

CD-PIM: A High-Bandwidth and Compute-Efficient LPDDR5-Based PIM for Low-Batch LLM Acceleration on Edge-Device

Ye Lin¹, Chao Fang¹, Xiaoyong Song², Qi Wu¹, Anying Jiang¹, Yichuan Bai¹, Li Du^{1,†}

¹School of Electronic Science and Engineering, Nanjing University, China ²China Mobile Research Institute, China
Email: {yelin, fantasysee, qiwu, jay, yicbai}@smail.nju.edu.cn, songxiaoyong@chinamobile.com, ldu@nju.edu.cn

Abstract—Edge deployment of low-batch large language models (LLMs) faces critical memory bandwidth bottlenecks when executing memory-intensive general matrix-vector multiplications (GEMV) operations. While digital processing-in-memory (PIM) architectures promise to accelerate GEMV operations, existing PIM-equipped edge devices still suffer from three key limitations: limited bandwidth improvement, component under-utilization in mixed workloads, and low compute capacity of computing units (CUs). In this paper, we propose CD-PIM to address these challenges through three key innovations. First, we introduce a high-bandwidth compute-efficient mode (HBCEM) that enhances bandwidth by dividing each bank into four pseudo-banks through segmented global bitlines. Second, we propose a low-batch interleaving mode (LBIM) to improve component utilization by overlapping GEMV operations with GEMM operations. Third, we design a compute-efficient CU that performs enhanced GEMV operations in a pipelined manner by serially feeding weight data into the computing core. Forth, we adopt a column-wise mapping for the key-cache matrix and row-wise mapping for the value-cache matrix, which fully utilizes CU resources. Our evaluation shows that compared to a GPU-only baseline and state-of-the-art PIM designs, our CD-PIM achieves $11.42\times$ and $4.25\times$ speedup on average within a single batch in HBCEM mode, respectively. Moreover, for low-batch sizes, the CD-PIM achieves an average speedup of $1.12\times$ in LBIM compared to HBCEM.

I. INTRODUCTION

Transformer-based large language models (LLMs) [1] have demonstrated exceptional capabilities across a wide range of applications, such as question answering [2] and code completion [3]. Compared to cloud services [4], which handle large-batch requests requiring high compute capacity, on-device LLM inference [5] typically processes low-batch requests and demands high bandwidth to achieve low-latency interactions in real-time applications [6].

The inference process in LLMs is typically divided into two stages: the prefill stage and the decode stage [7]. In the prefill stage, the LLM processes the entire input token sequence in a single step, primarily relying on compute-intensive general matrix-matrix multiplication (GEMM) operations [8]. In contrast, the decode stage generates one token at a time in an autoregressive manner, which relies on memory-intensive

general matrix-vector multiplication (GEMV) operations [9]. However, the limited memory bandwidth of the edge device leads to increased decode latency [10].

To mitigate memory bandwidth bottlenecks in the decode stage, digital processing-in-memory (PIM) architectures are increasingly adopted as hardware accelerators. Based on the placement of computing unit (CU), PIM designs can be categorized into subarray-level [11], bank-level [12], bankgroup (BG)-level [13], and I/O-level [14] architectures. Among these, bank-level PIM offers a favorable trade-off between area cost and internal memory bandwidth by exploiting bank-level parallelism, where CUs are integrated within the memory bank [15]. This approach is prevalent in recently released production-grade PIM devices [16]. In such architectures, PIM-equipped processors perform the GEMM operations on the host processor while offloading the GEMV operations to PIM, thereby improving the overall performance of LLM inference [17]–[20].

To achieve efficient on-device inference, the current PIM designs mainly face three challenges: (1) The previous work AttAcc [13] leverages the large number of high-bandwidth memory (HBM) banks to achieve an internal bandwidth of 242 TB/s for cloud services. However, edge devices have far fewer memory banks, making existing PIM architectures essential for efficiently accelerating LLM inference on the edge. This resource limitation prevents on-device inference from meeting strict low-latency requirements. In addition, existing PIM architectures employ CU with relatively low compute capacity for GEMV operations, as they operate at an internal memory clock frequency lower than the maximum supported digital clock frequency [15], and only one CU is integrated per bank to perform GEMV operations [21]. (2) Current PIM architectures operate in a blocked mode, where the processor and PIM cannot execute simultaneously, which limits the efficiency of PIM. In particular, for compute-intensive workloads with long input sequences and short output sequences, the prefill latency on the processor becomes the dominant factor in the end-to-end inference latency. A straightforward solution would be to use PIM as a stand-alone accelerator with dedicated memory, enabling simultaneous execution of GEMV operations in the PIM and GEMM operations in the processor [22] [23]. However, this approach is impractical for edge devices due to constrained memory resources. Therefore, PIM must be integrated into the existing memory system to enable concurrent

[†]Corresponding author. This work was funded in part by the National Key Research and Development Program of China under Grant 2022YFB4400900, in part by the Strategic Industries and Key Technologies Project of Jiangsu Province under Grant BE2023020-3, in part by the Basic Research Program of Jiangsu Province under Grant BK20243042, and in part by the Nanjing University-China Mobile Communications Group Co., Ltd. Joint Institute.

