

CINM (Cinnamon): A Compilation Infrastructure for Heterogeneous Compute In-Memory and Compute Near-Memory Paradigms

Asif Ali Khan

asif_ali.khan@tu-dresden.de
TU Dresden

Clément Fournier

clement.fournier@tu-dresden.de
TU Dresden

Hamid Farzaneh

hamid.farzaneh@tu-dresden.de
TU Dresden

Lorenzo Chelini

lorenzo.chelini@intel.com
Intel Switzerland

Karl F. A. Friebe

karl.friebe@tu-dresden.de
TU Dresden

Jeronimo Castrillon

jeronimo.castrillon@tu-dresden.de
TU Dresden, ScaDS.AI

Abstract

The rise of data-intensive applications exposed the limitations of conventional processor-centric von-Neumann architectures that struggle to meet the off-chip memory bandwidth demand. Therefore, recent innovations in computer architecture advocate compute-in-memory (CIM) and compute-near-memory (CNM), non-von-Neumann paradigms achieving orders-of-magnitude improvements in performance and energy consumption. Despite significant technological breakthroughs in the last few years, the programmability of these systems is still a serious challenge. Their programming models are too low-level and specific to particular system implementations. Since such future architectures are predicted to be highly heterogeneous, developing novel compiler abstractions and frameworks becomes necessary. To this end, we present *CINM* (*Cinnamon*), a first end-to-end compilation flow that leverages the hierarchical abstractions to generalize over different CIM and CNM devices and enable device-agnostic and device-aware optimizations. Cinnamon progressively lowers input programs and performs optimizations at each level in the lowering pipeline. To show its efficacy, we evaluate CINM on a set of benchmarks for a real CNM system (UPMEM) and the memristors-based CIM accelerators. We show that Cinnamon, supporting multiple hardware targets, generates high-performance code comparable to or better than state-of-the-art implementations.

Keywords: Hardware Emerging architectures, Hardware Emerging tools and methodologies, Hardware Emerging languages and compilers, Computing methodologies Parallel computing methodologies

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ASPLOS '24, April 27-May 1, 2024, La Jolla, CA, USA

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0391-1/24/04.

<https://doi.org/10.1145/3622781.3674189>

ACM Reference Format:

Asif Ali Khan, Hamid Farzaneh, Karl F. A. Friebe, Clément Fournier, Lorenzo Chelini, and Jeronimo Castrillon. 2024. CINM (Cinnamon): A Compilation Infrastructure for Heterogeneous Compute In-Memory and Compute Near-Memory Paradigms. In *29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 4 (ASPLOS '24), April 27-May 1, 2024, La Jolla, CA, USA*. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3622781.3674189>

1 Introduction

Application domains such as social and streaming media, internet-of-everything, communications and services, and virtual assistant technologies such as Alexa and Siri are generating data at a break-neck pace, i.e., in the *quintillion bytes* range every day. This mind-boggling data volume is mostly raw and requires processing and analysis [6]. In the conventional *processor centric* von-Neumann computing paradigm, these applications quickly hit hard performance and energy efficiency boundaries as data have to be moved between the CPU and the memory via a narrow memory channel. On a mobile device, the data movement alone consumes 62% of the total system energy [11]. To overcome this data movement and other challenges associated with the memory subsystem, computer architects are moving to *non-Von-Neumann* system models like *computing near memory* (CNM) [71] and *computing in memory* (CIM) [66].¹ The idea is to bring computations closer to the data. In CNM, dedicated CMOS logic is integrated into the memory chip to diminish the data movement problem. Conceptually, this tight coupling of the logic and memory devices can be applied at any level in the memory hierarchy with various memory technologies. For DRAM, both planar and stacked structures, such as Micron's hybrid memory cube [58], AMD's and SK Hynix's high bandwidth memory [49], and Samsung's wide IO [42] have been used to realize CNM systems [71]. While CNM greatly reduces the data movement on the CPU bus, it still requires communicating data between the memory and the compute

¹CIM and CNM are also referred to as (PIM, IMC) and (NMP, PNM), respectively, in the literature. We will use CIM and CNM in this paper.

units. The CIM model completely eliminates data movement to compute units by exploiting the physical properties of the memory devices to implement various logic and compute operations in-place [66]. CIM systems based on novel memory devices with inherent computing capabilities, such as phase change memory (PCM), resistive RAM (RRAM), magnetic RAM (MRAM), and spintronics-based racetrack memories (RTMs) have demonstrated orders of magnitude performance and energy gains for machine learning and other application domains [10, 15, 29, 30, 53].

Of late, several innovative CNM and CIM systems have been proposed, some of which are also commercially available [38]. These include domain-specific architectures such as the Neurocube [41], ISAAC [68], Microsoft Brainwave NPU [22], and several DNN accelerators [64] among others. These systems are orders of magnitude faster and more energy-efficient than general-purpose Von-Neumann machines, but only target specific application domains. UPMEM [74] has shown case studies of CNM in more general-purpose off-the-shelf systems. Recently, Samsung [45, 51] and SK hynix [50] proposed machine learning specific CNM systems based on the HBM2 and GDDR6 DRAM standards supporting TFLOPS. On the CIM front, in just the last couple of years, all major memory tech giants such as Samsung [34], TSMC [14, 40], Intel [77], GlobalFoundries [16], and IBM [37, 47] have fabricated CIM chips based on memristive and CMOS technologies that attain unparalleled performance and energy efficiency.

Even though various companies currently provide CIM and CNM systems for machine learning and other application domains (such as Axelera, d-Matrix, Synthara, UPMEM, etc.), their programmability remains a significant challenge. Most of these systems provide low-level device libraries and leave the mapping problem, synchronization, and optimizations to the programmer. This makes the programmability and operability of these devices extremely difficult. In isolated efforts, compilers have been proposed to automatically map compute primitives to devices and perform load balancing and technology-specific optimizations [28, 35, 70].

However, they target only homogeneous architectures and are application-specific, e.g., GEMM on memristive crossbars [70]. Since future architectures are predicted to be highly heterogeneous and general-purpose, there is a pressing need to develop novel compiler abstractions and compiler frameworks that enable device-agnostic and device-specific optimizations. The same is also highlighted by several recent articles including one from Meta (Facebook) which states: “We’ve investigated applying processing-in-memory (PIM) to our workloads and determined there are several challenges to using these approaches. Perhaps the biggest challenge of PIM is its programmability” [8].

To this end, our goal is to develop a high-level framework that abstracts over CIM and CNM devices, enabling their programming through high-level frameworks and domain-specific languages, and generating highly efficient code for

them. We present CINM, pronounced as *Cinnamon*, a novel framework based on the *multi-level intermediate representation* (MLIR) that empowers the progressive lowering of abstractions and allows reasoning about computational primitives and their memory behavior and operations at various abstractions. CINM supports PCM and RRAM-based CIM accelerators and the UPMEM CNM architecture.²

We are using a high-end real UPMEM system and the extended gem5 simulator [70] to evaluate our generated codes for CNM and CIM systems, respectively. For PCM and ReRAM-based accelerators, CINM implements and extends OCC [70], an automatic compilation flow for memristive crossbar arrays. The hierarchical lowering in the CINM enables identifying the most suitable target for each primitive in the input application and transformations at different abstractions to optimize for individual devices. For evaluation, we use the sets of benchmarks available for these systems, i.e., PrIM benchmarks for CNM [24] and the machine learning benchmarks from [70] for CIM. Concretely, we make the following contributions:

1. We investigate the landscape of CIM and CNM systems to understand their properties and compute the primitives they support (Section 2.4).
2. We present CINM, an end-to-end compilation framework based on MLIR that seamlessly maps computational motifs to different backend targets (Section 3.1).
3. CINM implements multiple hardware-oblivious and hardware-specific abstractions. Concretely, we introduce *cim/cnm* abstractions that implement abstract operations for CIM/CNM paradigms, which are subsequently lowered differently for different hardware targets in their respective device dialects (Section 3.2).
4. We introduce a high-level *cim* dialect that abstracts over all CINM devices and provides a placeholder for implementing cost models to automate the mapping of k kernels/regions onto d devices in a heterogeneous system setup (Section 3).
5. We implement device-specific abstractions to perform device-aware optimizations and mapping to their respective libraries.
6. Our evaluation shows that CINM can effectively reproduce or beat the performance of the hand-optimized codes in the selected benchmark suites (Section 4).

2 Background and motivation

This section presents MLIR as well as the CNM and CIM computing models using different memory technologies.

2.1 The MLIR compiler infrastructure

MLIR is a toolkit to represent and transform intermediate representation (IR) at different abstraction levels across different application domains and heterogeneous hardware targets [46].

²The selection of architectures is influenced by the availability of infrastructure where these systems can be evaluated.

It offers a nonopinionated IR with few builtins, leaving most of the IR customizable. MLIR allows compiler developers to plug in into the compiler their own abstraction and empowers them to optimize for a specific domain or target by matching at the appropriate abstraction levels.

MLIR implements a set of reusable abstractions modeled with *dialects*. A dialect is a logical group of custom types, operations, and attributes. Operations are building blocks of the IR and consume and produce new values. Each value in MLIR is associated with a type known at compile time. Attributes associate compile-time information to operations. Dialects in MLIR preserve transformation validity preconditions in their IR in order to minimize the cost and complexity of analysis passes. They are typically associated with domains (`linalg` with linear algebra, `TOSA` with tensor operations), representations (`affine` with the polyhedral model, `scf` with control flow), or targets (`gpu`, `cim`). Abstractions in MLIR can be progressively lowered (e.g., from high-level domain-specific to low-level platform-specific dialects) and raised [12].

2.2 Compute near memory

Compute near memory (CNM) is a data-centric paradigm aiming to process data in memory proximity. Compute units, e.g., CPU, GPU, FPGA, ASIC, or CGRA, are physically placed closer to the memory (in the memory controller, in peripheral circuitry, on the memory chip, or connected to the memory chip via a shared crossbar) to minimize data movement. The original idea of CNMs dates back to 70s [73]. In the 90s, architectures such as EXECUBE [43] and IRAM [57] demonstrated significant performance gains in a range of applications. However, design complexity and manufacturing costs hindered commercialization. Recent advances in manufacturing and stacking technologies alleviate these practicality concerns, paving the way for many novel CNM architectures.

Stacked DRAM structures such as the hybrid memory cube (HMC) [58] and the high bandwidth memory (HBM) [49] are considered the true enablers of CNM systems. These architectures stack multiple DRAM dies on top of a logic layer using through silicon vias (TSVs), where the logic layer can implement fixed function units. These stacked solutions deliver higher bandwidth and improved performance compared to other DRAM families but can lead to higher refresh power and limited capacities. UPMEM integrated co-processors with the DDR4 DRAM on the same DRAM die [74]. The co-processor, known as the data processing unit (DPU), is a general-purpose 32-bit RISC processor. Due to its massive local and cumulative bandwidth and parallelism, UPMEM demonstrated an order of magnitude gains in performance and energy consumption on different applications [24].

