

Sorted Weight Sectioning for Energy-Efficient Unstructured Sparse DNNs on Compute-in-Memory Crossbars

Matheus Farias
Harvard University
matheusfarias@g.harvard.edu

H. T. Kung
Harvard University
kung@harvard.edu

Abstract—We introduce *sorted weight sectioning* (SWS): a weight allocation algorithm that places sorted deep neural network (DNN) weight sections on bit-sliced compute-in-memory (CIM) crossbars to reduce analog-to-digital converter (ADC) energy consumption. Data conversions are the most energy-intensive process in crossbar operation. SWS effectively reduces this cost leveraging (1) small weights and (2) zero weights (weight sparsity).

DNN weights follow bell-shaped distributions, with most weights near zero. Using SWS, we only need low-order crossbar columns for sections with low-magnitude weights. This reduces the quantity and resolution of ADCs used, exponentially decreasing ADC energy costs without significantly degrading DNN accuracy.

Unstructured sparsification further sharpens the weight distribution with small accuracy loss. However, it presents challenges in hardware tracking of zeros: we cannot switch zero rows to other layer weights in unsorted crossbars without index matching. SWS efficiently addresses unstructured sparse models using offline remapping of zeros into earlier sections, which reveals full sparsity potential and maximizes energy efficiency.

Our method reduces ADC energy use by 89.5% on unstructured sparse BERT models. Overall, this paper introduces a novel algorithm to allow energy-efficient CIM crossbars for unstructured sparse DNN workloads.

Index Terms—resistive crossbars, computing in memory, software optimizations, sparsity.

I. INTRODUCTION

Resistive compute-in-memory (CIM) crossbars have emerged as promising deep neural network (DNN) accelerators by reducing costly data movement between memory and processing units. [1]–[3]. By integrating storage and computation directly within the memory array, crossbars perform multiply-accumulate (MAC) operations with higher speed and lower power consumption compared to conventional digital systems [4]–[8]. However, the power efficiency of CIM architectures is limited by the energy consumption of analog-to-digital converters (ADCs), which can account for up to 85% of the total energy and area in these systems [9]. This presents a critical challenge to the practical deployment of CIM-based accelerators for DNN workloads.

One way to improve the energy efficiency of CIM crossbars is by leveraging sparsity in DNNs. Sparsity, achieved through pruning techniques, reduces the number of weights, decreasing memory usage and MAC operations. Pruning can be either unstructured and structured. While unstructured pruning removes individual weights to increase sparsity with minimal

model accuracy loss, it complicates efficient tracking due to the irregular patterns of zero weights.

On the other hand, structured pruning removes entire neurons, filters, or channels, creating regular zero patterns that simplify hardware tracking. This method reduces crossbar mapping complexity, improving energy efficiency. However, structured pruning can result in significant model accuracy loss due to the aggressive reduction in the network’s capacity to learn and represent complex features [10]. Therefore, a key challenge is merging the accuracy retention of unstructured pruning with the hardware efficiency of structured pruning.

To address this challenge, we propose a novel technique called *sorted weight sectioning* (SWS), which optimizes unstructured sparse DNN implementation on CIM crossbars. SWS strategically organizes weights exploiting two critical DNN characteristics to minimize ADC energy consumption while maintaining model accuracy: (1) bell-shaped distribution, where a large proportion of weights are distributed around zero, and (2) weight sparsity. By sorting weights by magnitude and grouping them into sections, SWS maps smaller weights to crossbar columns with lower power-of-two multipliers, termed *low-order columns*. These sections require less precise and fewer ADCs, thereby reducing energy use without compromising model accuracy. Additionally, the technique remaps isolated zero weights, achieving both energy efficiency and robust model performance.

SWS involves four steps: (1) sorting the weight vector by magnitude, (2) partitioning the sorted vector into sections of rows, where each row corresponds to a weight value, (3) programming each section into a CIM crossbar, and (4) permuting the activation vector according to the sorted weight order to ensure dot-product correctness. This approach, summarized in Figure 1, enables fine-grained control of ADC usage and enhances energy efficiency of CIM architectures for sparse DNNs. The main contributions of this paper are:

- A mathematical analysis of how programming sections with sorted weights reduce the ADC energy consumption.
- A novel permutation strategy based on sorted sectioning leveraging DNN weight distribution and sparsity to decrease ADCs energy use without compromising accuracy.
- End-to-end experiments in the state-of-the-art simulator CiMLoop [11] and PyTorch to show energy consumption

