

High Area/Energy Efficiency RRAM CNN Accelerator with Pattern-Pruning-Based Weight Mapping Scheme

Songming Yu, Lu Zhang, Jingyu Wang, Jinshan Yue, Zhuqing Yuan,
Xueqing Li, Huazhong Yang, Yongpan Liu

Department of Electronic Engineering, Tsinghua University, Beijing, China
e-mail: ypliu@tsinghua.edu.cn

Abstract—Resistive random access memory (RRAM) is an emerging device for processing-in-memory (PIM) architecture to accelerate convolutional neural network (CNN). However, due to the highly coupled crossbar structure in the RRAM array, it is difficult to exploit the CNN sparsity feature to improve the performance in RRAM-based CNN accelerator. To optimize the weight mapping of sparse network in the RRAM array and improve area and energy efficiency, we propose a novel weight mapping scheme and corresponding RRAM-based CNN accelerator architecture based on pattern pruning and operation unit(OU) mechanism. Experimental results show that our work can achieve 4.16x-5.20x crossbar area efficiency, 1.98x-2.15x energy efficiency, and 1.15x-1.35x performance speedup in comparison with the traditional method.

I. INTRODUCTION

Deep convolutional neural networks have been widely used in many fields, such as face recognition, pose estimation, and image classification. CNN models have achieved better accuracy than humans in several cases [1]. However, as the sizes of modern CNN models becoming larger, every inference operation contains massive matrix-vector multiplications. The computational complexity and storage overhead of CNN models severely limit their application, such as on low-power embedded platform.

To overcome these drawbacks, dedicated hardware accelerator architectures have been proposed to accelerate the neural network. Among them, RRAM-based PIM accelerator [2]–[4] is one of the most promising architectures. It uses RRAM array to perform arithmetic operations as well as data storage. This kind of accelerator can achieve high performance and energy efficiency, because RRAM crossbar arrays can perform matrix-vector multiplication, which is the most time and energy consuming operation in CNN, by means of high degree of parallelism and high energy efficiency.

The RRAM has shown its great potential for accelerating the computation of the neural networks, but the architectural exploration in the context of different application scenarios is still in early stage. For example, the sparsity of the network has not been fully studied and optimized in RRAM-based CNN accelerator. Previous works [5], [6] have reported that there is significant weight redundancy in neural networks. By properly designed fine-grained pruning algorithms, more than 85% redundant weights in networks can be pruned with little accuracy loss [7]. However, this kind of pruning algorithm makes the network structure irregular, which severely affects the computing efficiency in RRAM-based architecture. Due

to the highly coupled RRAM crossbar structure, it is hard to optimize the weight mapping of irregular sparse networks to improve the area efficiency. On the contrary, structured pruning, which is hardware-friendly for RRAM-based architecture, causes more significant accuracy loss than the fine-grained pruning with the same weight sparsity [8].

Recently, researches [9], [10] proposed a new pruning algorithm—pattern pruning. Pattern pruning can be considered as an intermediate type between non-structured and structured pruning. It can achieve both high accuracy and high regularity level. So, it is a very promising method which contains the advantages of both non-structured and structured pruning. We find pattern pruning has the potential to be combined with RRAM-based CNN accelerator to achieve high energy and area efficiency, which, to our knowledge, has not been studied by other researchers before. Thus, we develop a novel weight mapping scheme for sparse network based on pattern pruning. Besides, paper [11] mentioned that due to the hardware limitation, the computation in the RRAM crossbar must be executed in smaller granularity, called operation unit(OU), instead of activating an entire crossbar array in one cycle. This mechanism also provides more possibility to optimize the weight mapping of sparse networks in RRAM array.

Our contributions mainly include:

- 1) A novel area-efficiency weight mapping scheme based on pattern pruning and OU mechanism to exploit the network weight sparsity in RRAM-based PIM architecture.
- 2) An RRAM-based energy-efficiency sparse CNN accelerator architecture supporting our novel mapping scheme.
- 3) Exploit both weight and activation sparsity to achieve high energy efficiency in our architecture.