GEMV and GEMM execution; (3) In the attention layer, fixed mapping of key (K)-cache matrix and value (V)-cache matrix leads to imbalanced CUs resource utilization due to differing computation flows [13]. In particular, if the K-cache matrix (H_{dim}, L) is partitioned row-wise, where H_{dim} denotes the embedding dimension and L represents the number of tokens accumulated from the prefill stage up to the current decoding step, the appended ($H_{\text{dim}}, 1$) column vector is processed by only one CU, leading to imbalanced PIM CUs utilization. However, for the V-cache matrix, row-wise data mapping can fully utilize CU resources.

To address this challenge, we propose cross-division-PIM (CD-PIM), a high-bandwidth and compute-efficient PIM accelerator optimized for low-batch LLM inference on edge devices. The CD-PIM architecture partitions each memory bank into four pseudo-banks (Pbanks), thereby improving internal bandwidth with minor changes to routing and peripheral circuitry. Furthermore, it seamlessly integrates PIM into the existing memory systems and supports two operation modes: high-bandwidth compute-efficient mode (HBCEM) and low-batch interleaving mode (LBIM). In HBCEM, CD-PIM maximizes internal memory bandwidth by simultaneously activating four Pbanks to accelerate GEMV operations, whereas in LBIM, it overlaps the GEMV latency with GEMM latency to reduce end-to-end inference latency. Furthermore, a compute-efficient CU is designed to perform GEMV operations in a pipelined manner.

The main contributions of our work are as follows:

- For challenge (1), we propose a high-bandwidth compute-efficient LPDDR5-based CD-PIM architecture that fully accelerates GEMV operations in HBCEM for memory-intensive workloads with short input sequences and long output sequences. The CD-PIM divides each bank into four Pbanks by segmenting the global bitline (GBL), achieving a $4 \times$ improvement in memory bandwidth. A lightweight high compute CU is designed to efficiently perform GEMV operations in a pipelined manner by serially feeding weight data from the separated buffers into the computing core, enabling higher compute capacity. (Sec. III-A)
- For challenge (2), we overlap the GEMV latency in PIM with GEMM latency in processor to reduce end-to-end inference latency in LBIM for compute-intensive workload by activating two Pbanks for GEMV operations and another two Pbanks for GEMM operations, enabling the simultaneous execution of PIM and the processor. (Sec. III-B)
- For challenge (3), we adopt column-wise mapping for the K-cache matrix and row-wise mapping for the V-cache matrix, which fully utilizes CU resources through inner-product and outer-product operations. (Sec. III-C)

We evaluate the proposed architecture on both compute-intensive and memory-intensive workloads using an NVIDIA Jetson AGX Orin 64GB [24] and an Apple iPhone 15 Pro [25]. Our evaluation demonstrates that compared to a GPU-only baseline and state-of-the-art (SOTA) PIM designs, the CD-PIM

achieves an average of $11.42 \times$ and $4.25 \times$ speedup in HBCEM. For compute-intensive workloads with low batch sizes, the CD-PIM achieves an average of $1.12 \times$ speedup in LBIM compared to HBCEM.

II. BACKGROUND

A. Transformer-based LLM Inference

Transformer-based LLM consists of multiple cascaded Transformer decoder layers [26]. For each request, inference begins with the prefill stage, in which multiple tokens from the input sequence are processed in parallel to generate the first output token [27]. The decode stage then generates subsequent tokens iteratively, each time using the previously generated token as input. This process continues until the entire output sequence is produced. Owing to this iterative nature, the input to each decoder layer during the decode stage is a single vector, leading to memory-intensive GEMV operations, in contrast to the GEMM operations performed in the prefill stage. Since the prefill stage is executed only once per request, whereas the decode stage must be repeated hundreds of times, the decode latency dominates the end-to-end inference latency.

B. Bank-level PIM-equipped Accelerators

Bank-level PIM-equipped LLM accelerators typically place the CU within the bank to exploit higher internal memory bandwidth than the external memory interface. To overcome the limited bandwidth improvements on device, FOLD-PIM [5] divides the GBL into upper and lower segments and employs a single CU operating at twice the internal memory clock frequency. Similarly, Pipe-PIM [15] divides the GBL into left and right segments and employs two CUs, each operating at the internal memory clock frequency. Although these PIM architectures highly efficiently accelerate the GEMV operation, they operate in blocked mode, thereby limiting PIM utilization.

A digital PIM integrates CU for performing GEMV operations and a vector buffer for storing and reusing the input vector and partial sums at the bank’s periphery. In bank-level digital PIM, the CU is placed within the bank region. However, the vector buffer may reside in the bank region [5], the BG region [13], or the I/O region [14]. Locating the vector buffer farther from the CU increases data transfer energy. To minimize transfer energy, this work places the vector buffer adjacent to the computing unit.