Each DPU has a small private scratchpad working RAM (WRAM) backed by the shared main RAM (MRAM). UPMEM provides an SDK and a set of tools that allow developers to adapt to the PIM programming. More recently, Samsung and SK Hynix presented their FIMDRAM [45, 51] and AiM [50]

architectures, respectively. Similar to UPMEM, these architectures integrate co-processors on the same DRAM die (using HBM2 and GDDR6 DRAMs). However, unlike UPMEM, the co-processors in both architectures are optimized explicitly for ML-specific workloads.

2.3 Compute in memory

The compute in memory paradigm radically departs from traditional architectures by implementing certain compute motifs *using* the physical attributes of the devices. Memristive devices such as PCM and RRAM cells can be programmed to different resistance states using external current/voltage, where each state represents some information. When organized in a crossbar configuration, these memristive devices allow for in-place fixed-size matrix-vector (MV) multiplication in constant time [33]. However, these computations are in the analog domain and require converters from the digital to the analog domain and back. In a different crossbar setup, memristors can be used to implement the entire set of logical operations [44] entirely in the digital domain. The write operation in these resistive technologies is typically very slow and reduces the device's lifetime. Therefore, the selection of an application for CIM acceleration requires careful consideration.

Magnetic memories such as MRAM and RTM can also be used to implement certain operations in place. The tunnel magnetoresistance in the magnetic tunnel junctions (MTJs) of MRAM cells is a natural implementation of the XOR operation, which can be exploited to implement other logic operations [30]. Similar to memristors, MRAM cells can also be organized in crossbars to realize MV operations [34]. RTM devices also use MTJs as access interfaces and can use the same basic principles to implement various logic operations [79]. They also offer novel access mechanisms that allow efficient implementation of population count and the majority operations [39, 55]. Conventional charge-based SRAM and DRAM technologies can also implement a series of logic and compute operations in-place [4, 67].

2.4 The need for CIM and CNM abstractions

Figure 1 presents a partial taxonomy of prominent CNM and CIM systems. On the CNM side, the figure shows only real-world systems. Similarly, only promising and mature CIM technologies are presented on the CIM side.




Figure 1. A partial taxonomy of CNM and CIM systems based on pronounced technologies.

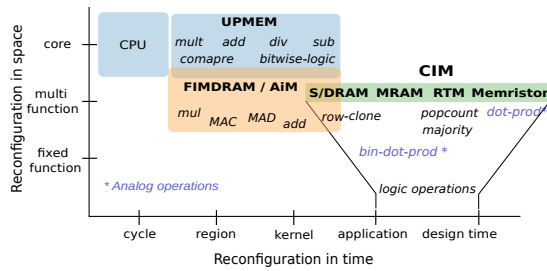


Figure 2. CNM and CIM programmability landscape.

These architectures are typically optimized for a specific domain or function. Figure 2 shows the landscape of these architectures with their supported operators and their specificity and flexibility to be reconfigured in time and space. For instance, the general-purpose CPU is programmed at the granularity of the core every new instruction cycle. On the contrary, application-specific ICs (ASICs) are optimized for a particular application and can not be reprogrammed. The near-memory logic in CNM systems can be general-purpose (UPMEM), or multi-function (AiM, FIMDRAM), and they are programmed at the kernel and region granularities where, in the latter case, a kernel is partitioned into regions before offloading it to the CNM devices. CIM systems are usually fixed-function (e.g., in dot-product), but they can also be multi-function (e.g., logic operations) and can be programmed at the application granularity.

Unfortunately, even for this limited set of systems, there is a lack of programming models that abstract over them and can be leveraged to program them. All CNM and CIM systems use low-level, architecture-specific libraries to expose their device traits. The radically different design decisions and architectures of these systems make their programmability a serious challenge. For instance, in UPMEM, the programmer is responsible for load balancing on thousands of DPUs, explicit data movement and bandwidth management between the CPUs and DPUs, MRAM and WRAM, and the data coherency [24, 74]. The AiM architecture has unique features that allow operations including row clone, element-wise multiplication, and addition on a set of banks with different granularities [50]. However, it is not clear how these systems are programmed. Samsung’s FIMDRAM has its own *closed-source* software stack [51]. The programming models and tools of different CIM technologies, e.g., for PCM [70], do not apply to other technologies (such as MRAM) as they have different properties.

Figure 3a shows an implementation of the GEMM kernel in UPMEM’s C interface. Performing simple read and write for a computation requires specific API calls and manual address translation. For other architectures, this code has to be completely rewritten using their device-specific function calls. With these programming models, it is extremely difficult to program and effectively utilize heterogeneous systems

integrating these technologies. Even for the same system, any device or system changes may lead to a considerable rewriting of the input applications. To enable the integration and exploration of these devices in heterogeneous setups, novel programming models that abstract from these devices to a higher level are needed. CINM’s device-agnostic abstraction is a step in that direction. The GEMM input to CINM is device independent, as shown in Figure 3b, and can be lowered to CIM or CNM device code. This also highlights the expressiveness and conciseness of CINM compared to the low-level device-specific programming models. CINM allows rigorous analysis and reasoning about individual kernels and supports a rich set of optimizations and device interfaces.

```

BARRIER_INIT(my_barrier, NR_TASKLETS);
int main() { ...
    barrier_wait(&my_barrier);
    int32_t point_per_tasklet = (ROWS*COLS)/NR_TASKLETS;
    uint32_t mram_base_addr_A = (uint32_t) (DPU_MRAM_HEAP_POINTER );
    uint32_t mram_base_addr_B = (uint32_t) (DPU_MRAM_HEAP_POINTER +
        ROWS * COLS * sizeof(T));
    uint32_t mram_base_addr_C = (uint32_t) (DPU_MRAM_HEAP_POINTER +
        2 * ROWS * COLS * sizeof(T));
    for(int i = (tasklet_id * point_per_tasklet) ;
        i < ( (tasklet_id+1)*point_per_tasklet) ; i++) {
        if( new_row != row ){ ...
            mram_read((__mram_ptr void const*) (mram_base_addr_A +
                mram_offset_A), cache_A, COLS * sizeof(T));
        }
        mram_read((__mram_ptr void const*) (mram_base_addr_B +
            mram_offset_B), cache_B, COLS * sizeof(T));
        dot_product(cache_C, cache_A, cache_B, number_of_dot_products);
        ... } ...
    mram_write( cache_C, (__mram_ptr void *) (mram_base_addr_C
        + mram_offset_C), point_per_tasklet * sizeof(T));
}
    
```

(a) Matmul on UPMEM. Each tasklet runs this code to generate a single element in the resultant matrix.

```

func.func @matmul(%A: tensor<64x64xi32>, %B: tensor<64x64xi32>, %C
    ↪ : tensor<64x64xi32>) -> tensor<64x64xi32> {
    %D = linalg.matmul ins(%A, %B : tensor<64x64xi32>, tensor<64
        ↪ x64xi32>) outs(%C: tensor<64x64xi32>)
    return %D : tensor<64x64xi32>
}
    
```

(b) GEMM code at the linalg abstraction in CINM.

Figure 3. Matmul code using the UPMEM programming model and the device-unaware linalg abstraction in CINM.

Compared to device libraries, compiler frameworks like CINM are interesting alternatives or complements because for most CIM and CNM systems, such libraries do not exist or are device-specific and thus not portable. In addition, libraries use kernels as-is, while compilers like ours, if the device supports it, can fuse operations to reduce the data movement. Compiler-optimized codes have also been shown to be on-par with the libraries [36]. CINM can be extended to map to optimized libraries when they become available (in the same way many DSLs map to, e.g., BLAS calls).

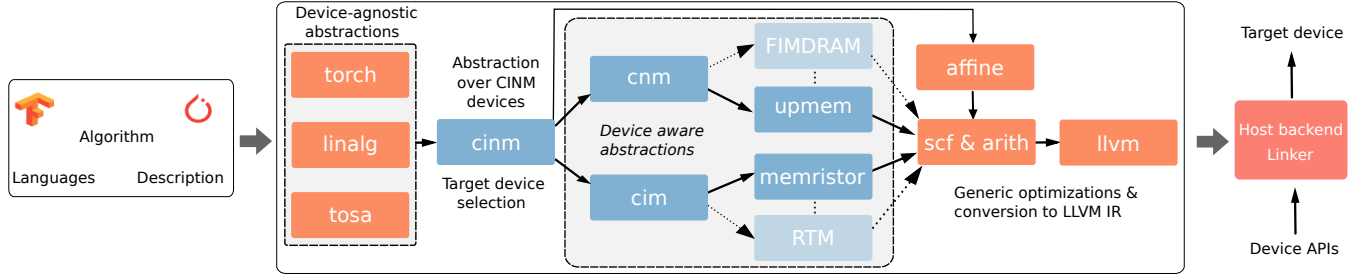


Figure 4. The CINM compiler. The abstraction lowers from left to right. Blue boxes show the new dialects introduced in CINM.

3 CINM (Cinnamon): Compilation for in and near memory computing

The explosion of CIM and CNM technologies and architectures has led to fragmented toolchains and device-specific low-level programming APIs. This imposes a high barrier of entry for programmers, leading to low adoption and potentially slower technology evolution cycles. Based on the taxonomy and operator pools described in Section 2.4, we present CINM – an end-to-end compilation flow that generates high-performance code for various target devices. CINM leverages MLIR to optimize input programs by progressively lowering from high-level domain abstractions to low-level device abstractions. Each device dialect supports device-aware transformations to ensure effective utilization of the underlying system. In the scope of this work, we target only memristors-based CIM systems and the UPMEM CNM systems, as highlighted in Figure 1. However, considering the ongoing device innovations, we design our abstractions with a special focus on extensibility. The restriction to selected architectures in this work is enforced by the lack of open-source tools for other architectures that can be used to evaluate the generated code for them.

3.1 The CINM lowering pipeline

Figure 4 presents a high-level overview of the CINM compilation flow. The entry point to the compilation flow is the `cinm` dialect, which currently accepts the IRs from the `linalg`, `TOSA` and `torch` abstractions. The `cinm` abstraction is a generalization over different CINM technologies and takes over the shared responsibilities of host-device interfacing and device mapping. The latter may also require the input program to be (re)written in CINM-supported operations (see Figure 2), which can then be processed by the low-level dialects. The `cinm` dialect is then lowered to the `cim`, `cnm` or `affine` dialects.

`cim` and `cnm` abstractions implement custom types and operations that are common to these architectures. For instance, in all CNM devices, the host allocates the grid of compute devices and transfers data before launching the kernel. The `cnm` dialect implements abstract prototypes of these functions using custom types that are contextually converted to the device

types. The concrete mapping from `cnm` operations to the target devices is then accomplished using device dialects. These dialects serve as interfaces to their respective accelerators and runtimes. All device dialects provide their lowering for code generation, which could mean emitting runtime library calls (e.g., for `upmem`) or CPU instructions (for devices embedded as ISA extensions).

3.2 Progressive lowering

This section describes the CINM pipeline, particularly focusing on our newly added dialects and their primitives.