Experimental results show that in our pattern-pruning-based mapping scheme, we can achieve 4.67x-5.20x crossbar area efficiency with almost no accuracy loss, which means we save 76.0%-80.8% crossbar area than the baseline mapping algorithm. Besides, we achieve a 1.98x-2.15x energy efficiency and 1.15x-1.35x performance speedup in our CNN accelerator design by utilizing both weight mapping scheme and activation sparsity.

II. BACKGROUND AND MOTIVATION

A. Weight Mapping in RRAM-based CNN Accelerator

In RRAM-based CNN accelerator, RRAM crossbars are used to perform matrix-vector multiplication, which is the most energy and time consuming operation in convolutional neural network. Fig. 1 shows a baseline weight mapping

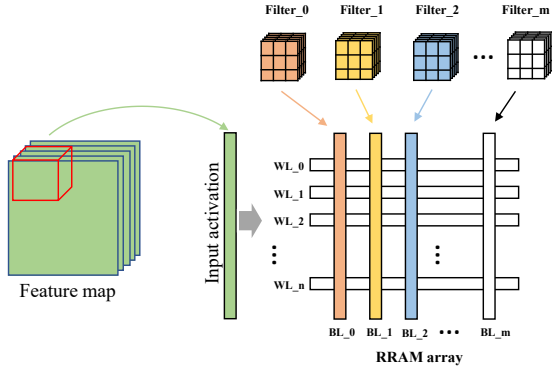


Fig. 1. A straightforward weight mapping scheme(as a motivation example)

method for convolution layer weights. In previous over-idealized RRAM-based CNN accelerator designs [2]–[4], the whole crossbar is assumed to be activated in one clock cycle. And all the weights in a filter are mapped to one column in the crossbar. When executing the matrix-vector multiplication, the input activations are converted to the voltage signals through DACs and fed into each wordline(WL) of the crossbar, then currents in each RRAM cell are accumulated in the bitlines. Finally, the output current signals are collected at the end of each bitline(BL) and converted to digital signals by ADCs.

However, this kind of mapping method and accelerator design makes it hard to exploit the weight sparsity in RRAM-based CNN accelerator. In this method, all the weights in the same row of the RRAM array will be multiplied with the same input, and the multiplication results of each RRAM cell in the same column will accumulate to the same output. So the position of the weights in the crossbar can't be easily changed. If a weight is zero, it still needs to occupy an RRAM cell. Only when all weights in a wordline or bitline are zeros can we remove these zero weights and save those wordlines/bitlines.

Furthermore, because of the conductance deviation per cell and the limitation of the ADC resources, the activated wordline and bitline number in one cycle is limited. Only a small block of the RRAM crossbar, called operation unit [11], can be executed per cycle. For example, in a recent state-of-the-art RRAM-based CNN accelerator design [12], only nine wordlines and eight bitlines can be activated in a cycle. This provides more possibility for us to exploit the weight sparsity of the convolutional neural network.

B. Previous Works on Sparse CNN Accelerator

Several studies have already realized the problem and proposed their solutions. [13] was aware of the problem of irregular sparsity, so they turned to the regular sparsity and apply regularization on filter and channel dimension to obtain structured sparse network which can be directly mapped to the RRAM crossbar. But this paper does not mention the data about how much RRAM crossbar they save in their architecture. [11] makes use of the feature we mentioned above that only a small block, called operation unit, can be executed per cycle. So they can exploit the weight sparsity in a small granularity. Compared to [13], [11] achieves higher performance speed and energy saving. But it does not mention the crossbar they saved either. [14] proposed a k-means



Fig. 2. From kernel to pattern

clustering which shuffles the column in the weights matrix to gather the zero weights and use the crossbar-grained pruning algorithm to prune the all-zero crossbars. However, only 6% to 22% of crossbar resources can be saved in this algorithm, which is still not efficient enough.

