the input matrix. Unlike traditional algebra kernels like Matrix Matrix Multiplication or LU decomposition, the elements of the matrix in SpMV are used only *once* [83, 84]. Second, due to the sparsity of the matrix, the matrix is stored in a compressed format (e.g., CSR) to avoid unnecessary computations and data accesses. Specifically, the non-zero elements of the matrix are stored contiguously in memory, while additional data structures assist in the proper traversal of the matrix, i.e., to discover the positions of the non-zero elements. For example, CSR uses the `rowptr[]` and `colind[]` arrays to discover the positions of the non-zero elements of the matrix. These additional data structures cause additional memory access operations, memory bandwidth pressure and contention with other requests in the memory subsystem. Third, due to the sparsity of the input matrix, SpMV causes irregular memory accesses to the elements of the input vector x . The memory accesses to the elements of the input vector are input driven, i.e., they follow the sparsity pattern of the input matrix. This irregularity results to poor data locality on the elements of the input vector and expensive data accesses, because it increases the average access latency due to a high number of cache misses on commodity systems with deep cache hierarchies [83, 84]. As a result, memory-centric near-bank PIM systems constitute a better fit for the widely used SpMV kernel, because they provide high levels of parallelism, large aggregate memory bandwidth and low memory access latency [16, 82, 94, 151, 241].

2.2 Near-Bank PIM Systems

Figure 3 shows the baseline organization of a near-bank PIM system that we assume in this work. The PIM system consists of a host CPU, standard DRAM memory modules, and PIM-enabled memory modules. PIM-enabled modules are connected to the host CPU using one or more memory channels, and include multiple PIM chips. A PIM chip (Figure 3 right) tightly integrates a low-area PIM core with a DRAM bank. We assume that each PIM core can additionally include a small private instruction memory and a small data (scratchpad or cache) memory. PIM cores can access data located on their local DRAM bank, and typically there is no direct communication channel among PIM cores. The DRAM banks of PIM chips are accessible by the host CPU for copying input data and retrieving results via the memory bus.

2.2.1 The UPMEM PIM Architecture

The UPMEM PIM system [55, 82, 94] includes the host CPU with standard main memory, and UPMEM PIM modules. An UPMEM PIM module is a standard DDR4-2400 DIMM [119] with 2 ranks.




Fig. 3. High-level organization of a near-bank PIM architecture.

Each rank contains 64 PIM cores, which are called DRAM Processing Units (DPUs). In the current UPMEM PIM system, there are 20 double-rank PIM DIMMs with 2560 DPUs.¹

DPU Architecture and Interface. Each DPU has exclusive access to a 24-KB instruction memory, called **IRAM**, a 64-KB scratchpad memory, called **WRAM**, and a 64-MB DRAM bank, called **MRAM**. A DPU is a multithreaded in-order 32-bit RISC core that can potentially reach 500 MHz [241]. The DPU has 24 hardware threads, each of which has 24 32-bit general purpose registers. The DPU pipeline has 14 stages, and supports a single cycle 8x8-bit multiplier. Multiplications on 64-bit integers, 32-bit floats and 64-bit floats are not supported in hardware, and require longer routines with a large number of operations [82, 94, 241]. Threads share the IRAM and WRAM, and can access the MRAM by executing transactions at 64-bit granularity via a DMA engine, i.e., data can be accessed from/to MRAM as a multiple of 8 bytes, up to 2048 bytes. MRAM transactions are serialized in the DMA engine. The ISA provides DMA instructions to move instructions from MRAM to IRAM, or data between MRAM and WRAM. The DPU accesses the WRAM through 8-, 16-, 32- and 64-bit load/store instructions. DPUs use the *Single Program Multiple Data* programming model, where software threads, called **tasklets**, execute the same code, but operate in different pieces of data, and can execute different control-flow paths during runtime. Tasklets can synchronize using mutexes, barriers, handshakes and semaphores provided by the UPMEM runtime library.

CPU-DPU Data Transfers. Standard main memory and PIM-enabled memory have different data layouts. The UPMEM SDK [242] has a transposition library to execute necessary data shuffling when moving data between main memory and MRAM banks of PIM-enabled memory modules via a programmer-transparent way. The CPU-DPU and DPU-CPU data transfers can be performed in parallel, i.e., concurrently across multiple MRAM banks, with the limitation that *the transfer sizes from/to all MRAM banks need to be the same*. The UPMEM SDK provides two options: (i) perform parallel transfers to all MRAM banks of all ranks, or (ii) iterate over each rank to perform parallel transfers to MRAM banks of the same rank, and serialize data transfers across ranks.

3 The SparseP Library

This section describes the parallelization techniques that we explore for SpMV on real PIM architectures, and presents the SpMV implementations of our *SparseP* package. Section 3.1 describes SpMV execution on a real PIM system. Section 3.2 presents an overview of the data partitioning techniques that we explore. Section 3.3 and Section 3.4 describe in detail the parallelization techniques across PIM cores, and across threads within a PIM core, respectively. Section 3.5 describes the kernel implementation for all compressed matrix storage formats.

¹There are thirty two faulty DPUs in the system where we run our experiments. They cannot be used and do not affect the correctness of our results, but take away from the system's full computational power of 2560 DPUs.

3.1 SpMV Execution on a PIM System

Figure 4 shows the SpMV execution on a real PIM system, which is broken down in four steps: (1) the time to load the input vector into DRAM banks of PIM-enabled memory (**load**), (2) the time to execute the SpMV kernel on PIM cores (**kernel**), (3) the time to retrieve from DRAM banks to the host CPU results for the output vector (**retrieve**), and (4) the time to merge partial results and assemble the final output vector on the host CPU (**merge**). In our analysis, we omit the time to load the matrix into PIM-enabled memory, since this step can typically be hidden in real-world applications (it can be overlapped with other computation performed by the application or amortized if the application performs multiple SpMV iterations on the same matrix).




Fig. 4. Execution of the SpMV kernel on a real PIM system.

3.2 Overview of Data Partitioning Techniques

To parallelize the SpMV kernel, we implement well-crafted data partitioning schemes to split the matrix across multiple DRAM banks of PIM cores. *SparseP* supports two general types of data partitioning techniques, shown in Figure 5.




Fig. 5. Data partitioning techniques of the *SparseP* package.

First, we provide an 1D partitioning technique (Figure 5a), where the matrix is horizontally partitioned across PIM cores, and the whole input vector is copied into the DRAM bank of each PIM core. With the 1D partitioning technique, almost the entire SpMV computation is performed using only PIM cores, since the merge step in the host CPU is negligible: a very small number of partial results is created, i.e., only for a few rows that are split across neighboring PIM cores. Thus, the number of partial elements of the output vector is at most equal to the number of PIM cores used. Second, we provide a 2D partitioning technique (Figure 5b), where the matrix is partitioned into 2D tiles, the number of which is equal to the number of PIM cores. With the 2D partitioning technique, we aim to strive a balance between computation and data transfer costs, since only a subset of

the elements of the input vector is copied into the DRAM bank of each PIM core. However, PIM cores assigned to tiles that horizontally overlap, i.e., tiles that share the same rows of the original matrix (rows that are split across multiple tiles), produce *many* partial results for the elements of the output vector. These partial results are transferred to the host CPU, and merged by CPU cores, which assemble the final output vector. In the *SparseP* library, the merge step performed by the CPU cores is parallelized using the OpenMP API [51].

In both data partitioning schemes, matrices are stored in a row-sorted way, i.e., the non-zero elements are sorted in increasing order of their row indices. Therefore, each PIM core computes results for a *continuous* subset of elements of the output vector. This way we minimize data transfer costs, since we only transfer necessary data to the host CPU, i.e., *the values* of the elements of the output vector produced at PIM cores. If each PIM core instead computed results for a *non-continuous* subset of elements of the output vector, an additional array *per core*, which would store *the indices* of the *non-continuous* elements within the output vector, would need to be transferred to the host CPU, causing additional data transfer overheads.

3.3 Parallelization Techniques Across PIM Cores

To parallelize SpMV across multiple PIM cores *SparseP* supports various parallelization schemes for both 1D and 2D partitioning techniques.

3.3.1 1D Partitioning Technique

To efficiently parallelize SpMV across multiple PIM cores via the 1D partitioning technique, *SparseP* provides various load balancing schemes for each supported compressed matrix format. Figure 6 presents an example of parallelizing SpMV across multiple PIM cores using load balancing schemes for the CSR and COO formats. For the CSR and COO formats, we balance either the rows, such that each PIM core processes almost the same number of rows, or the non-zero elements, such that each PIM core processes almost the same number of non-zero elements. In the CSR format, since the matrix is stored in row-order, i.e., the `rowptr[]` array stores the index pointers of the non-zero elements of *each row*, and thus balancing the non-zero elements across PIM cores is performed at row granularity. In the COO format, the matrix is stored in non-zero order using the `tuples[]` array, and thus balancing the non-zero elements can be performed either at row granularity, or by splitting a row across two neighboring PIM cores to provide a near-perfect non-zero element balance across cores. In the latter case, as mentioned, a small number of partial results for the output vector is merged by the host CPU: if the row is split between two neighboring PIM cores at most one element needs to be accumulated at the host CPU cores.




Fig. 6. Load balancing schemes across PIM cores for the CSR (left) and COO (right) formats with the 1D partitioning technique. The colored cells of the matrix represent non-zero elements.

Figure 7 presents an example of parallelizing SpMV across multiple PIM cores using load balancing schemes of the BCSR and BCOO formats. In Figure 7, the cells of the matrix represent sub-blocks of size 4x4: the *grey* cells represent sub-blocks that do not have *any* non-zero element, and the *colored* cells represent sub-blocks that have k non-zero elements, where k is the number shown inside the colored cell. In the BCSR and BCOO formats, since the matrix is stored in sub-blocks of non-zero elements, we balance either the blocks, such that each PIM core processes almost the same number of blocks, or the non-zero elements, such that each PIM core processes almost the same number of non-zero elements. Similarly to CSR, in the BCSR format, the matrix is stored in block-row-order, i.e., the `browptr[]` array stores the index pointers of the non-zero blocks of *each* block row (recall that a block row represents r consecutive rows of the original matrix, where r is the vertical dimension of the sub-block), and thus balancing the blocks or the non-zero elements across cores is limited to be performed at block-row granularity. In the BCOO format, given that a block-row might be split across two PIM cores, a small number of partial results for the output vector is merged by the host CPU: between two neighboring PIM cores at most block size r elements (r is the vertical dimension of the block size) might need to be accumulated at the host CPU cores.




Fig. 7. Load balancing schemes across PIM cores for the BCSR (left) and BCOO (left) formats with the 1D partitioning technique. The cells of the matrix represent sub-blocks of size 4x4. The colored cells of the matrix represent non-zero sub-blocks, and the number inside a colored cell describes the number of non-zero elements of the corresponding sub-block.

3.3.2 2D Partitioning Technique

SparseP includes three 2D partitioning techniques, shown in Figure 8:

- (1) **equally-sized** (Figure 8a): The 2D tiles are statically created to have the same height and width. This way the subsets of the elements for the input and output vectors have the same sizes across all PIM cores.
- (2) **equally-wide** (Figure 8b): The 2D tiles have the same width and variable height. This way the subset of the elements for the input vector has the same size across PIM cores, while the subset of the elements for the output vector varies across PIM cores. We balance the non-zero elements across the tiles of the *same* vertical partition, such that we can provide high non-zero element balance across PIM cores assigned to the same vertical partition.
- (3) **variable-sized** (Figure 8c): The 2D tiles have both variable width and height. We balance the non-zero elements both across the vertical partitions and across the tiles of the *same* vertical partition. This way we can provide high non-zero element balance across all PIM cores.

SparseP provides various load balancing schemes across PIM cores in the *equally-wide* and *variable-sized* techniques. In the *equally-wide* technique, for the CSR and COO formats, we balance




Fig. 8. The 2D partitioning techniques of *SparseP* package assuming 4 PIM cores and 2 vertical partitions.

the non-zero elements across the tiles of the same vertical partition. Load balancing in the CSR format is performed at row-granularity, i.e., splitting the `rowptr[]` array across PIM cores. For the BCSR and BCOO formats, we balance either the blocks or the non-zero elements across the tiles of the same vertical partition. Load balancing in the BCSR format is performed at block-row granularity, i.e., splitting the `browptr[]` array across PIM cores. In the *variable-sized* technique, we first balance the non-zero elements across the vertical partitions, such that the vertical partitions include the same number of non-zero elements. Then, across the tiles of the same vertical partition, we balance the non-zero elements for the CSR (at row-granularity) and COO formats, and either the blocks or the non-zero elements for the BCSR (at block-row granularity) and BCOO formats.

Table 1 summarizes the parallelization approaches across PIM cores. Please also see Appendix C for all SpMV kernels provided by the *SparseP* software package. All kernels support a wide range of data types, i.e., 8-bit integer (**int8**), 16-bit integer (**int16**), 32-bit integer (**int32**), 64-bit integer (**int64**), 32-bit float (**fp32**), and 64-bit float (**fp64**) data types.

3.4 Parallelization Techniques Across Threads within a PIM Core

PIM cores can support multiple hardware threads to exploit high memory bank bandwidth [82, 94]. To parallelize SpMV across multiple threads within a multithreaded PIM core *SparseP* supports various load balancing schemes for each compressed matrix format, and three synchronization approaches to ensure correctness among threads of a PIM core.

3.4.1 Load Balancing Approaches

In a similar way as explained in Figure 6, for the CSR and COO formats, we balance either the rows, such that each thread processes almost the same number of rows, or the non-zero elements, such that each thread processes almost the same number of non-zero elements. In the CSR format, matrix is stored in row-order, and thus load balancing across threads is performed at row granularity. In the UPMEM PIM system, elements of the output vector are accessed at 64-bit granularity in DRAM memory. Thus, when balancing is performed at row granularity, we assign rows to threads in chunks of $8/\text{sizeof}(\text{data_type})$ to ensure 8-byte alignment on the elements of the output vector. In the COO format, balancing the non-zero elements can be performed either at row granularity or by splitting the row between threads, i.e., providing an almost perfect non-zero balance across threads. In the latter case, synchronization among threads for write accesses on the elements of the output vector can be implemented with three synchronization approaches described in Section 3.4.2.

For the BCSR and BCOO formats, we balance either the blocks, such that each thread processes almost the same number of blocks, or the non-zero elements, such that each thread processes almost the same number of non-zero elements. In the BCSR format, the matrix is stored in block-row order, and thus load balancing across threads is performed at block row granularity. For both formats, the block sizes are *configurable* in *SparseP*. In our evaluation, we use block sizes of 4×4 , since these are

Partitioning Technique	Compressed Format	Load Balancing Across PIM Cores
1D	CSR	rows (CSR.row) nnz* (CSR.nnz)
	COO	rows (COO.row) nnz* (COO.nnz-rgrn) nnz (COO.nnz)
	BCSR	blocks [†] (BCSR.block) nnz [†] (BCSR.nnz)
	BCOO	blocks (BCOO.block) nnz (BCOO.nnz)
2D <i>equally-sized</i>	CSR (DCSR)	-
	COO (DCOO)	-
	BCSR (DBC SR)	-
	BCOO (DBCOO)	-
2D <i>equally-wide</i>	CSR (RBD CSR)	nnz*
	COO (RBD COO)	nnz
	BCSR	blocks [†] (RBD BCSR) nnz [†]
	BCOO	blocks (RBD BCOO) nnz
2D <i>variable-sized</i>	CSR (B DCSR)	nnz*
	COO (B DCOO)	nnz
	BCSR	blocks [†] (B DBCSR) nnz [†]
	BCOO	blocks (B DBCCOO) nnz

Table 1. Parallelization techniques across PIM cores of the *SparseP* library. *: row-granularity, [†]: block-row-granularity

the most common dimensions to cover various sparse matrices [14, 63, 123]. In the UPMEM PIM architecture, elements of the output vector are accessed at 64-bit granularity. Therefore, for the BCSR format, with an 8-bit integer data type and small block sizes (4x4 or smaller), threads use synchronization primitives to ensure correctness when writing the elements of the output vector. This is because different threads may write to the same 64-bit-aligned DRAM memory location. Synchronization among threads for writes to the elements of the output vector is necessary for all configurations of the BCOO format, and can be implemented with three approaches described next.

3.4.2 Synchronization Approaches

SparseP provides three synchronization approaches.

- (1) **Coarse-Grained Locking (lb-cg)**. One global mutex protects the elements of the entire output vector.
- (2) **Fine-Grained Locking (lb-fg)**. Multiple mutexes protect the elements of the output vector. *SparseP* associates mutexes to the elements of the output vector in a round-robin manner. The UPMEM API supports up to 56 mutexes [242]. In our evaluation, we use 32 mutexes such that we can find the corresponding mutex for a particular element of the output vector only with a shift operation on the MRAM address, avoiding costly division operations.
- (3) **Lock-Free (lf)**. Since the formats are row-sorted or block-row-sorted, race conditions in the elements of the output vector arise *only in a few elements*, i.e., either when a row (or a block

row for BCSR/BCOO) is split across threads, or when continuous elements of the output vector processed by different threads belong to the same 64-bit-aligned DRAM location in the UPMEM PIM system. In our proposed lock-free approach, threads temporarily store partial results for these few elements in the data (scratchpad) memory (i.e., WRAM in the UPMEM PIM system), and later one single thread merges the partial results, and writes the final result for the corresponding element of the output vector to the DRAM bank.

Table 2 summarizes the parallelization techniques across threads of a PIM core. All kernels support a wide range of data types, i.e., 8-bit integer (**int8**), 16-bit integer (**int16**), 32-bit integer (**int32**), 64-bit integer (**int64**), 32-bit float (**fp32**), and 64-bit float (**fp64**) data types.

Compressed Format	Load Balancing Across Threads	Synchronization Approach
CSR	rows (CSR.row) nnz* (CSR.nnz)	- -
COO	rows (COO.row) nnz* (COO.nnz-rgrn) nnz (COO.nnz)	- - lb-cg / lb-fg / lf
BCSR	blocks [†] (BCSR.block) nnz [†] (BCSR.nnz)	lb-cg / lb-fg (only for int8 and small block sizes) lb-cg / lb-fg (only for int8 and small block sizes)
BCOO	blocks (BCOO.block) nnz (BCOO.nnz)	lb-cg / lb-fg / lf lb-cg / lb-fg / lf

Table 2. Parallelization schemes across threads of a PIM core. *: row-granularity, †: block-row-granularity

3.5 Kernel Implementation

We briefly describe the *SparseP* implementations for all compressed matrix formats, i.e., the way that threads access data involved in the kernel from/to the local DRAM bank. The SpMV kernels include three types of data structures: (i) the arrays that store the non-zero elements, i.e., the values (values[]) and the positions of the non-zero elements (rowptr[], colind[] for CSR, tuples[] for COO, browptr[], bcolind[] for BCSR, browind[], bcolind[] for BCOO), (ii) the array that stores the elements of the input vector, and (iii) the array that stores the partial results created for the elements of the output vector.

First, SpMV performs streaming memory accesses to the arrays that store the non-zero elements and their positions. Therefore, to exploit spatial locality and immense bandwidth in data (scratchpad or cache) memory, each thread reads the non-zero elements by fetching large chunks of bytes in a coarse-grained manner from DRAM to data memory (i.e., WRAM in the UPMEM PIM system). Then, it accesses elements through data memory in a fine-grained manner. In the UPMEM PIM system, we fetch chunks of 256-byte data to discover the non-zero elements, as suggested by the UPMEM API [242], since 256-byte transfer sizes highly exploit the available local bandwidth of DRAM bank [82, 94]. For the BCSR and BCOO formats, only for the array that stores the values of the non-zero elements (i.e., bvalues[]), we fetch from DRAM to data memory block size chunks, i.e., chunks of $r \times c \times \text{sizeof}(\text{data_type})$ bytes, assuming that the matrix is stored in blocks of size $r \times c$.

Second, SpMV causes irregular memory accesses to the elements of the input vector. Specifically, the accesses to the elements of the input vector are input-driven, i.e., they are determined by the column positions (column indexes) of the non-zero elements of each particular matrix. Given that matrices involved in SpMV are very sparse [60, 63, 83, 104, 121, 152, 201, 226, 250], i.e., the column indexes of the non-zero elements significantly vary, memory accesses to the input vector incur poor data locality. Thus, in our SpMV implementations, threads of a PIM core directly access elements of the input vector through DRAM bank at fine-granularity [82, 94, 242], i.e., using the

smallest possible granularity: for the CSR and COO formats at 64-bit granularity, and for the BCSR and BCOO formats at the granularity of $c \times \text{sizeof}(\text{data_type})$ bytes, where c is the horizontal dimension of the block size.

