

CLSA-CIM: A Cross-Layer Scheduling Approach for Computing-in-Memory Architectures

Rebecca Pelke[✉], Jose Cubero-Cascante[✉], Nils Bosbach[✉], Felix Staudigl[✉], Rainer Leupers[✉], Jan Moritz Joseph[✉]

Institute for Communication Technologies and Embedded Systems

RWTH Aachen University, Germany

{pelke, cubero, bosbach, staudigl, leupers, joseph}@ice.rwth-aachen.de

Abstract—The demand for efficient machine learning (ML) accelerators is growing rapidly, driving the development of novel computing concepts such as resistive random access memory (RRAM)-based tiled computing-in-memory (CIM) architectures. CIM allows to compute within the memory unit, resulting in faster data processing and reduced power consumption. Efficient compiler algorithms are essential to exploit the potential of tiled CIM architectures. While conventional ML compilers focus on code generation for CPUs, GPUs, and other von Neumann architectures, adaptations are needed to cover CIM architectures. Cross-layer scheduling is a promising approach, as it enhances the utilization of CIM cores, thereby accelerating computations. Although similar concepts are implicitly used in previous work, there is a lack of clear and quantifiable algorithmic definitions for cross-layer scheduling for tiled CIM architectures.

To close this gap, we present CLSA-CIM, a cross-layer scheduling algorithm for tiled CIM architectures. We integrate CLSA-CIM with existing weight-mapping strategies and compare performance against state-of-the-art (SOTA) scheduling algorithms. CLSA-CIM improves the utilization by up to $17.9\times$, resulting in an overall speedup increase of up to $29.2\times$ compared to SOTA.

Index Terms—RRAM, CIM, compiler, cross-layer scheduling

I. INTRODUCTION

The increasing demand for efficient computation of data-intensive machine learning (ML) applications has led to specialized architectures such as graphics processing units (GPUs) and tensor processing units (TPUs). However, a major performance limitation is the data movement between main memory and compute units, known as the von Neumann bottleneck [1]. Novel computing-in-memory (CIM) technologies, such as resistive random access memory (RRAM), tackle this bottleneck by unifying memory and computation unit [2]. These designs outperform their CMOS-based counterparts in memory capacity, device density, and power consumption [3].

In recent years, several CIM architectures have been introduced [4], [5], [6], [7]. These architectures adopt a tiled structure, as shown in Figure 1(a), wherein tiles are interconnected via a network on chip (NoC). To achieve high energy efficiency and inference performance, maximizing the utilization of the processing elements (PEs), located inside the tiles, is essential. This imposes a special challenge for CIM architectures since the neural network (NN)'s weights are statically assigned to the PEs and remain there during inference. To increase the PEs's utilization, the compiler needs to exploit

This work was funded by the Federal Ministry of Education and Research (BMBF, Germany) in the project NeuroSys (Project Nos. 03ZU1106CA).

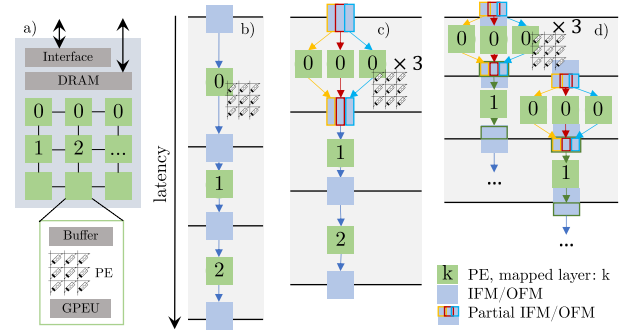


Fig. 1: NN inference on (a) tiled CIM architectures: (b) Layer-by-layer scheduling, (c) weight duplication mapping, and (d) cross-layer scheduling

both *intra-* and *cross-layer scheduling* of the workload [8]. Previous research has mainly focused on intra-layer scheduling techniques [9], [10], [11], [12], which only consider parallel execution of individual layers. *Weight duplication*, a mapping method proposed in [13], [14], [15], involves assigning the same weights to multiple PEs to divide input data between them. This approach is restricted by the limited number of PEs and can only accelerate individual layers. It does not increase the utilization of the PEs. In contrast, *cross-layer scheduling* accomplishes this by considering optimizations across layer boundaries [16]. Cross-layer scheduling forwards parts of a layer's output feature map (OFM) to subsequent layers before the entire OFM has been computed (Figure 1(d)). While previous research has addressed the development of tiled CIM architectures that enable cross-layer scheduling in hardware [6], [13], [17], there is a lack of a software methodology to fully exploit this feature. We want to close this gap by presenting the following contributions:

- We extend the existing weight duplication approaches by developing an algorithm that decides which parts of the NN are duplicated to achieve minimum inference latency. Details of the TensorFlow implementation are provided.
- We introduce CLSA-CIM, a novel approach for cross-layer scheduling on tiled CIM architectures. CLSA-CIM seamlessly integrates with existing mapping strategies, such as weight duplication, while leveraging established intra-layer scheduling techniques.
- In a case study, we show that CLSA-CIM has the potential to achieve up to $29.2\times$ speedup by massively improving the PE utilization.

II. BACKGROUND AND RELATED WORK

A. RRAM-based Tiled CIM Architectures

RRAM devices offer programmable conductance values and are used in crossbar structures for efficient in-memory matrix-vector multiplication (MVM) operations [18]. RRAM cells have a limited endurance [19]. It therefore makes sense to store all NN weights only once before inference. This also avoids costly rewriting processes [4]. Consequently, RRAM-based CIM architectures typically incorporate a large number of crossbars to store the entire NN [5]. Various CIM architectures have been proposed to enable efficient and parallel MVM execution [4], [5], [6], [20]. Drawing from these accelerators, we will define fundamental hardware requirements that must be fulfilled to support cross-layer scheduling (see Figure 1(a)):

- Tiles that exchange data with other tiles via a NoC.
- All tiles operate in parallel and independently.
- Within the tiles, there are buffers to store parts of the input and output data.
- Due to limited buffer memory, all tiles have fast access to a global DRAM for data exchange.
- Inside the tiles, there are crossbar(s), also called PE(s).
- The number of tiles and PEs is sufficient to store all weights of the NN at least once on the architecture.
- Each tile has general purpose execution unit (GPEU) to execute other operations than MVM (e.g., *pooling*).

While the majority of the tiled CIM accelerators meet these cross-layer scheduling requirements [5], [6], [20], they differ in GPEU logic, NoC structure, tile count, PE dimensions, or buffer capacities. From the cross-layer scheduling perspective, these differences are not relevant as long as the above-mentioned requirements are met.

B. Intra-Layer Scheduling and Layer-by-Layer Inference

The performance gains of tiled CIM accelerators stem from the fact that the PEs can efficiently perform MVMs in parallel. Previous works exploit intra-layer parallelism to speed up inference [9], [10], [11], [12]. However, the overall PE utilization remains low as only one layer's PE(s) are active at any time, which is called layer-by-layer inference.

C. Weight Duplication Mapping

To speed up layer-by-layer inference, weight duplication aims to further enhance the intra-layer parallelization capabilities of NNs by storing the same weight sets in two or more PEs. This creates parallel duplicated layer nodes in the NNs graph, which can be executed in parallel on CIM architectures (see Figure 1(c)). This idea has been proposed in previous works [13], [14], [15]. Weight duplication comes at the cost of increased resource requirements. It is mandatory to determine which layers should be duplicated and how often. We extend existing approaches by developing an algorithm that estimates the costs and values of duplicating certain weights. This is explained in more detail in Section III-C.

D. Cross-Layer Inference

The idea of cross-layer inference is that partial results can be passed to the next PEs even before an entire layer has been fully computed. This improves the inference latency and PE utilization. While cross-layer inference approaches have been proposed for tiled accelerators [8], they have not yet been applied to CIM systems. It is particularly well-suited for CIM architectures because of their weight-stationary data flow. Some architectures are specifically designed for this purpose. For instance, the authors of [13] presented a CIM architecture that includes synchronization mechanisms particularly tailored to support cross-layer inference. However, they do not provide a general approach on the software side. CLSA-CIM, our software-based scheduling approach, overcomes this limitation. Implementation details are presented in Section III and Section IV.

III. CROSS-LAYER SCHEDULING - PREPARATION

Preprocessing transforms the NN model (from TensorFlow) into a unified structure, which serves as input for the CLSA-CIM algorithm. It involves architecture-specific high-level optimizations, the refinement of existing mapping approaches including concepts like *im2col* and weight duplication, as well as the use of existing intra-layer scheduling concepts.