Recently, [9], [10] mentioned about a new kind of convolutional neural network pruning algorithm: pattern pruning. As Fig. 2 shows, pattern means a shape of kernels, and is defined as a boolean mask that indicates whether the weights are nonzero in each position. Taking the common 3×3 kernels as an example, in an irregular pruned network, weight in any position of the kernel could be zero, so the theoretical maximum pattern numbers is $2^{3 \times 3} = 512$. By following proper pruning algorithm, we can limit the total pattern numbers to a very low level, such as less than 8 patterns in each network layer, with little accuracy and sparsity loss. By pattern pruning, we can make the irregular sparse network regular in kernel dimension. Pattern pruning can achieve both accuracy and regularity at a higher level. So we develop a novel weight mapping scheme based on pattern pruning to optimize the mapping of the sparse network on the RRAM array and achieve high RRAM area efficiency and energy efficiency.

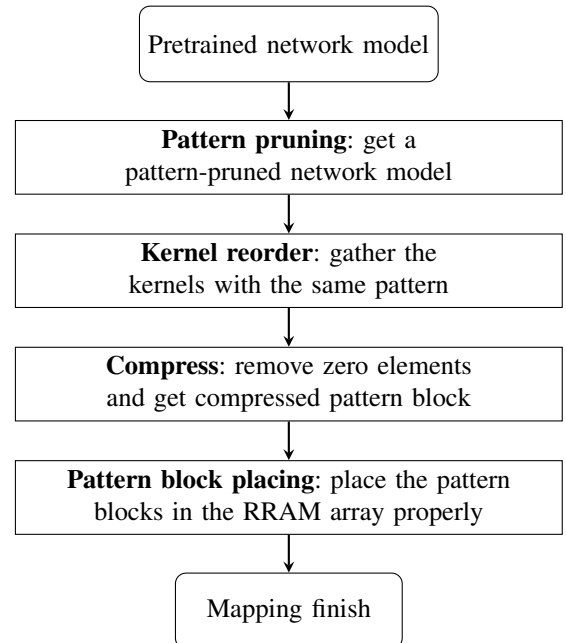


Fig. 3. The flowchart of our pattern-pruning-based weight mapping scheme

III. PATTERN-PRUNING-BASED WEIGHT MAPPING SCHEME

The flowchart in Fig. 3 shows an overview of our mapping scheme. We will explain each step in more detail in the following part of this section.

A. Pattern Pruning

Before we map the weights on the RRAM crossbars, we need to apply the pattern pruning algorithm to train the pattern-pruned network.

Here, we use an alternating direction method of multiplier (ADMM)-based pattern compression method [10]. For more detail, you can refer to the paper [10]. Experimental results shows that ADMM-based pruning method can achieve higher sparsity with less accuracy loss compared to other heuristic pruning algorithms [6].

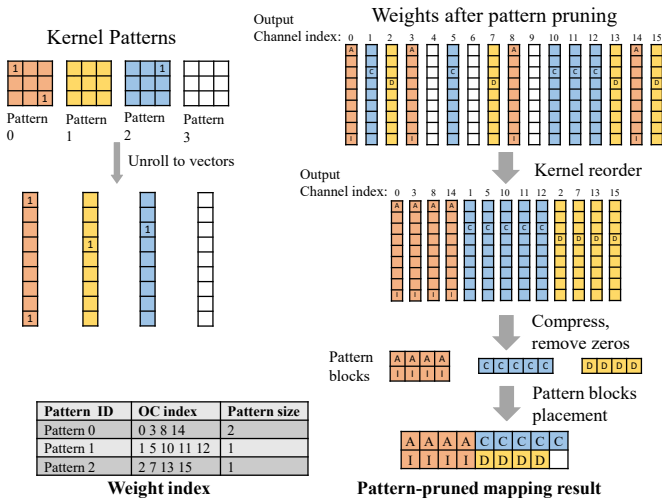


Fig. 4. A case study of our pattern-pruning-based mapping method. On the lower left is the weight index.