Third, regarding the output vector, threads temporarily store partial results for the same elements of the output vector in data (scratchpad or cache) memory to exploit data locality, until all the non-zero elements of the *same* row or the *same* block row have been traversed (recall matrices are stored in a row-sorted way). Then, the produced results are written to DRAM bank at fine-granularity [82, 94, 242]: for the CSR and COO formats at 64-bit granularity, and for the BCSR and BCOO formats at the granularity of $r \times \text{sizeof}(\text{data_type})$ bytes, where r is the vertical dimension of the block size.

4 Evaluation Methodology

We conduct our evaluation on an UPMEM PIM system that includes a 2-socket Intel Xeon Silver 4110 CPU [118] at 2.10 GHz (host CPU), standard main memory (DDR4-2400) [119] of 128 GB, and 20 UPMEM PIM DIMMs with 160 GB PIM-capable memory and 2560 DPUs.²

First, we evaluate SpMV execution using one single DPU and multiple tasklets (Section 5). Table 3 shows our evaluated small matrices that fit in the 64 MB DRAM memory of a single DPU. The evaluated matrices vary in sparsity (i.e., NNZ / (rows x columns)), standard deviation of non-zero elements among rows (NNZ-r-std) and columns (NNZ-c-std). The highlighted matrices in Table 3 with red color exhibit block pattern [63, 138], i.e., they include *a lot* of dense sub-blocks (almost all their non-zero elements fit in dense sub-blocks).

Matrix Name	Sparsity	NNZ-r-std	NNZ-c-std
delaunay_n13	7.32e-04	1.343	1.343
wing_nodal	1.26e-03	2.861	2.861
raefsky4	3.396e-03	15.956	15.956
pkustk08	0.006542	61.537	61.537

Table 3. Small Matrix Dataset.

Second, we evaluate SpMV execution using *multiple* DPUs of the UPMEM PIM system (Section 6). We evaluate SpMV execution using both 1D (Section 6.1) and 2D (Section 6.2) partitioning techniques, and compare them (Section 6.3) using a wide variety of sparse matrices with diverse sparsity patterns. We select 22 representative sparse matrices from the Sparse Suite Collection [53], the characteristics of which are shown in Table 4. As the values of the last two metrics increase (i.e., NNZ-r-std and NNZ-c-std), the matrix becomes very irregular [182, 236], and is referred to as *scale-free* matrix. In our evaluation, we refer to all matrices between hgc to bns matrices of Table 4 as *regular* matrices. The matrices in which NNZ-r-std is larger than 25, i.e., all matrices between wbs to ask in Table 4, we refer to as *scale-free* matrices. Please see Appendix D for a complete description of our dataset of large sparse matrices.

Third, we compare the performance and energy consumption of SpMV execution on the UPMEM PIM system to those on the Intel Xeon Silver 4110 CPU [118] and the NVIDIA Tesla V100 GPU [190] (Section 7).

In Section 8, we summarize our key takeaways and provide programming recommendations for software designers, and suggestions and hints for hardware and system designers of future PIM systems.

²There are thirty two faulty DPUs in the system where we run our experiments. They cannot be used and do not affect the correctness of our results, but take away from the system's full computational power of 2560 DPUs.

Matrix Name	Sparsity	NNZ-r-std	NNZ-c-std
hugetric-00020 (hgc)	4.21e-07	0.031	0.031
mc2depi (mc2)	7.59e-06	0.076	0.076
parabolic_fem (pfm)	1.33e-05	0.153	0.153
roadNet-TX (rtn)	1.98e-06	1.037	1.037
rajat31 (rjt)	9.24e-07	1.106	1.106
af_shell1 (ash)	6.90e-05	1.275	1.275
delaunay_n19 (del)	1.14e-05	1.338	1.338
thermomech_dK (tdk)	6.81e-05	1.431	1.431
memchip (mem)	2.02e-06	2.062	1.173
amazon0601 (amz)	2.08e-05	2.79	15.29
FEM_3D_thermal2 (fth)	1.59e-04	4.481	4.481
web-Google (wbg)	6.08e-06	6.557	38.366
ldoor (ldr)	5.13e-05	11.951	11.951
poisson3Db (psb)	3.24e-04	14.712	14.712
boneS10 (bns)	6.63e-05	20.374	20.374
webbase-1M (wbs)	3.106e-06	25.345	36.890
in-2004 (in)	8.846e-06	37.230	144.062
pkustk14 (pks)	6.428e-04	46.508	46.508
com-Youtube (cmb)	4.639e-06	50.754	50.754
as-Skitter (skt)	7.71e-06	136.861	136.861
sx-stackoverflow (sxw)	5.352e-06	137.849	65.367
ASIC_680k (ask)	8.303e-06	659.807	659.807

Table 4. Large Matrix Dataset. Matrices are sorted by NNZ-r-std, i.e., based on their irregular pattern. The highlighted matrices with red color exhibit block pattern [63, 138].

5 Analysis of SpMV Execution on One DPU

This section characterizes SpMV performance with various load balancing schemes and compressed matrix formats using multiple tasklets in a single DPU. Section 5.1 compares load balancing schemes of each compressed matrix format, and Section 5.2 compares the scalability of various compressed matrix formats.

5.1 Load Balancing Schemes Across Tasklets of One DPU

We compare the parallelization schemes of each compressed matrix format supported by *SparseP* library (presented in Table 2) across multiple threads of a multithreaded PIM core. Figure 9 compares the load balancing schemes of each compressed matrix format using 16 tasklets in a single DPU. For the BCSR and BCOO formats, we omit results for the fine-grained locking approach, since it performs similarly with the coarse-grained locking approach: as we explain in Appendix A.1, fine-grained locking does not increase parallelism over coarse-grained, since in the UPMEM PIM hardware, DRAM memory accesses of the critical section are serialized in the DMA engine of the DPU [82, 94, 242].

We draw four findings from Figure 9. First, we find that SpMV execution using int8, int16, and int32 data types achieves similar execution times across them. This is because the multiplication operation of these data types is sufficiently supported by hardware [82]. In contrast, execution time




Fig. 9. Execution time achieved by various load balancing schemes of each compressed matrix format using 16 tasklets of a single DPU.

sharply increases when using more heavyweight data types, i.e., int64 and floating point data types, in which multiplication is emulated in software using the 8x8-bit multiplier of the DPU [82, 94, 242].

Second, we observe that balancing the non-zero elements across tasklets typically outperforms balancing the rows for the CSR/COO formats or blocks for the BCSR/BCOO formats, since the non-zero element multiplications are computationally very expensive and can significantly affect load balance across tasklets. However, in `delaunay_n13` matrix, balancing the non-zero elements causes high row/block imbalance across tasklets, since one tasklet processes a significantly higher number of rows/blocks over the rest, thereby causing high operation imbalance across tasklets within the DPU core pipeline. As a result, balancing the rows/blocks outperforms balancing the non-zero elements due to the particular pattern of `delaunay_n13` matrix. In addition, performance benefits of balancing the blocks over balancing the non-zero elements are significant in the BCSR/BCOO formats, because they operate at block granularity and incur high loop control costs.

Third, we observe that the lock-free approach (COO.nnz-lf) outperforms the lock-based approaches (COO.nnz-lb-cg and COO.nnz-lb-fg) in `delaunay_n13` matrix, especially in data types where the multiplication operation is supported directly in hardware. In `delaunay_n13` matrix, one tasklet processes a much larger number of rows than the rest, i.e., it performs a much larger number of critical sections than the rest. In other words, one tasklet performs a much larger number of lock acquisitions/releases and memory instructions than the rest. Thus, lock-based approaches cause high operation imbalance in the DPU core pipeline with significant performance costs. Instead, lock-free and lock-based approaches in the BCOO format perform similarly, since lock acquisition/release costs can be hidden due to BCOO's higher loop control costs and larger critical sections. Overall, based on the second and the third findings, we conclude that in matrices and formats, where the load balancing and/or the synchronization scheme used cause *high* disparity in the number of non-zero elements/blocks/rows processed across tasklets or the number of lock acquisitions/lock releases/memory accesses performed across tasklets, the DPU core pipeline can incur significant performance overheads.

OBSERVATION 1:

High operation imbalance in computation, control, synchronization, or memory instructions executed by multiple threads of a PIM core can cause *high performance overheads* in the compute-bound and area-limited PIM cores.

Fourth, we find that the fine-grained locking approach (COO.nnz-lb-fg) performs similarly with the coarse-grained locking approach (COO.nnz-lb-cg). This is because the critical section includes memory accesses to the local DRAM bank, which, in the UPMEM PIM hardware, are serialized in the DMA engine of the DPU. Therefore, fine-grained locking does not increase execution parallelism over coarse-grained locking, since concurrent accesses to MRAM bank are not supported in the UPMEM PIM hardware. Fine-grained locking does not improve performance over coarse-grained locking, also when using block-based formats (e.g., BCSR/BCOO formats), as we demonstrate in Appendix A.1. Therefore, we recommend PIM hardware designers to provide lightweight synchronization mechanisms [81] for PIM cores, and/or enable concurrent accesses to local DRAM memory, e.g., supporting sub-array level parallelism [41, 42, 96, 134, 217–220] or multiple DRAM banks per PIM core.

OBSERVATION 2:

Fine-grained locking approaches to parallelizing critical sections that perform memory accesses to different DRAM memory locations cannot improve performance over *coarse-grained* locking, when the PIM hardware does not support *concurrent accesses to a DRAM bank*.

5.2 Analysis of Compressed Matrix Formats on One DPU

We compare the scalability and the performance achieved by various compressed matrix formats. Figure 10 compares the supported compressed formats for the int8 (top graphs) and fp64 (bottom graphs) data types when balancing the non-zero elements across tasklets of a DPU.

We draw three findings. First, we find that even though a DPU supports 24 tasklets, SpMV execution typically scales up to 16 tasklets, since the DPU pipeline is fully utilized. In delaunay_n13 matrix, CSR.nnz scales up to 24 tasklets. In this matrix, when using 16 tasklets, performance of the CSR.nnz scheme is limited by memory accesses: *only* one tasklet processes $6 \times$ more rows than the rest, i.e., it performs $6 \times$ more memory accesses to fetch elements from the rowptr[] array. Thus, as we increase the number of tasklets from 16 to 24, the disparity in the number of rows across tasklets decreases, and the performance of the CSR.nnz scheme improves. Second, we observe that for the data types with hardware-supported multiplication operation (e.g., int8 data type), CSR achieves the highest scalability, since it provides a better balance between memory access and computation. In contrast, in the floating point data types (e.g., fp64 data type), the DPU is significantly bottlenecked by the expensive software-emulated multiplication operations, and thus all formats scale similarly. Third, we observe that the BCSR and BCOO formats outperform the CSR and COO formats in matrices that exhibit block pattern (i.e., raefsky4 and pkustk08 matrices), only when multiplication is supported by hardware (e.g., int8 data type). This is because they exploit spatial and temporal locality in data memory (i.e., WRAM) in the accesses of the elements of the input vector. Instead, in the fp64 data type, performance is severely bottlenecked by computation, thus the BCSR/BCOO formats perform worse than the CSR/COO formats, since they incur higher indexing costs to discover the positions of the non-zero elements [14, 121].




Fig. 10. Scalability of all compressed formats for the int8 (top graphs) and fp64 (bottom graphs) data types as the number of tasklets of a single DPU increases.

OBSERVATION 3:

Block-based formats (e.g., BCSR/BCOO) can provide high performance gains over non-block-based formats (e.g., CSR/COO) in matrices that exhibit block pattern, if the multiplication operation is supported by hardware. Otherwise, the state-of-the-art CSR and COO formats can provide high performance and scalability.

6 Analysis of SpMV Execution on Multiple DPUs

This section analyzes SpMV execution using multiple DPUs in the UPMEM PIM system using the large matrix data set of Table 4.

Section 6.1 evaluates the 1D partitioning schemes. Section 6.1.1 evaluates the actual kernel time of SpMV by comparing (a) all load balancing schemes of each compressed matrix format, and (b) the performance of all compressed matrix formats. Section 6.1.2 characterizes end-to-end SpMV execution time of the 1D partitioning technique including the data transfer costs for the input and output vectors.

Section 6.2 evaluates the 2D partitioning techniques. Section 6.2.1 presents three characterization studies on (a) performing fine-grained data transfers to transfer the elements of the input and output vectors to/from PIM-enabled memory, (b) the scalability of 2D partitioning techniques to thousands of DPUs, and (c) the number of vertical partitions to perform on the matrix. Section 6.2.2 compares the end-to-end performance of all compressed matrix formats for each of the three types of 2D partitioning techniques. Section 6.2.3 compares the best-performing SpMV implementations of all three types of 2D partitioning techniques.

Section 6.3 compares the best-performing (on average across all matrices and data types) SpMV implementations of the 1D and 2D partitioning techniques.

6.1 Analysis of SpMV Execution Using 1D Partitioning Techniques

We evaluate the 1D partitioning schemes highlighted in bold in Table 1. Specifically, for C00.nnz, we present the coarse-grained locking (C00.nnz-lb) and lock-free (C00.nnz-lf) approaches, since the fine-grained locking approach performs similarly with the coarse-grained locking approach, as shown in the previous section (Section 5.1). Similarly, for the BCSR (int8 data type) and BCOO formats, we present only the coarse-grained locking approach, since all synchronization approaches perform similarly (Section 5.1). Finally, in all experiments presented henceforth, we use 16 tasklets and load-balance the non-zero elements across tasklets within the DPU, since this load balancing scheme provides the highest performance benefits on average across all matrices and data types, according to our evaluation shown in Section 5.

6.1.1 Analysis of Kernel Time

We compare the kernel time of SpMV achieved by various load balancing schemes for each particular compressed matrix format, and then we compare the kernel time of the compressed matrix formats.

Analysis of Load Balancing Schemes Across DPUs. Figure 11 compares load balancing techniques for each compressed matrix format using 2048 DPUs and the int32 data type.

We draw four findings. First, we observe that CSR.nnz and C00.nnz-rgrn, i.e., balancing the non-zero elements across DPUs (at row granularity), either outperform or perform similarly to CSR.row and C00.row, respectively, i.e., balancing the rows across DPUs, except for hgc and del matrices. In these two matrices, CSR.nnz and C00.nnz-rowgrn incur a high disparity in rows assigned to DPUs, i.e., only one DPU processes 4× and 11× more rows than the rest, for hgc and del matrices, respectively. This in turn creates a high disparity in the elements of the output vector processed across DPUs, causing performance to be limited by the DPU that processes the largest number of rows. Thus, we find that adaptive load balancing approaches and selection methods based on the characteristics of each input matrix need to be developed to achieve high performance across all matrices.

OBSERVATION 4:

Adaptive load balancing schemes and selection methods for the balancing scheme on rows/blocks/non-zero elements based on the characteristics of each input matrix need to be developed to provide best performance across all matrices.

Second, we find that C00.nnz-lb and C00.nnz-lf, which provide an almost perfect non-zero element balance across DPUs, significantly outperform C00.row and C00.nnz-rgrn in *scale-free* matrices (i.e., from wbs to ask matrices) by on average 6.73×. Scale-free matrices have only a few rows, that include a much larger number of non-zero elements compared to the remaining rows of the matrix. Therefore, perfectly balancing the non-zero elements across DPUs provides high performance gains.

OBSERVATION 5:

Perfectly balancing the non-zero elements across PIM cores can provide significant performance benefits in highly irregular, scale-free matrices.

Third, we find that the lock-free C00.nnz-lf scheme outperforms the lock-based C00.nnz-lb scheme by 1.34× on average, and provides high performance benefits when there is a high row imbalance across tasklets within the DPU. When one tasklet processes a much larger number of




Fig. 11. Performance comparison of load balancing techniques for each particular compressed format using 2048 DPUs and the int32 data type.

rows versus the rest, it executes a much larger number of critical sections. As a result, the core pipeline incurs high imbalance in lock acquisitions/releases, causing the lock-based approach to incur high performance overheads in relatively compute-bound DPUs [82, 94].

OBSERVATION 6:

Lock-free approaches can provide high performance benefits over *lock-based* approaches in PIM architectures, because they minimize synchronization overheads in PIM cores.

Finally, in the BCSR and BCOO formats, balancing the blocks across DPUs performs similarly (on average across all matrices) to balancing the non-zero elements across DPUs.

To further investigate the performance of the various load balancing schemes, Figure 12 compares them using all the data types. We present the geometric mean of all matrices using 2048 DPUs. In the CSR and COO formats, balancing the non-zero elements across DPUs on average outperforms balancing the rows across DPUs by 1.18 $\times$ and 1.20 $\times$, respectively. We observe that in the COO format almost perfectly balancing the non-zero elements across DPUs provides significant performance

benefits (2.55 $\times$, averaged across all the data types), compared to balancing the rows, especially when multiplication is not supported by hardware (e.g., for the floating point data types). In contrast, in the BCSR and BCOO formats, balancing the blocks across DPUs performs only slightly better (on average 2.7% across all the data types) than balancing the non-zero elements.




Fig. 12. Performance comparison of load balancing techniques for each data type using 2048 DPUs.

Comparison of Compressed Matrix Formats. Figures 13 and 14 compare the throughput (in GOperations per second) and the performance, respectively, achieved by various compressed formats using 2048 DPUs and the int32 data type. For the CSR and COO formats, we select balancing the non-zero elements across DPUs, and for the BCSR and BCOO formats, we select balancing the blocks across DPUs, since these are the best-performing schemes for each format averaged across all matrices and data types (Figure 12).

We draw four findings. First, matrices that exhibit block pattern (almost all non-zero elements of the matrix fit in dense sub-blocks), i.e., ash, ldr, bns, pks matrices, have the highest throughput, since they leverage higher data locality compared to matrices with non-block pattern. Second, in scale-free matrices, the COO and BCOO formats significantly outperform the CSR and BCSR formats by 6.94 $\times$ and 13.90 $\times$, respectively. This is because they provide better non-zero element balance across DPUs. In the CSR and BCSR formats, the non-zero element balance is limited to be performed at row and block-row granularity, respectively, causing performance to be limited by the DPU that processes the largest number of non-zero elements. Third, we observe that the BCOO format can outperform the CSR format even in *non-blocked* scale-free matrices. Fourth, we find that when the CSR and BCSR formats provide sufficient non-zero element balance across DPUs, i.e., in many regular matrices such as rtn, tdk, amz, and fth, they can outperform the COO and BCOO formats, respectively.




Fig. 13. Throughput of various compressed formats using 2048 DPUs and the int32 data type.




Fig. 14. Performance comparison of various compressed formats using 2048 DPUs and the int32 data type. Performance is normalized to that of CSR.nnz.

OBSERVATION 7:

In *scale-free* matrices, the COO and BCOO formats significantly outperform the CSR and BCSR formats, because they provide higher non-zero element balance across PIM cores.

6.1.2 Analysis of End-To-End SpMV Execution

Figure 15 shows the end-to-end execution time of 1D-partitioned kernels using 2048 DPUs and the int32 data type. The times are broken down into (i) the time for CPU to DPU transfer to load the input vector into DRAM banks (load), (ii) the kernel time on DPUs (kernel), (iii) the time for DPU to CPU transfer to retrieve the results for the output vector (retrieve), and (iv) the time to merge partial results on the host CPU cores (merge).




Fig. 15. Total execution time when using 2048 DPUs and the int32 data type for CR: CSR.nnz, CO: COO.nnz-1f, BR: BCSR.block and BO: BCOO.block kernels.

We draw four findings. First, the load data transfers constitute more than 90% of the total execution time, because the input vector is replicated and broadcast into each DPU, causing a large number of bytes to be transferred through the narrow off-chip memory bus. An exception is in the CSR and BCSR formats for sxw, ask matrices, which include one very dense row, and thus kernel time is highly bottlenecked by one DPU that processes a significantly larger number of non-zero elements than the rest. Second, the kernel time constitutes on average only 4.3% of the total execution time, since SpMV is effectively parallelized to thousands of DPUs. Third, the retrieve data transfers constitute on average 3.4% of the total execution time, because the output vector is split across DPUs. Fourth, the merge time on the host CPU is negligible (less than 1% of the total execution time), since only a few partial results for the elements of the output vector are merged by the host CPU cores in the 1D partitioning techniques.

OBSERVATION 8:

The end-to-end performance of the 1D partitioning techniques is severely bottlenecked by the data transfer costs to replicate and broadcast the whole input vector into *each* DRAM bank of PIM cores, which takes place through the narrow off-chip memory bus.

To further investigate on the costs to the load input vector into all DRAM banks of PIM-enabled memory, we present in Figure 16 the total execution time achieved by COO.nnz-1f when varying (a) the data type using 2048 DPUs (normalized to the experiment for the int8 data type), and (b) the number of DPUs for the int32 data type (normalized to 64 DPUs).