On-device LLM inference services typically exhibit two key characteristics: handling only a small number of requests concurrently and requiring low-latency interaction. For example, AIOS [28] is presented as an LLM-embedded management system that supports six interfaces. However, current PIM architectures operate in blocked mode, preventing concurrent operation of PIM and processor computations [29]. This serialization inherently under-utilizes resources. In particular, for low-batch scenarios, the blocked computation mode leads to longer inference latency, constraining PIM adoption in edge devices. Although prior work [30] enables PIM-NPU co-processing via dual row buffer banks, this approach incurs considerable area overhead.

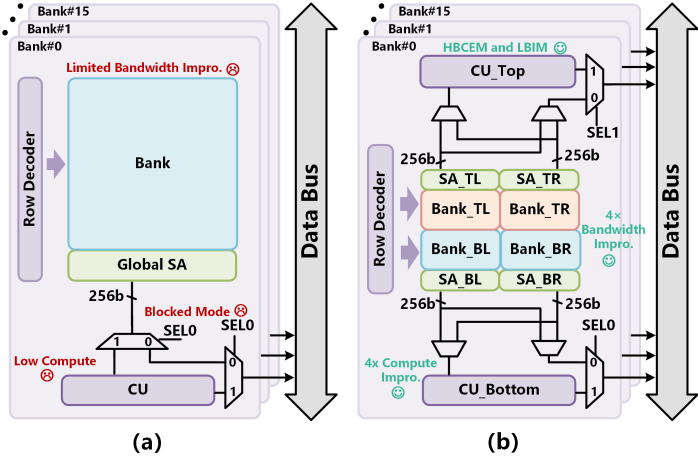


TABLE I
TWO EDGE DEVICE SPECIFICATIONS

Device	DRAM	Throughput (TFLOPS)	Bandwidth (GB/s)
Jetson AGX Orin	LPDDR5	42.5	204.8
iPhone 15 Pro	LPDDR5	4.29	51.2

digital PIM designs [34] execute GEMV operations at the internal clock frequency, thereby limiting CU compute capacity. This reduced compute capacity prevents such PIM architectures from fully accelerating low-batch LLM inference.

III. CD-PIM: HIGH-BANDWIDTH COMPUTE-EFFICIENT PIM ARCHITECTURE

In this section, we introduce CD-PIM, a high-bandwidth compute-efficient PIM architecture that accelerates GEMV operations by exploiting bank-level parallelism at low implementation cost. CD-PIM leverages its high memory bandwidth to accelerate the decode stage in HBCEM. Moreover, to support low-batch, compute-intensive workloads, a CD-PIM-equipped processor overlaps decode latency by simultaneously executing CD-PIM for the decode stage and the processor for the prefill stage in LBIM. Notably, both the input and weight data in this paper are represented with 8-bit precision, which does not lead to any noticeable degradation in LLM inference accuracy. In addition, we introduce three new instructions to support PIM operation.

A. The Proposed CD-PIM

Compared with conventional DRAM-PIM architectures, as shown in Fig. 1(a), CD-PIM integrates additional CU directly into each bank. Fig. 1(b) illustrates the overall architecture of the CD-PIM with two CUs for accelerating GEMV operation. The top and bottom CU perform INT8 multiplication-and-accumulation (MAC) operations in a pipelined manner by receiving weight data from banks through the left and right global BLSAs (SA_TL, SA_TR, SA_BL, SA_BR). Each CU can handle a 32 B data width per compute cycle. To enhance compute capacity, the CU is designed to operate at 400 MHz, which is twice the internal clock frequency of LPDDR5 [5]. As a result, the CU can pipelined process (64, 1) vector data from both the left-SA and right-SA within a single memory cycle. Each CU is equipped with a 64 B input buffer for storing the input vector and a 128 B output buffer for storing partial sums. By employing two pipelined CUs in each bank, CD-PIM provides a 2 \times computational advantage over DH-PIM [34] equipped with the same DRAM die. In addition, owing to the dedicated mapping strategies for the K-cache and V-cache matrix, the CU is capable of supporting both inner-product and outer-product operations.

Fig. 2 illustrates the details of the CD-PIM architecture. First, each bank is divided into a left-bank and a right-bank by equally splitting the global bitline (GBL), the bitline sense amplifier (BLSA), and the global BLSA into left-GBL/right-GBL, BLSA-left/BLSA-right, and global BLSA-

Fig. 1. (a) Conventional PIM architecture with a single CU; (b) CD-PIM architecture with two CUs for performing the GEMV operation.