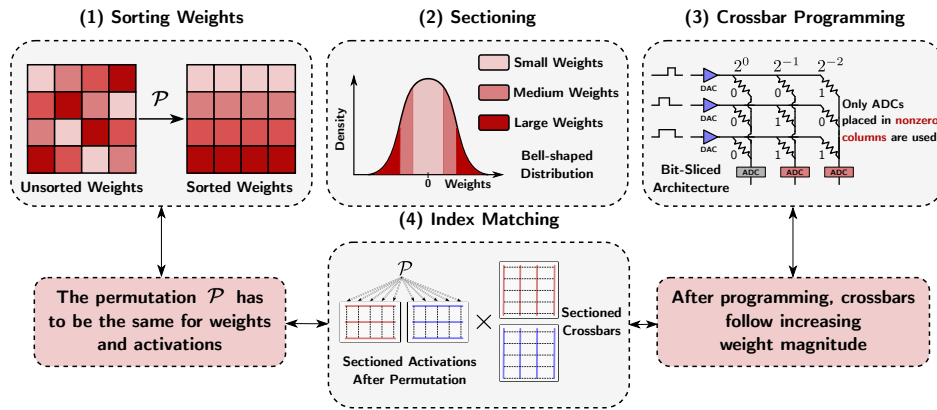


Fig. 1: Summary of the approach. (1) sort the weights, (2) partition into sections of increasing weight magnitude, (3) program each crossbar section, and (4) permute the activation vector to ensure dot-product correctness.

reduction after applying sorted weight sectioning.

II. BACKGROUND

In resistive crossbars, weights are stored in the conductance of programmable resistors called memristors [12], [13] and grouped in 1T1R cells to simplify programming [14]. This work employs a bit-sliced design, where each crossbar row represents a bit-sliced weight value, and each column acts as a negative power-of-two multiplier [5], [15]. Bit-sliced architectures are commonly used for precision-demanding applications [7]. Despite using more memristors than multi-level implementations [8], [16], bit-sliced architectures fit signed weights without additional digital logic and with fewer nonidealities [4].

The dot product is computed by accumulating currents obtained by multiplying electric potentials with conductances across each column. Finally, ADCs convert currents back to digital. Digitized column outputs are shifted according to their column index (i.e., multiplied by a power of two). Past works address digital partial sum accumulation to increase accuracy [5], [17], [18]. Small crossbars reduce nonidealities but require more ADCs: one for each partial sum accumulation section. Combining sectioning with bit-slicing is discussed in [15].

ADCs are the main energy bottleneck, consuming up to 85% of the area and energy for mixed-signal tasks [9]. A b -bit Flash ADC requires 2^b active comparators for conversion. Moreover, one of the most common ADC architectures, the SAR-ADC, has a built-in DAC module. Thus, we target analog-to-digital conversions to minimize crossbar energy consumption.

III. SORTED WEIGHT SECTIONING

In crossbar computing, dividing the matrix-matrix multiplication (MMM) into sections generalizes the conventional approach by segmenting the operation based on the architecture's constraints, such as size or precision limits. Each section corresponds to a subset of the matrix, and the results are aggregated to compute the final output (see Figure 2). This sectioning not only overcomes hardware limitations but also enables new optimizations.




Fig. 2: (a) The dot product of a 15-element row vector (red) of A and the 15-element column vector (red) of B to compute final elements of C , and (b) the dot product of a colored section, such as the section 1 (blue) of A and the corresponding section (blue) of B computes partial results for C . Accumulating dot products for partitioned A and B in Figure 2b generalize conventional dot products in Figure 2a by using matrix sections as matrix elements.




Fig. 3: (a) Pretrained weights follow a bell-shaped distribution around zero. (b) This example compares the number of ADCs required for sections s_1 , s_2 , and s_3 using unsorted and sorted approaches. As most frequent weights are near zero, we require only low-order columns, thereby, using fewer ADCs.

A. Mathematical Formulation of SWS

Now we justify why SWS provides a weight allocation strategy that enhances energy efficiency.