Fig. 16. End-to-end execution time breakdown achieved by COO.nnz-1f when varying (a) the data type using 2048 DPUs (normalized to the experiment for the int8 data type), and (b) the number of DPUs for the int32 data type (normalized to 64 DPUs).

We draw two conclusions. First, the load data transfer costs increase proportionally to the number of bytes of the data type, and still dominate performance even for the data type with the smallest memory footprint (int8). Second, the load data transfer costs and the associated memory footprint for the input vector increase proportionally to the number of DPUs used, and thus the best end-to-end performance is achieved using only a small portion of the available DPUs on the system.

OBSERVATION 9:

SpMV execution of the 1D-partitioned schemes cannot scale up to a large number of PIM cores due to high data transfer overheads to copy the input vector into *each* DRAM bank of PIM-enabled memory.

6.2 Analysis of SpMV Execution Using 2D Partitioning Techniques

We evaluate the 2D-partitioned kernels highlighted in bold in Table 1. Specifically, for the COO format we use the lock-free approach, and for the BCSR (in the int8 data type) and BCOO formats we use the coarse-grained locking approach. In the *equally-wide* and *variable-sized* techniques, for the BCSR and BCOO formats we balance the blocks across DPUs of the same vertical partition, since doing so performs slightly better than balancing the non-zero elements, as explained in Section 6.1.1. In all experiments, we balance the non-zero elements across 16 tasklets within a single DPU.

6.2.1 Sensitivity Studies on 2D Partitioning Techniques

We present three characterization studies on the 2D partitioning techniques. First, we evaluate the performance of fine-grained data transfers from/to PIM-enabled memory for the input and output vectors. Second, we evaluate the scalability of the 2D partitioning techniques to thousands of DPUs. Finally, we explore performance implications on the number of vertical partitions used in the 2D-partitioned kernels.

Analysis of Fine-Grained Data Transfers. The UPMEM API [242] has the limitation that *the transfer sizes from/to all DRAM banks involved in the same parallel transfer need to be the same*. The UPMEM API provides *parallel data transfers* either to all DPUs of all ranks (henceforth referred to as *coarse-grained* transfers), or at rank granularity, i.e., to 64 DPUs of the same rank (henceforth referred to as *fine-grained* transfers). In the first case, parallel data transfers are performed to all DPUs used at once, padding with empty bytes at the granularity of *all* DPUs used, e.g., 2048 DPUs in Figure 17. In the latter case, programmers iterate over the ranks of PIM-enabled DIMMs, and for *each* rank perform parallel data transfers to the 64 DPUs of the same rank padding with empty bytes at the granularity of 64 DPUs.

In SpMV execution, for the *equally-wide* and *variable-sized* techniques the heights and widths of 2D tiles vary, and thus padding with empty bytes is necessary for the load and retrieve data transfers of the elements of the input and output vector, respectively. Figure 17 compares coarse-grained data transfers, i.e., performing parallel transfers to all 2048 DPUs at once, with fine-grained data transfers, i.e., iterating over the ranks and for each rank performing parallel transfers to the 64 DPUs of the same rank. We evaluate both the *equally-wide* and *variable-sized* techniques using the COO format and with 2 and 32 vertical partitions. Please see Appendix A.2 for all matrices.

We draw two findings. First, when the number of vertical partitions is small, e.g., 2 vertical partitions, the disparity in widths across tiles in the *variable-sized* scheme is low. Thus, BT only slightly outperforms BY by 1% on average, since in BY *only* a small amount of padding is added on the load data transfers of the input vector. In contrast, the disparity in heights across tiles in the *equally-wide* and *variable-sized* schemes is high. Thus, RY and BY significantly outperform RC and BC by an average of 1.68 $\times$ and 1.60 $\times$, respectively. This is because fine-grained transfers to retrieve the elements of the output vector significantly decrease the amount of bytes transferred from PIM-enabled memory to host CPU over coarse-grained transfers. Second, when the number of vertical partitions is large, e.g., 32 vertical partitions, the disparity in heights across tiles in the *equally-wide* and *variable-sized* schemes is lower compared to when the number of vertical partitions is small. Thus, RY and BY provide smaller performance benefits over RC and BC (on average 1.24 $\times$ and 1.22 $\times$, respectively), respectively, compared to a small number of vertical partitions. In contrast, the disparity in heights across tiles in the *equally-wide* and *variable-sized* schemes is higher compared to when the number of vertical partitions is small. Thus, BT outperforms BY by 4.7% on average. Overall, we conclude that fine-grained data transfers (i.e., at rank granularity in the UPMEM PIM system) can significantly improve performance in the *equally-wide* and *variable-sized* schemes.




Fig. 17. Performance comparison of RC: RBDC00 with coarse-grained transfers, RY: RBDC00 with fine-grained transfers in the output vector, BC: BDC00 with coarse-grained transfers, BY: BDC00 with fine-grained transfers only in the output vector, and BT: BDC00 with fine-grained transfers in both the input and the output vector using the int32 data type, 2048 DPUs and having 2 (left) and 32 (right) vertical partitions. Performance is normalized to that of the RC scheme.

OBSERVATION 10:

Fine-grained parallel transfers in the equally-wide and variable-sized 2D partitioning techniques, i.e., minimizing the amount of padding with empty bytes in parallel data transfers to/from PIM-enabled memory, can provide large performance gains.

Scalability of the 2D Partitioning Techniques. We analyze scalability with the number of DPUs for the 2D partitioning techniques. Figures 18, 19 and 20 compare the performance of the *equally-sized*, *equally-wide* and *variable-sized* schemes, respectively, using the COO format and the int32 data type, as the number of DPUs increases.




Fig. 18. Execution time breakdown of *equally-sized* partitioning technique of the COO format using 4 (left) and 16 (right) vertical partitions when varying the number of DPUs used for the int32 data type. Performance is normalized to that with 256 DPUs.




Fig. 20. Execution time breakdown of *variable-sized* partitioning technique of the COO format using 4 (left) and 16 (right) vertical partitions when varying the number of DPUs used for the int32 data type. Performance is normalized to that with 256 DPUs.




Fig. 19. Execution time breakdown of *equally-wide* partitioning technique of the COO format using 4 (left) and 16 (right) vertical partitions when varying the number of DPUs used for the int32 data type. Performance is normalized to that with 256 DPUs.

We draw two findings. First, the *equally-sized* scheme (i.e., DC00) achieves high scalability with a large number of vertical partitions. The kernel time of *equally-sized* scheme is mainly limited by the DPU (or a few DPUs) that processes the largest number of non-zero elements. With a large number of *static* vertical partitions, the non-zero element disparity across DPUs is high, i.e., the kernel time is highly bottlenecked by the DPU that processes the largest number of non-zero elements. As a result, increasing the number of DPUs improves performance by decreasing the kernel time via better non-zero element balance across DPUs.

OBSERVATION 11:

The kernel time in the *equally-sized* schemes is limited by the PIM core (or a few PIM cores) assigned to the 2D tile with the largest number of non-zero elements.

Second, we observe that the *equally-wide* and *variable-sized* schemes (i.e., RBDC00 and BDC00) are severely bottlenecked by retrieve data transfer costs (a large number of partial results is created on PIM cores), and thus they are difficult to scale up to thousands of DPUs. Moreover, when the number of vertical partitions is high, the disparity in heights of the tiles is high. Thus, as the number of DPUs increases, the amount of padding needed in retrieve data transfers becomes very large, causing significant performance degradation.

OBSERVATION 12:

The scalability of the *equally-wide* and *variable-sized* schemes to a large number of PIM cores is severely limited by large data transfer overheads to retrieve partial results for the elements of the output vector from the DRAM banks of PIM-enabled memory to the host CPU via the narrow memory bus.

Effect of the Number of Vertical Partitions. In all experiments presented henceforth, we perform fine-grained data transfers (at rank granularity, i.e., 64 DPUs in the UPMEM PIM system) in the 2D partitioning schemes. Figure 21 evaluates performance implications on the number of vertical partitions performed in 2D-partitioned kernels. We use the COO format and vary the number of vertical partitions from 1 to 32, in steps of multiple of 2. We draw four findings.




Fig. 21. Execution time breakdown of 2D partitioning schemes using the COO format and 2048 DPUs when varying the number of vertical partitions from 1 to 32 for the int8 and fp64 data types. Performance is normalized to the performance of the experiment with 1 vertical partition.

First, in the *equally-sized* scheme, as the number of vertical partitions increases, kernel time increases, if there is *no* dense row in the matrix. This is because the disparity in the non-zero elements across 2D tiles increases as the number of vertical partitions increases. Thus, performance is limited by one DPU or a few DPUs that process the largest number of non-zero elements.

OBSERVATION 13:

As the number of vertical partitions increases, the *equally-sized* 2D partitioning scheme typically increases the non-zero element disparity across PIM cores (unless there is one dense row on the matrix), thereby increasing the kernel time.

Second, as the number of vertical partitions increases, retrieve data transfer costs and merge time increase. This is because the partial results created for the output vector increase proportionally with the number of vertical partitions. The performance overheads of retrieve data transfer costs are highly affected by the characteristics of the underlying hardware (e.g., the bandwidth provided

on I/O channels of the memory bus between host CPU and PIM-enabled DIMMs). Similarly, the performance cost of the merge step depends on the hardware characteristics of the host CPU (e.g., the number of the CPU cores, the available hardware threads, microarchitecture of CPU cores). We refer the reader to Appendix A.3 for a comparison of SpMV execution using two different UPMEM PIM systems with different hardware characteristics (Table 6).

Third, we find that in the *equally-wide* and *variable-sized* schemes, there is high disparity in heights of 2D tiles, and as a result on the number of partial results created across DPUs. Even with fine-grained parallel retrieve data transfers at rank granularity, the amount of padding needed in the *equally-wide* and *variable-sized* schemes is at 88.6% and 88.0%, respectively, causing high bottlenecks in the narrow memory bus. Therefore, in PIM systems that do not support very fine-grained parallel transfers to gather results from PIM-enabled memory to the host CPU *at DRAM bank granularity*, execution is highly limited by the amount of padding performed in retrieve data transfers, which can be very large in irregular workloads [22, 56, 60, 63, 80, 82, 83, 94, 104, 121, 152, 167, 194, 201, 226, 230, 250] like the SpMV kernel.

OBSERVATION 14:

The *equally-wide* and *variable-sized* 2D partitioning schemes require fine-grained parallel transfers *at DRAM bank granularity* to be supported by the PIM system, i.e., *zero padding in parallel* retrieve data transfers from PIM-enabled memory to the host CPU, to achieve high performance.

Fourth, we find that the number of vertical partitions that provides the best performance depends on the sparsity pattern of the input matrix, the data type, and the underlying hardware parameters (e.g., number of PIM cores, off-chip memory bus bandwidth, transfer latency costs between main memory and PIM-enabled memory, characteristics and microarchitecture of the host CPU cores that perform the merge step). For example, with the int8 data type, DC00 performs best for hgc and mem matrices with 8 and 4 vertical partitions, respectively. Instead, with the fp64 data type, DC00 performs best for hgc and mem matrices with 16 and 8 vertical partitions, respectively. We refer the reader to Appendix A.3 for a characterization study on the number of vertical partitions to perform in the 2D-partitioned kernels using two UPMEM PIM systems with different hardware characteristics. As we demonstrate in Appendix A.3, the number of vertical partitions that provides best performance on SpMV varies across the two different UPMEM PIM platforms. In this work, we leave for future work the exploration of selection methods for the number of vertical partitions that provide best SpMV execution. Overall, based on our analysis we conclude that the parallelization scheme that achieves the best performance in SpMV depends on both the input sparse matrix and the hardware characteristics of the PIM system.

OBSERVATION 15:

There is *no one-size-fits-all* parallelization approach for SpMV in PIM systems, since the performance of each parallelization scheme depends on the characteristics of the input matrix and the underlying PIM hardware.

6.2.2 Analysis of Compressed Formats

We compare the performance achieved by various compressed matrix formats for each of the three types of the 2D partitioning technique. The goal of this experiment is to find the best-performing compressed format for each 2D partitioning technique. Figures 22, 23, and 24 compare the performance of compressed matrix formats for the *equally-sized*, *equally-wide* and *variable-sized*

2D partitioning techniques, respectively. We use 2048 DPUs and the int32 data type having 4 vertical partitions. See Appendix A.4 for the complete evaluation on all large sparse matrices.




Fig. 22. End-to-end execution time breakdown of the *equally-sized* 2D partitioning technique for CR: DCSR, CO: DC00, BR: DBCSR and BO: DBC00 schemes using 4 vertical partitions and the int32 data type. Performance is normalized to that of DCSR.




Fig. 23. End-to-end execution time breakdown of the *equally-wide* 2D partitioning technique for CR: RBCSR, CO: RBDC00, BR: RDBCSR and BO: RDBC00 schemes using 4 vertical partitions and the int32 data type. Performance is normalized to that of RBCSR.




Fig. 24. End-to-end execution time breakdown of the *variable-sized* 2D partitioning technique for CR: BDCSR, CO: BDC00, BR: BDBCSR and BO: BDBC00 schemes using 4 vertical partitions and the int32 data type. Performance is normalized to that of BDCSR.

We draw two findings. First, as already explained, kernel time of the *equally-sized* scheme is limited by the DPU (or a few DPUs) assigned to the 2D tile with the largest number of non-zero elements. In scale-free matrices (e.g., in and ask), the disparity in the non-zero elements across 2D tiles is higher than in regular matrices (e.g., pfg and bns), causing kernel time to be a larger portion of the total execution time. Second, we find that the CSR and BCSR formats perform worse than the COO and BCOO formats, especially in the *equally-wide* and *variable-sized* schemes, due to higher kernel times. In the CSR and BCSR formats, data partitioning across DPUs and/or across

tasklets within a DPU is performed at row and block-row granularity, respectively. Thus, the CSR and BCSR formats can cause higher non-zero element imbalance across processing units compared to the COO and BCOO formats. Overall, the COO and BCOO formats outperform the CSR and BCSR formats by $1.59 \times$ and $1.53 \times$ (averaged across all three types of 2D partitioning techniques), respectively.

OBSERVATION 16:

The compressed matrix format used to store the input matrix determines the data partitioning across DRAM banks of PIM-enabled memory. Thus, it affects the load balance across PIM cores with corresponding performance implications. Overall, the COO and BCOO formats outperform the CSR and BCSR formats, because they provide higher non-zero element balance across PIM cores.

6.2.3 Comparison of 2D Partitioning Techniques

We compare the best-performing SpMV implementations of all 2D partitioning schemes, i.e., using the COO and BCOO formats. Figures 25 and 26 compare the throughput (in GOperations per second) and the performance, respectively, of DCOO, DBCOO, RBDCOO, RBDBCOO, BDCOO, BDBCOO schemes using 2048 DPUs and the int32 data type. For each implementation, we vary the number of vertical partitions from 2 to 32, in steps of multiple of 2, and select the best-performing execution throughput.




Fig. 25. Throughput of 2D partitioning techniques using the COO and BCOO formats, 2048 DPUs and the int32 type.




Fig. 26. Performance comparison of 2D partitioning techniques using the COO and BCOO formats, 2048 DPUs and the int32 type. Performance is normalized to that of DCOO.

We draw two conclusions. First, similarly to 1D-partitioned kernels, matrices that exhibit block pattern (e.g., ash, ldr, bns, pks) have the highest throughput (Figure 25). Second, the *equally-wide* and *variable-sized* schemes perform similarly, i.e., their performance varies only by $\pm 1.1\%$ on average. Even though the *variable-sized* technique can improve the non-zero element balance across DPUs, and thus kernel time, compared to the *equally-wide* technique, the total execution time does

not improve. In the UPMEM PIM system, performance of both techniques is severely bottlenecked by data transfer overheads due to a large amount of padding needed to retrieve results from PIM-enabled memory to the host CPU. Third, we find that the *equally-sized* technique outperforms the *equally-wide* and *variable-sized* techniques by $3.71\times$ on average, because it achieves lower data transfer overheads. The *equally-wide* and *variable-sized* techniques provide near-perfect non-zero element balance across DPUs, but they significantly increase the retrieve data transfer costs due to the large amount of padding with empty bytes performed. As a result, we recommend software designers to explore *relaxed* load balancing schemes, i.e., schemes that trade off computation balance across PIM cores for lower amounts of data transfer.

6.3 Comparison of 1D and 2D Partitioning Techniques

We compare the throughput (in GOperations per second) and the performance of the best-performing 1D- and 2D-partitioned kernels in Figures 27 and 28, respectively. For 1D partitioning, we use the lock-free COO (COO.nnz-1f) and coarse-grained locking BCOO (BCOO.block) kernels. For each matrix, we vary the number of DPUs from 64 to 2528, and select the best-performing end-to-end execution throughput. For 2D partitioning, we use the *equally-sized* COO (DCOO) and BCOO (DBC00) kernels with 2528 DPUs. For each matrix, we vary the number of vertical partitions from 2 to 32 (in steps of multiple of 2), and select the best-performing end-to-end execution throughput. The numbers shown over each bar of Figure 27 present the number of DPUs that provide the best-performing end-to-end execution throughput for each input-scheme combination. Please see Appendix A.5 for a performance comparison of the best-performing SpMV kernels on two UPMEM PIM systems with different hardware characteristics.




Fig. 27. Throughput of the best-performing 1D- and 2D-partitioned kernels for the fp32 data type.




Fig. 28. Performance comparison of the best-performing 1D- and 2D-partitioned kernels for the fp32 data type. Performance is normalized to that of COO.nnz-1f.

We draw two conclusions. First, we find that best performance is achieved using a much smaller number of DPUs than the available DPUs on the system. In the 1D-partitioned kernels (i.e.,

C00.nnz-1f and BC00.block), replicating the input vector into a large number of DPUs significantly increases the load data transfer costs. Thus, best performance is achieved using 253 DPUs on average across all matrices. In the 2D-partitioned kernels (i.e., DC00 and DBC00), creating *equally-sized* 2D tiles leads to a large disparity in non-zero element count across tiles, causing many tiles to be empty, i.e., without *any* non-zero element. Thus, best performance is achieved using 1329 DPUs on average across all matrices, since DPUs associated with empty tiles are idle.

OBSERVATION 17:

Expensive data transfers to PIM-enabled memory performed via the narrow memory bus impose significant performance overhead to end-to-end SpMV execution. Thus, it is hard to fully exploit all available PIM cores of the system.

Second, we observe that in regular matrices, the 2D-partitioned kernels outperform the 1D-partitioned kernels by $1.45\times$ on average. This is because the 2D-partitioned kernels use a larger number of DPUs, and thus their kernel times are lower. In contrast, in scale-free matrices, the 1D-partitioned kernels outperform the 2D-partitioned kernels by $1.41\times$ on average. This because the *equally-sized* 2D technique significantly increases the non-zero element disparity across DPUs, i.e., kernel time is bottlenecked by only one DPU or a few DPUs that process a much larger number of non-zero elements compared to the rest.

OBSERVATION 18:

In *regular* matrices, 2D-partitioned kernels outperform 1D-partitioned kernels, since the former provide a better trade-off between computation and data transfer overheads. In contrast, in *scale-free* matrices, 2D-partitioned kernels perform worse than 1D-partitioned kernels, since the former's performance is limited by one DPU or a few DPUs that process the largest number of non-zero elements.

7 Comparison with CPUs and GPUs

We compare SpMV execution on the UPMEM PIM architecture to a state-of-the-art CPU and a state-of-the-art GPU in terms of performance and energy consumption. Our goal is to quantify the potential of the UPMEM PIM architecture on the widely used memory-bound SpMV kernel.

We compare the UPMEM PIM system with 2528 DPUs to an Intel Xeon CPU [118] and an NVIDIA Tesla V100 GPU [190], the characteristics of which are shown in Table 5. We use peakperf [200] and stream [231] for CPU and GPU systems to calculate the peak performance, memory bandwidth, and Thermal Design Power (TDP). For the UPMEM PIM system, we estimate the peak performance as $Total_DPUs * AT$, where the arithmetic throughput (AT) is calculated for the multiplication operation in Appendix B (Figure 42), the total bandwidth as $Total_DPUs * Bandwidth_DPU$, where the $Bandwidth_DPU$ is 700 MB/s [55, 82, 94], and TDP as $(Total_DPUs / DPUs_per_chip) * 1.2W / chip$ from prior work [55, 82, 94].

System	Process Node	Total Cores	Frequency	Peak Performance	Memory Capacity	Total Bandwidth	TDP
Intel Xeon 4110 CPU [118]	14 nm	2x8 x86 cores (2x16 threads)	2.1 GHz	660 GFLOPS	128 GB	23.1 GB/s	2x85 W
NVIDIA Tesla V100 [190]	12 nm	5120 CUDA cores	1.25 GHz	14.13 TFLOPS	32 GB	897 GB/s	300 W
PIM System	2x nm	2528 DPUs	350 MHz	4.66 GFLOPS	159 GB	1.77 TB/s	379 W

Table 5. Evaluated CPU, GPU, and UPMEM PIM Systems.