A. High-Level Optimizations

BN folding: Batch normalization (BN) layers enhance training stability and convergence speed by normalizing the input distributions. For inference, the BN layer can be merged with the previous operation, known as BN folding [21]. It improves computational efficiency and memory utilization by adjusting the kernel weights w and the bias b in the Conv2D operation.

Partitioning: The NN is divided into *base layers*, i.e., operations executed on the PE (like convolutions and dense layers), and *non-base layers* (all remaining layers). In the NN graph, padding and bias addition are decoupled from the base layer, eliminating redundancy in the graph representation.

Quantization: Base layers need to be quantized due to the limited resolution of PE (RRAM) cells. For existing PEs, this resolution can be up to 4 Bits [4]. The preprocessing steps are summarized in Figure 2 using a minimal example.




Fig. 2: Partitioning, quantization (Q), and BN folding. The resulting canonical NN representation is split into *base* (green) and *non-base* (blue) layers

B. Im2col and Intra-Layer Scheduling

Base layers must be translated into MVMs to execute them on the PEs. One widely used technique for convolutions is converting them into general matrix multiplies (GEMMs). This can be accomplished through the use of *im2col* [9]. Figure 3 illustrates the *im2col* algorithm, which unrolls the individual kernels and arranges them in columns which leads to a $(K_W \cdot K_H \cdot K_I) \times K_O$ kernel matrix. The kernel matrix is subdivided

into submatrices of size $M \times N$, which are statically mapped to the accelerators PEs [12]. Previous research shows that all PEs within a layer can operate in parallel with minimal latency overhead, a concept known as intra-layer scheduling [22].

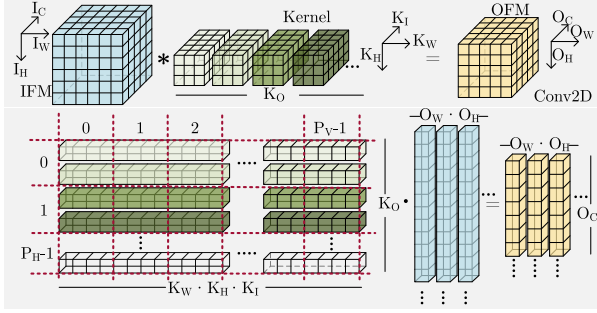


Fig. 3: Conv2D to GEMM transformation using im2col

Therefore, we simplify by assuming that the calculation of an $(1 \times 1 \times O_C)$ OFM vector takes place within t_{MVM} , which represents the MVM latency of a PE. Accordingly, the total latency to compute the OFM of a single layer using intra-layer scheduling is $t_{OFM} = O_H \cdot O_W \cdot t_{MVM}$.

C. Weight Duplication Mapping

Weight duplication reduces the inference latency of a single layer since the work, i.e., the input vectors, is evenly distributed among the duplicates. The latency of the Conv2D operation reduces to $t_{OFM} = \frac{1}{D} \cdot O_H \cdot O_W \cdot t_{MVM}$, where D is the number of duplicates. Duplicating the kernel matrix comes at the cost of requiring more PEs to store all weights. This means that weight duplication is rather beneficial for layers with high calculation latency (large $O_H \cdot O_W$ factor) and a small number of required PEs. As discussed in Section II-A, it is assumed that the architecture has a sufficient number of PEs to store all weights without rewriting. If the architecture has F PEs and the NN needs C_{num} PEs, with $C_{num} < F$, weight duplication can be applied to further reduce the inference latency (see Section II-C). The solution vector \mathbf{d} of Optimization Problem 1 determines which layers should be duplicated to achieve the best inference latency:

Optimization Problem 1 Weight Duplication

$$\begin{aligned} &\text{minimize: } \sum_i \frac{t_i}{d_i} \\ &\text{subject to: } \mathbf{c}^T \cdot \mathbf{d} \leq F, \\ &\quad \mathbf{d} \geq \mathbf{1}, \\ &\quad \mathbf{d} \in \mathbb{Z}_+^N \end{aligned}$$