Let W be a random variable following a normal distribution $N(0, \sigma)$ (approximation of bell-shaped curve of pretrained weights). Define the symmetric regions:

$$S_k = (-w_k, -w_{k-1}) \cup (w_{k-1}, w_k) \quad \text{and} \quad (1)$$

$$S_{k+1} = (-w_{k+1}, -w_k) \cup (w_k, w_{k+1}),$$

where $0 < w_{k-1} < w_k < w_{k+1}$ (sorted weight assumption).

Suppose $w \in S_k$ and $w' \in S_{k+1}$ are represented in binary form up to b bits:

$$w = \sum_{i=0}^{b-1} a_i 2^{-i} \quad \text{and} \quad w' = \sum_{i=0}^{b-1} a'_i 2^{-i}, \quad (2)$$

where $a_i, a'_i \in \{0, 1\}$ for $i = 0, 1, \dots, b-1$.

Goal: Show that for any i with $0 \leq i \leq b-1$:

$$P(a_i = 0) > P(a'_i = 0), \quad (3)$$

where $P(a_i = 0)$ is the probability of $a_i = 0$. In crossbars, if $a_i = 0$, it means the weight will not contribute to placing an ADC at the column end.

Intuitively, because S_k is closer to zero than S_{k+1} , and the normal distribution is symmetric and decreasing as $|w|$ increases, the probability density $f(w) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-w^2/2\sigma^2}$ is higher and decreases more steeply in S_k . This causes a greater imbalance between the probabilities of $a_n = 0$ and $a_n = 1$ in S_k compared to S_{k+1} . As a result, the chance that $a_n = 0$ is greater for w in S_k than for w' in S_{k+1} .

Mathematically, For each bit position n , the bits a_0, \dots, a_{n-1} define an interval $[L, U]$ for $|w|$:

$$L = \sum_{i=0}^{n-1} a_i 2^{-i}, \quad U = L + 2^{-n}. \quad (4)$$

The bit a_n divides this interval into two subintervals:

- If $a_n = 0$, $|w| \in [L, M]$
- If $a_n = 1$, $|w| \in [M, U]$

where $M = \frac{L+U}{2}$. Equivalently, for $[L', U']$, the midpoint M' .

Since $\frac{\int_L^M f(w) dw}{\int_L^U f(w) dw}$ monotonically decreases with L , the probabilities follow:

$$P(a_n = 0) = \frac{\int_L^M f(w) dw}{\int_L^U f(w) dw} > \frac{\int_{L'}^{M'} f(w) dw}{\int_{L'}^{U'} f(w) dw} = P(a'_n = 0). \quad (5)$$

Therefore, earlier sections require fewer ADCs.

B. Qualitative Analysis of SWS

In the previous subsection, we reduced the ADC count for a fixed section size using *sorted weight sectioning* (SWS). The method consists of two offline steps. We first sort crossbar rows by their magnitudes, defining a section-specific index matching function (see \mathcal{P} in Figure 1). Then, we partition the MMM into sections ordered by increasing magnitude.

Low-magnitude weight sections only activate low-order columns, and the ADC count is determined by columns that accumulated nonzero values (see Figure 3). Thus, earlier sections are filled with zero-entry columns, which do not need data conversions. Consequently, these sections need fewer ADCs. Note that zero activations are not fetched into crossbars because they output zeros in all columns. We also do not program zero weights since they are rows filled with zeros.

This approach has the side effect of remapping isolated zeros in unstructured sparse models to early sections, which makes it a perfect use case for SWS. We leverage this hardware-friendly

sparsity mapping while having a small accuracy drop compared to structured sparsity. With SWS, sparsity reduces the number of sections, since sections full of zeros are not programmed.

We reuse the input permutation function on the fly for every inference task, ensuring index matching. This process’s memory/latency depends on the feature size, since we must allocate memory to hold and permute the fetched data while maintaining constant computation throughput.

To minimize permutation latency, each of the f inputs is connected to a dedicated multiplexer controlled by a precomputed index. This setup achieves a full permutation in one cycle. For L crossbars, the space overhead is $\mathcal{O}(f)$, and sorting to compute these indices costs $\mathcal{O}(fL \log f)$. To mitigate sorting overhead, we reuse the data buffer to save memory and crossbars to reduce latency.

C. Impact on Energy and Accuracy

The impact of SWS on the accuracy and energy consumption of the model depends on the DNN weight distribution. DNN weights follow a bell-shaped distribution with long tails [19]–[22] due to normalization layers during training. This ensures more sections of small weights than of large.