7.1 Performance Comparison

For the CPU system, we use the optimized CSR kernel from the TACO library [136]. For the GPU system, we use the CSR5 CUDA [49, 164] for the int32 data type and cuSparse [50] for the other data types. For the UPMEM PIM system, we use the lock-free COO 1D-partitioned kernel (**COO.nnz-1f**) and the *equally-sized* COO 2D-partitioned kernel (**DCOO**). In the former, we run experiments from 64 to 2528 DPUs, and in the latter, we use 2528 DPUs, and vary the number of vertical partitions from 2 to 32, in steps of multiple of 2. In both schemes, we select the best-performing end-to-end execution throughput. We also include the lock-free COO 1D-partitioned kernel using 2528 DPUs, named **COO.k1**, to evaluate SpMV execution using *all* available DPUs of the system.

Figure 29 shows the throughput of SpMV (in GOperations per second) in all systems, comparing both the end-to-end execution throughput (i.e., including the load and retrieve data transfer costs for the input and output vectors in case of the UPMEM PIM and GPU systems), and only the actual kernel throughput (i.e., including the kernel time in DPUs and the merge time in host CPU for the UPMEM PIM system).




Fig. 29. Performance comparison between the UPMEM PIM system, Intel Xeon CPU and Tesla V100 GPU on SpMV execution.

We draw three conclusions. First, when data transfer costs to/from host CPU are included, CPU outperforms both the GPU and UPMEM PIM systems, since data transfers impose high overhead. When only the actual kernel time is considered, GPU performs best, since it is the system that provides the highest computation throughput, e.g., 14.13 TFlops for the fp32 data type. Second, we evaluate the portion of the machine’s peak performance achieved on SpMV in all systems, and observe that SpMV execution on the UPMEM PIM system achieves a much higher fraction of the peak performance compared to CPU and GPU systems. For the fp32 data type, SpMV achieves on average 0.51% and 0.21% of the peak performance in CPU and GPU, respectively, while it achieves 51.7% of the peak performance in the UPMEM PIM system using the C00.k1 scheme. Achieving a high portion of machine’s peak performance is highly desirable, since the software highly exploits the computation capabilities of the underlying hardware. This way, it improves the processor/resource utilization, and the cost of ownership of the underlying hardware. Third, we observe that when all DPUs are used, as in C00.k1, SpMV execution on the UPMEM PIM outperforms SpMV execution on the CPU by 1.09 $\times$ and 1.25 $\times$ for the int8 and int32 data types, respectively, the multiplication of which is supported by hardware. In contrast, SpMV execution on the UPMEM PIM performs 1.27 $\times$ and 2.39 $\times$ worse than SpMV execution on the CPU for the fp32 and fp64 data types, the multiplication of which is software emulated in the DPUs of the UPMEM PIM system.

OBSERVATION 19:

SpMV execution can achieve a *significantly higher* fraction of the peak performance on real memory-centric PIM architectures compared to that on processor-centric CPU and GPU systems, since PIM architectures greatly mitigate data movement costs.

7.2 Energy Comparison

For energy measurements, we consider only the actual kernel time in all systems (in the UPMEM PIM we consider the kernel and merge steps of SpMV execution). We use Intel RAPL [128] on the CPU, and NVIDIA SMI [189] on the GPU. For the UPMEM PIM system, we measure the number of cycles, instructions, WRAM accesses and MRAM accesses of each DPU, and estimate energy with energy weights provided by the UPMEM company [241]. Figure 30 shows the energy consumption (in Joules) and performance per energy (in (GOp/s)/W) for all systems.

We draw three findings. First, GPU provides the lowest energy on SpMV over the other two systems, since the energy results typically follow the performance results. Second, we find that the 2D-partitioned kernel, i.e., DC00, consumes more energy than the 1D-partitioned kernels, i.e., C00 and C00.k1, due to the energy consumed in the host CPU cores. CPU cores merge a large number of partial results in the 2D-partitioned kernels to assemble the final output vector, thereby increasing the energy consumption. Finally, we find that the 1D-partitioned kernels provide better energy efficiency on SpMV over the CPU system, when the multiplication operation is supported by hardware. Specifically, 1D-partitioned kernels provide 3.16 $\times$ and 4.52 $\times$ less energy consumption, and 1.74 $\times$ and 1.14 $\times$ better performance per energy over the CPU system for the int8 and int32 data types, respectively.

OBSERVATION 20:

Real PIM architectures can provide high energy efficiency on SpMV execution.




Fig. 30. Energy comparison between the UPMEM PIM system, Intel Xeon CPU and Tesla V100 GPU on SpMV execution.

7.3 Discussion

These evaluations are useful for programmers to anticipate how much performance and energy savings memory-centric PIM systems can provide on SpMV over commodity processor-centric CPU and GPU systems. However, our evaluated SpMV kernels do not constitute the best-performing approaches for *all* matrices. Designing methods to select the best-performing SpMV parallelization scheme depending on the particular characteristics of the input matrix would further improve performance and energy savings of SpMV execution on memory-centric PIM systems. Moreover, the UPMEM PIM hardware is still maturing and is expected to run at a higher frequency in the near future (500 MHz instead of 350 MHz) [82, 241]. Hence, SpMV execution on the UPMEM PIM architecture might achieve even higher performance and energy benefits over the results we report in this comparison. Finally, note that our proposed *SparseP* kernels can be adapted and evaluated on other current and future real PIM systems with potentially higher computation capabilities and energy efficiency than the UPMEM PIM system.

8 Key Takeaways and Recommendations

This section summarizes our key takeaways in the form of recommendations to improve multiple aspects of PIM hardware and software.

Recommendation #1. Design algorithms that provide high load balance across threads of a PIM core in terms of computations, loop control iterations, synchronization points and memory accesses. Section 5 shows that in matrices and formats where the parallelization scheme used causes *high disparity* in the non-zero elements/blocks/rows processed across threads of a PIM core, or the

number of lock acquisitions/lock releases/DRAM memory accesses performed across threads, SpMV performance severely degrades in compute-bound DPUs [82, 94]. Therefore, from a programmer's perspective, providing high operation balance across parallel threads is of vital importance in low-area and low-power PIM cores with relatively low computation capabilities [82, 94].

Recommendation #2. *Design compressed data structures that can be effectively partitioned across DRAM banks, with the goal of providing high computation balance across PIM cores.* Sections 6.1.1 and 6.2.2 demonstrate that (i) the compressed matrix format used to store the input matrix determines the data partitioning across DRAM banks of PIM-enabled memory, and (ii) SpMV execution using the CSR and BCSR formats performs significantly worse than SpMV execution using the COO and BCOO formats. This is because the matrix is stored in row- or block-row-order for the CSR and BCSR formats, respectively, and thus data partitioning across DRAM banks is limited to be performed at row or block-row granularity, respectively, leading to high non-zero element imbalance across PIM cores. Therefore, we recommend that programmers design compressed data structures that can provide effective data partitioning schemes with high computation balance across thousands of PIM cores.

Recommendation #3. *Design adaptive algorithms that (i) trade off computation balance across PIM cores for lower data transfer costs to PIM-enabled memory, and (ii) adapt their configuration to the particular patterns of each input given, as well as the characteristics of the PIM hardware.* Our analysis in Sections 6.1.1, 6.2.1 and 6.2.3 demonstrates that the best-performing SpMV execution on the UPMEM PIM system can be achieved using algorithms that (i) trade off computation for lower data transfer costs, and (ii) select the load balancing strategy and data partitioning policy based on the particular sparsity pattern of the input matrix. In addition, the performance of each balancing scheme and data partitioning technique for SpMV execution highly depends on the characteristics of the underlying PIM hardware, as we explain in Section 6.2.1 and Appendix A.3. To this end, we recommend that software designers implement heuristics and selection methods for their algorithms to adapt their configuration to the underlying hardware characteristics of the PIM system and the input data given.

Recommendation #4. *Provide low-cost synchronization support and hardware support to enable concurrent memory accesses by multiple threads to the local DRAM bank to increase parallelism in a multithreaded PIM core.* Section 5 shows that (i) lock acquisitions/releases can cause high overheads in the DPU pipeline, and (ii) fine-grained locking approaches to increase parallelism in critical sections do not improve performance over coarse-grained approaches in the UPMEM PIM hardware. This is because the DMA engine of the DPU serializes DRAM memory accesses included in the critical sections. Based on these key takeaways, we recommend that hardware designers provide lightweight synchronization mechanisms for multithreaded PIM cores [81], and enable concurrent access to local DRAM memory arrays to increase execution parallelism. For example, sub-array level parallelism [42, 134] or multiple DRAM banks per PIM core could be supported in the PIM hardware to improve parallelism.

Recommendation #5. *Optimize the broadcast collective operation in data transfers from main memory to PIM-enabled memory to minimize overheads of copying the input data into all DRAM banks in the PIM system.* Figures 15 and 16 show that SpMV execution using the 1D partitioning technique cannot scale up to a large number of PIM cores. This is because it is severely limited by data transfer costs to broadcast the input vector into *each* DRAM bank of PIM-enabled DIMMs via the narrow off-chip memory bus. To this end, we suggest that hardware and system designers provide a fast broadcast collective primitive to DRAM banks of PIM-enabled memory modules [233].

Recommendation #6. *Optimize the gather collective operation at DRAM bank granularity for data transfers from PIM-enabled memory to the host CPU to minimize overheads of retrieving the output*

results. Figures 19, 20 and 21 demonstrate that SpMV execution using the *equally-wide* and *variable-sized* 2D partitioning schemes is severely limited by data transfers to retrieve results for the output vector from DRAM banks of PIM-enabled DIMMs. This is due to two reasons: (i) 2D-partitioned kernels create a large number of partial results that need to be transferred from PIM-enabled memory to the host CPU via the narrow memory bus in order to assemble the final output vector, and (ii) the UPMEM PIM system has the limitation that the transfer sizes from/to all DRAM banks involved in the same parallel transfer need to be the same, and therefore a large amount of padding with empty bytes is performed in the *equally-wide* and *variable-sized* schemes. To this end, we suggest that hardware and system designers provide an optimized *gather* primitive to efficiently collect results from multiple DRAM banks to host CPU [233], and support parallel fine-grained data transfers from PIM-enabled memory to host CPU *at DRAM bank granularity* to avoid padding with empty bytes.

Recommendation #7. *Design high-speed communication channels and optimized libraries for data transfers to/from thousands of DRAM banks of PIM-enabled memory.* Section 7 demonstrates that SpMV execution on the memory-centric UPMEM PIM system achieves a much higher fraction of the machine’s peak performance (on average 51.7% for the 32-bit float data type), compared to that on processor-centric CPU and GPU systems. However, the end-to-end performance of both 1D- and 2D-partitioned kernels is significantly limited by data transfer overheads on the narrow memory bus. To this end, we recommend that the hardware architecture and the software stack of real PIM systems be enhanced with low-cost and fast data transfers to/from PIM-enabled memory modules, and/or with support for efficient direct communication among PIM cores [41, 209, 217–219, 249].

9 Related Work

To our knowledge, this is the first work that (i) extensively characterizes the Sparse Matrix Vector Multiplication (SpMV) kernel in a real PIM system, and (ii) presents an open-source SpMV library for real-world PIM systems. We briefly discuss closely related prior work.

Processing-In-Memory (PIM). A large body of prior work examines Processing-Near-Memory (PNM) [3, 4, 8, 9, 16, 30–32, 39, 47, 52, 57, 66, 68, 76–78, 81, 89, 90, 101, 102, 109, 110, 112, 126, 129, 130, 144, 151, 166, 167, 176, 177, 179–181, 191, 194, 199, 206, 213, 224, 225, 233, 241, 269, 271, 280, 281]. PNM integrates processing units near or inside the memory via a 3D PNM configuration (i.e., processing units are located at the logic layer of 3D-stacked memories) [3, 30–32, 47, 57, 76, 166, 180, 181, 206, 269, 271, 281], a 2.5D PNM configuration (i.e., processing units are located in the same package as the CPU connected via silicon interposers) [68, 81, 224], a 2D PNM configuration (i.e., processing units are placed inside DDRX DIMMs) [9, 16, 44, 89, 90, 126, 143, 147, 148, 179, 185, 199, 213, 282], or at the memory controller of CPU systems [101, 102, 167]. These works propose hardware designs for irregular applications like graph processing [3, 4, 31, 32, 52, 180, 281], bioinformatics [39, 81, 130, 147, 148], neural networks [29, 30, 48, 68, 78, 89, 129, 224], pointer-chasing workloads [47, 81, 110, 166], and databases [57]. However, *none* of these works examines the SpMV kernel in such systems.

Several prior works enable Processing-Using-Memory (PUM) [2, 6, 11–13, 23, 41, 43, 54, 58, 69, 73–75, 96–98, 122, 131, 141, 142, 153, 156, 157, 209, 217, 218, 221, 227, 249, 253, 255, 258, 260, 267, 268, 278]. PUM exploits the operational principles of memory cells to perform computation within the memory chip. Prior works propose PUM designs using SRAM [2, 58, 73, 122], DRAM [6, 11, 41, 54, 69, 75, 96, 131, 156, 209, 217, 218, 249, 253, 260], PCM [157] or RRAM/memristive memory technologies [12, 13, 23, 43, 74, 97, 98, 141, 142, 153, 221, 227, 255, 258, 267, 268, 278]. A few PUM works [2, 43, 54, 58, 96, 156, 221] enable the multiplication operation inside memory cells with the goal of performing efficient matrix vector multiplication at low cost within the memory chip. These works design hardware-based solutions to accelerate the *dense* matrix vector multiplication (GEMV)

kernel via PUM. However, there is *no* prior work that leverages PUM to accelerate the *Sparse* Matrix Vector Multiplication (SpMV) kernel using state-of-the-art compressed matrix storage formats.

Sparse Matrix Kernels in PIM Systems. Xie et al. [259] design heterogenous PIM units to accelerate SpMV via a 3D PNM configuration, i.e., in HMC-based PIM systems. Sun et al. [233] leverage the buffer device space of DIMM modules to add one processing unit per each DIMM module, and design low-cost inter-DIMM broadcast collectives to minimize data transfer overheads on irregular workloads, like SpMV and graph processing, executed in 2D PNM configurations. Zhu et al. [280] propose a PIM accelerator for Sparse Matrix Matrix Multiplication via a 3D PNM configuration. Fujiki et al. [72] enhance the memory controllers of GPUs with PIM cores to transform the matrix from the CSR to the DCSR format [106] on the fly to minimize memory traffic on SpMV execution. These works propose hardware designs for sparse matrix kernels. In contrast, our work studies software optimizations and strategies to efficiently map compressed matrix storage formats on real near-bank PIM systems, and accelerate SpMV execution on such systems.

SpMV in Commodity Systems. Numerous prior works propose optimized SpMV algorithms for CPUs [5, 37, 59, 60, 62, 63, 108, 136, 165, 171, 172, 182, 193, 204, 210, 236–238, 246, 248, 251, 252, 256, 257, 274], GPUs [18, 27, 48, 61, 70, 91, 107, 162, 203, 228, 232, 234, 244, 254, 261, 265], heterogeneous CPU-GPU systems [10, 19, 34, 116, 117, 202, 262, 264], and distributed CPU systems [24, 28, 38, 40, 85, 125, 150, 161, 183, 196, 201, 243]. Optimized SpMV kernels for processor-centric CPU and GPU systems exploit the shared memory model of these systems and data locality in deep cache hierarchies. However, these kernels cannot be directly mapped to most near-bank PIM systems, which have a distributed memory model and a shallow cache hierarchy. Most well-tuned SpMV kernels for distributed CPU and CPU-GPU systems improve performance by overlapping computation with communication among processing units, and exploiting data locality in large cache memories. In contrast, real near-bank PIM architectures are fundamentally different from CPU-GPU systems, since they are *highly distributed*, i.e., there is no direct communication among PIM cores, and include a shallow memory hierarchy. Therefore, SpMV kernels designed for common processor-centric systems cannot be directly used in near-bank PIM systems.

Hardware Accelerators for SpMV. Recent works design accelerators for SpMV [71, 86, 121, 158, 175, 187, 212, 239] or other sparse kernels [15, 103, 113, 173, 188, 197, 198, 207, 208, 270, 272, 273, 279]. In contrast, our work proposes software optimizations and provides the first characterization study of SpMV on a real PIM system.

Compressed Matrix Storage Formats. Prior works propose a range of compressed matrix storage formats [17, 25, 36, 106, 114, 135, 137–139, 146, 163, 165, 169, 170, 174, 204, 205, 211, 214, 216, 248, 261, 263] and selection methods to find the most efficient compressed format [14, 20, 21, 154, 155, 168, 186, 215, 232, 235, 276, 277]. In this work, we extensively explore the four most widely used *general* compressed matrix formats, and observe that the compressed format (i) needs to provide good balance between computation and memory accesses inside the core pipeline, and (ii) affects load balancing across PIM cores, with corresponding performance implications. Therefore, some compressed formats designed for commodity processor-centric systems might not be suitable or efficient for real PIM systems. We leave the exploration of other PIM-suitable compressed matrix storage formats for future work.

10 Conclusion

We present *SparseP*, the first open-source SpMV library for real Processing-In-Memory (PIM) systems, and conduct the first comprehensive characterization analysis of the widely-used SpMV kernel on a real-world PIM architecture.

First, we design efficient SpMV kernels for real PIM systems. Our proposed *SparseP* software package supports (1) a wide range of data types, (2) two types of well-crafted data partitioning

techniques of the sparse matrix to DRAM banks of PIM-enabled memory, (3) the most popular compressed matrix formats, (4) a wide variety of load balancing schemes across PIM cores, (5) several load balancing schemes across threads of a multithreaded PIM core, and (6) three synchronization approaches among threads within PIM core.

Second, we conduct an extensive characterization study of *SparseP* kernels on the state-of-the-art UPMEM PIM system. We analyze SpMV execution on one single multithreaded PIM core and thousands of PIM cores using 26 sparse matrices with diverse sparsity patterns. We also compare the performance and energy consumption of SpMV on the UPMEM PIM system with those of state-of-the-art CPU and GPU systems to quantify the potential of a real memory-centric PIM architecture on the widely used SpMV kernel over conventional processor-centric architectures. Our analysis of *SparseP* kernels provides programming recommendations for software designers, as well as suggestions and hints for hardware and system designers of future PIM systems.

We believe and hope that our work will provide valuable insights to programmers in the development of efficient sparse linear algebra kernels and other irregular kernels from different application domains tailored for real PIM systems, as well as to architects and system designers in the development of future memory-centric computing systems.

Acknowledgments

We thank the UPMEM company for valuable support. We thank the anonymous reviewers from SIGMETRICS 2022, and our shepherd, Bhuvan Urgaonkar, for their comments and suggestions. We thank the SAFARI Research Group members for feedback and the stimulating, scholarly and collaborative intellectual environment they provide. We thank the CSLAB Research Group members for continued and undivided support, insightful comments and valuable feedback. We acknowledge the support of SAFARI Research Group’s industrial partners, especially ASML, Facebook, Google, Huawei, Intel, Microsoft, VMware, the Semiconductor Research Corporation and the ETH Future Computing Laboratory. Christina Giannoula is funded for her postgraduate studies from the Foundation for Education and European Culture. The *SparseP* software package is publicly available at <https://github.com/CMU-SAFARI/SparseP>.