3.2.1 CINM front-ends. The entry point to the CINM framework is the `cinm` abstraction. To arrive at this abstraction, CINM presently starts from the high-level abstractions `linalg`, `tosa`, or `torch`, which serve as widely adopted targets for several DSLs and frameworks, particularly in the ML domain, including TensorFlow [19], ONNX [32], OpenEarth Compiler [25], CFDlang [72], and others. In terms of applicability, they handle a wider set of applications from the ML and other domains, as demonstrated in Section 4, but can not necessarily handle *any* input application. It is important to note that the CINM framework itself is not tied to these front-ends. In general, the program representation can be in any high-level or domain-specific language that already provides a front-end to the `linalg` abstraction, e.g., Tensor Comprehension (via `teckyl`), TorchScript (via `torch-mlir`), and a subset of C/C++ (via `Polygeist` [54]), among others. For non-`linalg` operations or fine-grained control over the target selection process, front-ends must provide direct lowering into `cinm`. More advanced front-ends, such as `Mojo` [2], which handle optimizations in a compiler-as-library approach, would need to entirely sidestep `cinm` and directly emit `cnm` and `cim`.

3.2.2 The `cinm` dialect. The `cinm` dialect is responsible for target selection, i.e., delegating the implementation of k kernels or regions in an input application to the most suitable d devices. This is not arbitrary, as it requires precise cost models for individual devices and an exhaustive search mechanism to evaluate tradeoffs before making mapping decisions. For some operations such as `cinm.gemm`, estimating the execution time (or other metric) on a specific architecture might not

Table 1. The `cinm` dialect operations. T marks any shaped type, S a scalar type, and E an appropriate enumeration type.

Operation	Type	Description	CIM	CNM
<code>cinm.{add,sub,mul,div,min,max}{%lhs, %rhs}</code>	$T \times T \rightarrow T$	Element-wise arithmetic	✓	✓
<code>cinm.{and,or,xor,not}{%lhs, %rhs}</code>	$T \times T \rightarrow T$	Element-wise bit-wise logic	✓	✓
<code>cinm.gemv{%lhs, %rhs}</code>	$S^{m \times n} \times S^n \rightarrow S^m$	Matrix-vector product	✓	✓
<code>cinm.gemm{%lhs, %rhs}</code>	$S^{m \times k} \times S^{k \times n} \rightarrow S^{m \times n}$	Matrix-matrix product	✓	✓
<code>cinm.transpose{%in, %perms}</code>	$S^n \times \mathbb{N}^n \rightarrow S'$	Transposition	✗	✓
<code>cinm.{histogram,majority}{%in}</code>	$S^n \rightarrow S^k$	Histogram and bit-wise majority	✗	✓
<code>cinm.topk{%in, %k}</code>	$S^n \times \mathbb{N} \rightarrow S^k \times \mathbb{N}^k$	Finds k largest values & their indices	✗	✓
<code>cinm.simSearch #E, #k {%in1, %in2}</code>	$E \times \mathbb{N}^k \times S^n \times S^n \times \mathbb{N} \rightarrow S^k$	Finds k most similar values & their indices with metric E	✓	✓
<code>cinm.mergePartial #op #dir {%lhs, %rhs}</code>	$E \times D \times T \times T \rightarrow T$	Hardware-defined operation that merges partial results of E	✓	✓
<code>cinm.popCount{%in}</code>	$T \rightarrow \mathbb{N}$	Counts 1s in a bit vector	✓	✗
<code>cinm.reduce #op {%in}</code>	$E \times S^n \rightarrow S$	Performs reduction in group (S, E)	✗	✓
<code>cinm.scan #op {%in}</code>	$E \times S^n \rightarrow S^n$	Performs inclusive scan in group (S, E)	✗	✓

be that difficult. However, for more complex operations or a combination of the easy-to-estimate operations, additional analysis and rewriting passes becomes necessary. For instance, none of the discussed CINM architectures are optimized for convolution and contraction operations. In the absence of rewriting, the subsequent lowering from `cinm` will have no choice but to map these operations to more general compute-capable devices, like, e.g., UPMEM or the host CPU. However, having detected these kernels and identified them as profitable, they can be rewritten as `matmul`, amenable to both CIM and CNM acceleration. Therefore, the `cinm` abstraction must evaluate all code variants before making any mapping decision.

CINM provides a mechanism to implement and integrate device cost models that can be used for target selection and offloading decisions. Due to the absence of cost models for our targeted devices (see Section 3.3), presently, either the user specifies the target device via the command line or the analysis passes at `cinm` automatically select a target. For instance, `matmul`-like operations or those that can be rewritten as `matmul` are greedily offloaded to the CIM crossbar if the tensor dimensions are greater than a user-defined threshold. Convolutions and contractions are identified using the analysis algorithm from OCC [70]. Similarly, search operations suited to content-addressable memories (CAMs) can be detected using the analysis algorithm from C4CAM [20]. For all other `cinm` operations, CINM presently selects UPMEM as a target.

CINM is designed to support offloading at finer granularities. For instance, within a kernel, if some operations can be accelerated on one device and others on another, the framework should seamlessly perform such offloading. However, this requires extensive analysis, which is only possible when the cost models become available.

For the cost model to work at this abstraction, it must be aware of all the primitives supported by the underlying devices, as shown in Figure 2. `cinm`, therefore, exposes the set of operations listed in Table 1 to be targeted during offloading. The `cinm` IR is emitted by a conversion pass that applies to the `linalg` input. The `linalg` input may have been produced by a front-end (see Sec. 3.2.1) for a high-level framework or another higher-level front-end dialect like `tosa` or `torch`.

`linalg` operators that do not have a direct counterpart in the `cinm` dialect are first canonicalized into a sequence of lower-level operators that are mappable to `cinm`. For instance, in an MLP application represented at the `tosa` abstraction, the `tosa.fully_connected` operation is first decomposed into a sequence of `linalg` operations (`transpose`, `matmul`, and `bias addition` using a `generic` operation). During `linalg` to `cinm` conversion, the `generic` operation responsible for adding the bias is rewritten with a `cinm.add` operation in the canonicalization pass. Operators that still cannot be converted are run on the host CPU, as they cannot participate in the `cinm` target selection.

Figure 5 shows a convolution kernel at the `linalg` (Figure 5a) and `cinm` (Figure 5b) abstraction levels. The conversion process for this convolution kernel from `linalg` to `cinm` involves rewriting convolutions with an `im2col` (lines 6-7) operation followed by `gemm` (lines 8-9) and `expand` (lines 10-11) operations. The `cinm` IR is subsequently lowered to `cim`, `cnm`, `affine` or a combination thereof.

```

1 %conv = linalg.conv2d_nhwc_hwc ins(%img, %flt:
2   tensor<1x128x128x3xi16>, tensor<3x3x3x8xi16>) outs
3   (%bias: tensor<1x126x126x8xi16>->tensor<1x126x126x8xi16>

```

(a) `linalg` IR for 2D convolution.

```

1 %cbuf= linalg.generic #im2col_traits
2   ins(%img: tensor<1x128x128x3xi16>)
3   outs(%init: tensor<1x126x126x3x3xi16>) {
4     ^bb0(%arg0: i16, %arg1: i16): linalg.yield %arg0 : i16
5   } -> tensor<1x126x126x3x3xi16>
6 %im2col= tensor.collapse_shape %cbuf [[0,1,2],[3,4,5]]
7   : tensor<1x126x126x3x3xi16> into tensor<15876x27xi16>
8 %rbuf= cinm.op.gemm %arg0, %flt :
9   tensor<15876x27xi16>, tensor<27x8xi16>
10 %conv= tensor.expand_shape %rbuf [[0,1,2],[3,4,5]]
11   : tensor<15876x8xi16> into tensor<1x126x126x8xi16>

```

(b) `cinm` IR for the 2D convolution kernel, rewritten as GEMM.**Figure 5.** Convolution kernel at different abstractions.

3.2.3 The `cnm` dialect. Every CNM architecture shown in Figure 1 has unique features that make it superior to others for certain workloads. In general, the performance efficiency is strictly determined by the parallel processing units (PUs)

```

1 #scatter_map = affine_map<(d0, d1) ->
2   (d0 floordiv 16, d1 floordiv 16, d0 mod 16, d1 mod 16)>
3   ...
4 %C_pad = scf.for %i0 = %cst0_i to %cst15888_i step %cst128
5   iter_args(%in_result = %in) -> tensor<15888x16xi16> {
6   %A_tile = tensor.extract_slice %A_pad[%i0, %cst0][128,32][1,1]
7   : tensor<15888x32xi16> to tensor<128x32xi16>
8   %wg = cnm.workgroup [8 2] {cnm.physical_dims=["dpu","thread"]}
9   %A_buf = cnm.alloc() for %wg { cnm.physical_space = "global" }
10  : !cnm.buffer<16x16xi16, level 0> for !cnm.workgroup<8x2>
11  ...
12  %sc_a_tok = cnm.scatter %A_tile into %A_buf[#scatter_map]
13  of %wg : tensor<128x32xi16> into
14  !cnm.buffer<16x16xi16, level 0> of !cnm.workgroup<8x2>
15  ...
16  %e_tok = cnm.launch %wg (%A_buf, %B_buf, %C_buf:
17  !cnm.buffer<16x16xi16, level 0>, ...) {
18  ^bb0(%A_space: memref<16x16xi16>, %B_space: ...):
19  scf.for %arg2 = 0 to %cst16 {
20  %0 = memref.load %A_space[%arg0, %arg1] : memref<16x16xi16>
21  %1 = memref.load %B_space[%arg1, %arg2] : memref<16x16xi16>
22  ... }
23  } : !cnm.workgroup<8x2>
24  %C_tile, %g_tok = cnm.gather %C_buf[#scatter_map] of %wg
25  : !cnm.workgroup<8x2> into tensor<128x32xi16>
26  ...
27  scf.yield %out_result : tensor<15888x16xi16> }
28  ...

```

(a) `cnm` IR for convolution.

```

1 ...
2 %C_pad = scf.for %i0 = %cst0_i to %cst15888_i step %cst128
3   iter_args(%in_result = %in) -> tensor<15888x16xi16> {
4   %bias = arith.constant dense<0.0> : tensor<16x16xi16>
5   %C_tile = scf.for %i1 = %cst0_i to %cst128 step %cst16
6   iter_args(%in_tile = %bias) -> tensor<16x16xi16> {
7   %A_tile = tensor.extract_slice %A_pad[%i0, %i1][16,16][1,1]
8   : tensor<15888x32xi16> to tensor<16x16xi16>
9   %B_tile = tensor.extract_slice %B_pad[%i1, 0][16, 16][1, 1]
10  : tensor<32x16xi16> to tensor<16x16xi16>
11  %id = cim.acquire : cim_id
12  %C_tile = cim.execute(%id: cim_id, %A_tile :
13  tensor<16x16xi16>, %B_tile: tensor<16x16xi16>){
14  %out = cim.gemm %A_tile, %B_tile :
15  tensor<16x16xi16>, tensor<16x16xi16>
16  cim.yield %out: tensor<16x16xi16> }
17  cim.release %id : cim_id
18  %out_tile = arith.addf %in_tile, %C_tile
19  : tensor<16x16xi16>
20  scf.yield %out_tile : tensor<16x16xi16> }
21  %out_result = tensor.insert_slice
22  %C_tile, %in_result[%i0, 0][16, 16][1, 1]
23  : tensor<16x16xi16> into tensor<15888x16xi16>
24  scf.yield %out_result : tensor<15888x16xi16> }
25  ...