High-order columns are filled with zeros in sections with small magnitude weights, which makes these sections require fewer ADCs. Furthermore, low-order column outputs are scaled by smaller power-of-two multipliers, motivating the use of different ADC resolutions per column. We can use lower-precision conversions for columns farther from the crossbar. As a result, we can use both a smaller number of ADCs and a lower resolution to reduce energy costs without significantly degrading the DNN accuracy. We note that energy savings are more pronounced for models with sharper weight distributions around zero and longer tails – models with smaller σ will require fewer ADCs due to the higher occurrence of zero bits.

IV. EXPERIMENTS

We simulate crossbar computation in PyTorch and CiMLoop on ImageNet-1K [23], CIFAR-10 [24], and MNIST [25] datasets. We apply our method on all layers of models from PyTorch (ResNets and VGGs), timm (ViTs and DeITs patch 16 224), and transformers (BERTs) libraries trained in 32-bit floating point. In all cases, the crossbars have eight columns. We consider two scenarios for evaluation:

- **Sorted (our approach):** crossbar rows are sorted by their magnitudes. This scenario can have either a fixed resolution (ADCs with the same resolution) or an unfixed resolution.
- **Unsorted (state-of-the-art approach considered in [5], [15]):** crossbar rows are placed randomly.

A. Results

We fix the ADC resolution to 10 bits and use 128-row sectioning for various DNN models on ImageNet-1K. We reduced the number of ADCs used, thereby lowering energy consumption with SWS (see Figure 4). Because we only apply SWS fixing ADC resolutions at this point, we did not perceive




Fig. 4: ADC energy savings applying SWS with fixed ADC resolution on convolutional (red) and vision transformer (blue) models for ImageNet-1K.

TABLE I: Energy breakdown for language models in CIMLoop.

Baseline Models	Drivers + DAC (pJ/MAC)	CIM Unit (pJ/MAC)	ADC (pJ/MAC)	Total (pJ/MAC)	Parameters
MobileBERT	1.169801	0.956911	14.1433	16.27	25M
Microsoft Phi	0.3078242	0.265546	3.9607	4.5340702	1.3B
GPT-2 Medium	0.346854	0.299687	2.4511	3.097641	355M
BERT 80% Pruned	0.30253	0.02443	0.5132	0.84016	68M
BERT 85% Pruned	0.32412	0.02171	0.4588	0.80463	51M
BERT 90% Pruned	0.31555	0.01929	0.3343	0.66914	34M
Sorted Models					
MobileBERT	1.549332 (↑ 32.44%)	0.935127 (↓ 2.27%)	3.3212 (↓ 76.52%)	5.805659 (↓ 64.32%)	25M
Microsoft Phi	0.4174333 (↑ 35.61%)	0.245217 (↓ 7.65%)	0.9827 (↓ 75.19%)	1.64535 (↓ 63.71%)	1.3B
GPT-2 Medium	0.444351 (↑ 28.11%)	0.287684 (↓ 4.00%)	0.6313 (↓ 74.24%)	1.36335 (↓ 55.90%)	355M
BERT 80% Pruned	0.36512 (↑ 20.69%)	0.01832 (↓ 25.01%)	0.0733 (↓ 85.72%)	0.45674 (↓ 45.64%)	68M
BERT 85% Pruned	0.39834 (↑ 22.90%)	0.01593 (↓ 26.62%)	0.0550 (↓ 88.01%)	0.46927 (↓ 41.68%)	51M
BERT 90% Pruned	0.38848 (↑ 23.11%)	0.01311 (↓ 32.04%)	0.0352 (↓ 89.47%)	0.43679 (↓ 34.72%)	34M

a significant accuracy drop. In fact, the largest drop was 0.09% in DeIT-Tiny. SWS achieves better energy savings for CNNs because transformer models have (1) wider weight distributions (noted in Section III) and (2) relevant outliers, which makes quantization harder [22], [26]–[28]. Furthermore, we used CIMLoop to get energy values and account for the sorting hardware driver (see Table I). We highlight the reduction in CIM Unit and ADC energy consumption. This occurs with SWS because sparsity decreases the number of sections and allows ADCs to be skipped.

We further reduced energy consumption by lowering the ADC resolutions in each column, as discussed in Section III (see the unfixed resolution results in Table II). Table II presents the ADC resolution from the highest- to the lowest-order column, selected to balance energy savings and accuracy. We could specifically prune the ADC of the lowest-order column in LeNet-5 to enhance energy efficiency without compromising accuracy.