B. Weight Mapping Algorithm

After we get the pattern-pruned network, we can perform the weight mapping algorithm. Fig. 4 shows our mapping algorithm workflow. We take a small layer with only one input channel and 16 output channels for example. After pattern pruning, all 16 3×3 kernels have only four patterns, including an all-zero pattern. We unroll the patterns and kernels to one dimension vectors and mark different patterns in different colors. Firstly, we reorder the kernels according to the pattern types and gather the kernels with the same pattern. Then, we can compress the kernels by removing all the zero elements. After the compression, the adjacent kernels with the same pattern form a pattern block. Inside each pattern block, the matrix-vector multiplication can be computed in parallel due to the fact that the weights in the same position of original kernels are still in the same row in the crossbar. Finally, we place the pattern blocks on the crossbar following a proper strategy. We will explain this strategy in the next paragraph. In previous weight mapping method, the weights matrix is directly mapped to the RRAM crossbar and all those weights (16 kernels with the size of 3×3) will take up a 9×16

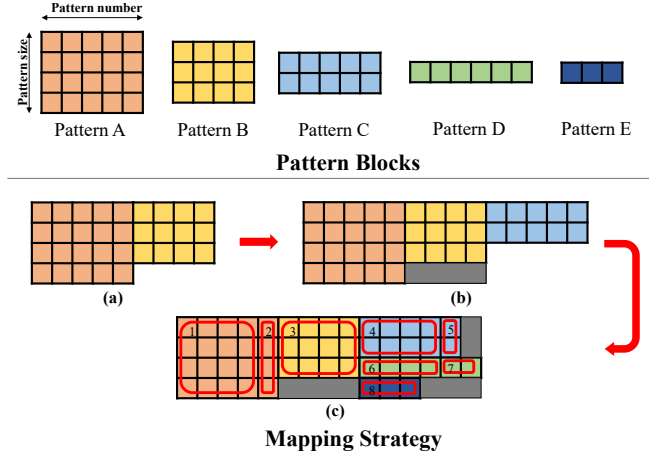


Fig. 5. A more specific example to explain our mapping strategy. The red boxes means the operation unit(OU).

crossbar array. However, we optimize the pattern-pruning-based mapping scheme by storing all the weights in a 2×9 crossbar array.

To explain how we place each pattern block on the crossbar (the last step in Fig. 4), we use a more specific example. As shown in Fig. 5, after getting each pattern block, we reorder all the blocks according to the pattern size (the number of nonzero elements in that pattern). First, we place the pattern block with the biggest pattern size, and put the next pattern block on the left, aligning it to the top of the former block. Then, if the number of rows behind the current block is enough for the next block, we can place it there and align it left. Otherwise, we place the next block on the left and is also aligned to the top of the former block. There is only one row left behind the current block, as shown in Fig. 5(a), which is not enough for the next pattern block with a pattern size of two, so the next pattern block is placed in new columns. And one row marked in grey colors is wasted, as shown in Fig. 5(b). The next two blocks with only one row can be placed behind the former block and are left-aligned, and a little more area marked in grey is wasted. And the red boxes in Fig. 5(c) show the OU organization for 4×4 OU size.

For a practical network layer with more than one input channel, we apply all those operations for every input channel and store all the weights channel by channel. We also explain our algorithm by pseudocode in Algorithm 1.

Besides, we need to store the indexes of the kernels, since we have reordered the kernels. We store the indexes pattern by pattern in the same order as mapping the pattern blocks to the crossbar, and for each pattern, we store the corresponding output channel index of each kernel and the pattern shape (including pattern size). We will explain how we can get the placement information of the weights from the indexes in next section. And the index overhead will be analyzed in section V.

IV. ARCHITECTURE

In our mapping scheme, the weights in the RRAM crossbar are compressed and no longer in sequential order. So it cannot be deployed in traditional RRAM-based CNN accelerator without hardware design modification. Referencing former