Vector \mathbf{t} contains the latencies needed to calculate the OFM of the base layers with $\mathbf{t}^T = (t_{OFM_0}, t_{OFM_1}, \dots, t_{OFM_{N-1}})$. Vector \mathbf{c} contains the number of required $M \times N$ PEs for every base layer, e.g., convolutions (see Figure 3):

$$c_i = \underbrace{\left\lceil \frac{K_{W,l} \cdot K_{H,l} \cdot K_{l,l}}{N} \right\rceil}_{=: P_{V,i}} \cdot \underbrace{\left\lceil \frac{K_{O,l}}{M} \right\rceil}_{=: P_{H,i}}, \quad \sum_i c_i = C_{num} \quad (1)$$

The vector \mathbf{d} also specifies the number of base layer duplicates to be created. The latency values t_i for calculating one layer i using intra-layer scheduling are set according to Section III-B. Note that the solution determines which weights to duplicate, but it does not determine how to distribute the work, i.e., the input feature map (IFM), among the duplicates. Keeping in mind the intra-layer scheduling algorithm in Section III-B, the IFMs and OFMs should be cut along the I_W/O_W and/or I_H/O_H dimensions, as shown in Figure 4.




Fig. 4: Implementation of weight duplication using three duplicates

The example provides details of the TensorFlow-specific implementation of weight duplication. The OFM is divided into $2 \times 2 \times 1$ disjoint parts. In the NN graph, this is realized by applying one `tf.slice` operation on the IFM for each duplicate. The IFM slices may overlap depending on the kernel shape and stride. After the distributed calculations, the OFMs are concatenated using `tf.keras.layers.Concatenate`. The depth of the concatenated tree corresponds to the number of dimensions along which it has been cut. The influence of weight duplication on the inference latency will be discussed in Section V-A.

IV. CROSS-LAYER SCHEDULING - CLSA-CIM

CLSA-CIM builds upon the mapping and intra-layer scheduling concepts from Section III. It aims to minimize the inference latency by maximizing the utilization of the tiles (see Section II-A). The algorithm comprises two preprocessing stages for creating the necessary data structures, *determine sets* and *determine dependencies* (Figure 5(a)-(b)), followed by two scheduling stages: First, intra-layer scheduling is applied, followed by the actual cross-layer scheduling (Figure 5(c)).

1) *Stage 1 - Determine Sets*: The OFMs is divided into disjoint sets, which are the minimum scheduling units. This means that all elements within this set must be processed before elements from another set of the same OFM can be calculated. The sets should ideally contain a similar number of elements; otherwise, the execution time for each set may vary. Additionally, a hyperrectangle shape allows to identify the set's location and size using two coordinates. Increasing the number of sets provides a more detailed scheduling granularity. The sets should be sufficiently large to facilitate the execution of non-base layer operations, such as pooling. In the example in Figure 5(a), the sets must contain at least 2×2 values to accommodate $(2, 2)$ pooling with a stride of $(2, 2)$. Next, the intra-layer dependencies are determined. For each set of the OFM, the corresponding set of the IFM is calculated. When adding new base layers to the algorithms, this dependency has to be specified.

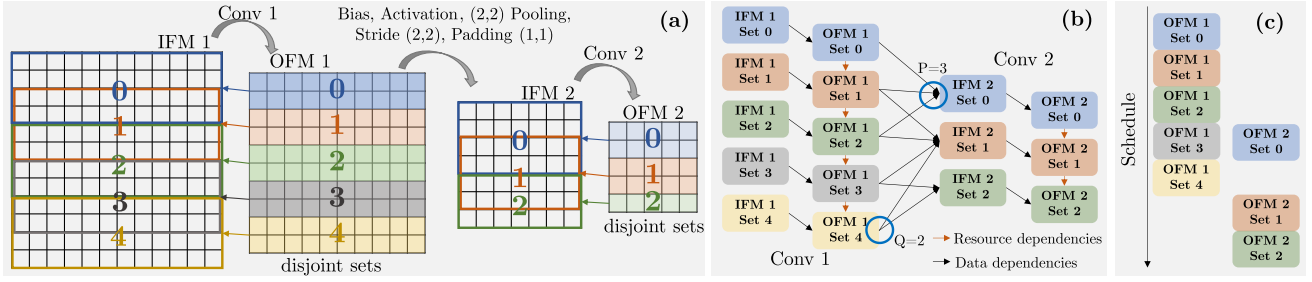


Fig. 5: Minimal example for CLSA-CIM using two consecutive Conv2D layers and a non-base layer path including bias, activation, pooling, and padding: Determine sets (a), determine dependencies (b), and cross-layer scheduling (c)

2) *Stage II - Determine Dependencies*: This stage calculates the dependencies between consecutive base layers. The two points specifying the location and size of the OFM set of a predecessor are propagated along the non-base layer path to determine which IFM sets are affected. In Figure 5(b), it is evident that each OFM set can influence multiple IFM sets (denoted as Q), and likewise, each IFM set can be affected by multiple OFM sets (denoted as P).