```


(b) `cim` IR for convolution.**Figure 6.** `cnm` and `cim` IRs for 2D convolution.

and data locality, i.e., in WRAM for UPMEM and in GRF for FIMDRAM. These aspects are abstractly represented by the `cnm` abstraction. This intermediate dialect abstracts over the common features of CNM architectures, which use a tightly coupled hierarchy of memory and compute elements. It aims to provide a low-complexity transition from parallel workloads to memory-distributed programs that can guarantee access patterns and mappings.

Table 2 presents the set of operations supported by the `cnm` abstraction. We separate the host and the device codes and represent device resources (memory/compute) with workgroups. Figure 6a shows the `cnm` IR of the convolution example in Figure 5a, after applying tiling at the `cnm` abstraction considering a tile size of 16×16 , demonstrating workgroup creation, its

mapping onto the CNM system resources, and the launching of its execution.

A workgroup is a logical address space that arranges memory resources in a tree, with PUs at its leaves, see Figure 7. The dots in the left grid show logical PUs, while the different memory spaces on the right correspond to different memory hierarchy levels. For instance, the (i, j, k) space could represent the local-private memory, the (i, j) column, a shared memory shared by some PUs, and the (i) space another memory shared by a larger set of PUs, all backed by the global memory.

**Figure 7.** A `cnm` workgroup with dims (i, j, k) . The red PU on the left lives in the 4 non-aliasing memory spaces shown on the right. The red path shows the tree hierarchy of accesses.

PUs in a workgroup run in parallel, and lowering will eventually map (slices of) them to concurrent execution units of the target device. This is similar to CUDA, except that the workgroup dimensions can be chosen freely and that memory is intrinsically bound to levels in the tree. A PU can only access memory along its path to the root (the global memory), and `cnm` forbids the user from manipulating any memory directly. This is achieved by the abstract `allocate` operation that obtains an opaque buffer which is subsequently used by the `scatter` and `gather` operations for memory accesses. The buffer is then converted into regular `memrefs` on device memory inside the workgroup's launch operation. This pattern ensures that memory can be moved in the tree without invalidating the program. Note that physical memory hierarchy of devices can be different and this will work with any device as a buffer can always be moved safely to a parent level (at a latency cost). For instance, in UPMEM, MRAM is technically global but is partitioned for exclusive access by the code generator.

`cnm` allows us to map parallelism inherent in an algorithm to concurrency on the device without relying on implementation details, like thread IDs (which are emulated if needed as with UPMEM) and implicit memory sharing. Suppose we want to implement the Einsteinian tensor expression $x_{ijk} = A_{ir}B_{rjk} + C_{jk}$ with $A[M, P]$, $B[P, N, O]$, $C[N, O]$. We can immediately assume i, j, k as our parallel workgroup domain over $[M, N, O]$. We can also tell that each PU will have to have access to some slices $A' [P]$, $B' [P]$, and $C' []$

Table 2. The `cnm` dialect operations.

Operation	Description
<code>cnm.allocate(%arg, %arg2)</code>	Allocate workgroup on the CNM device.
<code>cnm.launch(%arg, %arg)</code>	Launch the workgroup execution.
<code>cnm.scatter(%arg, %arg2)</code>	Move specified elements' indices of the input tensor to the destination tensor.
<code>cnm.gather(%arg, %arg2)</code>	Symmetrical to scatter, copy back.
<code>cnm.wait(%arg, %arg2)</code>	Wait to synchronize.

to avoid reduction dependencies across PUs. We can interchange, coalesce, and split dimensions of the workgroup freely since the PUs are independent, and this will not change their working set A' , B' , C' either. When doing so, the compute remains unchanged, but the buffers commute, changing the total device memory required and the amount of scalars copied. Figure 8 shows an example of such a transform, which changes the memory footprint from $M(P + NO(P + 1))$ to $NO(MP + P + 1)$, which is advantageous for large M . From here, the resulting 2D workgroup can be tiled as shown in Figure 9 to utilize the device and obey physical memory restrictions.




Figure 8. An example workgroup transform for $x_{ijk} = A_{ir}B_{rjk} + C_{jk}$ from (i, j, k) to (h, i) , showing the buffer nodes storing the working set.

3.2.4 The `cim` dialect. The `cim` abstraction serves the same purpose as `cnm` but for CIM targets. Similar to the `cnm` abstraction, it implements functions for acquiring/releasing the CIM accelerators, the data transfers, and the launching of the CIM kernel execution. Note that the CIM and CNM systems are fundamentally different, and so is the process of their resource allocation and management. Since most CIM devices are nonvolatile, this abstraction also implements device locking to ensure consistent and permanent NVM states. Table 3 lists the set of operations supported by this dialect while Figure 6b shows its IR for the tiled version of the running example with the tile size set as 16×16 .

The rationale for separating the CIM abstraction from the CNM abstraction is based on the architectural differences between the two paradigms they serve. However, within the CIM paradigm itself, despite varying operations and configurations, there are common properties in different setups, e.g., content-addressable memory (CAM)-based CIM, logic CIM,

Table 3. The `cim` dialect operations.

Operation	Description
<code>cim.acquire()</code>	Acquire a CIM device, returns ID.
<code>cim.write(%arg, %arg2)</code>	Write the input tensor to the CIM device.
<code>cim.execute(%arg, %arg)</code>	Launch execution on the CIM device.
<code>cim.read(%arg)</code>	Read data from the acquired CIM device.
<code>cim.barrier(%arg, %arg2)</code>	Wait to synchronize or finish executing.
<code>cim.release(%arg)</code>	Release the device.

and crossbars. For instance, all these CIM types require a device setup before executing any meaningful operations, and the series of required operations are mostly similar. The exact implementation of the device setup depends on the CIM target. For example, with crossbars, the controller orchestrates data flow and issues instructions in the correct sequence. The device setup involves configuring the controller and determining parameters such as device levels (bits per cell), the number of shared ADCs, and the write mode (open-loop or write-verify), among others. In the case of RTM, this entails configuring the data storage mode (sequential or bit-interleaved) and setting up the device for transverse read operations [39], considering the transverse read distance. The `acquire` function in CIM acquires a CIM device by first setting it up. The `read/write` operations facilitate data transfers and accessing/updating the data in the crossbar arrays. The `execute` function launches the execution in the CIM accelerator. Note that this supports both fixed single-function (e.g., crossbars) and multi-functions (e.g., CIM logic) accelerators. The exact lowering is performed at the device abstraction level and varies depending on the specific technologies involved.

Since CIM array sizes are fixed, if the kernel size exceeds the array size, `cim` triggers the *tiling* transformation to partition the input tensors in order to fit them into the array. The `cim` dialect also implements CIM-specific optimizations. Most CIM devices are NVM-based, and the write operations in almost all NVM technologies are costly in terms of both performance and device lifespan. After applying the tiling pass, the order of the new loop nest (i, j) can be interchanged to (j, i) without encountering any dependency violations, as they are independent. In general, the *loop interchange* is applied to enhance the memory access pattern [80]; `cim` leverages this to minimize the number of writes.

3.2.5 Device dialects. Device dialects in CINM expose a set of device-specific concepts, including: the set of supported operations, device attributes e.g., memory array or tile sizes, the memory hierarchy (buffers, private and shared memories). They apply conversion patterns to translate the `cnm` operators and provide an interface to the device libraries.

Memristors: The `memristor` dialect implements the transformation passes from OCC [70] for memristive devices. We extend the OCC flow to adapt the `cnm` operations and transformations. The compulsory *tiling* transformation is applied at the `cim` abstraction to ensure large input buffers are divided

into blocks that can be mapped to the crossbar tiles. The implementation details of the tiling transformation are given in Section 3.2.6. The tile size in the transformation, and hence the number of tiles are determined by the crossbar tile size. This dialect materializes the `cinm` operations using the memristors' specific operations such as `read`, `write` and `read to perform computations`, program crossbars, and support data communication between the host and the device, respectively.

To enable parallel execution across multiple CIM tiles, `memristor` applies the *loop unrolling* transformation on the innermost loop of the matmul kernel. The transformation pass takes an unroll factor, and modifies the body and loop variable of the innermost loop accordingly. The partial results of individual tiles are accumulated using `cinm.mergePartial` as soon as they are ready. Finally, the `memristor` dialect maps all operations to the device function calls. All `memristor` operators have a one-to-one mapping with the device function calls exposed by the memristor devices' API. All other operations are lowered to the host instructions.

UPMEM: The `upmem` dialect features UPMEM device-specific transformation and optimization passes. Each UPMEM DRAM bank has an integrated DPU, complemented by a 4 kB instructions memory (IRAM), a 64 kB WRAM (scratchpad) and a 64 MB main memory (MRAM). The DPUs communicate with other DPUs via the host. An UPMEM DIMM module integrates M chips, each with N DPUs.

The `upmem` abstraction also enables configuring the number of tasklets per DPU and allocating buffers in both the private WRAM and the MRAM. The number of tasklets can be configured by the user. By default, CINM uses values that are empirically extracted for different operations and different tensor sizes. For instance, for the matmul operation, the best-performing results for large-size tensors were achieved by setting the tasklets to 16. For synchronization, `upmem` introduces operations that can be ultimately mapped to the UPMEM `barrier_wait` function calls for all threads.

Adding new devices Adding a new hardware target to CINM requires defining a new dialect (capturing device intrinsics) and implementing the necessary conversions. If it supports operations that are not in the `cinm` registered operations, which is not very common in these kinds of architectures, the `cinm` dialect would need to be updated to achieve automation. However, in general, the abstractions in `cinm`, `cim` and `cnm` are reused as-is for other architectures.

For example, to support Samsung's FIMDRAM, a new device dialect needs to be added to CINM that contains device-specific operations such as `ADD`, `MAD`, `MUL`, and `MAC` computing operands from different memory sources, e.g., (register file(s), bank), control operations such as `JUMP`, `EXIT`, and `barrier` operation. A new conversion pass needs to be implemented from the `cnm` abstraction to the new device abstraction. Since all of the operations for this target are




Figure 9. Various tiling shapes CINM implements.

already supported by `cinm`, *no further changes are needed* to the higher abstractions, i.e, `cinm`, and `cnm`.

Low-level dialects The optimized device-aware IRs are lowered to the low-level dialects common to various compilation paths. The `scf` dialect provides standard control flow primitives, i.e., `scf.for`, `scf.while` and `scf.if`. This is then lowered to the `llvm` dialect, which closely mirrors the LLVM IR, and can be translated to machine code.

3.2.6 Tiling and partitioning passes. CINM implements a generic tiling transformation using an MLIR interface that can be triggered by the lower-level dialects. Operations in device-aware dialects select tile sizes based on architectural properties (e.g., crossbar size) and call this interface. CINM architectures perform tiling for either parallelism, improving the local memory locality, or compulsory tiling to fit operands onto CIM devices. Although the transformation is the same, it has different impacts and produces different results. For instance, for the matmul operation with operands of sizes $A : m \times k$ and $B : k \times n$, tiling on different dimensions produces different results and offers different tradeoffs. For instance, the box and rectangular tilings in Figure 9 produce different partial results and have different localities. For CIM systems, the strategy for tiling and partitioning kernels must align with hardware constraints such as reliability, endurance, and other unique hardware-specific attributes. While these characteristics are inherently device-specific, they can be regarded as variables that dictate how to tile.