Deeper networks are more likely to achieve better energy savings since the distribution of later layers is sharper. This is evident when comparing the results for CIFAR-10 and MNIST. ImageNet required higher-resolution ADCs. Although deeper networks are expected to have better energy-saving results,

TABLE II: Energy and accuracy of various sectioning methods.

ResNet-50 (ImageNet-1K)	Accuracy	Energy Consumption	ADCs
Unsorted Sectioning	78.10%	–	[10-10-10-10-10-10-10]
Sorted Sectioning (Fixed Resolution)	78.03% (↓0.07%)	↓75.70%	[10-10-10-10-10-10-10]
Sorted Sectioning (Unfixed Resolution)	77.17% (↓0.92%)	↓81.64%	[10-10-10-10-10-9-8]
ViT-Base (ImageNet-1K)			
Unsorted Sectioning	76.21%	–	[10-10-10-10-10-10-10]
Sorted Sectioning (Fixed Resolution)	76.13% (↓0.08%)	↓74.54%	[10-10-10-10-10-10-10]
Sorted Sectioning (Unfixed Resolution)	75.08% (↓1.13%)	↓79.05%	[10-10-10-10-10-9-8]
VGG-11 (CIFAR-10)			
Unsorted Sectioning	88.78%	–	[8-8-8-8-8-8-8]
Sorted Sectioning (Fixed Resolution)	88.71% (↓0.07%)	↓73.46%	[8-8-8-8-8-8-8]
Sorted Sectioning (Unfixed Resolution)	87.98% (↓0.8%)	↓88.50%	[8-8-8-8-8-7-7-6]
LeNet-5 (MNIST)			
Unsorted Sectioning	98.34%	–	[6-6-6-6-6-6-6]
Sorted Sectioning (Fixed Resolution)	98.37% (↑0.03%)	↓14.8%	[6-6-6-6-6-6-6]
Sorted Sectioning (Unfixed Resolution)	98.36% (↑0.02%)	↓74.00%	[6-6-6-6-5-4-0]

complex datasets are more sensitive to quantization. This explains why even small modifications to ADC resolutions resulted in a 1.13% drop in accuracy on ViT-Base.

V. DISCUSSION

ISAAC [5] and CASCADE [15] used fixed ADC resolutions in their architectures. Section IV-A showed that DNNs have different active column distributions and may require less ADC resolution. In our case, reconfigurable ADCs are attached to each column. We can set the ADC precision for active columns and disable it for inactive ones.

Recent studies propose weight allocation algorithms to enhance crossbar energy efficiency. McDanel et al. [29] use term quantization to increase sparsity in crossbars, reducing ADC resolutions. Han et al. [30] maximize data reuse on convolutional neural network (CNN) layers to reduce ADC use. Huang et al. [31] manage a hybrid in/near memory system allocating weights on each system based on their sensitivity. Our work is orthogonal to these as it is (1) encoding-agnostic, (2) model-agnostic (supporting, e.g., CNNs and transformers), and (3) performed entirely within memory.

ISAAC’s flipping encoding reduces ADC resolution by one bit. However, it places sample-and-hold circuits on each crossbar column, significantly increasing latency. SWS reduces ADC resolution without considerable latency overhead in Section III. CASCADE uses pure analog accumulation with bit-sliced architecture. Sectioning is preferred for finer quantization and fewer analog nonidealities.

For future work, we can model ADC sampling frequency like in CASCADE, using lower frequencies for higher-order columns to save more energy. We can also consider non-idealities (e.g., sneak paths) when designing efficient weight allocations.

VI. CONCLUSION

We showed how sorting and sectioning pretrained DNN weight for CIM crossbars yields significant savings on ADC energy consumption. The approach succeeds by leveraging bell-shaped DNN weight distributions (see Figure 3) and unstructured sparsity. To the best of our knowledge, we are the first to observe this.

Our strategy applies to any computation formulated as matrix multiplication. Its effectiveness depends on the application’s tolerance to approximate computing and how values are distributed: the sharper the distribution, the fewer ADCs are used (and the lower the ADC resolution). For larger models, we can (1) increase weight bitwidths by adding more columns to the crossbar and (2) enhance output precision by raising ADC resolution for each column output. This enables flexible trading of accuracy for energy savings. SWS reduces ADC energy use by 89.5% in pruned BERT models.