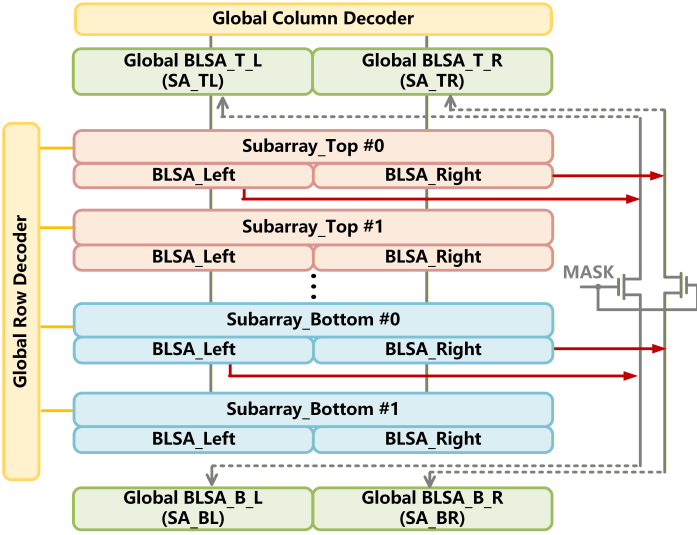


Fig. 2. Details of the CD-PIM architecture, where each bank is divided into four Pbanks by separating the GBL and isolation transistors.

Compared with HBM-based PIM [13], which integrates four dies per rank and over 128 banks per die for cloud devices, LPDDR5-PIM provides only a limited number of banks, restricting bank-level parallelism to enhance internal memory bandwidth. Therefore, expanding the internal memory bandwidth is crucial for efficiently accelerating GEMV operations on edge devices. Moreover, the variable workload distribution [31] and the diversity of compute and memory resources across edge devices [32] largely determine the proportion of time-to-first-token (TTFT) in the overall inference latency. In particular, for compute-intensive workloads with long input sequences and short output sequences, constrained compute resources cause TTFT to dominate the end-to-end inference latency.

In conventional DRAM dies, the bandwidth of the external memory interface exceeds the internal clock frequency due to the burst-transfer scheme [33]. However, prior bank-level

TABLE II
COMMAND SELECTION TABLE

Command Description	SEL0	SEL1
PIM_MAC_FM	1	1
MACT_LDB	0	1
MACB_LDT	1	0

left/global BLSA-right, respectively. Owing to this separation, the left-bank and right-bank can be accessed simultaneously by the column decoder. Furthermore, each left-GBL and right-GBL is further partitioned into an upper-GBL and a lower-GBL through isolation transistors, which divides both the left-bank and the right-bank into an upper-bank and a lower-bank. These upper and lower banks can be accessed simultaneously by the row decoder. As a result, CD-PIM partitions each bank into four Pbanks (Bank_TL, Bank_TR, Bank_BL, Bank_BR), all of which can be activated concurrently to perform GEMV operations, therefore achieving a higher memory bandwidth compared to FOLD-PIM [5].

B. HBCEM and LBIM

Table II summarizes the three PIM instructions by controlling the SEL0 and SEL1 signals. The PIM_MAC_FM instruction activates four Pbanks simultaneously to fully accelerate the GEMV operation. The MACT_LDB instruction activates the top CU to execute the GEMV operation, while the processor performs the GEMM operation by accessing data from the bottom bank. Similarly, the MACB_LDT instruction activates the bottom CU to execute the GEMV operation while accessing data from the top bank. Therefore, by simultaneously performing GEMV in PIM and GEMM in the processor, the GEMV latency is overlapped with the GEMM latency.

For conventional edge devices, the prefill and decode stages of LLM inference are executed on the processor. However, for memory-intensive workloads, the limited memory bandwidth causes significant inference latency, as shown in Fig. 4(a). For $(L_{in}, L_{out}) = (128, 2048)$ with LLAMA-1B on the NVIDIA Jetson AGX Orin 64 GB, the end-to-end inference latency reaches 35.7s, while the processor compute utilization is 85%. To accelerate decode stage, we offload the decode stage to PIM in HBCEM, activating four Pbanks to fully accelerate the GEMV operation using PIM_MAC_FM instruction as illustrated in Fig. 4(b). By exploiting the high internal memory bandwidth of PIM, the decode stage latency is reduced by 90.2% compared with processor execution, leading to a total inference latency of 3.53s.

For compute-intensive workloads, although PIM accelerates GEMV operation, the TTFT dominates the end-to-end inference latency, particularly on low compute capacity edge devices, as shown in Fig. 4(b). For $(L_{in}, L_{out}) = (2048, 128)$ with LLAMA-13B, offloading the decode stage to PIM causes the TTFT portion of the latency to increase from 53.7% on the NVIDIA Jetson AGX Orin 64 GB to 65.2% on the Apple iPhone 15 Pro. To address this issue, CD-PIM executes GEMV and GEMM operations concurrently, activating two Pbanks for the GEMV operations and another two Pbanks for