Algorithm 1 Weight mapping algorithm pseudocode

```
1:  $height\_free \leftarrow \max \text{ pattern length}$ 
2:  $i, j, max\_num, i\_pre \leftarrow 0$ 
3: for each  $in\_channel \in In\_channels$  do
4:   for each  $pattern \in Pattern \text{ Set}$  do
5:      $L \leftarrow \text{current pattern length}$ 
6:      $N \leftarrow \text{current pattern numbers}$ 
7:     if  $L < height\_free$  then
8:        $Array[i : i + L][j : j + N] \leftarrow patt\_block$ 
9:        $height\_free \leftarrow height\_free - L$ 
10:       $max\_num \leftarrow \max(max\_num, N)$ 
11:       $i \leftarrow i + L$ 
12:     else
13:        $i \leftarrow i\_pre$ 
14:        $height\_free \leftarrow \max \text{ pattern length} - L$ 
15:        $j \leftarrow j + max\_num$ 
16:       if  $j \geq \text{array width}$  then
17:          $j \leftarrow 0$ 
18:          $i, i\_pre \leftarrow i\_pre + max\_patt\_length$ 
19:       end if
20:        $Array[i : i + L][j : j + N] \leftarrow patt\_block$ 
21:        $max\_num \leftarrow 0$ 
22:     end if
23:   end for
24: end for
```

RRAM-based CNN accelerator architecture design [2], [11], [15], we design a new architecture to support our mapping scheme. Fig. 6 shows our architecture design, and the red arrow shows the dataflow from input to output. The weights are mapped to the RRAM crossbar by using our mapping scheme, and the indexes are stored in the weight index buffer. The main difference between the former designs [3], [14] and our architecture are explained in detail as follows.

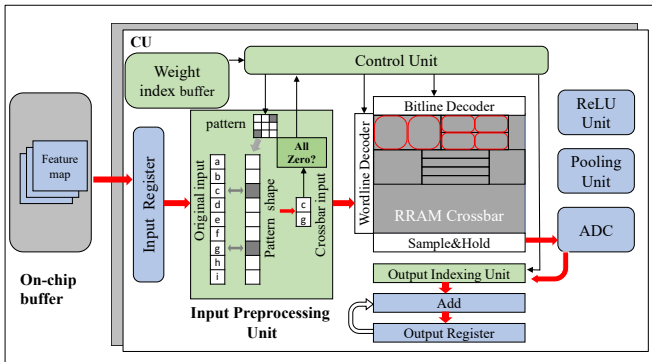


Fig. 6. An overview of our RRAM-based CNN accelerator architecture. The red arrow shows the dataflow from input to output.

A. Input Preprocessing Unit

In the RRAM crossbar, we only store the nonzero weights. For 3*3 kernel, only nonzero elements are stored together. When we send the inputs to crossbar, we only send the input activations corresponding to the nonzero weights. So we need to select the correct inputs according to the pattern of the current weights. The Input Preprocessing Unit is designed to

implement this function. It will get the pattern information from the control unit to send the inputs to the RRAM crossbar.

We also notice that because of the ReLU activation function, there is also considerable sparsity in the input activations. If we can utilize it, we can make further improvement in energy efficiency. So we add an all-zero detection module in the Input Preprocessing Unit. If the inputs are all zeros, a signal is sent to the control unit and all the operations will not be done to avoid useless computation and save energy.

B. Output Indexing Unit

In our mapping method, the weight kernels in the RRAM crossbar are no longer sequentially stored in each bitline of the crossbar, as we have explained in previous section. And the outputs collected in each bitline are also out of sequence. So before we store the outputs into the output register, we need to reorder those outputs. In the output indexing unit, the outputs collected in every cycle are reordered and stored into the right address according to the indexes stored in the weight index buffer.

C. Operation Unit Organization

As we have mentioned before, in every cycle we can only activate a small block of the RRAM crossbar (called operation unit) to perform the computation, instead of activating the whole crossbar. In the pattern-pruning-based weight mapping scheme, every operation unit must be limited inside a pattern block, because for different patterns, the weights stored in the same wordline correspond to different inputs, and cannot be computed in parallel. In Fig. 5(c), the red boxes show an example of OU Organization for an OU size of 4×4 .

Another problem is how to get the placement information of the weights in the crossbar from the indexes. Actually, the procedures are similar to the mapping strategy shown in Fig. 5. First, we get the index of the pattern with the biggest pattern size, the width of this pattern block is the number of input channels of this pattern, and the height of this pattern block is the pattern size. The next pattern block is next to the current block, since this is how the blocks are placed. The width and the height of this pattern can be got from the input channel indexes and the pattern size. Then we get the pattern size of the next pattern block. If there are enough rows behind the current block for next block, then we know it is placed there, otherwise we know that it will be placed in new columns. Repeat those steps until we get all the weights' placement.