The idea is simple to implement as it merely involves proper weight placements on the crossbar. We have provided end-to-end validation of these results via simulations for DNNs on well-known datasets in the literature. This paper suggests a fruitful new direction for future research on CIM implementations for DNNs based on sorted weight sectioning.

REFERENCES

- [1] M. Horowitz, "1.1 Computing's energy problem (and what we can do about it)," in *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, pp. 10–14, 2014.
- [2] X. Huang, C. Liu, Y.-G. Jiang, and P. Zhou, "In-memory computing to break the memory wall," *Chinese Physics B*, vol. 29, p. 078504, jul 2020.
- [3] M. E. Fouda, H. E. Yantir, A. M. Eltawil, and F. Kurdahi, "In-Memory Associative Processors: Tutorial, Potential, and Challenges," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 69, no. 6, pp. 2641–2647, 2022.
- [4] I. Chakraborty, M. Ali, A. Ankit, S. Jain, S. Roy, S. Sridharan, A. Agrawal, A. Raghunathan, and K. Roy, "Resistive Crossbars as Approximate Hardware Building Blocks for Machine Learning: Opportunities and Challenges," *Proceedings of the IEEE*, vol. 108, no. 12, pp. 2276–2310, 2020.
- [5] A. Shafiee, A. Nag, N. Muralimanohar, R. Balasubramonian, J. P. Strachan, M. Hu, R. S. Williams, and V. Srikumar, "ISAAC: A Convolutional Neural Network Accelerator with In-Situ Analog Arithmetic in Crossbars," in *2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA)*, pp. 14–26, 2016.
- [6] S. Jung, H. Lee, S. Myung, H. Kim, S. K. Yoon, S.-W. Kwon, Y. Ju, M. Kim, W. Yi, S. Han, B. Kwon, B. Seo, K. Lee, G.-H. Koh, K. Lee, Y. Song, C. Choi, D. Ham, and S. J. Kim, "A crossbar array of magnetoresistive memory devices for in-memory computing," *Nature*, vol. 601, no. 7892, pp. 211–216, 2022.
- [7] A. Sebastian, M. Le Gallo, R. Khaddam-Aljameh, and E. Eleftheriou, "Memory devices and applications for in-memory computing," *Nature Nanotechnology*, vol. 15, no. 7, pp. 529–544, 2020.
- [8] Q. Huo, Y. Yang, Y. Wang, D. Lei, X. Fu, Q. Ren, X. Xu, Q. Luo, G. Xing, C. Chen, X. Si, H. Wu, Y. Yuan, Q. Li, X. Li, X. Wang, M.-F. Chang, F. Zhang, and M. Liu, "A computing-in-memory macro based on three-dimensional resistive random-access memory," *Nature Electronics*, vol. 5, no. 7, pp. 469–477, 2022.
- [9] B. Li, L. Xia, P. Gu, Y. Wang, and H. Yang, "MERging the Interface: Power, area and accuracy co-optimization for RRAM crossbar-based mixed-signal computing system," in *2015 52nd ACM/EDAC/IEEE Design Automation Conference (DAC)*, pp. 1–6, 2015.
- [10] D. Blalock, J. J. G. Ortiz, J. Frankle, and J. Guttag, "What is the State of Neural Network Pruning?," *Proceedings of Machine Learning and Systems (MLSys)*, vol. 2, pp. 129–146, 2020.
- [11] T. Andrulis, J. S. Emer, and V. Sze, "CiMLoop: A Flexible, Accurate, and Fast Compute-In-Memory Modeling Tool," in *2024 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, 2024.
- [12] L. Wang, C. Yang, J. Wen, S. Gai, and Y. Peng, "Overview of emerging memristor families from resistive memristor to spintronic memristor," *Journal of Materials Science: Materials in Electronics*, vol. 26, pp. 4618–4628, Jul 2015.
- [13] Y. Liao, H. Wu, W. Wan, W. Zhang, B. Gao, H.-S. Philip Wong, and H. Qian, "Novel In-Memory Matrix-Matrix Multiplication with Resistive Cross-Point Arrays," in *2018 IEEE Symposium on VLSI Technology*, pp. 31–32, 2018.
- [14] M. Zangeneh and A. Joshi, "Design and Optimization of Nonvolatile Multibit 1T1R Resistive RAM," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 22, no. 8, pp. 1815–1828, 2014.