3.3 Device cost models

For the `cinm` dialect to make optimal mapping decisions, it must compare the performance of different implementations on different architectures considering the device constraints. This requires a cost models based on comparable metrics across devices, which is an outstanding research question. In this work, we do not focus on the development of cost models and search mechanisms as this is orthogonal to CINM and is left to future research when more reference points for comparison will be available.

In our proposed flow, we designed a mechanism that can be used to leverage such models when available. The `cinm` dialect declares an interface [1], whose implementations can be registered by a device dialect at the time it is loaded.

Considering the target hardware constraints, a `cinn` lowering conversion can be delegated to these interfaces to evaluate the cost model. When available, the appropriate selection algorithm, e.g., comparing the estimated ranges, will automate the mapping decision at the `cinn` level (current offloading mechanism is explained in Section 3). In this scenario, the `cinn` dialect will provide the advantage that the cost model can work on the constrained subset of interface operations defined by `cinn` instead of arbitrary programs.

In addition to target selection, cost models can also be used to guide optimizations in CINM, similar to previous compilers for conventional architectures and accelerators [9, 17, 61, 76]. They include identifying tile size and shape for a given architecture, data layout transformations [76], and other loop transformations [9].

3.4 Heterogeneous system setups

In this paper, we evaluate systems with a host CPU and a CIM/CNM device (CPU+UPMEM, CPU+memristor-crossbar) because we don't have a heterogeneous setup for evaluation. However, it's worth noting that heterogeneous systems utilizing conventional technologies already exist, such as Prodigy processors (combining CPU, GPU, and TPU) and Nvidia's A100 tensor core integrating multiple SMPs with GPUs for various data types and precisions. These systems are anticipated to become more prevalent with the emergence of CIM/CNM architectures. For instance, nothing prevents integrating a GPU into a UPMEM machine (CPU, GPU, UPMEM). Furthermore, advancements in System-on-Chips (SoCs) and Chiplet structures make it increasingly plausible to combine CPUs, GPUs, DPUs, TPUs, etc., once the necessary interfaces and interconnect support are in place.

The modular structure of CINM enables it to support such heterogeneous setups by allowing adding new device dialects, as detailed in Section 3.3, and making target selection and offloading decisions at finer-granularities, as explained in Section 3.2.2.

4 Evaluation

This section presents our experimental setup, describes our evaluated benchmarks and gives a detailed evaluation and analysis of our generated codes and optimizations.

4.1 Experimental setup

All experiments are run on an Intel Xeon CPU E5-2630 v2 @ 2.60GHz CPU having a maximum clock frequency of 3.1 GHz, 2 CPU sockets, 6 cores/socket, private L1 and L2 data caches sizing 384 kB and 3 MB per core, respectively, and a shared L3 cache of 30 MB (2 instances). The machine has 128 GB of main memory (DRAM) with Linux Ubuntu (22.04).

For the UPMEM backend, we use a high-end real UPMEM machine with 16 DIMMs. Each UPMEM DDR4-2400 DIMM consists of 16 PIM-enabled chips, integrating 128 DPUs. Each DPU runs at 350 MHz, comprises a 64 MB of main RAM

(MRAM) and a 64 kB of working RAM (WRAM). All data transfers to and between the DPUs are directed by and routed through the host. For CIM results, we use the same setup as in OCC [70]. All results reported in this paper are the geometric mean of ten execution runs. The simulation environment is based on the full-system gem5 simulator [52] that supports memristors' based CIM accelerators. For comparison, we use the same baseline as in OCC, i.e., an in-order host ARMv8-A core with 32 kB and 64 kB instruction and data caches, respectively, and a unified 2 MB L2 cache (cf. [70] for details). The ARM core orchestrates data/instructions to the CIM accelerator and executes all non-matmul-like operations. The execution pipeline comprises an array of CIM tiles, where each tile contains a memristor crossbar for executing analog vector-matrix multiplication. We simulated a PCM-based four-tile accelerator, each sizing 64×64 . The read/write latency and energy values for the PCM devices and associated periphery are extracted from [68] and [48]. For high-precision values, we use bit slicing by distributing computations across multiple columns, each representing a single bit. The bits are then weighted at the column output using a shift and add block.

4.1.1 Benchmarks. We evaluate CINM on two sets of benchmarks. For memristor-based CIM accelerators, we are not aware of any hand-optimized benchmark suite that we could use as a baseline. Therefore, we use OCC [70] as our baseline and compare the CINM generated codes with the OCC results. For CNM systems, we use the PriM benchmarks suite [24], the hand-optimized publicly available benchmarks for CNM systems, to evaluate CINM on UPMEM. PriM is a collection of memory-bound workloads from various application domains, including linear algebra, database, data analytics, graph processing, bioinformatics, and image processing. For the non-idiomatic PriM benchmarks, no existing front-end supports translating them into an MLIR representation. Therefore, we resort to manual translation to handle them. To reduce our engineering efforts, we opt not to translate and evaluate all benchmarks. Instead, we've chosen one benchmark from each of the 7 out of 9 categories to showcase the capabilities of our framework and the efficiency of our optimizations across a diverse range of domains. For the remaining two categories, they contain a single benchmark each in bioinformatics and sparse matrix multiplication. The translation effort for these was substantial, and we didn't observe any notable distinctions between them and other workloads. For all non-PriM benchmarks, we start from PyTorch and use its front-end (torch-mlir) to enter MLIR and, subsequently, CINM. All workloads in all configurations use INT32 data type.

mm, *2mm*, *3mm*: Generalized matrix-matrix multiplication *mm*, two consecutive matmuls *2mm*, and two matmuls and multiplication of their results *3mm*.

Convolution (conv) is a compute-bound kernel dominating the execution time of most of the ML models.




Figure 10. Performance comparison of different CIM configurations. All results are normalized to the ARM cpu.

Contraction is a generalization of matmul to N-dimensional tensors and is used in different shapes in different application domains. We use the examples from OCC [70], which were given by the indices involved in equivalent Einstein summation notation. These are a larger contraction *contr1* $C_{abcd} = A_{aebf}B_{dfce}$, performing two reductions, and two small contractions, *contrs1* $C_{ab} = A_{acd}B_{dbc}$ and *contrs2* $C_{abc} = A_{acd}B_{db}$, performing one reduction each.

MLP is a fully connected feed-forward neural network with three layers, each consisting of a matmul followed by addition.

In addition to these benchmarks, we use benchmarks from the PRIM benchmark suite [24]. These include vector addition (*va*), matrix-vector multiplication (*mv*), image histogram long (*hst-1*), breadth-first search (*bfs*), select kernel from databases (*se1*) and time series analysis (*ts*).

4.1.2 Evaluated configurations. For evaluation in this section, we compare the following configurations.

- *cpu-opt*: The host CPU with specification described in Sec. 4.1. All benchmarks are compiled with the Intel oneAPI DPC++/C++ Compiler 2023.1.0 on a Ubuntu host with loop-unrolled, loop-tiling, vectorization, and parallelization enabled.
- *cinm-nd*: Kernels executed in parallel on the UPMEM DPUs (having n DIMMs). Each DPU uses 16 tasklets (threads). The code is generated with CINM flow that tiles compute kernels before offloading.
- *cinm-opt-nd*: Code generated by CINM where the kernels are tiled based on WRAM size assigned to the threads, and the tiled loops are interchanged to improve the WRAM locality in the DPUs.
- *prim-nd*: DPU-code from the PRIM [24].
- *cim*: Code generated by the CINM flow for the memristive CIM target. The *cim* configuration applies the mandatory tiling transformation to fit compute kernels on the CIM crossbars. The *cim-min-writes* configuration interchanges the tiled loops to minimize the number of write operations on the CIM devices. *cim-parallel* unrolls the inner loop dimension to run multiple tiles in parallel while *cim-opt* simultaneously enables all optimizations.

For CIM results, we primarily target workloads that are similar or can be rewritten as *mm*. This is because CIM devices are particularly good for *mm*-like kernels (see Figure 10, cf. [66, 70]) as they can execute them in constant time.

4.2 CIM performance comparison

For PCM and ReRAM based CIM accelerators, the performance of the code generated by CINM matches the performance results presented in OCC. For completeness, we include these performance metrics obtained through the CINM workflow. On average (geomean), *cim* outperforms the arm CPU baseline by an order of magnitude (see Figure 10). Note that kernels that are not *mm* and can not be rewritten as *mm* are not shown in the figure. The *cim-min-writes* configuration reduces the number of writes by 7×, leading to an average performance improvement of 12.4×. The *cim-opt* delivers the best performance, i.e., 30× performance gain, by combining the loop interchange and loop unrolling transformations.

In terms of energy consumption, the *cim-opt* reduces the energy consumption by 5× (geomean), compared to the host cpu. However, for some benchmarks such as *mv*, *conv*, the slower operands movement and little reuse on the CIM array increases the energy consumption by 30% and 40%, respectively compared to the baseline cpu.

4.2.1 CINM evaluation on UPMEM. UPMEM systems have consistently outperformed other platforms such as CPUs, GPUs, and FPGAs, as demonstrated in numerous research studies [3]. In this section, we extend our comparison to include a state-of-the-art CPU with all optimizations enabled. However, our primary focus here is twofold: (1) to showcase the impact of our optimizations implemented in CINM, and (2) to conduct a performance evaluation of CINM-generated code in comparison to the best-available hand-optimized codes (PRIM suite) on a state-of-the-art UPMEM machine (detailed specifications in Sec. 4.1). Given that UPMEM is a general-purpose system capable of accelerating both ML workloads and PRIM workloads, we utilize both sets of benchmarks for our UPMEM evaluation. We employ ML workloads to showcase the effectiveness of our optimizations and the PRIM suite for comparative analysis.




Figure 11. Impact of CINM optimizations on performance.

4.2.2 Effectiveness of our device-aware optimizations.

Figure 11 shows the execution time (in ms) of the CINM generated code for the cinm and cinm-opt configurations (see Sec. 4.1.2). In all benchmarks, the device-aware CINM optimizations show considerable performance gains. On average (geometric mean) across all benchmarks, the cinm-opt-4d, cinm-opt-8d and cinm-opt-16d configurations are 47%, 42% and 40% faster than their respective baseline cinm-nd configurations. The speedup for 3mm benchmark compared to 2mm is relatively small due to the data dependencies of the third GEMM operation on the first two GEMM operations in 3mm. The host puts the synchronization barrier after the first two multiplications in order to get both operands for the third multiplication before offloading it to the DPUs.

4.3 CINM comparison to PriM

The PriM paper [24] extensively compares CPU and DPU systems using microbenchmarks from various domains, assessing metrics such as WRAM and MRAM bandwidth, along with DPUs’ arithmetic throughput on various operations. In this section, we evaluate our generated code against their hand-optimized versions.

In comparison to the baseline optimized CPU configuration, both cinm-nd and prim-nd demonstrate significant performance improvements, as illustrated in Figure 12. On average, the DPU configurations prime-4d, prime-8d, and prime-16d take 1.9 \times , 3.1 \times , and 5.1 \times , less execution time compared to the cpu-opt configuration respectively. This is primarily due to the higher number of compute units within the UPMEM systems compared to the CPU.