References

- [1] Shaizeen Aga, Nuwan Jayasena, and Mike Ignatowski, *Co-ML: A Case for Collaborative ML Acceleration Using near-Data Processing*, MEMSYS, 2019.
- [2] Shaizeen Aga, Supreet Jeloka, Arun Subramaniyan, Satish Narayanasamy, David Blaauw, and Reetuparna Das, *Compute Caches*, HPCA, 2017.
- [3] Junwhan Ahn, Sungpack Hong, Sungjoo Yoo, Onur Mutlu, and Kiyoungh Choi, *A Scalable Processing-In-Memory Accelerator for Parallel Graph Processing*, ISCA, 2015.
- [4] Junwhan Ahn, Sungjoo Yoo, Onur Mutlu, and Kiyoungh Choi, *PIM-Enabled Instructions: A Low-overhead, Locality-Aware Processing-in-Memory Architecture*, ISCA, 2015.
- [5] Baris Aktemur, *A Sparse Matrix-Vector Multiplication Method with Low Preprocessing Cost, Concurrency and Computation: Practice and Experience*, 2018.
- [6] Mustafa F. Ali, Akhilesh Jaiswal, and Kaushik Roy, *In-Memory Low-Cost Bit-Serial Addition Using Commodity DRAM Technology*, IEEE TCSE, 2020.
- [7] Mohammad Alian and Nam Sung Kim, *NetDIMM: Low-Latency Near-Memory Network Interface Architecture*, MICRO, 2019.
- [8] Mohammed Alser, Zülal Bingöl, Damla Senol Cali, Jeremie Kim, Saugata Ghose, Can Alkan, and Onur Mutlu, *Accelerating Genome Analysis: A Primer on an Ongoing Journey*, IEEE Micro, 2020.
- [9] Marco A. Z. Alves, Paulo C. Santos, Matthias Diener, and Luigi Carro, *Opportunities and Challenges of Performing Vector Operations inside the DRAM*, MEMSYS, 2015.
- [10] Petros Anastasiadis, Nikela Papadopoulos, Georgios Goumas, and Nectarios Koziris, *CoCoPeLia: Communication-Computation Overlap Prediction for Efficient Linear Algebra on GPUs*, ISPASS, 2021.
- [11] Shaahin Angizi and Deliang Fan, *GraphiDe: A Graph Processing Accelerator Leveraging In-DRAM-Computing*, GLSVLSI, 2019.
- [12] Aayush Ankit, Izzat El Hajj, Sai Rahul Chalamalasetti, Sapan Agarwal, Matthew Marinella, Martin Foltin, John Paul Strachan, Dejan Milojicic, Wen-Mei Hwu, and Kaushik Roy, *PANTHER: A Programmable Architecture for Neural Network Training Harnessing Energy-Efficient ReRAM*, IEEE TC, 2020.
- [13] Aayush Ankit, Izzat El Hajj, Sai Rahul Chalamalasetti, Geoffrey Ndu, Martin Foltin, R. Stanley Williams, Paolo Faraboschi, Wen-mei W Hwu, John Paul Strachan, Kaushik Roy, and Dejan S. Milojicic, *PUMA: A Programmable Ultra-Efficient Memristor-Based Accelerator for Machine Learning Inference*, ASPLOS, 2019.
- [14] Bahar Asgari, Ramyad Hadidi, Joshua Dierberger, Charlotte Steinichen, and Hyesoon Kim, *Copernicus: Characterizing the Performance Implications of Compression Formats Used in Sparse Workloads*, CoRR, 2020.
- [15] Bahar Asgari, Ramyad Hadidi, Tushar Krishna, Hyesoon Kim, and Sudhakar Yalamanchili, *ALRESCHA: A Lightweight Reconfigurable Sparse-Computation Accelerator*, HPCA, 2020.
- [16] Hadi Asghari-Moghaddam, Young Hoon Son, Jung Ho Ahn, and Nam Sung Kim, *Chameleon: Versatile and Practical Near-DRAM Acceleration Architecture for Large Memory Systems*, MICRO, 2016.
- [17] Mehmet Belgin, Godmar Back, and Calvin J. Ribbens, *Pattern-Based Sparse Matrix Representation for Memory-Efficient SVM Kernels*, ICS, 2009.
- [18] Nathan Bell and Michael Garland, *Implementing Sparse Matrix-Vector Multiplication on Throughput-Oriented Processors*, SC, 2009.
- [19] Akrem Benatia, Weixing Ji, and Yizhuo Wang, *Sparse Matrix Partitioning for Optimizing SpMV on CPU-GPU Heterogeneous Platforms*, IJHPCA, 2019.
- [20] Akrem Benatia, Weixing Ji, Yizhuo Wang, and Feng Shi, *Sparse Matrix Format Selection with Multiclass SVM for SpMV on GPU*, ICPP, 2016.
- [21] Akrem Benatia, Weixing Ji, Yizhuo Wang, and Feng Shi, *BestSF: A Sparse Meta-Format for Optimizing SpMV on GPU*, TACO, 2018.
- [22] Maciej Besta, Florian Marending, Edgar Solomonik, and Torsten Hoefer, *SlimSell: A Vectorizable Graph Representation for Breadth-First Search*, IPDPS, 2017.
- [23] Debjyoti Bhattacharjee, Rajeswari Devadoss, and Anupam Chattopadhyay, *ReVAMP: ReRAM based VLIW Architecture for In-Memory Computing*, DATE, 2017.
- [24] Rob H. Bisseling and Wouter Meesen, *Communication Balancing in Parallel Sparse Matrix-Vector Multiplication*, ETNA. Electronic Transactions on Numerical Analysis, 2005.
- [25] Åke Björck, *Numerical Methods for Least Squares Problems*, SIAM, 1996.
- [26] Jeff Bolz, Ian Farmer, Eitan Grinspun, and Peter Schröder, *Sparse Matrix Solvers on the GPU: Conjugate Gradients and Multigrid*, SIGGRAPH, 2003.
- [27] Jeff Bolz, Ian Farmer, Eitan Grinspun, and Peter Schröder, *Sparse Matrix Solvers on the GPU: Conjugate Gradients and Multigrid*, ACM Transactions on Graphics, 2003.

- [28] Erik G. Boman, Karen D. Devine, and Sivasankaran Rajamanickam, *Scalable Matrix Computations on Large Scale-Free Graphs Using 2D Graph Partitioning*, SC, 2013.
- [29] Amirali Boroumand, Saugata Ghose, Berkin Akin, Ravi Narayanaswami, Geraldo F. Oliveira, Xiaoyu Ma, Eric Shiu, and Onur Mutlu, *Google Neural Network Models for Edge Devices: Analyzing and Mitigating Machine Learning Inference Bottlenecks*, PACT, 2021.
- [30] Amirali Boroumand, Saugata Ghose, Youngsok Kim, Rachata Ausavarungnirun, Eric Shiu, Rahul Thakur, Daehyun Kim, Aki Kuusela, Allan Knies, Parthasarathy Ranganathan, and Onur Mutlu, *Google Workloads for Consumer Devices: Mitigating Data Movement Bottlenecks*, ASPLOS, 2018.
- [31] Amirali Boroumand, Saugata Ghose, Minesh Patel, Hasan Hassan, Brandon Lucia, Rachata Ausavarungnirun, Kevin Hsieh, Nastaran Hajinazar, Krishna T. Malladi, Hongzhong Zheng, and Onur Mutlu, *CoNDA: Efficient Cache Coherence Support for Near-data Accelerators*, ISCA, 2019.
- [32] Amirali Boroumand, Saugata Ghose, Minesh Patel, Hasan Hassan, Brandon Lucia, Kevin Hsieh, Krishna T. Malladi, Hongzhong Zheng, and Onur Mutlu, *LazyPIM: An Efficient Cache Coherence Mechanism for Processing-in-Memory*, CAL, 2017.
- [33] F. Nisa Bostanci, Ataberk Olgun, Lois Orosa, A. Giray Yağlıkcı, Jeremie S. Kim, Hasan Hassan, Oğuz Ergin, and Onur Mutlu, *DR-STRaNGe: End-to-End System Design for DRAM-based True Random Number Generators*, HPCA, 2022.
- [34] Brice Boyer, Jean-Guillaume Dumas, and Pascal Giorgi, *Exact Sparse Matrix-Vector Multiplication on GPU's and Multicore Architectures*, PASC, 2010.
- [35] Sergey Brin and Lawrence Page, *The Anatomy of a Large-scale Hypertextual Web Search Engine*, WWW, 1998.
- [36] Aydin Buluç, Jeremy T. Fineman, Matteo Frigo, John R. Gilbert, and Charles E. Leiserson, *Parallel Sparse Matrix-Vector and Matrix-Transpose-Vector Multiplication Using Compressed Sparse Blocks*, SPAA, 2009.
- [37] Aydin Buluç, Samuel Williams, Leonid Oliker, and James Demmel, *Reduced-Bandwidth Multithreaded Algorithms for Sparse Matrix-Vector Multiplication*, IPDPS, 2011.
- [38] Beata Bylina, Jarosław Bylina, Przemysław Stpiczynski, and Dominik Szalkowski, *Performance Analysis of Multicore and Multinodal Implementation of SpMV Operation*, FedCSIS, 2014.
- [39] D. S. Cali, G. S. Kalsi, Z. Bingöl, C. Firtina, L. Subramanian, J. S. Kim, R. Ausavarungnirun, M. Alser, J. Gomez-Luna, A. Boroumand, A. Norion, A. Scibisz, S. Subramoneyon, C. Alkan, S. Ghose, and O. Mutlu, *GenASM: A High-Performance, Low-Power Approximate String Matching Acceleration Framework for Genome Sequence Analysis*, MICRO, 2020.
- [40] U.V. Catalyurek and C. Aykanat, *Hypergraph-Partitioning-Based Decomposition for Parallel Sparse-Matrix Vector Multiplication*, IEEE TPDS, 1999.
- [41] Kevin K. Chang, Prashant J. Nair, Donghyuk Lee, Saugata Ghose, Moinuddin K. Qureshi, and Onur Mutlu, *Low-Cost Inter-Linked Subarrays (LISA): Enabling Fast Inter-Subarray Data Movement in DRAM*, HPCA, 2016.
- [42] Kevin Kai-Wei Chang, Donghyuk Lee, Zeshan Chishti, Alaa R. Alameldeen, Chris Wilkerson, Yoongu Kim, and Onur Mutlu, *Improving DRAM performance by parallelizing refreshes with accesses*, HPCA, 2014.
- [43] Ping Chi, Shuangchen Li, Cong Xu, Tao Zhang, Jishen Zhao, Yongpan Liu, Yu Wang, and Yuan Xie, *PRIME: A Novel Processing-in-Memory Architecture for Neural Network Computation in ReRAM-Based Main Memory*, ISCA, 2016.
- [44] Benjamin Y. Cho, Jeageun Jung, and Mattan Erez, *Accelerating Bandwidth-Bound Deep Learning Inference with Main-Memory Accelerators*, SC, 2021.
- [45] Benjamin Y. Cho, Yongkee Kwon, Sangkug Lym, and Mattan Erez, *Near Data Acceleration with Concurrent Host Access*, ISCA, 2020.
- [46] Seunghwan Cho, Haerang Choi, Eunhyeok Park, Hyunsung Shin, and Sungjoo Yoo, *McDRAM v2: In-Dynamic Random Access Memory Systolic Array Accelerator to Address the Large Model Problem in Deep Neural Networks on the Edge*, IEEE Access, 2020.
- [47] Jiwon Choe, Amy Huang, Tali Moreshet, Maurice Herlihy, and R. Iris Bahar, *Concurrent Data Structures with Near-Data-Processing: An Architecture-Aware Implementation*, SPAA, 2019.
- [48] Jee W. Choi, Amik Singh, and Richard W. Vuduc, *Model-Driven Autotuning of Sparse Matrix-Vector Multiply on GPUs*, PpopP, 2010.
- [49] CSR5, CSR5 Cuda, https://github.com/weifengliu-ssslab/Benchmark_SpMV_using_CSR5, 2015.
- [50] cuSparse, cuSparse, <https://docs.nvidia.com/cuda/cusparse/index.html>, 2021.
- [51] Leonardo Dagum and Ramesh Menon, *OpenMP: An Industry-Standard API for Shared-Memory Programming*, IEEE Comput. Sci. Eng., 1998.
- [52] Guohao Dai, Tianhao Huang, Yuze Chi, Jishen Zhao, Guangyu Sun, Yongpan Liu, Yu Wang, Yuan Xie, and Huazhong Yang, *GraphH: A Processing-in-Memory Architecture for Large-Scale Graph Processing*, IEEE TCAD, 2018.
- [53] Timothy A. Davis and Yifan Hu, *The University of Florida Sparse Matrix Collection*, TOMS, 2011.
- [54] Quan Deng, Lei Jiang, Youtao Zhang, Minxuan Zhang, and Jun Yang, *DrAcc: a DRAM based Accelerator for Accurate CNN Inference*, DAC, 2018.
- [55] F. Devaux, *The True Processing In Memory Accelerator*, Hot Chips, 2019.

- [56] Jack Dongarra, Andrew Lumsdaine, Xinhui Niu, Roldan Pozoz, and Karin Remington, *Sparse Matrix Libraries in C++ for High Performance Architectures*, Mathematics, 1994.
- [57] Mario Drumond, Alexandros Daglis, Nooshin Mirzadeh, Dmitrii Ustiugov, Javier Picorel, Babak Falsafi, Boris Grot, and Dionisios Pnevmatikatos, *The Mondrian Data Engine*, ISCA, 2017.
- [58] Charles Eckert, Xiaowei Wang, Jingcheng Wang, Arun Subramaniyan, Ravi Iyer, Dennis Sylvester, David Blaauw, and Reetuparna Das, *Neural Cache: Bit-Serial in-Cache Acceleration of Deep Neural Networks*, ISCA, 2018.
- [59] Athena Elafrou, Georgios Goumas, and Nectarios Koziris, *Performance Analysis and Optimization of Sparse Matrix-Vector Multiplication on Intel Xeon Phi*, IPDPSW, 2017.
- [60] Athena Elafrou, G. Goumas, and N. Koziris, *Performance Analysis and Optimization of Sparse Matrix-Vector Multiplication on Modern Multi- and Many-Core Processors*, ICPP, 2017.
- [61] Athena Elafrou, Georgios Goumas, and Nectarios Koziris, *BASMAT: Bottleneck-Aware Sparse Matrix-Vector Multiplication Auto-Tuning on GPGPUs*, PPOPP, 2019.
- [62] Athena Elafrou, Georgios Goumas, and Nectarios Koziris, *Conflict-Free Symmetric Sparse Matrix-Vector Multiplication on Multicore Architectures*, SC, 2019.
- [63] Athena Elafrou, Vasileios Karakasis, Theodoros Gkountouvas, Kornilios Kourtis, Georgios Goumas, and Nectarios Koziris, *SparseX: A Library for High-Performance Sparse Matrix-Vector Multiplication on Multicore Platforms*, ACM TOMS, 2018.
- [64] R. D. Falgout, *An Introduction to Algebraic Multigrid*, Computing in Science Engineering, 2006.
- [65] Robert D Falgout and Ulrike Meier Yang, *hypr: A Library of High Performance Preconditioners*, ICCS, 2002.
- [66] Amin Farmahini-Farahani, Jung Ho Ahn, Katherine Morrow, and Nam Sung Kim, *NDA: Near-DRAM Acceleration Architecture Leveraging Commodity DRAM Devices and Standard Memory Modules*, HPCA, 2015.
- [67] Amin Farmahini-Farahani, Jung Ho Ahn, Katherine Morrow, and Nam Sung Kim, *DRAMA: An Architecture for Accelerated Processing Near Memory*, IEEE CAL, 2015.
- [68] Ivan Fernandez, Ricardo Quislan, Christina Giannoula, Mohammed Alser, Juan Gómez-Luna, Eladio Gutiérrez, Oscar Plata, and Onur Mutlu, *NATSA: A Near-Data Processing Accelerator for Time Series Analysis*, ICCD, 2020.
- [69] João Dinis Ferreira, Gabriel Falcao, Juan Gómez-Luna, Mohammed Alser, Lois Orosa, Mohammad Sadrosadati, Jeremie S. Kim, Geraldo F. Oliveira, Taha Shahroodi, Anant Nori, and Onur Mutlu, *pLUTo: Enabling Massively Parallel Computation In DRAM via Lookup Tables*, CoRR, 2021.
- [70] Salvatore Filippone, Valeria Cardellini, Davide Barbieri, and Alessandro Fanfarillo, *Sparse Matrix-Vector Multiplication on GPGPUs*, ACM TOMS, 2017.
- [71] Jeremy Fowers, Kalin Ovtcharov, Karin Strauss, Eric S. Chung, and Greg Stitt, *A High Memory Bandwidth FPGA Accelerator for Sparse Matrix-Vector Multiplication*, FCCM, 2014.
- [72] Daichi Fujiki, Niladri Chatterjee, Donghyuk Lee, and Mike O'Connor, *Near-Memory Data Transformation for Efficient Sparse Matrix Multi-Vector Multiplication*, SC, 2019.
- [73] Daichi Fujiki, Scott Mahlke, and Reetuparna Das, *Duality Cache for Data Parallel Acceleration*, ISCA, 2019.
- [74] Pierre-Emmanuel Gaillardon, Luca Amarú, Anne Simon, Eike Linn, Rainer Waser, Anupam Chattopadhyay, and Giovanni De Micheli, *The Programmable Logic-in-Memory (PLiM) computer*, DATE, 2016.
- [75] Fei Gao, Georgios Tziantzioulis, and David Wentzlaff, *ComputeDRAM: In-Memory Compute Using Off-the-Shelf DRAMs*, MICRO, 2019.
- [76] Mingyu Gao, Grant Ayers, and Christos Kozyrakis, *Practical Near-Data Processing for In-Memory Analytics Frameworks*, PACT, 2015.
- [77] Mingyu Gao and Christos Kozyrakis, *HRL: Efficient and Flexible Reconfigurable Logic for Near-Data Processing*, HPCA, 2016.
- [78] Mingyu Gao, Jing Pu, Xuan Yang, Mark Horowitz, and Christos Kozyrakis, *TETRIS: Scalable and Efficient Neural Network Acceleration with 3D Memory*, ASPLOS, 2017.
- [79] Saugata Ghose, Amirali Boroumand, Jeremie Kim, Juan Gómez-Luna, and Onur Mutlu, *Processing-in-Memory: A Workload-Driven Perspective*, IBM JRD, 2019.
- [80] Christina Giannoula, Georgios Goumas, and Nectarios Koziris, *Combining HTM with RCU to Speed up Graph Coloring on Multicore Platforms*, ISC HPC, 2018.
- [81] Christina Giannoula, Nandita Vijaykumar, Nikela Papadopoulou, Vasileios Karakostas, Ivan Fernandez, Juan Gómez-Luna, Lois Orosa, Nectarios Koziris, Georgios I. Goumas, and Onur Mutlu, *SynCron: Efficient Synchronization Support for Near-Data-Processing Architectures*, HPCA, 2021.
- [82] Juan Gómez-Luna, Izzat El Hajj, Ivan Fernandez, Christina Giannoula, Geraldo F. Oliveira, and Onur Mutlu, *Benchmarking a New Paradigm: An Experimental Analysis of a Real Processing-in-Memory Architecture*, CoRR, 2021.
- [83] Georgios Goumas, Kornilios Kourtis, Nikos Anastopoulos, Vasileios Karakasis, and Nectarios Koziris, *Understanding the Performance of Sparse Matrix-Vector Multiplication*, PDP, 2008.