- [15] T. Chou, W. Tang, J. Botimer, and Z. Zhang, "CASCADE: Connecting RRAMs to Extend Analog Dataflow In An End-To-End In-Memory Processing Paradigm," in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture, MICRO '52*, p. 114–125, 2019.
- [16] F. Cai, J. M. Correll, S. H. Lee, Y. Lim, V. Bothra, Z. Zhang, M. P. Flynn, and W. D. Lu, "A fully integrated reprogrammable memristor-CMOS system for efficient multiply-accumulate operations," *Nature Electronics*, vol. 2, pp. 290–299, Jul 2019.
- [17] X. Sun, S. Yin, X. Peng, R. Liu, J.-s. Seo, and S. Yu, "XNOR-RRAM: A scalable and parallel resistive synaptic architecture for binary neural networks," in *2018 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pp. 1423–1428, 2018.
- [18] M. A. Zidan, H. A. H. Fahmy, M. M. Hussain, and K. N. Salama, "Memristor-based memory: The sneak paths problem and solutions," *Microelectronics Journal*, vol. 44, no. 2, pp. 176–183, 2013.
- [19] S. Han, H. Mao, and W. J. Dally, "Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding," in *International Conference in Learning Representations (ICLR)*, 2016.
- [20] J. Fang, A. Shafiee, H. Abdel-Aziz, D. Thorsley, G. Georgiadis, and J. Hassoun, "Post-Training Piecewise Linear Quantization for Deep Neural Networks," in *The European Conference on Computer Vision (ECCV)*, 2020.
- [21] M. Horton, Y. Jin, A. Farhadi, and M. Rastegari, "Layer-Wise Data-Free CNN Compression," in *International Conference on Pattern Recognition (ICPR)*, 2022.
- [22] T. Tambe, E.-Y. Yang, Z. Wan, Y. Deng, V. Janapa Reddi, A. Rush, D. Brooks, and G.-Y. Wei, "Algorithm-Hardware Co-Design of Adaptive Floating-Point Encodings for Resilient Deep Learning Inference," in *2020 57th ACM/IEEE Design Automation Conference (DAC)*, pp. 1–6, 2020.
- [23] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, 2009.
- [24] A. Krizhevsky and G. Hinton, "Learning Multiple Layers of Features from Tiny Images," 2009.
- [25] L. Deng, "The MNIST database of handwritten digit images for machine learning research," *IEEE Signal Processing Magazine*, vol. 29, no. 6, pp. 141–142, 2012.
- [26] Z. Yao, R. Y. Aminabadi, M. Zhang, X. Wu, C. Li, and Y. He, "ZeroQuant: Efficient and Affordable Post-Training Quantization for Large-Scale Transformers," in *Advances in Neural Information Processing Systems (S. Koyejo et al., eds.)*, vol. 35, pp. 27168–27183, Curran Associates, Inc., 2022.
- [27] G. Xiao, J. Lin, M. Seznec, H. Wu, J. Demouth, and S. Han, "SmoothQuant: Accurate and Efficient Post-Training Quantization for Large Language Models," in *Proceedings of the 40th International Conference on Machine Learning*, 2023.
- [28] Y. Bondarenko, M. Nagel, and T. Blankevoort, "Quantizable Transformers: Removing Outliers by Helping Attention Heads Do Nothing," in *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [29] B. McDanel, S. Q. Zhang, and H. T. Kung, "Saturation RRAM Leveraging Bit-Level Sparsity Resulting from Term Quantization," in *2021 IEEE International Symposium on Circuits and Systems (ISCAS)*, pp. 1–5, 2021.
- [30] L. Han, P. Huang, Z. Zhou, Y. Chen, X. Liu, and J. Kang, "A Convolution Neural Network Accelerator Design with Weight Mapping and Pipeline Optimization," in *2023 60th ACM/IEEE Design Automation Conference (DAC)*, pp. 1–6, 2023.
- [31] C.-T. Huang, C.-Y. Chang, Y.-C. Chuang, and A.-Y. A. Wu, "BWA-NIMC: Budget-based Workload Allocation for Hybrid Near/In-Memory-Computing," in *2023 60th ACM/IEEE Design Automation Conference (DAC)*, pp. 1–6, 2023.