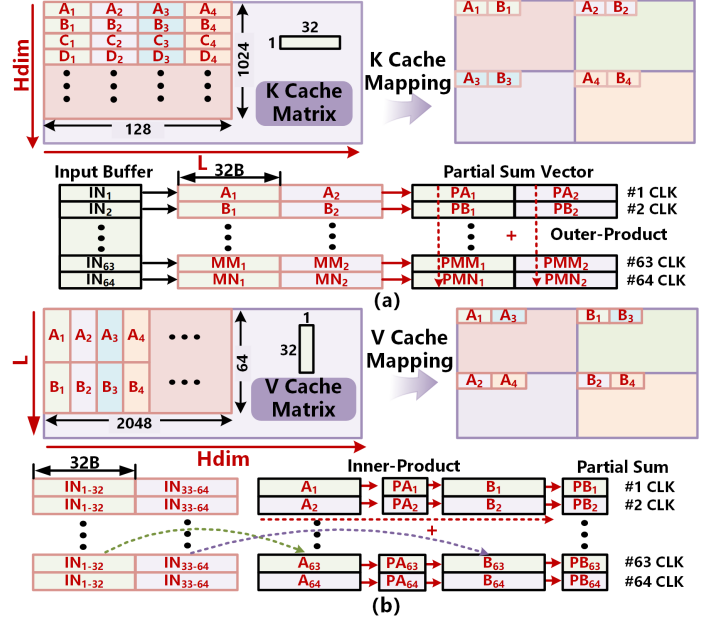


Fig. 3. Detailed data mapping of the CD-PIM: (a) K-cache matrix data mapping and outer-product compute flow; (b) V-cache matrix data mapping and inner-product compute flow.

the GEMM operations in LBIM, using the MACT_LDB and MACB_LDT instructions. Although the PIM compute capacity in LBIM is reduced by half compared to HBCEM, the GEMV latency is effectively overlapped with GEMM latency, as illustrated in Fig. 4(c).

C. Data Mapping

In the decode stage, the computation flow is composed of the QKV generation layer, the attention layer, and the feed-forward network (FFN). First, the input token is fed into the QKV generation layer to derive the query (Q), K, and V vectors of size $(1, H_{dim})$. The newly generated K and V vectors are concatenated with the K-cached matrix (H_{dim}, L) and V-cached matrix (L, H_{dim}) . Subsequently, the Q vector is forwarded to the attention layer, where a GEMV operation is performed with the K-cached matrix. The resulting vector is normalized using the softmax function to produce the attention weight vector $(1, L)$. This weight vector is then multiplied with the V-cached matrix to compute the attention output $(1, H_{dim})$. Finally, the attention output is passed to the FFN layer to generate the next token.

During the attention layer, the data mapping of the KV-cached matrix is adjusted according to the compute flow. **For the K-cache matrix mapping in column-wise, each data chunk corresponds to a vector of size (1×32) , which can be accessed by a Pbank in a single burst operation.** At the bank level, each bank is able to access a matrix organized as a (2×64) data chunk. Before the GEMV operation, the Q vector is loaded into the CU input buffer. The Q vector is partitioned into sub-vectors of size (1×1024) , which are sequentially transferred to the LPDDR5 dies. Each (1×1024) sub-vector is further divided into sixteen (1×64) sub-vectors, corresponding to the 16 banks of the LPDDR5 device. Within

each bank, a (1×64) sub-vector is broadcast to the input buffers of the two CUs.

During the GEMV operation, each CU performs an outer product between the 1B input data stored in the input buffer and a $(1 \times 32B)$ weight vector, producing a (1×32) partial-sum vector. The same input data is then multiplied with another $(1 \times 32B)$ weight vector to obtain a new (1×32) partial-sum vector. Consequently, within a single internal memory clock cycle, the CU multiplies the 1B input data with $(1 \times 64B)$ weight data, resulting in a (1×64) partial-sum vector. In the following cycle, the produced (1×64) partial-sum vector is accumulated with the results generated in the previous step, as illustrated in Fig. 3(a). For example, the input data IN_1 is multiplied with the weight vector A_1 to produce the partial-sum vector PA_1 . The same input data IN_1 is also multiplied with the weight vector A_2 , yielding the partial-sum vector PA_2 . Next, the input data IN_2 is multiplied with the weight vector B_1 and added to PA_1 to generate the updated partial-sum vector PB_1 . Similarly, IN_2 is multiplied with the weight vector B_2 and accumulated with PA_2 , producing the updated partial-sum vector PB_2 . Following this compute flow, each bank can perform a $((1, 1) \times (1, 128))$ GEMV operation within a single memory internal clock cycle. For each (1×1024) input-vector segment, the input buffer capacity supports the processing of a (64×128) matrix per bank. At the die level, a (1024×128) matrix can be processed. Subsequently, the next (1×1024) input-vector segment is transferred to the DRAM die, enabling continuous computation.

For the V-cache matrix mapping in row-wise, three main differences are identified compared to the K-cache matrix, as shown in Fig. 3(b). First, each data chunk corresponds to a (32×1) vector, and each bank can access a matrix organized as a (64×2) data chunk, whereas the K-cache matrix accesses a (2×64) data chunk. Second, the attention vector $(1, L)$ is divided into multiple (1×64) sub-vectors. Each sub-vector is broadcast to all input buffers of the CUs. Third, during the GEMV operation, each CU performs an inner product between a (1×32) attention vector IN_{1-32} and a (32×1) weight vector A_1 to generate a partial sum PA_1 . The subsequent $(1 \times 32B)$ attention data IN_{33-64} is then multiplied with a $(32 \times 1B)$ weight vector B_1 and added to PA_1 to produce a new partial sum PB_1 . Following this computation flow, each bank can execute a $((1, 64) \times (64, 2))$ GEMV operation within a single internal clock cycle. For each (1×64) attention vector, the output buffer capacity supports the processing of a (64×128) matrix within a bank, and at the die level, a (64×2048) matrix can be processed. The next (1×64) attention vector is then transferred to the DRAM die for continuous computation.