- [84] Georgios Goumas, Kornilios Kourtis, Nikos Anastopoulos, Vasileios Karakasis, and Nectarios Koziris, *Performance Evaluation of the Sparse Matrix-Vector Multiplication on Modern Architectures*, *J. Supercomput.*, 2009.
- [85] Anael Grandjean, Johannes Langguth, and Bora Uçar, *On Optimal and Balanced Sparse Matrix Partitioning Problems*, *CLUSTER*, 2012.
- [86] Paul Grigoras, Pavel Burovskiy, Eddie Hung, and Wayne Luk, *Accelerating SpMV on FPGAs by Compressing Nonzero Values*, *FCCM*, 2015.
- [87] SAFARI Research Group, *PrIM Benchmark Suite*, <https://github.com/CMUSAFARI/prim-benchmarks>, 2021.
- [88] SAFARI Research Group, *SparseP Software Package*, <https://github.com/CMU-SAFARI/SparseP>, 2022.
- [89] Peng Gu, Xinfeng Xie, Yufei Ding, Guoyang Chen, Weifeng Zhang, Dimin Niu, and Yuan Xie, *iPIM: Programmable In-Memory Image Processing Accelerator Using Near-Bank Architecture*, *ISCA*, 2020.
- [90] Peng Gu, Xinfeng Xie, Shuangchen Li, Dimin Niu, Hongzhong Zheng, Krishna T. Malladi, and Yuan Xie, *DLUX: A LUT-Based Near-Bank Accelerator for Data Center Deep Learning Training Workloads*, *IEEE TCAD*, 2021.
- [91] Ping Guo, Liqiang Wang, and Po Chen, *A Performance Modeling and Optimization Analysis Tool for Sparse Matrix-Vector Multiplication on GPUs*, *IEEE TPDS*, 2014.
- [92] Udit Gupta, Samuel Hsia, Vikram Saraph, Xiaodong Wang, Brandon Reagan, Gu-Yeon Wei, Hsien-Hsin S. Lee, David Brooks, and Carole-Jean Wu, *DeepRecSys: A System for Optimizing End-To-End At-Scale Neural Recommendation Inference*, *ISCA*, 2020.
- [93] Udit Gupta, Xiaodong Wang, Maxim Naumov, Carole-Jean Wu, Brandon Reagan, David Brooks, Bradford Cottel, Kim M. Hazelwood, Bill Jia, Hsien-Hsin S. Lee, Andrey Malevich, Dheevatsa Mudigere, Mikhail Smelyanskiy, Liang Xiong, and Xuan Zhang, *The Architectural Implications of Facebook’s DNN-based Personalized Recommendation*, *CoRR*, 2019.
- [94] Juan Gómez-Luna, Izzat El Hajj, Ivan Fernandez, Christina Giannoula, Geraldo F. Oliveira, and Onur Mutlu, *Benchmarking Memory-Centric Computing Systems: Analysis of Real Processing-In-Memory Hardware*, *IGSC*, 2021.
- [95] Ramyad Hadidi, Bahar Asgari, Burhan Ahmad Mudassar, Saibal Mukhopadhyay, Sudhakar Yalamanchili, and Hyesoon Kim, *Demystifying the Characteristics of 3D-stacked Memories: A case Study for Hybrid Memory Cube*, *IISWC*, 2017.
- [96] Nastaran Hajinazar, Geraldo F. Oliveira, Sven Gregorio, João Dinis Ferreira, Nika Mansouri Ghiasi, Minesh Patel, Mohammed Alser, Saugata Ghose, Juan Gómez-Luna, and Onur Mutlu, *SIMDRAM: A Framework for Bit-Serial SIMD Processing Using DRAM*, *ASPLOS*, 2021.
- [97] Said Hamdioui, Shahar Kvatinaky, Gert Cauwenberghs, Lei Xie, Nimrod Wald, Siddharth Joshi, Hesham Mostafa Elsayed, Henk Corporaal, and Koen Bertels, *Memristor for Computing: Myth or Reality?*, *DATE*, 2017.
- [98] Said Hamdioui, Lei Xie, Hoang Anh Du Nguyen, Mottaqiallah Taozil, Koen Bertels, Henk Corporaal, Hailong Jiao, Francky Catthoor, Dirk Wouters, Linn Eike, and Jan van Lunteren, *Memristor Based Computation-In-Memory Architecture for Data-Intensive Applications*, *DATE*, 2015.
- [99] Song Han, Xingyu Liu, Huizi Mao, Jing Pu, Ardavan Pedram, Mark A. Horowitz, and William J. Dally, *EIE: Efficient Inference Engine on Compressed Deep Neural Network*, *ISCA*, 2016.
- [100] Song Han, Jeff Pool, John Tran, and William J. Dally, *Learning Both Weights and Connections for Efficient Neural Networks*, *NIPS*, 2015.
- [101] Milad Hashemi, Eiman Ebrahimi, Onur Mutlu, Yale N Patt, et al., *Accelerating Dependent Cache Misses with an Enhanced Memory Controller*, *ISCA*, 2016.
- [102] Milad Hashemi, Onur Mutlu, and Yale N Patt, *Continuous Runahead: Transparent Hardware Acceleration for Memory Intensive Workloads*, *MICRO*, 2016.
- [103] Kartik Hegde, Hadi Asghari-Moghaddam, Michael Pellauer, Neal Crago, Aamer Jaleel, Edgar Solomonik, Joel Emer, and Christopher W. Fletcher, *ExTensor: An Accelerator for Sparse Tensor Algebra*, *MICRO*, 2019.
- [104] Ahmed E. Helal, Jan Laukemann, Fabio Checconi, Jesmin Jahan Tithi, Teresa Ranadive, Fabrizio Petrini, and Jeewhan Choi, *ALTO: Adaptive Linearized Storage of Sparse Tensors*, *ICS*, 2021.
- [105] Pascal Hénon, Pierre Ramet, and Jean Roman, *PASTIX: A High-Performance Parallel Direct Solver for Sparse Symmetric Positive Definite Systems*, *PMAA*, 2002.
- [106] Changwan Hong, Aravind Sukumaran-Rajam, Bortik Bandyopadhyay, Jinsung Kim, Süreyya Emre Kurt, Israt Nisa, Shivani Sabhlok, Ümit V. Çatalyürek, Srinivasan Parthasarathy, and P. Sadayappan, *Efficient Sparse-Matrix Multi-Vector Product on GPUs*, *HPDC*, 2018.
- [107] Changwan Hong, Aravind Sukumaran-Rajam, Bortik Bandyopadhyay, Jinsung Kim, Süreyya Emre Kurt, Israt Nisa, Shivani Sabhlok, Ümit V. Çatalyürek, Srinivasan Parthasarathy, and P. Sadayappan, *Efficient Sparse-Matrix Multi-Vector Product on GPUs*, *HPDC*, 2018.
- [108] Kaixi Hou, Wu-chun Feng, and Shuai Che, *Auto-Tuning Strategies for Parallelizing Sparse Matrix-Vector (SpMV) Multiplication on Multi- and Many-Core Processors*, *IPDPSW*, 2017.
- [109] Kevin Hsieh, Eiman Ebrahimi, Gwangsun Kim, Niladri Chatterjee, Mike O’Connor, Nandita Vijaykumar, Onur Mutlu, and Stephen W. Keckler, *Transparent Offloading and Mapping: Enabling Programmer-Transparent Near-Data*

- Processing in GPU Systems*, ISCA, 2016.
- [110] Kevin Hsieh, Samira Khan, Nandita Vijaykumar, Kevin Chang, Amirali Boroumand, Saugata Ghose, and Onur Mutlu, *Accelerating Pointer Chasing in 3D-stacked Memory: Challenges, Mechanisms, Evaluation*, ICCD, 2016.
 - [111] Jiayi Huang, Ramprakash Reddy Puli, Pritam Majumder, Sungkeun Kim, Rahul Boyapati, Ki Hwan Yum, and Eun Jung Kim, *Active-Routing: Compute on the Way for Near-Data Processing*, HPCA, 2019.
 - [112] Wenqin Huangfu, Xueqi Li, Shuangchen Li, Xing Hu, Peng Gu, and Yuan Xie, *Medal: Scalable DIMM based Near Data Processing Accelerator for DNA Seeding Algorithm*, MICRO, 2019.
 - [113] Ranggi Hwang, Taehun Kim, Youngeun Kwon, and Minsoo Rhu, *Centaur: A Chiplet-based, Hybrid Sparse-Dense Accelerator for Personalized Recommendations*, ISCA, 2020.
 - [114] Eun-Jin Im and Katherine A. Yelick, *Optimizing Sparse Matrix Vector Multiplication on SMP*, PPSC, 1999.
 - [115] Eun-Jin Im, Katherine Yelick, and Richard Vuduc, *Sparsity: Optimization Framework for Sparse Matrix Kernels*, IJHPCA, 2004.
 - [116] Sivaramakrishna Bharadwaj Indarapu, Manoj Maramreddy, and Kishore Kothapalli, *Architecture- and Workload-Aware Heterogeneous Algorithms for Sparse Matrix Vector Multiplication*, ICPADS, 2013.
 - [117] Sivaramakrishna Bharadwaj Indarapu, Manoj Maramreddy, and Kishore Kothapalli, *Architecture- and Workload-Aware Heterogeneous Algorithms for Sparse Matrix Vector Multiplication*, COMPUTE, 2014.
 - [118] Intel, *Intel Xeon Silver 4110 Processor*, <https://ark.intel.com/content/www/us/en/ark/products/123547/intel-xeon-silver-4110-processor-11m-cache-2-10-ghz.html>, 2017.
 - [119] JEDEC, *JESD79-4 DDR4 SDRAM standard*, JEDEC, 2012.
 - [120] Hongshin Jun, Jinhee Cho, Kangseol Lee, Ho-Young Son, Kwiwook Kim, Hanho Jin, and Keith Kim, *HBM DRAM Technology and Architecture*, IMW, 2017.
 - [121] Konstantinos Kanellopoulos, Nandita Vijaykumar, Christina Giannoula, Roknoddin Azizi, Skanda Koppula, Nika Mansouri Ghiasi, Taha Shahroodi, Juan Gomez Luna, and Onur Mutlu, *SMASH: Co-Designing Software Compression and Hardware-Accelerated Indexing for Efficient Sparse Matrix Operations*, MICRO, 2019.
 - [122] Mingu Kang, Min-Sun Keel, Naresh R. Shanbhag, Sean Eilert, and Ken Curewitz, *An Energy-Efficient VLSI Architecture for Pattern Recognition via Deep Embedding of Computation in SRAM*, ICASSP, 2014.
 - [123] Vasileios Karakasis, Georgios Goumas, and Nectarios Koziris, *Performance Models for Blocked Sparse Matrix-Vector Multiplication Kernels*, ICPP, 2009.
 - [124] W. H. Kautz, *Cellular Logic-in-Memory Arrays*, IEEE TC, 1969.
 - [125] Enver Kayaaslan, Bora Uçar, and Cevdet Aykanat, *Semi-Two-Dimensional Partitioning for Parallel Sparse Matrix-Vector Multiplication*, IPDPS Workshop, 2015.
 - [126] Liu Ke, Udit Gupta, Carole-Jean Wu, Benjamin Youngjae Cho, Mark Hempstead, Brandon Reagen, Xuan Zhang, David Brooks, Vikas Chandra, Utku Diril, et al., *RecNMP: Accelerating Personalized Recommendation with Near-Memory Processing*, ISCA, 2020.
 - [127] Kamil Khan, Sudeep Pasricha, and Ryan Gary Kim, *A Survey of Resource Management for Processing-In-Memory and Near-Memory Processing Architectures*, Journal of Low Power Electronics and Applications, 2020.
 - [128] Kashif Nizam Khan, Mikael Hirki, Tapio Niemi, Jukka K Nurminen, and Zhonghong Ou, *Rapl in Action: Experiences in Using RAPL for Power Measurements*, TOMPECS, 2018.
 - [129] Duckhwan Kim, Jaeha Kung, Sek Chai, Sudhakar Yalamanchili, and Saibal Mukhopadhyay, *Neurocube: A Programmable Digital Neuromorphic Architecture with High-Density 3D Memory*, ISCA, 2016.
 - [130] Jeremie Kim, Damla Senol Cali, Hongyi Xin, Donghyuk Lee, Saugata Ghose, Mohammed Alser, Hasan Hassan, Oguz Ergin, Can Alkan, and Onur Mutlu, *GRIM-Filter: Fast Seed Location Filtering in DNA Read Mapping Using Processing-in-Memory Technologies*, BMC Genomics, 2018.
 - [131] Jeremie S. Kim, Minesh Patel, Hasan Hassan, and Onur Mutlu, *The DRAM Latency PUF: Quickly Evaluating Physical Unclonable Functions by Exploiting the Latency-Reliability Tradeoff in Modern Commodity DRAM Devices*, HPCA, 2018.
 - [132] Jeremie S. Kim, Minesh Patel, Hasan Hassan, Lois Orosa, and Onur Mutlu, *D-RaNGe: Using Commodity DRAM Devices to Generate True Random Numbers with Low Latency and High Throughput*, HPCA, 2019.
 - [133] Min-Jae Kim, Jeong-Geun Kim, Su-Kyung Yoon, and Shin-Dug Kim, *Functionality-Based Processing-in-Memory Accelerator for Deep Convolutional Neural Networks*, IEEE Access, 2021.
 - [134] Yoongu Kim, Vivek Seshadri, Donghyuk Lee, Jamie Liu, and Onur Mutlu, *A Case for Exploiting Subarray-Level Parallelism (SALP) in DRAM*, ISCA, 2012.
 - [135] David R Kincaid, Thomas C Oppe, and David M Young, *Itpackv 2D User's Guide*, , 1989.
 - [136] Fredrik Kjolstad, Stephen Chou, David Lugato, Shoaib Kamil, and Saman Amarasinghe, *TACO: A Tool to Generate Tensor Algebra Kernels*, ASE, 2017.
 - [137] Kornilios Kourtis, Georgios Goumas, and Nectarios Koziris, *Optimizing Sparse Matrix-Vector Multiplication Using Index and Value Compression*, CF, 2008.

- [138] Kornilios Kourtis, Vasileios Karakasis, Georgios Goumas, and Nectarios Koziris, *CSX: An Extended Compression Format for Spmv on Shared Memory Systems*, PPOPP, 2011.
- [139] Moritz Kreutzer, Georg Hager, Gerhard Wellein, Holger Fehske, Achim Basermann, and Alan R. Bishop, *Sparse Matrix-vector Multiplication on GPGPU Clusters: A New Storage Format and a Scalable Implementation*, IPDPSW, 2012.
- [140] Pranith Kumar and Hyesoon Kim, *Parallel Hash Table Design for NDP Systems*, MEMSYS, 2020.
- [141] Shahar Kvatinsky, Dmitry Belousov, Slavik Liman, Guy Satat, Nimrod Wald, Eby G. Friedman, Avinoam Kolodny, and Uri C. Weiser, *MAGIC—Memristor-Aided Logic*, IEEE TTCSII, 2014.
- [142] Shahar Kvatinsky, Avinoam Kolodny, Uri C. Weiser, and Eby G. Friedman, *Memristor-Based IMPLY Logic Design Procedure*, ICCD, 2011.
- [143] Youngeun Kwon, Yunjae Lee, and Minsoo Rhu, *TensorDIMM: A Practical Near-Memory Processing Architecture for Embeddings and Tensor Operations in Deep Learning*, MICRO, 2019.
- [144] Youngeun Kwon, Yunjae Lee, and Minsoo Rhu, *TensorDIMM: A Practical Near-Memory Processing Architecture for Embeddings and Tensor Operations in Deep Learning*, MICRO, 2019.
- [145] Young-Cheon Kwon, Suk Han Lee, Jaehoon Lee, Sang-Hyuk Kwon, Je Min Ryu, Jong-Pil Son, O Seongil, Hak-Soo Yu, Haesuk Lee, Soo Young Kim, Youngmin Cho, Jin Guk Kim, Jongyoon Choi, Hyun-Sung Shin, Jin Kim, BengSeng Phuah, HyoungMin Kim, Myeong Jun Song, Ahn Choi, Daeho Kim, SooYoung Kim, Eun-Bong Kim, David Wang, Shinhaeng Kang, Yuhwan Ro, Seungwoo Seo, JoonHo Song, Jaeyoun Youn, Kyomin Sohn, and Nam Sung 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*, ISSCC, 2021.
- [146] Daniel Langr and Pavel Tvrđik, *Evaluation Criteria for Sparse Matrix Storage Formats*, TPDS, 2016.
- [147] Dominique Lavenier, Remy Cimadomo, and Romaric Jodin, *Variant Calling Parallelization on Processor-in-Memory Architecture*, BIBM, 2020.
- [148] Dominique Lavenier, Jean-Francois Roy, and David Furodet, *DNA Mapping Using Processor-in-Memory Architecture*, BIBM, 2016.
- [149] Donghyuk Lee, Saugata Ghose, Gennady Pekhimenko, Samira Khan, and Onur Mutlu, *Simultaneous Multi-Layer Access: Improving 3D-Stacked Memory Bandwidth at Low Cost*, TACO, 2016.
- [150] Seyong Lee and Rudolf Eigenmann, *Adaptive Runtime Tuning of Parallel Sparse Matrix-Vector Multiplication on Distributed Memory Systems*, ICS, 2008.
- [151] Sukhan Lee, Shin-Haeng Kang, Jaehoon Lee, H. Kim, Eojin Lee, Seung young Seo, H. Yoon, Seungwon Lee, K. Lim, Hyunsung Shin, Jinhyun Kim, O. Seongil, Anand Iyer, David Wang, K. Sohn, and N. Kim, *Hardware Architecture and Software Stack for PIM Based on Commercial DRAM Technology: Industrial Product*, ISCA, 2021.
- [152] J. Leskovec and R. Sosič, *SNAP: A General-Purpose Network Analysis and Graph-Mining Library*, TIST, 2016.
- [153] Yifat Levy, Jehoshua Bruck, Yuval Cassuto, Eby G. Friedman, Avinoam Kolodny, Eitan Yaakobi, and Shahar Kvatinsky, *Logic Operations in Memory Using a Memristive Akers Array*, Microelectronics Journal, 2014.
- [154] Jiajia Li, Guangming Tan, Mingyu Chen, and Ninghui Sun, *SMAT: An Input Adaptive Auto-Tuner for Sparse Matrix-Vector Multiplication*, PLDI, 2013.
- [155] Kenli Li, Wangdong Yang, and Keqin Li, *Performance Analysis and Optimization for SpMV on GPU Using Probabilistic Modeling*, IEEE TPDS, 2015.
- [156] Shuangchen Li, Dimin Niu, Krishna T Malladi, Hongzhong Zheng, Bob Brennan, and Yuan Xie, *DRISA: A DRAM-based Reconfigurable In-Situ Accelerator*, MICRO, 2017.
- [157] Shuangchen Li, Cong Xu, Qiaosha Zou, Jishen Zhao, Yu Lu, and Yuan Xie, *Pinatubo: A Processing-In-Memory Architecture for Bulk Bitwise Operations in Emerging Non-Volatile Memories*, DAC, 2016.
- [158] Colin Yu Lin, Zheng Zhang, Ngai Wong, and Hayden Kwok-Hay So, *Design Space Exploration for Sparse Matrix-Matrix Multiplication on FPGAs*, FPT, 2010.
- [159] Greg Linden, Brent Smith, and Jeremy York, *Amazon.com Recommendations: Item-to-Item Collaborative Filtering*, IC, 2003.
- [160] Baoyuan Liu, Min Wang, Hassan Foroosh, Marshall Tappen, and Marianna Pensky, *Sparse Convolutional Neural Networks*, CVPR, 2015.
- [161] Changxi Liu, Biwei Xie, Xin Liu, Wei Xue, Hailong Yang, and Xu Liu, *Towards Efficient SpMV on Sunway Manycore Architectures*, ICS, 2018.
- [162] Weifeng Liu and Brian Vinter, *An Efficient GPU General Sparse Matrix-Matrix Multiplication for Irregular Data*, IPDPS, 2014.
- [163] Weifeng Liu and Brian Vinter, *CSR5: An Efficient Storage Format for Cross-Platform Sparse Matrix-Vector Multiplication*, ICS, 2015.
- [164] Weifeng Liu and Brian Vinter, *CSR5: An Efficient Storage Format for Cross-Platform Sparse Matrix-Vector Multiplication*, ICS, 2015.