Comparing the prime-nd configuration to the CINM’s generated code (cinm-nd), the latter consistently outperforms the former except in the mv kernel where the performances are comparable and the ts kernel where prime exhibits a marginal advantage. On average, cinm-4d, cinm-8d, and cinm-16d configurations take 1.6 \times , 1.9 \times , and 2 \times , less execution time compared to the prime-4d, prime-8d, prime-16d configurations respectively.

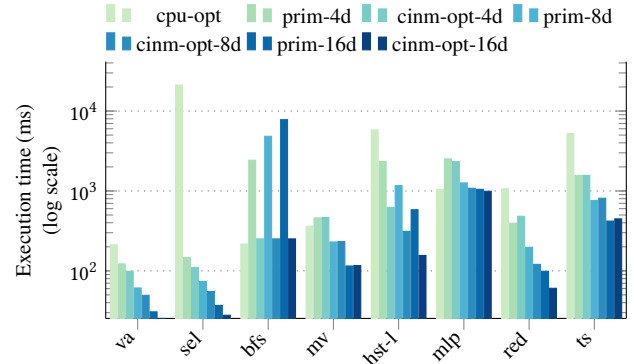


Figure 12. CPU vs cinm-opt-nd and prime-nd.

In most cases, the performance gains of cinm-nd over prime-nd are significant but not overwhelming. For instance, in the va kernel which has no data dependencies, the prime configuration takes 122.214 ms, 61.107 ms and 30.7 ms (absolute), respectively to compute the kernel, resulting in a speedup of up to $> 7\times$ (16d) compared to the cpu-opt configuration. The cinm-nd configurations is on average 1.23 \times better compared to the prime-nd configuration.

However, in workloads such as hst-1, the CINM generated codes demonstrate significantly reduced execution times compared to the prime configuration. cinm-nd takes 0.623 s, 0.311 s and 0.155 s for 4d, 8d and 16d configuration which are on average 3.7 \times less compared to their respective prime-nd configurations. The substantial gain in such benchmarks come from the better exploitation of WRAM by CINM. Overall, our analysis of the generated code and its comparison to the PriM implementations suggests that CINM’s improvements are derived from its efficient management of partial results (also dependent on the tiling size and shape, see Sec. 3.2.6) and their accumulation.

Table 4. Comparing lines of code in CINM vs baseline.

Application	CINM (MLIR)	UPMEM (C/C++)	Reduction (times)
mm / 2mm / 3mm	7 / 19 / 27	180 / 184 / 218	26 / 10 / 8
bfs / ts	29 / 25	315 / 172	11 / 7
contrs2 / contrs1 / contrl	14 / 14 / 16	200 / 197 / 197	14 / 14 / 14
conv / hst-1	5 / 6	203 / 134	40 / 22
mv / mlp	7 / 58	179 / 109	26 / 4
va / sel / red	7 / 12 / 13	101 / 145 / 119	14 / 12 / 9

Table 4 compares the applications’ representation in terms of lines of code (LoC). The *UPMEM (C/C++)* column shows the number of lines of code required to implement an application for both the Host and the DPUs in C/C++. The *CINM (MLIR)* column shows the same application represented using the cinm abstraction. Although comparing LoCs across different programming models may be misleading, it highlights the programmability and productivity aspects of CINM. On average, idiomatic CINM is 15 \times more concise compared to the low-level device codes.

5 Related work

Numerous software frameworks and compilation flows exist that aim to improve the programmability of CINM systems. However, most often, they target a specific architecture (memristive crossbars) or a specific domain (often machine learning). For instance, XLA-NDP [56], built on XLA [65] with a hierarchical flow, is a compiler and runtime solution designed specifically for the NDPX [26] architecture and ML training. PIMFlow [69] targets a GPU-based CNM system (the GPU memory is PIM-enabled) for convolutional neural networks. It identifies the candidate CNN ONNX graph, transforms it to create inter-node parallelism across GPU and PIM, and explores different scenarios to optimize for execution time. CHOPPER [60] focuses on bit-serial SIMD processing using DRAM, leveraging bit-slicing techniques. Considering the CINM unit as a vector processor or SIMD element within the logic layer of the 3D stacked memories, PRIMO [5] seeks patterns of vector operations in the assembly code with suitable operand sizes and substitutes them with the device APIs.

Compared to CNM systems, CIM systems have received relatively more attention. Ambrosi et al. [7] present a software stack for a memristor-based accelerator, which converts ONNX models, enabling optimization at higher levels of abstraction. Han et al. [27]’s compiler operates on the Polyhedral IR to identify operators suitable for offloading onto memristor accelerators. The same objective, i.e., automatic offloading to CIM accelerators, is achieved in many other frameworks using different IRs for different CIM targets [18, 31, 62, 70, 75, 81]. ComPRIME [23] translates boolean and inverter-graph-based functions into CIM-Logic instructions to offload them onto ReRAM crossbars. Sherlock [21] analyses the data flow graph and efficiently maps operands to RRAM-based CIM-logic array to jointly optimize for performance and reliability.

Table 5 compares CINM with existing compilers using various metrics. The CIM-* and CNM metrics indicate the supported device sets by the compiler. The Cost Model metric determines whether the compiler incorporates or allows for adding a cost model for target selection or optimization. The device-agnostic representation metric indicates whether the compiler requires device-specific details (such as API-calls or annotations) in the input application. The domain/device-specific optimization metric indicates if the compiler can apply optimizations specific to the domain or device. The reusable metric indicates whether the compiler can accommodate a range of devices, while the hierarchical metric denotes if it employs any form of hierarchy in its workflow.

Besides compilers, software frameworks exist that aim to enhance these systems’ programmability. SimplePIM [13] offers a high-level C++ library tailored for the UPMEM system. For CAM-based accelerators, frameworks like DT2CAM [63] and X-TIME [59] have been developed to map decision trees onto TCAMs and ACAMs, respectively. The closest framework to CINM having a similar goal is the LLVM-IR based

Table 5. Comparison to existing frameworks

Metric	[56]	[31]	[5]	[27]	[21, 81]	[75]	[7]	[18]	[69]	[78]	[60]	[62, cinm 70]
CIM-Logic	x	✓	✓	x	✓	✓	x	x	x	✓	✓	✓
CIM-Crossbar	x	✓	x	✓	x	✓	✓	✓	✓	x	x	✓
CIM-CAM	x	x	x	x	x	✓	x	x	x	x	x	✓
CNM	✓	x	x	x	x	x	x	x	✓	✓	x	✓
Cost model	✓	✓	x	x	x	x	x	x	✓	✓	x	✓
Device-agnostic input	✓	✓	✓	✓	x	✓	✓	✓	✓	✓	✓	✓
Domain-specific optimization	✓	x	x	✓	x	x	x	✓	x	✓	✓	✓
Device-specific optimization	✓	x	✓	✓	✓	x	✓	✓	x	✓	✓	✓
Reusable	x	✓	✓	✓	x	✓	✓	✓	✓	x	✓	✓
Hierarchical	✓	x	x	x	x	x	✓	✓	✓	✓	x	✓

Infinity Stream [78] framework that extracts the data flow graph from the C program and applies optimization and transformation passes on it for hybrid in/near memory execution. However, unlike Infinity Stream, CINM has multiple levels of IRs, allowing implementation of domain and device-specific optimizations and making analysis much easier compared to the LLVM IR. In terms of modularity and reusability, adding a new device in CINM does not require any changes to the higher level of abstractions in the CINM flow.

6 Conclusions

We presented CINM, a compilation infrastructure for heterogeneous compute-in-memory and compute-near-memory devices. CINM uses MLIR rewriting and introduces reusable abstractions and components that can be leveraged to generalize it to other hardware targets. We investigated the landscape of CIM and CNM systems and presented a partial taxonomy of architectures along with their supported operators. We introduced the `cinm` abstraction that generalizes over all CIM and CNM devices and provides mechanisms to select a hardware target for the input kernel. The `cnm` and `cim` dialects implement custom functions and types common to their respective CNM and CIM devices. As concrete use cases, we presented optimizations and code generation for memristor-based accelerators (CIM) and the UPMEM system (CNM). We demonstrated that by using existing and our novel reusable abstractions, CINM generates code that is comparable (and sometimes) better than the best-available open-source hand-optimized implementation. Given the emergence of CIM and CNM systems, we see CINM as a timely framework that empowers users to fully leverage these systems.

Acknowledgments

This work was partially funded by the Center for Advancing Electronics Dresden (cfaed) and the German Research Council (DFG) through the HetCIM project (502388442) under the Special Priority Program (SPP) 2377, the CO4RTM project (450944241), and the AI competence center ScaDS.AI Dresden/Leipzig in Germany (01IS18026A-D).