3) *Stage III - Intra-Layer Scheduling*: In the third stage, the scheduling order of the OFM sets is determined for each base layer individually. The execution order of the OFM sets can be seen in Figure 5(b). The orange-colored connections between OFM sets of the same layer indicate resource dependencies, which means that the same crossbars are needed to calculate those sets. In the example, $OFM1_{set0}$ has to be scheduled before the other sets of $Conv1$ since it allocates the resources first. The dependencies marked in black are data dependencies. To generate the IFM set of $Conv2$ from the OFM sets of $Conv1$, non-base layer operations, e.g., pooling, are applied.

4) *Stage IV - Cross-Layer Scheduling*: Figure 5(c) shows the resulting schedule. Note that the non-base layer operations are not illustrated due to simplicity. CLSA-CIM ascertains the earliest feasible starting point for computing each OFM set in the NN. In other words, an OFM set is scheduled once all the required IFM sets of its predecessors have been scheduled.

A. Combine Weight Duplication and Cross-Layer Scheduling

Weight duplication is a mapping technique, whereas cross-layer inference is a scheduling technique. These concepts can be used independently. Combining them can further reduce inference latency. The weight duplication algorithm is applied first, resulting in a non-sequential NN graph where each layer can have multiple predecessors and successors. CLSA-CIM is applied after that. It is designed to handle non-sequential models in a generic manner, requiring no additional modifications or adjustments. This allows for the seamless integration of weight duplication and CLSA-CIM.

V. EVALUATION

This chapter evaluates the performance of CLSA-CIM. We distinguish between three approaches: weight duplication mapping combined with layer-by-layer inference ($wdup$), cross-layer inference ($xinf$), and the combination of weight duplication mapping and cross-layer inference ($wdup+xinf$).

All speedup measurements are referenced to the layer-by-layer inference (see Section II-B). As there are currently no commercially available CIM chips, we use a custom system-level simulator, similar to previous works [13], [14], [23]. We calculate the maximum achievable utilization and minimum inference latency achievable with CLSA-CIM. For the simulation, three core parameters are required: the number of PEs, the dimensions of a PE, and the MVM latency. In a case study, we assume a 256×256 crossbar and an MVM latency of $t_{MVM} = 1,400$ ns [4], which we call a *cycle*. The number of CIM cores is kept variable in the simulation to investigate its impact on the latency. If future CIM architectures meet the prerequisites outlined in Section II-A, CLSA-CIM can be used. CLSA-CIM increases the architecture utilization Ut , which is defined as the mean over the ratio of the active cycle time $t_{p,active_cycles}$ to the total inference time of the NN t_{NN_cycles} for PE p :

$$Ut := \frac{1}{\#PE} \left(\sum_{p \in PE} \frac{t_{p,active_cycles}}{t_{NN_cycles}} \right) \quad (2)$$

The number of PEs is varied for each benchmark to enable weight duplication. The notation “ $wdup_{+x}$ ”, e.g., “ $wdup_{+32}$ ”, means that the architecture has 32 PEs more than needed to store all NN weights exactly once.

A. CLSA-CIM - A Case Study

We analyze our scheduling approach (CLSA-CIM) with a TinyYOLOv4 case study. TinyYOLOv4 is a non-sequential NN for object detection and classification. Table I shows an extract of the base layer structure of TinyYOLOv4.

TABLE I: Extract of the base layer structure of TinyYOLOv4

Layer	IFM shape (HWC)	OFM shape (HWC)	#PE 256×256	Cycles t_{init}
conv2d	(417, 417, 3)	(208, 208, 32)	1	43264
conv2d_1	(209, 209, 32)	(104, 104, 64)	2	10816
conv2d_2	(106, 106, 64)	(104, 104, 64)	3	10816
...
conv2d_16	(15, 15, 256)	(13, 13, 512)	18	169
conv2d_20	(26, 26, 256)	(26, 26, 255)	1	676
conv2d_17	(13, 13, 512)	(13, 13, 255)	2	169

TinyYOLOv4 has 18 Conv2D layers. The minimum number of PEs required to store all weights at least once, PE_{min} , is 117. The time t_{init} is the duration of executing the layer (in cycles) using only intra-layer scheduling (see Section III-B).