- [165] Xing Liu, Mikhail Smelyanskiy, Edmond Chow, and Pradeep Dubey, *Efficient Sparse Matrix-Vector Multiplication on X86-Based Many-Core Processors*, ICS, 2013.
- [166] Zhiyu Liu, Irina Calciu, Maurice Herlihy, and Onur Mutlu, *Concurrent Data Structures for Near-Memory Computing*, SPAA, 2017.
- [167] Elliot Lockerman, Axel Feldmann, Mohammad Bakhshalipour, Alexandru Stanescu, Shashwat Gupta, Daniel Sanchez, and Nathan Beckmann, *Livia: Data-Centric Computing Throughout the Memory Hierarchy*, ASPLOS, 2020.
- [168] Marco Maggioni and Tanya Berger-Wolf, *AdELL: An Adaptive Warp-Balancing ELL Format for Efficient Sparse Matrix-Vector Multiplication on GPUs*, ICPP, 2013.
- [169] Michele Martone, *Efficient Multithreaded Untransposed, Transposed or Symmetric Sparse Matrix-Vector Multiplication with the Recursive Sparse Blocks Format*, Parallel Comput., 2014.
- [170] Michele Martone, Salvatore Filippone, Marcin Paprzycki, and Salvatore Tucci, *On BLAS Operations with Recursively Stored Sparse Matrices*, SYNASC, 2010.
- [171] John Mellor-Crummey and John Garvin, *Optimizing Sparse Matrix-Vector Product Computations Using Unroll and Jam*, IJHPCA, 2004.
- [172] Duane Merrill and Michael Garland, *Merge-Based Parallel Sparse Matrix-Vector Multiplication*, SC, 2016.
- [173] Asit K. Mishra, Eriko Nurvitadhi, Ganesh Venkatesh, Jonathan Pearce, and Debbie Marr, *Fine-grained Accelerators for Sparse Machine Learning Workloads*, ASP-DAC, 2017.
- [174] Alexander Monakov, Anton Lokhmotov, and Arutyun Avetisyan, *Automatically Tuning Sparse Matrix-Vector Multiplication for GPU Architectures*, HiPEAC, 2010.
- [175] Anurag Mukkara, Nathan Beckmann, Maleen Abeydeera, Xiaosong Ma, and Daniel Sanchez, *Exploiting Locality in Graph Analytics through Hardware-Accelerated Traversal Scheduling*, MICRO, 2018.
- [176] Onur Mutlu, Saugata Ghose, Juan Gómez-Luna, and Rachata Ausavarungnirun, *Enabling Practical Processing In and Near Memory for Data-Intensive Computing*, DAC, 2019.
- [177] Onur Mutlu, Saugata Ghose, Juan Gómez-Luna, and Rachata Ausavarungnirun, *Processing Data Where It Makes Sense: Enabling In-Memory Computation*, MICPRO, 2019.
- [178] Onur Mutlu, Saugata Ghose, Juan Gómez-Luna, and Rachata Ausavarungnirun, *A Modern Primer on Processing in Memory, Emerging Computing: From Devices to Systems - Looking Beyond Moore and Von Neumann*, 2021.
- [179] Anirban Nag and Rajeev Balasubramonian, *OrderLight: Lightweight Memory-Ordering Primitive for Efficient Fine-Grained PIM Computations*, MICRO, 2021.
- [180] L. Nai, R. Hadidi, J. Sim, H. Kim, P. Kumar, and H. Kim, *GraphPIM: Enabling Instruction-Level PIM Offloading in Graph Computing Frameworks*, HPCA, 2017.
- [181] R. Nair, S. F. Antao, C. Bertolli, P. Bose, J. R. Brunheroto, T. Chen, C.-Y. Cher, C. H. A. Costa, J. Doi, C. Evangelinos, and et al., *Active Memory Cube: A Processing-in-Memory Architecture for Exascale Systems*, IBM JRD, 2015.
- [182] Naveen Namashivayam, Sanyam Mehta, and Pen-Chung Yew, *Variable-Sized Blocks for Locality-Aware SpMV*, CGO, 2021.
- [183] S.G. Nastea, O. Frieder, and T. El-Ghazawi, *Load-Balancing in Sparse Matrix-Vector Multiplication*, SPDP, 1996.
- [184] Maxim Naumov, Dheevatsa Mudigere, Hao-Jun Michael Shi, Jianyu Huang, Narayanan Sundaraman, Jongsoo Park, Xiaodong Wang, Udit Gupta, Carole-Jean Wu, Alisson G. Azzolini, Dmytro Dzhulgakov, Andrey Mallevich, Ilia Cherniavskii, Yinghai Lu, Raghuraman Krishnamoorthi, Ansha Yu, Volodymyr Kondratenko, Stephanie Pereira, Xianjie Chen, Wenlin Chen, Vijay Rao, Bill Jia, Liang Xiong, and Misha Smelyanskiy, *Deep Learning Recommendation Model for Personalization and Recommendation Systems*, CoRR, 2019.
- [185] Joel Nider, Craig Mustard, Andrada Zoltan, and Alexandra Fedorova, *Processing in Storage Class Memory*, HotStorage, 2020.
- [186] Yuyao Niu, Zhengyang Lu, Meichen Dong, Zhou Jin, Weifeng Liu, and Guangming Tan, *TileSpMV: A Tiled Algorithm for Sparse Matrix-Vector Multiplication on GPUs*, IPDPS, 2021.
- [187] Eriko Nurvitadhi, Asit Mishra, and Debbie Marr, *A Sparse Matrix Vector Multiply Accelerator for Support Vector Machine*, CASES, 2015.
- [188] Eriko Nurvitadhi, Asit Mishra, Yu Wang, Ganesh Venkatesh, and Debbie Marr, *Hardware Accelerator for Analytics of Sparse Data*, DAC, 2016.
- [189] NVIDIA, *NVIDIA System Management Interface Program*, <http://developer.download.nvidia.com/compute/DCGM/docs/nvidia-smi-367.38.pdf>, 2016.
- [190] NVIDIA, *NVIDIA Tesla V100 GPU Architecture*, <https://images.nvidia.com/content/volta-architecture/pdf/volta-architecture-whitepaper.pdf>, 2017.
- [191] Ataberk Olgun, Juan Gómez Luna, Konstantinos Kanellopoulos, Behzad Salami, Hasan Hassan, Oğuz Ergin, and Onur Mutlu, *PiDRAM: A Holistic End-to-end FPGA-based Framework for Processing-in-DRAM*, CoRR, 2021.
- [192] Ataberk Olgun, Minesh Patel, A. Giray Yağlikçi, Haocong Luo, Jeremie S. Kim, F. Nisa Bostanci, Nandita Vijaykumar, Oğuz Ergin, and Onur Mutlu, *QUAC-TRNG: High-Throughput True Random Number Generation Using Quadruple Row*

- Activation in Commodity DRAM Chips*, ISCA, 2021.
- [193] Leonid Oliker, Xiaoye Li, Parry Husbands, and Rupak Biswas, *Effects of Ordering Strategies and Programming Paradigms on Sparse Matrix Computations*, *SIAM Rev.*, 2002.
 - [194] Geraldo F. Oliveira, Juan Gómez-Luna, Lois Orosa, Saugata Ghose, Nandita Vijaykumar, Ivan Fernandez, Mohammad Sadrosadati, and Onur Mutlu, *DAMOV: A New Methodology and Benchmark Suite for Evaluating Data Movement Bottlenecks*, *IEEE Access*, 2021.
 - [195] Lois Orosa, Yaohua Wang, Mohammad Sadrosadati, Jeremie S. Kim, Minesh Patel, Ivan Puddu, Haocong Luo, Kaveh Razavi, Juan Gómez-Luna, Hasan Hassan, Nika Mansouri-Ghiasi, Saugata Ghose, and Onur Mutlu, *CODIC: A Low-Cost Substrate for Enabling Custom in-DRAM Functionalities and Optimizations*, *ISCA*, 2021.
 - [196] Brian A. Page and Peter M. Kogge, *Scalability of Hybrid Sparse Matrix Dense Vector (SpMV) Multiplication*, *HPCS*, 2018.
 - [197] Subhankar Pal, Jonathan Beaumont, Dong-Hyeon Park, Aporva Amarnath, Siying Feng, Chaitali Chakrabarti, Hun-Seok Kim, David Blaauw, Trevor Mudge, and Ronald Dreslinski, *OuterSPACE: An Outer Product Based Sparse Matrix Multiplication Accelerator*, *HPCA*, 2018.
 - [198] Angshuman Parashar, Minsoo Rhu, Anurag Mukkara, Antonio Puglielli, Rangharajan Venkatesan, Brucek Khailany, Joel Emer, Stephen W. Keckler, and William J. Dally, *SCNN: An Accelerator for Compressed-sparse Convolutional Neural Networks*, *ISCA*, 2017.
 - [199] Jaehyun Park, Byeongho Kim, Sungmin Yun, Eojin Lee, Minsoo Rhu, and Jung Ho Ahn, *TRiM: Enhancing Processor-Memory Interfaces with Scalable Tensor Reduction in Memory*, *MICRO*, 2021.
 - [200] peakperf, *peakperf*, <https://github.com/Dr-Noob/peakperf.git>, 2021.
 - [201] Daniël M. Pelt and Rob H. Bisseling, *A Medium-Grain Method for Fast 2D Bipartitioning of Sparse Matrices*, *IPDPS*, 2014.
 - [202] Juan C. Pichel and Francisco F. Rivera, *Sparse Matrix-Vector Multiplication on the Single-Chip Cloud Computer Many-Core Processor*, *J. Parallel Distrib. Comput.*, 2013.
 - [203] Juan C. Pichel, Francisco F. Rivera, Marcos Fernández, and Aurelio Rodríguez, *Optimization of Sparse Matrix-Vector Multiplication Using Reordering Techniques on GPUs*, *Microprocess. Microsyst.*, 2012.
 - [204] Ali Pinar and Michael T. Heath, *Improving Performance of Sparse Matrix-Vector Multiplication*, *SC*, 1999.
 - [205] Udo W. Pooch and Al Nieder, *A Survey of Indexing Techniques for Sparse Matrices*, *ACM Comput. Surv.*, 1973.
 - [206] Seth H Pugsley, Jeffrey Jestes, Huihui Zhang, Rajeev Balasubramonian, Vijayalakshmi Srinivasan, Alper Buyukto-sunoglu, Al Davis, and Feifei Li, *NDC: Analyzing the Impact of 3D-Stacked Memory+Logic Devices on MapReduce Workloads*, *ISPASS*, 2014.
 - [207] Eric Qin, Ananda Samajdar, Hyoukjun Kwon, Vineet Nadella, Sudarshan Srinivasan, Dipankar Das, Bharat Kaul, and Tushar Krishna, *SIGMA: A Sparse and Irregular GEMM Accelerator with Flexible Interconnects for DNN Training*, *HPCA*, 2020.
 - [208] Eric Qin, Ananda Samajdar, Hyoukjun Kwon, Vineet Nadella, Sudarshan Srinivasan, Dipankar Das, Bharat Kaul, and Tushar Krishna, *Sigma: A Sparse and Irregular GEMM Accelerator with Flexible Interconnects for DNN Training*, *HPCA*, 2020.
 - [209] Seyyed Hossein SeyyedAghaei Rezaei, Mehdi Modarressi, Rachata Ausavarungnirun, Mohammad Sadrosadati, Onur Mutlu, and Masoud Daneshdatab, *NoM: Network-on-Memory for Inter-Bank Data Transfer in Highly-Banked Memories*, *IEEE CAL*, 2020.
 - [210] Hongbo Rong, Jongsoo Park, Lingxiang Xiang, Todd A. Anderson, and Mikhail Smelyanskiy, *Sparso: Context-Driven Optimizations of Sparse Linear Algebra*, *PACT*, 2016.
 - [211] Youcef Saad, *Krylov Subspace Methods on Supercomputers*, *SIAM J. Sci. Stat. Comput.*, 1989.
 - [212] Fazle Sadi, Joe Sweeney, Tze Meng Low, James C. Hoe, Larry Pileggi, and Franz Franchetti, *Efficient SpMV Operation for Large and Highly Sparse Matrices Using Scalable Multi-Way Merge Parallelization*, *MICRO*, 2019.
 - [213] Elaheh Sadredini, Reza Rahimi, Mohsen Imani, and Kevin Skadron, *Sunder: Enabling Low-Overhead and Scalable Near-Data Pattern Matching Acceleration*, *MICRO*, 2021.
 - [214] SciPy, *List-of-list Sparse Matrix*, *SciPy*, 2021.
 - [215] Naser Sedaghati, Te Mu, Louis-Noël Pouchet, Srinivasan Parthasarathy, and P. Sadayappan, *Automatic Selection of Sparse Matrix Representation on GPUs*, *ICS*, 2015.
 - [216] Shubhabrata Sengupta, Mark Harris, Yao Zhang, and John D. Owens, *Scan Primitives for GPU Computing*, *GH*, 2007.
 - [217] Vivek Seshadri, Yoongu Kim, Chris Fallin, Donghyuk Lee, Rachata Ausavarungnirun, Gennady Pekhimenko, Yixin Luo, Onur Mutlu, Phillip B. Gibbons, Michael A. Kozuch, and Todd C. Mowry, *RowClone: Fast and energy-efficient in-DRAM bulk data copy and initialization*, *MICRO*, 2013.
 - [218] Vivek Seshadri, Donghyuk Lee, Thomas Mullins, Hasan Hassan, Amirali Boroumand, Jeremie Kim, Michael A. Kozuch, Onur Mutlu, Phillip B. Gibbons, and Todd C. Mowry, *Ambit: In-Memory Accelerator for Bulk Bitwise Operations Using Commodity DRAM Technology*, *MICRO*, 2017.

- [219] Vivek Seshadri and Onur Mutlu, *Simple Operations in Memory to Reduce Data Movement*, *Advances in Computers*, 2017.
- [220] Vivek Seshadri and Onur Mutlu, *In-DRAM Bulk Bitwise Execution Engine*, *CoRR*, 2020.
- [221] Ali Shafiee, Anirban Nag, Naveen Muralimanohar, Rajeev Balasubramonian, John Paul Strachan, Miao Hu, R. Stanley Williams, and Vivek Srikumar, *ISAAC: A Convolutional Neural Network Accelerator with In-Situ Analog Arithmetic in Crossbars*, *ISCA*, 2016.
- [222] Hyunsung Shin, Dongyoung Kim, Eunhyeok Park, Sungho Park, Yongsik Park, and Sungjoo Yoo, *McDRAM: Low Latency and Energy-Efficient Matrix Computations in DRAM*, *IEEE TCAD*, 2018.
- [223] Gagandeep Singh, Lorenzo Chelini, Stefano Corda, Ahsan Javed Awan, Sander Stuijk, Roel Jordans, Henk Corporaal, and Albert-Jan Boonstra, *Near-Memory Computing: Past, Present, and Future*, *MICPRO*, 2019.
- [224] Gagandeep Singh, Dionysios Diamantopoulos, Christoph Hagleitner, Juan Gómez-Luna, Sander Stuijk, Onur Mutlu, and Henk Corporaal, *NERO: A Near High-Bandwidth Memory Stencil Accelerator for Weather Prediction Modeling*, *FPL*, 2020.
- [225] Gagandeep Singh, Juan Gómez-Luna, Giovanni Mariani, Geraldo F. Oliveira, Stefano Corda, Sander Stuijk, Onur Mutlu, and Henk Corporaal, *NAPEL: Near-Memory Computing Application Performance Prediction via Ensemble Learning*, *DAC*, 2019.
- [226] A. Smith, *6 New Facts About Facebook*, <http://mediashift.org>, 2019.
- [227] Linghao Song, Youwei Zhuo, Xuehai Qian, Hai Li, and Yiran Chen, *GraphR: Accelerating Graph Processing Using ReRAM*, *HPCA*, 2018.
- [228] Markus Steinberger, Rhaleb Zayer, and Hans-Peter Seidel, *Globally Homogeneous, Locally Adaptive Sparse Matrix-Vector Multiplication on the GPU*, *ICS*, 2017.
- [229] Harold S. Stone, *A Logic-in-Memory Computer*, *IEEE TC*, 1970.
- [230] Foteini Strati, Christina Giannoula, Dimitrios Siakavaras, Georgios Goumas, and Nectarios Koziris, *An Adaptive Concurrent Priority Queue for NUMA Architectures*, *ACM CF*, 2019.
- [231] stream, *stream*, <https://github.com/jeffhammond/STREAM.git>, 2021.
- [232] Bor-Yiing Su and Kurt Keutzer, *ClSPMV: A Cross-Platform OpenCL SpMV Framework on GPUs*, *ICS*, 2012.
- [233] Weiwei Sun, Zhaoshi Li, Shouyi Yin, Shaojun Wei, and Leibo Liu, *ABC-DIMM: Alleviating the Bottleneck of Communication in DIMM-based Near-Memory Processing with Inter-DIMM Broadcast*, *ISCA*, 2021.
- [234] Xiangzheng Sun, Yunquan Zhang, Ting Wang, Xianyi Zhang, Liang Yuan, and Li Rao, *Optimizing SpMV for Diagonal Sparse Matrices on GPU*, *ICPP*, 2011.
- [235] Guangming Tan, Junhong Liu, and Jiajia Li, *Design and Implementation of Adaptive SpMV Library for Multicore and Many-Core Architecture*, *ACM Trans. Math. Softw.*, 2018.
- [236] Wai Teng Tang, Ruizhe Zhao, Mian Lu, Yun Liang, Huynh Phung Huyng, Xibai Li, and Rick Siow Mong Goh, *Optimizing and Auto-Tuning Scale-Free Sparse Matrix-Vector Multiplication on Intel Xeon Phi*, *CGO*, 2015.
- [237] O. Temam and W. Jalby, *Characterizing the Behavior of Sparse Algorithms on Caches*, *SC*, 1992.
- [238] S. Toledo, *Improving the Memory-System Performance of Sparse-Matrix Vector Multiplication*, *IBM Journal of Research and Development*, 1997.
- [239] Yaman Umuroglu and Magnus Jahre, *An Energy Efficient Column-Major Backend for FPGA SpMV Accelerators*, *ICCD*, 2014.
- [240] UPMEM, *Introduction to UPMEM PIM. Processing-in-memory (PIM) on DRAM Accelerator (White Paper)*, , 2018.
- [241] UPMEM, *UPMEM Website*, <https://www.upmem.com>, 2020.
- [242] UPMEM, *UPMEM User Manual. Version 2021.3*, UPMEM, 2021.
- [243] Brendan Vastenhouw and Rob H. Bisseling, *A Two-Dimensional Data Distribution Method for Parallel Sparse Matrix-Vector Multiplication*, *SIAM Rev.*, 2005.
- [244] F. Vázquez, J. J. Fernández, and E. M. Garzón, *A New Approach for Sparse Matrix Vector Product on NVIDIA GPUs, Concurrency and Computation: Practice and Experience*, 2011.
- [245] R. Vuduc, J.W. Demmel, K.A. Yelick, S. Kamil, R. Nishtala, and B. Lee, *Performance Optimizations and Bounds for Sparse Matrix-Vector Multiply*, *SC*, 2002.
- [246] Richard Vuduc, James W Demmel, and Katherine A Yelick, *OSKI: A Library of Automatically Tuned Sparse Matrix Kernels*, *Journal of Physics: Conference Series*, 2005.
- [247] Richard Wilson Vuduc and James W. Demmel, *Automatic Performance Tuning of Sparse Matrix Kernels*, *PhD Thesis*, 2003.
- [248] Richard W. Vuduc and Hyun-Jin Moon, *Fast Sparse Matrix-Vector Multiplication by Exploiting Variable Block Structure*, *HPCC*, 2005.
- [249] Yaohua Wang, Lois Orosa, Xiangjun Peng, Yang Guo, Saugata Ghose, Minesh Patel, Jeremie S. Kim, Juan Gómez-Luna, Mohammad Sadrosadati, Nika Mansouri-Ghiasi, and Onur Mutlu, *FIGARO: Improving System Performance via Fine-Grained In-DRAM Data Relocation and Caching*, *MICRO*, 2020.

- [250] J.B. White and P. Sadayappan, *On Improving the Performance of Sparse Matrix-Vector Multiplication*, HIPC, 1997.
- [251] Jeremiah Willcock and Andrew Lumsdaine, *Accelerating Sparse Matrix Computations via Data Compression*, ICS, 2006.
- [252] Samuel Williams, Leonid Oliker, Richard Vuduc, John Shalf, Katherine Yelick, and James Demmel, *Optimization of Sparse Matrix-Vector Multiplication on Emerging Multicore Platforms*, SC, 2007.
- [253] Lingxi Wu, Rasool Sharifi, Marzieh Lenjani, Kevin Skadron, and Ashish Venkat, *Sieve: Scalable In-situ DRAM-based Accelerator Designs for Massively Parallel k-mer Matching*, ISCA, 2021.
- [254] Tianji Wu, Bo Wang, Yi Shan, Feng Yan, Yu Wang, and Ningyi Xu, *Efficient PageRank and SpMV Computation on AMD GPUs*, ICPP, 2010.
- [255] Yue Xi, Bin Gao, Jianshi Tang, An Chen, Meng-Fan Chang, Xiaobo Sharon Hu, Jan Van Der Spiegel, He Qian, and Huaqiang Wu, *In-memory Learning with Analog Resistive Switching Memory: A Review and Perspective*, Proc. IEEE, 2021.
- [256] Guoqing Xiao, Kenli Li, Yuedan Chen, Wangquan He, Albert Y. Zomaya, and Tao Li, *CASpMV: A Customized and Accelerative SpMV Framework for the Sunway TaihuLight*, IEEE TPDS, 2021.
- [257] Biwei Xie, Jianfeng Zhan, Xu Liu, Wanling Gao, Zhen Jia, Xiwen He, and Lixin Zhang, *CVR: Efficient Vectorization of SpMV on X86 Processors*, CGO, 2018.
- [258] Lei Xie, Hoang Anh Du Nguyen, Mottaqiallah Taouil, Koen Bertels, and Said Hamdioui, *Fast Boolean Logic Mapped on Memristor Crossbar*, ICCD, 2015.
- [259] Xinfeng Xie, Zheng Liang, Peng Gu, Abanti Basak, Lei Deng, Ling Liang, Xing Hu, and Yuan Xie, *SpaceA: Sparse Matrix Vector Multiplication on Processing-in-Memory Accelerator*, HPCA, 2021.
- [260] Xin Xin, Youtao Zhang, and Jun Yang, *ELP2IM: Efficient and Low Power Bitwise Operation Processing in DRAM*, HPCA, 2020.
- [261] Shengen Yan, Chao Li, Yunquan Zhang, and Huiyang Zhou, *YaSpMV: Yet Another SpMV Framework on GPUs*, PPOPP, 2014.
- [262] Wangdong Yang, Kenli Li, and Keqin Li, *A Hybrid Computing Method of SpMV on CPU-GPU Heterogeneous Computing Systems*, JPDC, 2017.
- [263] Wangdong Yang, Kenli Li, Yan Liu, Lin Shi, and Lanjun Wan, *Optimization of Quasi-Diagonal Matrix-Vector Multiplication on GPU*, Int. J. High Perform. Comput. Appl., 2014.
- [264] Wangdong Yang, Kenli Li, Zeyao Mo, and Keqin Li, *Performance Optimization Using Partitioned SpMV on GPUs and Multicore CPUs*, IEEE Trans. Comput., 2015.
- [265] Xintian Yang, Srinivasan Parthasarathy, and P. Sadayappan, *Fast Sparse Matrix-Vector Multiplication on GPUs: Implications for Graph Mining*, Proc. VLDB Endow., 2011.
- [266] Amir Yazdanbakhsh, Choungki Song, Jacob Sacks, Pejman Lotfi-Kamran, Hadi Esmaeilzadeh, and Nam Sung Kim, *In-DRAM near-Data Approximate Acceleration for GPUs*, PACT, 2018.
- [267] Jintao Yu, Hoang Anh Du Nguyen, Lei Xie, Mottaqiallah Taouil, and Said Hamdioui, *Memristive Devices for Computation-In-Memory*, DATE, 2018.
- [268] Geng Yuan, Payman Behnam, Zhengang Li, Ali Shafiee, Sheng Lin, Xiaolong Ma, Hang Liu, Xuehai Qian, Mahdi Nazm Bojnordi, Yanzhi Wang, and Caiwen Ding, *FORMS: Fine-grained Polarized ReRAM-based In-situ Computation for Mixed-signal DNN Accelerator*, ISCA, 2021.
- [269] Dongping Zhang, Nuwan Jayasena, Alexander Lyashevsky, Joseph L. Greathouse, Lifan Xu, and Michael Ignatowski, *TOP-PIM: Throughput-Oriented Programmable Processing in Memory*, HPDC, 2014.
- [270] Guowei Zhang, Nithya Attaluri, Joel S. Emer, and Daniel Sanchez, *Gamma: Leveraging Gustavson's Algorithm to Accelerate Sparse Matrix Multiplication*, ASPLOS, 2021.
- [271] Mingxing Zhang, Youwei Zhuo, Chao Wang, Mingyu Gao, Yongwei Wu, Kang Chen, Christos Kozyrakis, and Xuehai Qian, *GraphP: Reducing Communication for PIM-Based Graph Processing with Efficient Data Partition*, HPCA, 2018.
- [272] Shijin Zhang, Zidong Du, Lei Zhang, Huiying Lan, Shaoli Liu, Ling Li, Qi Guo, Tianshi Chen, and Yunji Chen, *Cambricon-X: An Accelerator for Sparse Neural Networks*, MICRO, 2016.
- [273] Zhehai Zhang, Hanrui Wang, Song Han, and William J. Dally, *SpArch: Efficient Architecture for Sparse Matrix Multiplication*, HPCA, 2020.
- [274] Haoran Zhao, Tian Xia, Chenyang Li, Wenzhe Zhao, Nanning Zheng, and Pengju Ren, *Exploring Better Speculation and Data Locality in Sparse Matrix-Vector Multiplication on Intel Xeon*, ICCD, 2020.
- [275] Yue Zhao, Jiajia Li, Chunhua Liao, and Xipeng Shen, *Bridging the Gap between Deep Learning and Sparse Matrix Format Selection*, PPOPP, 2018.
- [276] Yue Zhao, Jiajia Li, Chunhua Liao, and Xipeng Shen, *Bridging the Gap between Deep Learning and Sparse Matrix Format Selection*, PPOPP, 2018.
- [277] Yue Zhao, Weijie Zhou, Xipeng Shen, and Graham Yiu, *Overhead-Conscious Format Selection for SpMV-Based Applications*, IPDPS, 2018.