References

- [1] Mlir interface. <https://mlir.llvm.org/docs/Interfaces/>, 2023. Accessed: 2023-08-25.
- [2] Mojo, the programming language. <https://www.modular.com/max/mojo>, 2023. Accessed: 2024-04-24.
- [3] Upmem use cases. <https://www.upmem.com/ressources/>, 2023. Accessed: 2023-09-01.
- [4] A. Agrawal et al. X-sram: Enabling in-memory boolean computations in cmos static random access memories. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 65(12):4219–4232, 2018.
- [5] Hameeza Ahmed, Paulo C. Santos, João P. C. Lima, Rafael F. Moura, Marco A. Z. Alves, Antônio C. S. Beck, and Luigi Carro. A compiler for automatic selection of suitable processing-in-memory instructions. pages 564–569, 2019.
- [6] Kerem Akarvardar and H-S Philip Wong. Technology prospects for data-intensive computing. *Proc. of the IEEE*, 111(1):92–112, 2023.
- [7] Joao Ambrosi, Aayush Ankit, Rodrigo Antunes, Sai Rahul Chalamalasetti, Soumitra Chatterjee, Izzat El Hajj, Guilherme Fachini, Paolo Faraboschi, Martin Foltin, Sitao Huang, Wen-Mei Hwu, Gustavo Knuppe, Sunil Vishwanathpur Lakshminarasimha, Dejan Milojicic, Mohan Parthasarathy, Filipe Ribeiro, Lucas Rosa, Kaushik Roy, Plinio Silveira, and John Paul Strachan. Hardware-software co-design for an analog-digital accelerator for machine learning. pages 1–13, 2018.
- [8] Michael Anderson, Benny Chen, Stephen Chen, Summer Deng, Jordan Fix, Michael Gschwind, Aravind Kalaiah, Changkyu Kim, Jaewon Lee, Jason Liang, et al. First-generation inference accelerator deployment at facebook. *arXiv preprint arXiv:2107.04140*, 2021.
- [9] Riyadh Baghdadi, Massinissa Merouani, Mohamed-Hicham Leghetas, Kamel Abdous, Taha Arbaoui, Karima Benatchba, and Saman P. Amarasinghe. A deep learning based cost model for automatic code optimization. *CoRR*, abs/2104.04955, 2021.
- [10] Robin Bläsing, Asif Ali Khan, Panagiotis Ch Filippou, Chirag Garg, Fazal Hameed, Jeronimo Castrillon, and Stuart SP Parkin. Magnetic racetrack memory: From physics to the cusp of applications within a decade. *Proceedings of the IEEE*, 108(8):1303–1321, 2020.
- [11] Amiralí 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. In *Proceedings of the Twenty-Third International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS '18, page 316–331, New York, NY, USA, 2018. Association for Computing Machinery.
- [12] Lorenzo Chelini, Andi Drebes, Oleksandr Zinenko, Albert Cohen, Nicolas Vasilache, Tobias Grosser, and Henk Corporaal. Progressive raising in multi-level ir. In *2021 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*, pages 15–26, 2021.
- [13] J. Chen, J. Gomez-Luna, I. El Hajj, Y. Guo, and O. Mutlu. Simplepim: A software framework for productive and efficient processing-in-memory. In *2023 32nd International Conference on Parallel Architectures and Compilation Techniques (PACT)*, pages 99–111, Los Alamitos, CA, USA, oct 2023. IEEE Computer Society.
- [14] Yen-Cheng Chiu, Chia-Sheng Yang, Shih-Hsin Teng, Hsiao-Yu Huang, Fu-Chun Chang, Yuan Wu, Yu-An Chien, Fang-Ling Hsieh, Chung-Yuan Li, Guan-Yi Lin, et al. A 22nm 4mb stt-mram data-encrypted near-memory computation macro with a 192gb/s read-and-decryption bandwidth and 25.1-55.1 tops/w 8b mac for ai operations. In *2022 IEEE International Solid-State Circuits Conference (ISSCC)*, volume 65, pages 178–180. IEEE, 2022.
- [15] L. Chua. Memristor-the missing circuit element. *IEEE Transactions on Circuit Theory*, 18(5):507–519, 1971.
- [16] Sourav De, Franz Mueller, Nellie Laleni, Maximilian Lederer, Yannick Raffel, Shaown Mojumder, Alptekin Vardar, Sukhrob Abdulazhanov, Tarek Ali, Stefan Dünkel, et al. Demonstration of multiply-accumulate operation with 28 nm fefet crossbar array. *IEEE Electron Device Letters*, 43(12):2081–2084, 2022.
- [17] Jialin Dou and Marcelo Cintra. A compiler cost model for speculative parallelization. *ACM Trans. Archit. Code Optim.*, 4(2):12–es, jun 2007.
- [18] Andi Drebes, Lorenzo Chelini, Oleksandr Zinenko, Albert Cohen, Henk Corporaal, Tobias Grosser, Kanishkan Vadivel, and Nicolas Vasilache. Tc-cim: Empowering tensor comprehensions for computing-in-memory. In *10th Int. Workshop on Polyhedral Compilation Techniques*, 2020.
- [19] Abadi Martín et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [20] Hamid Farzaneh, João Paulo Cardoso de Lima, Mengyuan Li, Asif Ali Khan, Xiaobo Sharon Hu, and Jeronimo Castrillon. C4cam: A compiler for cam-based in-memory accelerators. In *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3*, ASPLOS 2024, New York, NY, USA, 2024. Association for Computing Machinery.
- [21] Hamid Farzaneh, João Paulo Cardoso De Lima, Ali Nezhadi Khelejani, Asif Ali Khan, Mahta Mayahinia, Mehdi Tahoori, and Jeronimo Castrillon. SHERLOCK: Scheduling efficient and reliable bulk bit-wise operations in NVMs. In *Proceedings of the 61th Annual Design Automation Conference (DAC'24)*, DAC '24. ACM, June 2024.
- [22] Jeremy Fowers, Kalin Ovtcharov, Michael Papamichael, Todd Massengill, Ming Liu, Daniel Lo, Shlomi Alkalay, Michael Haselman, Logan Adams, Mahdi Ghandi, et al. A configurable cloud-scale dnn processor for real-time ai. In *2018 ACM/IEEE 45th Annual International Symposium on Computer Architecture (ISCA)*, pages 1–14. IEEE, 2018.
- [23] Steffen Frerix, Saeideh Shirinzadeh, Saman Fröhlich, and Rolf Drechsler. Comprime: A compiler for parallel and scalable rram-based in-memory computing. pages 1–6, 2019.
- [24] Juan Gómez-Luna, Izzat El Hajj, Ivan Fernandez, Christina Giannoula, Geraldo F. Oliveira, and Onur Mutlu. Benchmarking a new paradigm: Experimental analysis and characterization of a real processing-in-memory system. *IEEE Access*, 10:52565–52608, 2022.
- [25] Tobias Gysi, Christoph Müller, Oleksandr Zinenko, Stephan Herhut, Eddie Davis, Tobias Wicky, Oliver Fuhrer, Torsten Hoefler, and Tobias Grosser. Domain-specific multi-level ir rewriting for gpu: The open earth compiler for gpu-accelerated climate simulation. *Trans. Archit. Code Optim.*, 18(4), 2021.
- [26] Hyungkyu Ham, Hyunuk Cho, Minjae Kim, Jueon Park, Jeongmin Hong, Hoyjin Sung, Eunhyeok Park, Euicheol Lim, and Gwangsun Kim. Near-data processing in memory expander for dnn acceleration on gpus. *IEEE Computer Architecture Letters*, 20(2):171–174, 2021.
- [27] Jianhui Han, Xiang Fei, Zhaolin Li, and Youhui Zhang. Polyhedral-based compilation framework for in-memory neural network accelerators. *J. Emerg. Technol. Comput. Syst.*, 18(1), sep 2021.
- [28] Kevin Hsieh, Eiman Ebrahimi, Gwangsun Kim, Niladrish Chatterjee, Mike O'Connor, Nandita Vijaykumar, Onur Mutlu, and Stephen W Keckler. Transparent offloading and mapping (tom) enabling programmer-transparent near-data processing in gpu systems. *ACM SIGARCH Computer Architecture News*, 44(3):204–216, 2016.
- [29] Daniele Ielmini and H-S Philip Wong. In-memory computing with resistive switching devices. *Nature electronics*, 1(6):333–343, 2018.
- [30] Shubham Jain, Ashish Ranjan, Kaushik Roy, and Anand Raghunathan. Computing in memory with spin-transfer torque magnetic ram. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 26(3):470–483, 2017.
- [31] Hai Jin, Bo Lei, Haikun Liu, Xiaofei Liao, Zhuohui Duan, Chencheng Ye, and Yu Zhang. A compilation tool for computation offloading in rram-based cim architectures. 20(4), oct 2023.
- [32] Tian Jin et al. Compiling onnx neural network models using mlir. *arXiv preprint arXiv:2008.08272*, 2020.
- [33] V. Joshi et al. Accurate deep neural network inference using computational phase-change memory. *Nature communications*, 11(1):1–13, 2020.