Fig. 6: Visualization of weight duplication mapping (a) and CLSA-CIM (b) using $x = 16$ additional PEs, speedup and PE utilization for different weight mapping (weight duplication) and scheduling (layer-by-layer, CLSA-CIM) combinations (c)

Since the $O_H \cdot O_W$ factor is higher for the first layers, they are more time-consuming to compute. The first layers need fewer PEs, which makes them a good choice for weight duplication.

Figure 6 provides an illustration of the weight duplication mapping (wdup) of the TinyYOLOv4 benchmark combined with layer-by-layer scheduling (Figure 6a) and CLSA-CIM (Figure 6b). The solution of Algorithm 1 in Section III-C reveals that for $x = 16$ additional PEs, the first 6 Conv2D layers need to be duplicated according to the table in Figure 6a. Figure 6c confirms what is visible in Figure 6b: CLSA-CIM (xinf) increases the utilization of the PEs to a total of 4.1%. In combination with weight duplication and $x = 32$ additional PEs (wdup₊₃₂), in total 117+32 PEs, the utilization increases up to 28.4%. This corresponds to an inference speedup of up to 21.9 \times . The relationship between speedup S and utilization Ut for $+x$ PEs and configuration c is

$$S_{x,c} \approx \frac{Ut_{x,c} \cdot (PE_{min} + x)}{Ut_{layer_by_layer} \cdot PE_{min}}. \quad (3)$$

B. Performance Results

We further evaluate benchmarks that have a higher demand for PEs than TinyYOLOv4. This includes sequential models like VGG16 and VGG19, and non-sequential models like TinyYOLOv3, ResNet50, ResNet101, and ResNet152. Since the latency and utilization depend on the dimensions of the IFM, the dimensions are listed in Table II.

TABLE II: List of benchmarks

Benchmark	Input shape (HWC)	Base layers (number)	Min. # required 256 \times 256 PEs
TinyYOLOv3	(416, 416, 3)	13	142
VGG16	(224, 224, 3)	13	233
VGG19	(224, 224, 3)	16	314
ResNet50	(224, 224, 3)	53	390
ResNet101	(224, 224, 3)	104	679
ResNet152	(224, 224, 3)	155	936

The inference latency speedups and PE utilizations are compared for different combinations of mapping (wdup_{+x}) and scheduling (layer-by-layer, xinf). For additional PEs,

we consider the setups $x \in \{4, 8, 16, 32\}$, i.e., for VGG19, 314 to 346 PEs. This enables a better comparison across different benchmarks. In the tested configurations in Figure 7a, pure weight duplication yields a modest speedup for large models, from 1.1 \times to 1.9 \times . This is because the number of additional PEs (up to 32) is small compared to the minimum required PEs to store the entire NN on the accelerator. CLSA-CIM (xinf) achieves a speedup of up to 4.4 \times for large models compared to layer-by-layer scheduling. The best results are achieved by combining CLSA-CIM and weight duplication. This approach yields the highest speedup of 29.2 \times for the TinyYOLOv3 benchmark. Of particular interest is that only $x = 4$ additional PEs are sufficient to outperform the pure xinf configuration by a factor of almost 2 \times . We observe this even for ResNet152, where $x = 4$ PEs is very small compared to the minimum PE requirement of 936. This can be attributed to the fact that, as demonstrated in Figure 6a, the first layer is relatively computation-intensive. Figure 7b illustrates the PE utilizations for each benchmark. The utilization is increased by CLSA-CIM across all benchmarks, surpassing the impact of pure weight duplication. Again, the combination of weight duplication and CLSA-CIM delivers the best performance values. For smaller models, higher utilization rates can be achieved, with TinyYOLOv3 reaching a maximum utilization of 20.1%. This represents an improvement of 17.9 \times compared to layer-by-layer scheduling. Since the final layers often require many PEs (see Table I), but at the same time are less computationally intensive, the utilization of the architecture for a single NN inference usually remains below 10%. As the model depth increases, the utilization decreases, as observed in the ResNet benchmarks. This is due to the limited parallelization capabilities between layers which are far apart in the NN graph.