V. EVALUATION AND RESULTS

A. Evaluation Setup

In the evaluation section, a simulator is built in Python to implement our mapping algorithm and simulate the weight mapping workflow in the crossbars and computation workflow to get the crossbar area and energy efficiency, as well as the computation speedup. For comparison, we use the mapping method in Fig. 1 as the baseline.

For hardware energy model settings, according to ISAAC [3], RRAM related components (crossbars, ADCs, and DACs) consume more than 80% energy of the total chip, so we focus on those components when evaluating the energy efficiency. Table I shows our hardware configuration. For ADCs and

DACs energy, we use the data from [16]. And the RRAM crossbar array energy model is based on [17]. The operation unit size is set to 9×8 , the same as [12], which means that we can activate up to 9 wordlines and 8 bitlines per cycle.

TABLE I
HARDWARE PARAMETERS

Components	Parameters	Spec	energy
ADC	Precision	8 bits	1.67 pJ/op
	Frequency	1.2 GSps	
DAC	Precision	4 bits	0.0182 pJ/op
	Frequency	18 MSps	
RRAM Array	OU size	9×8	4.8 pJ/OU/op
	bits per cell	4	
	size	512×512	

We use a modified VGG16 network as our benchmark. The convolution layers in our network are the same as [18], but our network only contains one full-connected layer. By modification, we significantly reduced the parameters in FC layers, so we can focus on the results of the convolution layers. The datasets we use include CIFAR-10, CIFAR-100 [19] and ImageNet [20].

B. Pattern Pruning Result

The baseline VGG16 networks are trained on CIFAR-10, CIFAR-100 and ImageNet, and are irregularly pruned. The network trained on CIFAR-10 is of 81.95% sparsity, 419 mean pattern numbers per layer and 91.72% top-1 accuracy. The network trained on CIFAR-100 is of 81.95% sparsity, 450 mean pattern numbers per layer and 72.72%/91.44% top-1/top-5 accuracy. The networks trained on ImageNet is of 83.38% sparsity, 461 mean pattern numbers per layer and 71.90%/90.49% top-1/top-5 accuracy.

By pattern pruning, we can achieve 2-12 patterns per convolution layer, and more than 80% sparsity in convolution layers. Table II shows the pattern pruning results. After pattern pruning, the sparsity of the networks trained on CIFAR-10 and CIFAR-100 is even higher than the baseline networks, with little or no accuracy loss. And the pattern numbers in each layer are no more than 12, while on CIFAR-10 and CIFAR-100 the numbers are no more than 8, which makes the networks more structured.

C. Evaluation Results

Crossbar Area Efficiency: Fig. 7 shows the results of RRAM crossbar area efficiency. In our pattern pruned mapping algorithm, we achieve an area efficiency improvement of 4.67x/5.20x/4.16x for networks trained on CIFAR-10, CIFAR-100, and ImageNet, respectively, which means we save 78.5%/80.8%/76.0% RRAM crossbar array comparing to the baseline mapping method. This is very close to the theoretical best results (86.03%/85.23%/82.48%, the sparsity of the networks). This means our mapping algorithm has utilized most of the sparsity of the networks. The sparsity of the network trained on ImageNet is lower than the other two, and the pattern numbers are relatively higher, which means the network structure is more irregular. Lower sparsity and




Fig. 7. The results of RRAM crossbar array area efficiency on different datasets.