- [34] Seungchul Jung, Hyungwoo Lee, Sungmeen Myung, Hyunsoo Kim, Seung Keun Yoon, Soon-Wan Kwon, Yongmin Ju, Minje Kim, Wooseok Yi, Shinhee Han, et al. A crossbar array of magnetoresistive memory devices for in-memory computing. *Nature*, 601(7892):211–216, 2022.
- [35] Mahmut Taylan Kandemir, Jihyun Ryoo, Xulong Tang, and Mustafa Karakoy. Compiler support for near data computing. In *Proc. of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 90–104, 2021.
- [36] Navdeep Katel, Vivek Khandelwal, and Uday Bondhugula. High performance gpu code generation for matrix-matrix multiplication using mlir: Some early results, 2021.
- [37] Riduan Khaddam-Aljameh, Milos Stanisavljevic, Jordi Fornet Mas, Geethan Karunaratne, Matthias Brändli, Feng Liu, Abhairaj Singh, Silvia M Müller, Urs Egger, Anastasios Petropoulos, et al. Hermes-core—a 1.59-tops/mm² pcm on 14-nm cmos in-memory compute core using 300-ps/lsb linearized cco-based adcs. *IEEE Journal of Solid-State Circuits*, 57(4):1027–1038, 2022.
- [38] Asif Ali Khan, João Paulo C De Lima, Hamid Farzaneh, and Jeronimo Castrillon. The landscape of compute-near-memory and compute-in-memory: A research and commercial overview. *arXiv preprint arXiv:2401.14428*, 2024.
- [39] Asif Ali Khan, Sébastien Ollivier, Stephen Longofono, Gerald Hempel, Jeronimo Castrillon, and Alex K Jones. Brain-inspired cognition in next generation racetrack memories. *ACM Transactions on Embedded Computing Systems (TECS)*, 2022.
- [40] Win-San Khwa, Yen-Cheng Chiu, Chuan-Jia Jhang, Sheng-Po Huang, Chun-Ying Lee, Tai-Hao Wen, Fu-Chun Chang, Shao-Ming Yu, Tung-Yin Lee, and Meng-Fan Chang. A 40-nm, 2m-cell, 8b-precision, hybrid slc-mlc pcm computing-in-memory macro with 20.5-65.0 tops/w for tiny-al edge devices. In *2022 IEEE International Solid-State Circuits Conference (ISSCC)*, volume 65, pages 1–3. IEEE, 2022.
- [41] D. Kim, J. Kung, S. Chai, S. Yalamanchili, and S. Mukhopadhyay. Neurocube: A programmable digital neuromorphic architecture with high-density 3d memory. *ACM SIGARCH Computer Architecture News*, 44(3):380–392, 2016.
- [42] Jung-Sik Kim, Chi Sung Oh, Hocheol Lee, Donghyuk Lee, Hyong Ryol Hwang, Sooman Hwang, Byongwook Na, Joungwook Moon, Jin-Guk Kim, Hanna Park, et al. A 1.2v 12.8gb/s 2gb mobile wide-i/o dram with 4×128 i/os using tsv based stacking. *IEEE Journal of Solid-State Circuits*, 47(1):107–116, 2011.
- [43] Peter M Kogge. Execube—a new architecture for scaleable mpps. In *1994 International Conference on Parallel Processing*, volume 1, pages 77–84. IEEE, 1994.
- [44] S. Kvatinsky et al. Memristor-based material implication (imply) logic: Design principles and methodologies. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 22(10):2054–2066, 2013.
- [45] Young-Cheon Kwon, Suk Han Lee, Jaehoon Lee, Sang-Hyuk Kwon, Je Min Ryu, Jong-Pil Son, O Seongil, Hak-Soo Yu, Haesuk Lee, Soo Young Kim, et al. 25.4 a 20nm 6gb function-in-memory dram, based on hbm2 with a 1.2tflops programmable computing unit using bank-level parallelism, for machine learning applications. In *2021 IEEE International Solid-State Circuits Conference (ISSCC)*, volume 64, pages 350–352, 2021.
- [46] Chris Lattner, Mehdi Amini, Uday Bondhugula, Albert Cohen, Andy Davis, Jacques Pienaar, River Riddle, Tatiana Shpeisman, Nicolas Vasilech, and Oleksandr Zinenko. Mlir: Scaling compiler infrastructure for domain specific computation. In *IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*, pages 2–14, 2021.
- [47] Manuel Le Gallo, Riduan Khaddam-Aljameh, Milos Stanisavljevic, Athanasios Vasilopoulos, Benedikt Kersting, Martino Dazzi, Geethan Karunaratne, Matthias Brändli, Abhairaj Singh, Silvia M Mueller, et al. A 64-core mixed-signal in-memory compute chip based on phase-change memory for deep neural network inference. *Nature Electronics*, pages 1–14, 2023.
- [48] Manuel Le Gallo, Abu Sebastian, Giovanni Cherubini, Heiner Giefers, and Evangelos Eleftheriou. Compressed sensing with approximate message passing using in-memory computing. *IEEE Transactions on Electron Devices*, 65(10):4304–4312, 2018.
- [49] Dong Uk Lee, Kyung Whan Kim, Kwan Weon Kim, Hongjung Kim, Ju Young Kim, Young Jun Park, Jae Hwan Kim, Dae Suk Kim, Heat Bit Park, Jin Wook Shin, et al. A 1.2 v 8gb 8-channel 128gb/s high-bandwidth memory (hbm) stacked dram with effective microbump i/o test methods using 29nm process and tsv. In *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, pages 432–433. IEEE, 2014.
- [50] Seongju Lee, Kyuyoung Kim, Sanghoon Oh, Joonhong Park, Gimoon Hong, Dongyoon Ka, Kyudong Hwang, Jeongje Park, Kyeongpil Kang, Jungyeon Kim, Junyeol Jeon, Nahsung Kim, Yongkee Kwon, Kornijcuk Vladimir, Woojae Shin, Jongsoon Won, Minkyu Lee, Hyunha Joo, Haerang Choi, Jaewook Lee, Donguc Ko, Younggun Jun, Keewon Cho, Ilwoong Kim, Choungki Song, Chunseok Jeong, Daehan Kwon, Jieun Jang, Il Park, Junhyun Chun, and Joohwan Cho. A 1.25v 8gb, 16gb/s/pin gddr6-based accelerator-in-memory supporting 1tflops mac operation and various activation functions for deep-learning applications. In *2022 IEEE International Solid-State Circuits Conference (ISSCC)*, volume 65, pages 1–3, 2022.
- [51] Sukhan Lee, Shin-haeng Kang, Jaehoon Lee, Hyeonsu Kim, Eojin Lee, Seungwoo Seo, Hosang Yoon, Seungwon Lee, Kyoungwan Lim, Hyunsung Shin, Jinhyun Kim, O Seongil, Anand Iyer, David Wang, Kyomin Sohn, and Nam Sung Kim. Hardware architecture and software stack for pim based on commercial dram technology : Industrial product. In *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*, pages 43–56, 2021.
- [52] Jason Lowe-Power, Abdul Mutaal Ahmad, Ayaz Akram, Mohammad Alian, Rico Amslinger, Matteo Andreozzi, Adria Armejach, Nils Asmussen, Brad Beckmann, Srikanth Bharadwaj, et al. The gem5 simulator: Version 20.0+. *arXiv preprint arXiv:2007.03152*, 2020.
- [53] A. Mehonic et al. Memristors—from in-memory computing, deep learning acceleration, spiking neural networks, to the future of neuromorphic and bio-inspired computing. *arXiv preprint arXiv:2004.14942*, 2020.
- [54] William S. Moses, Lorenzo Chelini, Ruizhe Zhao, and Oleksandr Zinenko. Polygeist: Raising c to polyhedral mlir. In *2021 30th International Conference on Parallel Architectures and Compilation Techniques (PACT)*, pages 45–59, 2021.
- [55] Sébastien Ollivier, Stephen Longofono, Prayash Dutta, Jingtong Hu, Sanjukta Bhanja, and Alex K Jones. Pirm: Processing in racetrack memories. *arXiv preprint arXiv:2108.01202*, 2021.
- [56] Jueon Park and Hyojin Sung. Xla-ndp: Efficient scheduling and code generation for deep learning model training on near-data processing memory. *IEEE Computer Architecture Letters*, 22(1):61–64, 2023.
- [57] David Patterson, Thomas Anderson, Neal Cardwell, Richard Fromm, Kimberley Keeton, Christoforos Kozyrakis, Randi Thomas, and Katherine Yelick. Intelligent ram (iram): Chips that remember and compute. In *IEEE Int. Solids-State Circuits Conference. Digest of Technical Papers*, pages 224–225. IEEE, 1997.
- [58] J Thomas Pawlowski. Hybrid memory cube (hmc). In *2011 IEEE Hot chips 23 symposium (HCS)*, pages 1–24. IEEE, 2011.
- [59] Giacomo Pedretti et al. X-time: An in-memory engine for accelerating machine learning on tabular data with cams. *arXiv preprint arXiv:2304.01285*, 2023.
- [60] Xiangjun Peng, Yaohua Wang, and Ming-Chang Yang. Chopper: A compiler infrastructure for programmable bit-serial simd processing using memory in dram. In *International Symposium on High-Performance Computer Architecture (HPCA)*, pages 1275–1288. IEEE, 2023.
- [61] Angela Pohl, Biagio Cosenza, and Ben Juurlink. Portable cost modeling for auto-vectorizers. In *2019 IEEE 27th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)*, pages 359–369. IEEE, 2019.

- [62] Songyun Qu, Shixin Zhao, Bing Li, Yintao He, Xuyi Cai, Lei Zhang, and Ying Wang. Cim-mlc: A multi-level compilation stack for computing-in-memory accelerators. In *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, pages 185–200, 2024.
- [63] Mariam Rakka et al. Dt2cam: A decision tree to content addressable memory framework. *IEEE Transactions on Emerging Topics in Computing*, 2023.
- [64] Albert Reuther, Peter Michaleas, Michael Jones, Vijay Gadepally, Siddharth Samsi, and Jeremy Kepner. Survey and benchmarking of machine learning accelerators. In *2019 IEEE high performance extreme computing conference (HPEC)*, pages 1–9. IEEE, 2019.
- [65] Amit Sabne. Xla: Compiling machine learning for peak performance. *Google Res*, 2020.
- [66] Abu Sebastian, Manuel Le Gallo, Riduan Khaddam-Aljameh, and Evangelos Eleftheriou. Memory devices and applications for in-memory computing. *Nature nanotechnology*, 15(7):529–544, 2020.
- [67] V. Seshadri et al. Ambit: In-memory accelerator for bulk bitwise operations using commodity dram technology. In *Proc. of the IEEE/ACM International Symposium on Microarchitecture*, pages 273–287, 2017.
- [68] Ali Shafiee, Anirban Nag, Naveen Muralimanohar, Rajeev Balasubramanian, John Paul Strachan, Miao Hu, R Stanley Williams, and Vivek Srikumar. Isaac: A convolutional neural network accelerator with in-situ analog arithmetic in crossbars. *ACM SIGARCH Computer Architecture News*, 44(3):14–26, 2016.
- [69] Yongwon Shin, Juseong Park, Sungjun Cho, and Hyojin Sung. Pimflow: Compiler and runtime support for cnn models on processing-in-memory dram. page 249–262, 2023.
- [70] Adam Siemienuk, Lorenzo Chelini, Asif Ali Khan, Jeronimo Castrillon, Andi Drebes, Henk Corporaal, Tobias Grosser, and Martin Kong. Occ: An automated end-to-end machine learning optimizing compiler for computing-in-memory. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2021.
- [71] Gagandeep Singh, Lorenzo Chelini, Stefano Corda, Ahsan Javed Awan, Sander Stuijk, Roel Jordans, Henk Corporaal, and Albert-Jan Boonstra. A review of near-memory computing architectures: Opportunities and challenges. In *2018 21st Euromicro Conference on Digital System Design (DSD)*, pages 608–617. IEEE, 2018.
- [72] Stephanie Soldavini, Karl Friebe, Mattia Tibaldi, Gerald Hempel, Jeronimo Castrillon, and Christian Pilato. Automatic creation of high-bandwidth memory architectures from domain-specific languages: The case of computational fluid dynamics. *ACM Trans. Reconfigurable Technol. Syst.*, 16(2), mar 2023.
- [73] Harold S Stone. A logic-in-memory computer. *IEEE Transactions on Computers*, 100(1):73–78, 1970.
- [74] Upmem. Upmem processing in-memory (pim): Ultra-efficient acceleration for data-intensive applications. In *2022 UPMEM PIM Tech paper v2.7*, pages 1–22, 2022.
- [75] Kanishkan Vadivel, Lorenzo Chelini, Ali BanaGozar, Gagandeep Singh, Stefano Corda, Roel Jordans, and Henk Corporaal. Tdo-cim: transparent detection and offloading for computation in-memory. In *2020 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 1602–1605. IEEE, 2020.
- [76] Xavier Vera, Jaume Abella, Josep Llosa, and Antonio González. An accurate cost model for guiding data locality transformations. *ACM Trans. Program. Lang. Syst.*, 27(5):946–987, sep 2005.
- [77] H. Wang, R. Liu, R. Dorrance, D. Dasalukunte, D. Lake, and B. Carlton. A charge domain sram compute-in-memory macro with c-2c ladder-based 8-bit mac unit in 22-nm finfet process for edge inference. *IEEE Journal of Solid-State Circuits*, 58(4):1037–1050, 2023.
- [78] Zhengrong Wang, Christopher Liu, Aman Arora, Lizy John, and Tony Nowatzki. Infinity stream: Portable and programmer-friendly in-/near-memory fusion. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and*

Operating Systems, Volume 3, ASPLOS 2023, page 359–375, 2023.

- [79] Y. Wang et al. An ultralow-power memory-based big-data computing platform by nonvolatile domain-wall nanowire devices. In *Int. Symp. on Low Power Electronics and Design (ISLPED)*, pages 329–334, 2013.
- [80] Michael E Wolf and Monica S Lam. A loop transformation theory and an algorithm to maximize parallelism. *IEEE Transactions on Parallel & Distributed Systems*, 2(04):452–471, 1991.
- [81] Jintao Yu, Tom Hogervorst, and Razvan Nane. A domain-specific language and compiler for computation-in-memory skeletons. page 71–76, 2017.

A Artifact Appendix

A.1 Abstract

The artifact demonstrates the functionality of the CINM (Cinnamon) compilation flow. The core idea of the paper is to propose separate domain-specific (c_{inm}) and device-specific (c_{im}, c_{nm}) abstractions. The artifact includes the sources for these abstractions and the necessary transformations and conversion passes to progressively lower them. It focuses on conversions to illustrate different IRs, as shown in Fig. 5 and Fig. 6.

A.2 Description

A.2.1 How to access. The artifacts are available at GitHub on the following link.

<https://github.com/tud-ccc/Cinnamon.git>

A.2.2 Hardware dependencies. UPMEM hardware or SDK.

A.2.3 Software dependencies. The CINM framework depends on a patched version of LLVM 18.1.6. Additionally, a number of software packages are required to build it, like CMake. The repository contains a script, *build.sh* that installs all needed dependencies and builds the sources.

A.2.4 Data sets. All benchmarks at the CINM abstraction are already in the Git repository under *cinnamon/benchmarks/*.

A.3 Experiment workflow

The Git repository contains different scripts for building the CINM flow and compiling different benches for the UPMEM target.

- *build.sh* to build the sources;
- *compile-benches.sh* to compile all the benches using Cinnamon.

After running the *compile-benches.sh*, the generated code and the intermediate IRs for each bench can be found under *cinnamon/benchmarks/*.

The user can also try running individual benchmarks by manually trying individual conversions. The benchmark files have a comment at the top giving the command used to lower them to the upmem IR.