IV. EVALUATION

A. Methodology

We implemented CD-PIM on top of 4GB LPDDR5 by modifying Ramulator2 [35], an open-source DRAM simulator. Each LPDDR5 die provides 16 data pins, operating at a rate of 6.4 Gbps per pin. In our experimental setup, we employ the Apple iPhone 15 Pro and the NVIDIA Jetson AGX Orin as

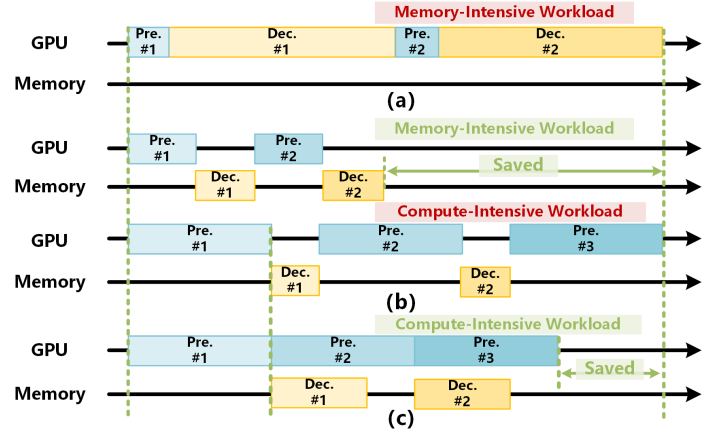


Fig. 4. Timing diagrams of LLM inference: (a) GPU with memory-intensive workload; (b) CD-PIM-equipped GPU in HBCEM for memory- and compute-intensive workload; (c) CD-PIM-equipped GPU in LBIM for compute-intensive workload.

edge devices to evaluate the performance of CD-PIM, which are equipped with 4 and 16 LPDDR5 dies, respectively, providing total memory capacities of 16 GB and 64 GB. For performance evaluation, we compare CD-PIM with GPU-only baseline and AttAcc [13] baseline execution on the Apple iPhone 15 Pro and the NVIDIA Jetson AGX Orin. We further assess CD-PIM under batch sizes of one and four, considering both HBCEM and LBIM across different LLM parameters.

B. Performance Results

Fig. 5 shows the normalized performance evaluation of the GPU and AttAcc compared with CD-PIM in HBCEM and LBIM under a single-batch scenario. On the NVIDIA Jetson AGX Orin, CD-PIM in HBCEM achieves a speedup of $4.48 \times - 10.51 \times$ for LLAMA-1B, $6.71 \times - 13.74 \times$ for LLAMA-7B, and $7.47 \times - 14.6 \times$ for LLAMA-13B relative to the GPU-only baseline [36]. These increases result from exploiting the higher internal memory bandwidth of PIM compared to the external memory bandwidth of the GPU. Furthermore, as the model size increases from LLAMA-1B to LLAMA-13B, the decode stage incurs higher latency. By accelerating the GEMV operation with PIM, the performance improvements become more pronounced.

For the Apple iPhone 15 Pro, CD-PIM in HBCEM achieves a higher speedup compared to the NVIDIA Jetson AGX Orin. Owing to the iPhone's lower memory bandwidth relative to the Jetson, the decode stage incurs longer latency, making the workload more memory-intensive. Consequently, accelerating the GEMV operation with PIM yields greater performance improvements on the iPhone. In particular, for the memory-intensive case with $(Lin, Lout) = (128, 2048)$, the speedup increases from $10.1 \times$ on the Jetson AGX Orin to $18.6 \times$ for LLAMA-1B on the iPhone 15 Pro.

However, due to limited compute capacity, CD-PIM in HBCEM is less efficient for compute-intensive workloads compared to memory-intensive workloads. Fig. 6 presents the performance evaluation of CD-PIM in both HBCEM and LBIM under a four-batch scenario on the NVIDIA Jetson AGX




Fig. 5. Normalized performance of LLaMA-1B, -7B, and -13B (batch size = 1) under various Lin and Lout on the NVIDIA Jetson AGX Orin and Apple iPhone 15 Pro.




Fig. 6. Normalized performance of the NVIDIA Jetson AGX Orin equipped with CD-PIM under HBCEM and LBIM (batch size = 4).




Fig. 7. Normalized performance of the Apple iPhone 15 Pro equipped with CD-PIM under HBCEM and LBIM (batch size = 4).