C. Limitations and Future Work

As mentioned in Section II-A, our current work focuses on cases where the number of crossbars is sufficient to accommodate complete NNs on the architecture. However, in future research, we aim to explore more general scenarios. CLSA-

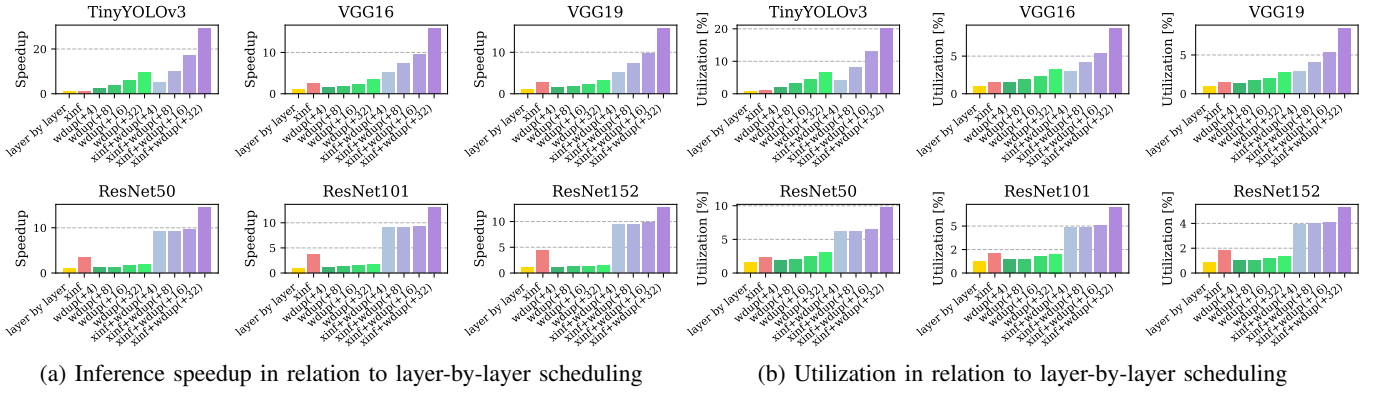


Fig. 7: Combinations of mapping and scheduling in contrast to layer-by-layer scheduling without weight duplication: layer-by-layer scheduling with weight duplication (wdup), CLSA-CIM (xinf), and weight duplication with CLSA-CIM (wdup+xinf)

CIM is already designed to accept the crossbar dimensions as an input parameter, allowing for adaptability to arbitrary sizes. It is important to acknowledge that the speedup values presented in this study represent peak performance. There may be architecture-dependent factors that could potentially impact latency. For example, the costs associated with data movement have not been differentiated yet. Depending on the topology, forwarding partial results may incur varying costs. Furthermore, it is possible for cores to share resources such as adders, further imposing constraints on scheduling algorithms. Our future work will involve extending our abstract architecture description to account for these factors, enabling full architecture retargetability.

VI. CONCLUSION

The development of efficient scheduling algorithms for tiled CIM architectures is crucial to fully utilize the potential of CIM concepts. Our scheduling approach, CLSA-CIM, enables cross-layer inference on top of existing intra-layer scheduling and weight duplication mapping algorithms, which significantly enhances the utilization of PEs of up to $17.9\times$, resulting in an inference speedup of up to $29.2\times$. We conducted evaluations using state-of-the-art NNs, including a case study of the TinyYOLOv4 model to visualize the algorithms. In summary, our work contributes to the advancement of scheduling approaches and algorithms for CIM architectures. It sheds light on the benefits of combining cross-layer inference and weight duplication, paving the way for enhanced performance of ML applications on CIM architectures.