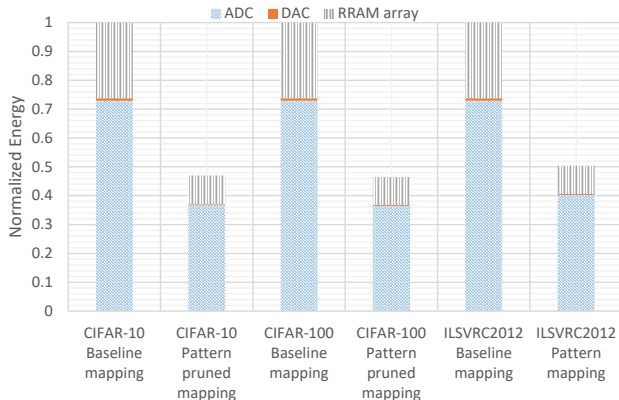


Fig. 8. The results of normalized energy on different datasets. The energy data are all normalized to the baseline results.

more irregular structure make the mapping efficiency lower than others.

Energy Efficiency: By pattern-pruning-based mapping algorithm, much fewer RRAM crossbar arrays are used, and in every cycle, less bitlines and wordlines, as well as the ADCs and DACs, are activated because of the pattern pruned compression, so the energy efficiency is also higher than the baseline mapping algorithm. Besides, all-zero detection module in the Input Preprocessing Unit also makes important contribution to the improvement of energy efficiency. Fig. 8 shows the energy efficiency results. We can see that the ADC energy is the main bottleneck. We achieve 2.13x/2.15x/1.98x energy efficiency on CIFAR-10, CIFAR-100, and ImageNet, respectively.

Performance Speedup: The speedup is achieved mainly by the deleted all-zero patterns which are neither stored in crossbars nor computed. Though the speedup ratio is relatively small, only 1.35x/1.15x/1.17x on CIFAR-10, CIFAR-100 and ImageNet respectively, it is acceptable for us since we achieved very high crossbar area efficiency and energy efficiency.

D. Index Overhead Analysis

As we have mentioned before, we need to store the kernels' output channel indexes because we reorder the kernels inside every input channel. We also store the pattern shapes for each

TABLE II
PATTERN PRUNING RESULTS

Dataset	Sparsity	Pattern Numbers in Each Conv layer	Total	top-1	top-5
CIFAR-10	86.03%(+4.08%)	[2, 2, 2, 6, 8, 8, 8, 6, 5, 4, 6, 6, 8]	71	92.63%(-0.09%)	/
CIFAR-100	85.23%(+3.28%)	[2, 2, 2, 2, 2, 8, 8, 8, 5, 6, 7, 6, 8]	66	72.73%(+0.01%)	92.23%(+0.79%)
ImageNet	82.48%(-0.90%)	[2, 2, 2, 2, 2, 9, 12, 12, 9, 10, 6, 4, 4]	76	71.15%(-0.75%)	89.98%(-0.51%)

layer, but this overhead can be ignored compared to the kernel indexes. For every kernel stored in the crossbars, we need an output channel index with no more than 9 bits (for 512 output channels).

In our evaluation, the total index overhead of the networks trained on CIFAR-10/CIFAR-100/ImageNet is 729.5KB/1013.5KB/990.6KB, respectively. The main factor that influences the size of the index overhead is the all-zero pattern ratio in each network. And in our results, the all-zero pattern ratio in each network is 40.9%/27.4%/28.5%. In our mapping method, all-zero patterns will not be stored in the crossbars, so their indexes will also be saved. Compared to the total network size, the overhead of the indexes is totally acceptable. For example, the size of the network trained on CIFAR-10 is 28.1MB before pruning, 6.0MB after pattern-pruning-based mapping(16 bits per weight), so the index overhead is only 12.2% of the network model size.

VI. CONCLUSIONS

RRAM is an emerging device for PIM architecture and has shown its great potential for accelerating the neural networks. To exploit the CNN sparsity in the RRAM-based accelerator, in this paper, we proposed a novel weight mapping scheme based on pattern pruning and a corresponding CNN accelerator architecture design to support the pattern-pruning-based weight mapping scheme. The results of our experiment show that in Our pattern-pruning-based mapping scheme, we can achieve 4.16x-5.20x crossbar area efficiency based on the pattern pruning results with almost no accuracy loss, which means we save 76.0%-80.8% crossbar area than the baseline mapping algorithm. And in our CNN accelerator design, we achieve a 1.98x-2.15x energy efficiency and 1.15x-1.35x performance speedup.