Orin platform, where Lin is fixed at 2048. For LLaMA-1B, when Lout increases from 2 to 128, CD-PIM in LBIM achieves a speedup of $1.01\times\text{--}1.41\times$ over HBCEM baseline by overlapping the latency of the decode stage. In contrast, for LLaMA-7B and LLaMA-13B, the performance speedup in LBIM initially grows with Lout from 2 to 32, but then decreases because at Lout = 128, the decode stage latency surpasses the TTFT latency.

For the Apple iPhone 15 Pro, CD-PIM in LBIM achieves a speedup of $1.01\times\text{--}1.23\times$ over HBCEM in LLaMA-1B, which is lower than the performance speedup on the NVIDIA Jetson AGX Orin, since for compute-intensive workloads, the reduced compute capacity has a stronger impact than the reduced memory bandwidth. Consequently, the TTFT latency contribution to end-to-end inference latency is higher on the iPhone 15 Pro than on the Jetson AGX Orin. For LLaMA-7B and LLaMA-13B, CD-PIM in LBIM achieves speedups of




Fig. 8. The area and power breakdown of CU.

$1.01\times\text{--}1.46\times$ and $1.01\times\text{--}1.44\times$ over HBCEM, respectively, when Lout ranges from 2 to 128, because the decode stage latency increases to nearly half of the total inference latency, as shown in Fig. 7.

C. Area and Power Overhead

Compared with CUs that only support inner-products, the proposed CU supports both inner- and outer-products. Fig. 8 illustrates the area and power breakdown of CU. We evaluated the CU area and power using Synopsys Design Compiler, which shows that each PU occupies $14,941\ \mu\text{m}^2$ and consumes 4.5 mW in TSMC 28 nm. This corresponds to 0.8% of a 32 Gb LPDDR5 die area and adds 144 mW in total, confirming the cost efficiency of CD-PIM.

V. CONCLUSION

In this paper, we propose CD-PIM to address the inefficiency of GEMV acceleration for memory-intensive workloads and the longer TTFT latency in compute-intensive workloads. The CD-PIM divides each bank into four Pbanks to enhance internal memory bandwidth and is equipped with two CUs to achieve high computational capacity. For memory-intensive workloads, CD-PIM in HBCEM fully accelerates GEMV operations; for compute-intensive workloads, CD-PIM in LBIM overlaps the decode latency by executing CD-PIM for the decode stage and the processor for the prefill stage simultaneously. A high compute capacity CU is designed to execute GEMV in a pipelined manner. We also analyze the area and power overhead of CD-PIM, showing that the proposed architecture delivers significant performance improvements with only minimal hardware cost.