REFERENCES

- [1] X. Zou, S. Xu, X. Chen, L. Yan, and Y. Han, "Breaking the von Neumann bottleneck: architecture-level processing-in-memory technology," *Science China Information Sciences*, 2021.
- [2] Y.-F. Chang *et al.*, "Memcomputing (Memristor + Computing) in Intrinsic SiOx-Based Resistive Switching Memory: Arithmetic Operations for Logic Applications," *IEEE (T-ED)*, 2017.
- [3] J. S. Vetter and S. Mittal, "Opportunities for Nonvolatile Memory Systems in Extreme-Scale High Performance Computing," *Computing in Science & Engineering*, 2015.
- [4] W. Wan *et al.*, "A compute-in-memory chip based on resistive random-access memory," *Nature*, 2022.
- [5] A. Ankit *et al.*, "PUMA: A Programmable Ultra-efficient Memristor-based Accelerator for Machine Learning Inference," in *ASPLOS XXIV*, 2019.
- [6] A. Shafiee *et al.*, "ISAAC: A convolutional neural network accelerator with in-situ analog arithmetic in crossbars," *ACM SIGARCH Computer Architecture News*, 2016.
- [7] P. Chi *et al.*, "Prime: A novel processing-in-memory architecture for neural network computation in reram-based main memory," *ACM SIGARCH Computer Architecture News*, 2016.
- [8] J. Cai, Y. Wei, Z. Wu, S. Peng, and K. Ma, "Inter-layer Scheduling Space Definition and Exploration for Tiled Accelerators," in *50th ISCA*, 2023.
- [9] K. Yanai, R. Tanno, and K. Okamoto, "Efficient Mobile Implementation of A CNN-based Object Recognition System," in *Proceedings of the 24th ACM international conference on Multimedia*, 2016.
- [10] X. Peng, R. Liu, and S. Yu, "Optimizing Weight Mapping and Data Flow for Convolutional Neural Networks on Processing-in-Memory Architectures," *IEEE Transactions on Circuits and Systems Is*, 2019.
- [11] S. Negi, I. Chakraborty, A. Ankit, and K. Roy, "NAX: neural architecture and memristive xbar based accelerator co-design," in *DAC*, 2022.
- [12] A. Agrawal, C. Lee, and K. Roy, "X-CHANGR: Changing Memristive Crossbar Mapping for Mitigating Line-Resistance Induced Accuracy Degradation in Deep Neural Networks," *arXiv preprint arXiv:1907.00285*, 2019.
- [13] X. Liu *et al.*, "FPRA: A Fine-grained Parallel RRAM Architecture," in *2021 IEEE/ACM ISLPED*. IEEE, 2021.
- [14] J. Rhe, S. Moon, and J. H. Ko, "VWC-SDK: Convolutional Weight Mapping Using Shifted and Duplicated Kernel with Variable Windows and Channels," *IEEE JETCAS*, 2022.
- [15] Z. Zhu *et al.*, "Mixed Size Crossbar based RRAM CNN Accelerator with Overlapped Mapping Method," in *IEEE/ACM ICCAD*, 2018.
- [16] V. Sze, Y.-H. Chen, T.-J. Yang, and J. S. Emer, "Efficient Processing of Deep Neural Networks: A Tutorial and Survey," *Proceedings of the IEEE*, 2017.
- [17] A. Symons, L. Mei, S. Coleman, P. Houshmand, S. Karl, and M. Verhelst, "Towards Heterogeneous Multi-core Accelerators Exploiting Fine-grained Scheduling of Layer-Fused Deep Neural Networks," *arXiv preprint arXiv:2212.10612*, 2022.
- [18] W. Cao, Y. Zhao, A. Bolor, Y. Han, X. Zhang, and L. Jiang, "Neural-PIM: Efficient Processing-In-Memory With Neural Approximation of Peripherals," *IEEE Transactions on Computers*, 2021.
- [19] C. Nail *et al.*, "Understanding rram endurance, retention and window margin trade-off using experimental results and simulations," in *IEEE IEDM*. IEEE, 2016.
- [20] L. Song, X. Qian, H. Li, and Y. Chen, "PipeLayer: A Pipelined ReRAM-Based Accelerator for Deep Learning," in *IEEE HPCA*, 2017.
- [21] B. Jacob *et al.*, "Quantization and Training of Neural Networks for Efficient Integer-Arithmetic-Only Inference," in *CVPR*, 2018.
- [22] R. Pelke, N. Bosbach, J. Cubero, F. Staudigl, R. Leupers, and J. M. Joseph, "Mapping of CNNs on multi-core RRAM-based CIM architectures," in *IFIP/IEEE VLSI-SoC*. IEEE, 2023.
- [23] A. Lu, X. Peng, W. Li, H. Jiang, and S. Yu, "NeuroSim Simulator for Compute-in-Memory Hardware Accelerator: Validation and Benchmark," *Frontiers in artificial intelligence*, 2021.