- [278] Le Zheng, Sangho Shin, Scott Lloyd, Maya Gokhale, Kyungmin Kim, and Sung-Mo Kang, *RRAM-based TCAMs for Pattern Search*, *ISCAS*, 2016.
- [279] Xuda Zhou, Zidong Du, Qi Guo, Shaoli Liu, Chengsi Liu, Chao Wang, Xuehai Zhou, Ling Li, Tianshi Chen, and Yunji Chen, *Cambricon-S: Addressing Irregularity in Sparse Neural Networks through A Cooperative Software/Hardware Approach*, *MICRO*, 2018.
- [280] Qiuling Zhu, Tobias Graf, H. Ekin Sumbul, Larry Pileggi, and Franz Franchetti, *Accelerating Sparse Matrix-Matrix Multiplication with 3D-Stacked Logic-In-Memory Hardware*, *HPEC*, 2013.
- [281] Youwei Zhuo, Chao Wang, Mingxing Zhang, Rui Wang, Dimin Niu, Yanzhi Wang, and Xuehai Qian, *GraphQ: Scalable PIM-Based Graph Processing*, *MICRO*, 2019.
- [282] Vasileios Zois, Divya Gupta, Vassilis J. Tsotras, Walid A. Najjar, and Jean-Francois Roy, *Massively Parallel Skyline Computation for Processing-in-Memory Architectures*, *PACT*, 2018.

APPENDIX

A Extended Results

A.1 Synchronization Approaches in Block-Based Compressed Matrix Formats

We compare the coarse-grained locking (*lb-cg*) and the fine-grained locking (*lb-fg*) approaches in the BCOO format. Figure 31 shows the performance achieved by the BCOO format for all the data types when balancing the blocks or the non-zero elements across 16 tasklets of one DPU. We evaluate all small matrices of Table 3, i.e., delaunay_n13 (**D**), wing_nodal (**W**), raefsky4 (**R**) and pkustk08 (**P**) matrices.




Fig. 31. Performance of the BCOO format with various load balancing schemes and synchronization approaches for all the data types and small matrices using 16 tasklets of one DPU.

Our key finding is that the fine-grained locking approach performs similarly with the coarse-grained locking approach. The fine-grained locking approach does not increase parallelism in the UPMEM PIM architecture, since memory accesses executed by multiple tasklets to the local DRAM bank are serialized in the DMA engine of the DPU. The same key finding holds independently of the compressed matrix format used.

A.2 Fine-Grained Data Transfers in 2D Partitioning Techniques

Figures 32 and 33 compare coarse-grained data transfers (i.e., performing parallel data transfers to all 2048 DPUs at once, padding with empty bytes at the granularity of 2048 DPUs) with fine-grained data transfers (i.e., iterating over the ranks and for each rank performing parallel data transfers to the 64 DPUs of the same rank, padding with empty bytes at the granularity of 64 DPUs) for all matrices of our large matrix dataset in the *equally-wide* and *variable-sized* schemes, respectively. The reported key findings of Figure 17 (Section 6.2.1) apply to all matrices with diverse sparsity patterns.




Fig. 32. Performance comparison of RC: RBDCOO with coarse-grained transfers, RY: RBDCOO with fine-grained transfers in the output vector, BC: BDCOO with coarse-grained transfers, BY: BDCOO with fine-grained transfers only in the output vector, and BT: BDCOO with fine-grained transfers in both the input and the output vector using the int32 data type, 2048 DPUs and having 2 vertical partitions. Performance is normalized to that of the RC scheme.




Fig. 33. Performance comparison of RC: RBDC00 with coarse-grained transfers, RY: RBDC00 with fine-grained transfers in the output vector, BC: BDC00 with coarse-grained transfers, BY: BDC00 with fine-grained transfers only in the output vector, and BT: BDC00 with fine-grained transfers in both the input and the output vector using the int32 data type, 2048 DPUs and having 32 vertical partitions. Performance is normalized to that of the RC scheme.

A.3 Effect of the Number of Vertical Partitions Using Two Different UPMEM PIM Systems

We compare SpMV execution in the two different UPMEM PIM systems using 2048 DPUs and 16 tasklets for each DPU. Table 6 shows the characteristics of two different UPMEM PIM systems. We calculate the available PIM peak performance and PIM bandwidth assuming 2048 DPUs for both PIM systems³. We estimate the PIM peak performance as $Total_DPUs * AT$, where the arithmetic throughput (AT) is calculated for the multiplication operation by running the arithmetic throughput microbenchmark of the PRIM benchmark suite [82, 94] in each of the two UPMEM PIM systems (See Appendix B). We estimate the PIM bandwidth as $Total_DPUs * Bandwidth_DPU$, where the $Bandwidth_DPU$ is calculated according to prior work [82, 94]. Specifically, the theoretical maximum MRAM bandwidth (i.e., $Bandwidth_DPU$) is 700 MB/s and 850 MB/s at a DPU frequency of 350 MHz (PIM system A) and 425 MHz (PIM system B), respectively.

System	Avail. DPUs	Frequency	PIM Peak Performance	PIM Bandwidth	Host CPU	CPU Peak Performance	Bus Bandwidth
PIM System A	2048 DPUs	350 MHz	3.78 GFLOPS	1.43 TB/s	Intel Xeon Silver 4110 @2.1 GHz	660 GFLOPS	23.1 GB/s
PIM System B	2048 DPUs	425 MHz	4.63 GFLOPS	1.74 TB/s	Intel Xeon Silver 4215 @2.5 GHz	1016 GFLOPS	21.8 GB/s

Table 6. Evaluated UPMEM PIM Systems.

Figures 34, 35 and 36 compare SpMV execution in the two different UPMEM PIM systems (2048 DPUs) using 2D-partitioned kernels with the COO format, when varying the number of vertical partitions from 1 to 32 (in steps of multiple of 2) for the int32 (left) and fp64 (right) data types.

We observe that the number of vertical partitions that provides the best performance on SpMV execution varies depending on the input matrix and the PIM system. For example, in PIM system B with the int32 data type, DC00 performs best for the hgc matrix with 16 vertical partitions, while in PIM system A, DC00 performs best for the same matrix with 8 vertical partitions. Similarly, in PIM system A with the fp64 data type, BDC00 performs best for the rjt matrix with 4 vertical partitions. Instead, in PIM system B with the fp64 data type, BDC00's performance does not improve for the rjt matrix when having more than 1 vertical partition (i.e., compared to when using the 1D partitioning technique). We conclude that the best-performing parallelization scheme that achieves the best performance in SpMV depends on the characteristics of both the input sparse matrix and the underlying PIM system.

³Both UPMEM PIM systems support 20 UPMEM PIM DIMMs with 2560 DPUs in total. However, both UPMEM-based PIM systems include multiple faulty DPUs. Thus, for a fair comparison between two systems we conduct our experiments using 2048 DPUs in both systems.




Fig. 34. Execution time breakdown of DCOO using 2048 DPUs when varying the number of vertical partitions from 1 to 32 for the int32 (left) and fp64 (right) data types on two different UPMEM PIM systems.




Fig. 35. Execution time breakdown of RBDCOO using 2048 DPUs when varying the number of vertical partitions from 1 to 32 for the int32 (left) and fp64 (right) data types on two different UPMEM PIM systems.




Fig. 36. Execution time breakdown of BDCOO using 2048 DPUs when varying the number of vertical partitions from 1 to 32 for the int32 (left) and fp64 (right) data types on two different UPMEM PIM systems.

A.4 Performance of Compressed Matrix Formats Using 2D Partitioning Techniques

Figures 37, 38, 39 compare the performance achieved by various compressed matrix formats for each of the three types of the 2D partitioning technique for all matrices of our large matrix dataset. The reported key findings explained in Section 6.2.2 apply to all matrices with diverse sparsity patterns.




Fig. 37. End-to-end execution time breakdown of the *equally-sized* 2D partitioning technique for CR: DCSR, CO: DC00, BR: DBCSR and BO: DBC00 schemes using 4 vertical partitions and the int32 data type. Performance is normalized to that of DCSR.




Fig. 38. End-to-end execution time breakdown of the *equally-wide* 2D partitioning technique for CR: RBDCSR, CO: RBDC00, BR: RBDBCSR and BO: RBDBC00 schemes using 4 vertical partitions and the int32 data type. Performance is normalized to that of RBDCSR.




Fig. 39. End-to-end execution time breakdown of the *variable-sized* 2D partitioning technique for CR: BDC00, CO: BDC00, BR: BDBCSR and BO: BDBC00 schemes using 4 vertical partitions and the int32 data type. Performance is normalized to that of BDCSR.

A.5 Analysis of 1D- and 2D-Partitioned Kernels in Two UPMEM PIM Systems

Figures 40 and 41 compare the throughput and the performance, respectively, achieved by the best-performing 1D- and 2D-partitioned kernels in two different UPMEM PIM systems (Table 6 presents the characteristics of the two UPMEM PIM systems). For 1D partitioning, we use the lock-free COO (COO.nnz-1f) and coarse-grained locking BCOO (BCOO.block) kernels. For each matrix, we vary the number of DPUs from 64 to 2048 DPUs, and select the best-performing end-to-end execution throughput. For 2D partitioning, we use the *equally-sized* COO (DC00) and BCOO (BC00) kernels with 2048 DPUs for both systems. For each matrix, we vary the number of vertical partitions from 2 to 32 (in steps of multiple of 2), and select the best-performing end-to-end execution throughput.

We draw three findings. First, we observe that in both systems the best performance is achieved using a smaller number of DPUs than 2048 DPUs. This is because SpMV execution in both UPMEM PIM systems is significantly bottlenecked by expensive data transfers performed via the narrow memory bus. As a result, the best-performing 1D- and 2D-partitioned kernels trade off computation with lower data transfer costs, thus causing many DPUs to be *idle*. Second, we find that in both systems the 2D-partitioned kernels outperform the 1D-partitioned kernels in regular matrices (i.e.,




Fig. 40. Throughput of 1D- and 2D-partitioned kernels for the fp32 data type using two different UPMEM PIM systems.




Fig. 41. Performance comparison of 1D- and 2D-partitioned kernels for the fp32 data type using two different UPMEM PIM systems. Performance is normalized to that of COO.nnz-lf (A).

from hgc to bns matrices on x axis), while the 1D-partitioned kernels outperform the 2D-partitioned kernels in scale-free matrices, i.e., in matrices that have high non-zero element disparity among rows and columns (i.e., from wbs to ask matrices on x axis). Third, we observe that PIM system B improves performance over PIM system A by $1.14\times$ (averaged across all matrices). This is because the DPUs of the PIM system B run at a higher frequency than that of PIM system A (425 MHz vs 350 MHz), providing higher peak performance on the system. Specifically, with 2048 DPUs, peak performance of the PIM system A and PIM system B is 3.78 GFlops and 4.63 GFlops, respectively, i.e., PIM system B provides $1.22\times$ higher computation throughput than PIM system A.

B Arithmetic Throughput of One DPU for the Multiplication Operation

We evaluate the arithmetic throughput of the DPU for the multiplication (MUL) operation. We use the arithmetic throughput microbenchmark of the PRIM benchmark suite [82, 94] and configure it for the all data types.

Figure 42 shows the measured arithmetic throughput (in MOperations per second) for the MUL operation varying the number of tasklets of one DPU at 350 MHz (PIM system A in Table 6) for all the data types. The arithmetic throughput for the MUL operation is 12.941 MOps, 10.524 MOps, 8.861 MOps, 2.381 MOps, 1.847 MOps, and 0.517 MOps for the int8, int16, int32, int64, fp32 and fp64 data types, respectively.




Fig. 42. Throughput of the MUL operation on one DPU at 350 MHz for all the data types.

Figure 43 shows the measured arithmetic throughput (in MOperations per second) for the MUL operation varying the number of tasklets of one DPU at 425 MHz (PIM system B in Table 6) for all the data types. The arithmetic throughput for the MUL operation is 15.656 MOps, 12.721 MOps, 10.732 MOps, 2.888 MOps, 2.259 MOps, and 0.631 MOps for the int8, int16, int32, int64, fp32 and fp64 data types, respectively.




Fig. 43. Throughput of the MUL operation on one DPU at 425 MHz for all the data types.

C The *SparseP* Software Package

Table 7 summarizes the SpMV PIM kernels provided by the *SparseP* library. All kernels support a wide range of data types, i.e., 8-bit integer, 16-bit integer, 32-bit integer, 64-bit integer, 32-bit float, and 64-bit float data types.

Partitioning Technique	Compressed Format	Balancing Across PIM Cores	Balancing Across Threads	Synchronization Approach
1D	CSR	rows nnz [★]	rows, nnz [★] rows, nnz [★]	- -
	COO	rows nnz [★] nnz	rows, nnz [★] rows, nnz [★] nnz	- - lb-cg / lb-fg / lf
	BCSR	blocks [†] nnz [†]	blocks [†] , nnz [†] blocks [†] , nnz [†]	lb-cg [‡] / lb-fg [‡] lb-cg [‡] / lb-fg [‡]
	BCOO	blocks nnz	blocks, nnz blocks, nnz	lb-cg / lb-fg / lf lb-cg / lb-fg / lf
2D <i>equally-sized</i>	CSR	-	rows, nnz [★]	-
	COO	-	nnz	lb-cg / lb-fg / lf
	BCSR	-	blocks [†] , nnz [†]	lb-cg [‡] / lb-fg [‡]
	BCOO	-	blocks, nnz	lb-cg / lb-fg
2D <i>equally-wide</i>	CSR	nnz [★]	rows, nnz [★]	-
	COO	nnz	nnz	lb-cg / lb-fg / lf
	BCSR	blocks [†] nnz [†]	blocks [†] , nnz [†] blocks [†] , nnz [†]	lb-cg [‡] / lb-fg [‡] lb-cg [‡] / lb-fg [‡]
	BCOO	blocks nnz	blocks, nnz blocks, nnz	lb-cg / lb-fg lb-cg / lb-fg
2D <i>variable-sized</i>	CSR	nnz [★]	rows, nnz [★]	-
	COO	nnz	nnz	lb-cg / lb-fg / lf
	BCSR	blocks [†] nnz [†]	blocks [†] , nnz [†] blocks [†] , nnz [†]	lb-cg [‡] / lb-fg [‡] lb-cg [‡] / lb-fg [‡]
	BCOO	blocks nnz	blocks, nnz blocks, nnz	lb-cg / lb-fg lb-cg / lb-fg

Table 7. The *SparseP* library. [★]: row-granularity, [†]: block-row-granularity, [‡]: (only for 8-bit integer and small block sizes)

D Large Matrix Dataset

We present the characteristics of the sparse matrices of our large matrix data set. Table 8 presents the sparsity of the matrix (i.e., $NNZ / (\text{rows} \times \text{columns})$), the standard deviation of non-zero elements among rows ($NNZ\text{-}r\text{-std}$) and columns ($NNZ\text{-}c\text{-std}$). Table 9 visualizes the sparsity patterns of each sparse matrix of our large matrix data set.

Matrix Name	Rows x Columns	NNZs	Sparsity	NNZ-r-std	NNZ-c-std
hugetric-00020	7122792 x 7122792	21361554	4.21e-07	0.031	0.031
mc2depi	525825 x 525825	2100225	7.59e-06	0.076	0.076
parabolic_fem	525825 x 525825	3674625	1.33e-05	0.153	0.153
roadNet-TX	1393383 x 1393383	3843320	1.98e-06	1.037	1.037
rajat31	4690002 x 4690002	20316253	9.24e-07	1.106	1.106
af_shell1	504855 x 504855	17588875	6.90e-05	1.275	1.275
delaunay_n19	524288 x 524288	3145646	1.14e-05	1.338	1.338
thermomech_dK	204316 x 204316	2846228	6.81e-05	1.431	1.431
memchip	2707524 x 2707524	14810202	2.02e-06	2.062	1.173
amazon0601	403394 x 403394	3387388	2.08e-05	2.79	15.29
FEM_3D_thermal2	147900 x 147900	3489300	1.59e-04	4.481	4.481
web-Google	916428 x 916428	5105039	6.08e-06	6.557	38.366
ldoor	952203 x 952203	46522475	5.13e-05	11.951	11.951
poisson3Db	85623 x 85623	2374949	3.24e-04	14.712	14.712
boneS10	914898 x 914898	55468422	6.63e-05	20.374	20.374
webbase-1M	1000005 x 1000005	3105536	3.106e-06	25.345	36.890
in-2004	1382908 x 1382908	16917053	8.846e-06	37.230	144.062
pkustk14	151926 x 151926	14836504	6.428e-04	46.508	46.508
com-Youtube	1134890 x 1134890	5975248	4.639e-06	50.754	50.754
as-Skitter	1696415 x 1696415	22190596	7.71e-06	136.861	136.861
sx-stackoverflow	2601977 x 2601977	36233450	5.352e-06	137.849	65.367
ASIC_680	682862 x 682862	3871773	8.303e-06	659.807	659.807

Table 8. Large Matrix Dataset. Matrices are sorted by $NNZ\text{-}r\text{-std}$, i.e., based on their irregular pattern.



















Matrix Name	Plot	Matrix Name	Plot
hugetric-00020		web-Google	
mc2depi		ldoor	
parabolic_fem		poisson3Db	
roadNet-TX		boneS10	
rajat31		webbase-1M	
af_shell1		in-2004	
delaunay_n19		pkustk14	
thermomech_dK		com-Youtube	
memchip		as-Skitter	
amazon0601		sx-stackoverflow	
FEM_3D_thermal2		ASIC_680	

Table 9. Sparsity patterns of the sparse matrices of our large matrix data set.