REFERENCES

- [1] C. Li *et al.*, “Specpim: Accelerating speculative inference on pim-enabled system via architecture-dataflow co-exploration,” in *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Volume 3, 2024, p. 950–965.
- [2] S. Zhang *et al.*, “Opt: Open pre-trained transformer language models,” 2022. [Online]. Available: <https://arxiv.org/abs/2205.01068>
- [3] E. Nijkamp *et al.*, “Codegen: An open large language model for code with multi-turn program synthesis,” 2023. [Online]. Available: <https://arxiv.org/abs/2203.13474>
- [4] L. Liu *et al.*, “Make llm inference affordable to everyone: Augmenting gpu memory with ndp-dimm,” in *IEEE International Symposium on High Performance Computer Architecture (HPCA)*, 2025, pp. 1751–1765.
- [5] K. Jeun *et al.*, “Fold-pim: A cost-efficient lpddr5-based pim for on-device slms,” *IEEE Computer Architecture Letters (CAL)*, vol. 24, no. 1, pp. 185–188, 2025.
- [6] Z. Zhou *et al.*, “Dimm-link: Enabling efficient inter-dimm communication for near-memory processing,” in *IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, 2023, pp. 302–316.
- [7] V. Pulavarthi *et al.*, “Are llms ready for practical adoption for assertion generation?” in *Design, Automation & Test in Europe Conference (DATE)*, 2025, pp. 1–7.
- [8] R. Qin *et al.*, “Nvcim-pt: An nvcim-assisted prompt tuning framework for edge llms,” in *Design, Automation & Test in Europe Conference (DATE)*, 2025, pp. 1–7.
- [9] Y. Hou *et al.*, “Nora: Noise-optimized rescaling of llms on analog compute-in-memory accelerators,” in *2025 Design, Automation & Test in Europe Conference (DATE)*, 2025, pp. 1–7.
- [10] F. Liu *et al.*, “Aster: Adaptive dynamic layer-skipping for efficient transformer inference via markov decision process,” in *Proceedings of the 33rd ACM International Conference on Multimedia (MM)*, 2025, pp. 11 853–11 861.
- [11] Q. Deng *et al.*, “Dracc: a dram based accelerator for accurate cnn inference,” in *55th ACM/ESDA/IEEE Design Automation Conference (DAC)*, 2018, pp. 1–6.
- [12] S. Cho *et al.*, “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*, vol. 8, pp. 135 223–135 243, 2020.
- [13] J. Park *et al.*, “Attacc! unleashing the power of pim for batched transformer-based generative model inference,” in *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Volume 2, 2024, p. 103–119.
- [14] S.-S. Park *et al.*, “An lpddr-based cxi-pnm platform for tco-efficient inference of transformer-based large language models,” in *IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, 2024, pp. 970–982.
- [15] T. Jeong *et al.*, “Pipepim: Maximizing computing unit utilization in ml-oriented digital pim by pipelining and dual buffering,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, vol. 43, no. 12, pp. 4585–4598, 2024.
- [16] S. Lee *et al.*, “A lynn 1.25v 8gb, 16gb/s/pin gddr6-based accelerator-in-memory supporting 1tflops mac operation and various activation functions for deep-learning applications,” in *IEEE International Solid-State Circuits Conference (ISSCC)*, vol. 65, 2022, pp. 1–3.
- [17] L. Wu *et al.*, “Pimoe: Towards efficient moe transformer deployment on npu-pim system through throttle-aware task offloading,” in *2025 62nd ACM/IEEE Design Automation Conference (DAC)*. IEEE, 2025, pp. 1–7.
- [18] Y. Hu *et al.*, “Plain: Leveraging high internal bandwidth in pim for accelerating large language model inference via mixed-precision quantization,” in *2025 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*. IEEE, 2025, pp. 1–9.
- [19] Y. Chen *et al.*, “P3-llm: An integrated npu-pim accelerator for llm inference using hybrid numerical formats,” *arXiv preprint arXiv:2511.06838*, 2025.
- [20] Y. Zhao *et al.*, “Um-pim: Dram-based pim with uniform & shared memory space,” in *2024 ACM/IEEE 51st Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2024, pp. 644–659.
- [21] H. Lee *et al.*, “Paise: Pim-accelerated inference scheduling engine for transformer-based llm,” in *IEEE International Symposium on High Performance Computer Architecture (HPCA)*, 2025, pp. 1707–1719.
- [22] Y. Kwon *et al.*, “Memory-centric computing with sk hynix’s domain-specific memory,” in *IEEE Hot Chips 35 Symposium (HCS)*, 2023, pp. 1–26.
- [23] M. Seo *et al.*, “Ianus: Integrated accelerator based on npu-pim unified memory system,” in *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Volume 3, 2024, p. 545–560.
- [24] NVIDIA, “Nvidia jetson agx orin series data sheet,” <https://www.diamondsystems.com/files/binaries/Jetson%20AGX%20Orin%20DS10662-001%20v1.2.pdf>, 2022, accessed: 2025-09-14.
- [25] R. Smith, “Apple announces m3 soc family: M3, m3 pro, and m3 max make their marks,” <https://www.anandtech.com/show/21116/apple-announces-m3-soc-family-m3-m3-pro-and-m3-max-make-their-marks>, 2023, accessed: 2025-09-14.
- [26] C. Fang *et al.*, “Anda: Unlocking efficient llm inference with a variable-length grouped activation data format,” in *IEEE International Symposium on High Performance Computer Architecture (HPCA)*, 2025, pp. 1467–1481.
- [27] S. Ma *et al.*, “Apt-llm: Exploiting arbitrary-precision tensor core computing for llm acceleration,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, 2025.
- [28] K. Mei *et al.*, “AIOS: LLM agent operating system,” in *Second Conference on Language Modeling (COLM)*, 2025. [Online]. Available: <https://openreview.net/forum?id=L4HHkCDz2x>
- [29] J. Ahn *et al.*, “Pim-enabled instructions: A low-overhead, locality-aware processing-in-memory architecture,” in *ACM/IEEE 42nd Annual International Symposium on Computer Architecture (ISCA)*, 2015, pp. 336–348.
- [30] G. Heo *et al.*, “Neupims: Npu-pim heterogeneous acceleration for batched llm inferencing,” in *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Volume 3, 2024, p. 722–737.
- [31] C. Li *et al.*, “H2-llm: Hardware-dataflow co-exploration for heterogeneous hybrid-bonding-based low-batch llm inference,” in *Proceedings of the 52nd Annual International Symposium on Computer Architecture (ISCA)*, 2025, p. 194–210.
- [32] S. H. Seo *et al.*, “Facil: Flexible dram address mapping for soc-pim cooperative on-device llm inference,” in *IEEE International Symposium on High Performance Computer Architecture (HPCA)*, 2025, pp. 1720–1733.
- [33] S.-Y. Cho *et al.*, “A 16-gb 37-gb/s gddr7 dram with pam3-optimized trx equalization and zq calibration,” *IEEE Journal of Solid-State Circuits (JSSC)*, vol. 60, no. 1, pp. 184–196, 2025.
- [34] B. J. Kim *et al.*, “Dh-pim: Maximizing computing unit utilization in digital pim by dual half mode extension,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, 2025.
- [35] H. Luo *et al.*, “Ramulator 2.0: A modern, modular, and extensible dram simulator,” *IEEE Computer Architecture Letters (RAL)*, vol. 23, no. 1, pp. 112–116, 2024.
- [36] H. Touvron *et al.*, “Llama: Open and efficient foundation language models,” 2023. [Online]. Available: <https://arxiv.org/abs/2302.13971>