REFERENCES

- [1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [2] P. Chi, S. Li, C. Xu, T. Zhang, J. Zhao, Y. Liu, Y. Wang, and Y. Xie, "Prime: A novel processing-in-memory architecture for neural network computation in reram-based main memory," in *Proceedings of the 43rd International Symposium on Computer Architecture, ISCA '16*, p. 27–39, IEEE Press, 2016.
- [3] 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 *Proceedings of the 43rd International Symposium on Computer Architecture, ISCA '16*, p. 14–26, IEEE Press, 2016.
- [4] L. Song, X. Qian, H. Li, and Y. Chen, "Pipelayer: A pipelined reram-based accelerator for deep learning," in *2017 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pp. 541–552, Feb 2017.
- [5] Y. He, X. Zhang, and J. Sun, "Channel pruning for accelerating very deep neural networks," in *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 1398–1406, Oct 2017.
- [6] A. Ren, T. Zhang, S. Ye, J. Li, W. Xu, X. Qian, X. Lin, and Y. Wang, "Admm-nn: An algorithm-hardware co-design framework of dnns using alternating direction methods of multipliers," in *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '19*, (New York, NY, USA), p. 925–938, Association for Computing Machinery, 2019.
- [7] T. Zhang, S. Ye, K. Zhang, X. Ma, N. Liu, L. Zhang, J. Tang, K. Ma, X. Lin, M. Fardad, *et al.*, "Structadmm: A systematic, high-efficiency framework of structured weight pruning for dnns," *arXiv preprint arXiv:1807.11091*, 2018.
- [8] H. Mao, S. Han, J. Pool, W. Li, X. Liu, Y. Wang, and W. J. Dally, "Exploring the regularity of sparse structure in convolutional neural networks," *arXiv preprint arXiv:1705.08922*, 2017.
- [9] X. Ma, F.-M. Guo, W. Niu, X. Lin, J. Tang, K. Ma, B. Ren, and Y. Wang, "Pconv: The missing but desirable sparsity in dnn weight pruning for real-time execution on mobile devices," *arXiv preprint arXiv:1909.05073*, 2019.
- [10] J. Wang, S. Yu, J. Yue, Z. Yuan, Z. Yuan, H. Yang, X. Li, and Y. Liu, "High pe utilization cnn accelerator with channel fusion supporting pattern-compressed sparse neural networks," in *2020 57th ACM/IEEE Design Automation Conference (DAC)*, pp. 1–6, 2020.
- [11] T.-H. Yang, H.-Y. Cheng, C.-L. Yang, I.-C. Tseng, H.-W. Hu, H.-S. Chang, and H.-P. Li, "Sparse reram engine: Joint exploration of activation and weight sparsity in compressed neural networks," in *Proceedings of the 46th International Symposium on Computer Architecture, ISCA '19*, (New York, NY, USA), p. 236–249, Association for Computing Machinery, 2019.
- [12] W. Chen, K. Li, W. Lin, K. Hsu, P. Li, C. Yang, C. Xue, E. Yang, Y. Chen, Y. Chang, T. Hsu, Y. King, C. Lin, R. Liu, C. Hsieh, K. Tang, and M. Chang, "A 65nm 1mb nonvolatile computing-in-memory reram macro with sub-16ns multiply-and-accumulate for binary dnn ai edge processors," in *2018 IEEE International Solid - State Circuits Conference - (ISSCC)*, pp. 494–496, Feb 2018.
- [13] H. Ji, L. Song, L. Jiang, H. Li, and Y. Chen, "Recom: An efficient resistive accelerator for compressed deep neural networks," in *2018 Design, Automation Test in Europe Conference Exhibition (DATE)*, pp. 237–240, March 2018.
- [14] J. Lin, Z. Zhu, Y. Wang, and Y. Xie, "Learning the sparsity for reram: Mapping and pruning sparse neural network for reram based accelerator," in *Proceedings of the 24th Asia and South Pacific Design Automation Conference, ASPDAC '19*, (New York, NY, USA), p. 639–644, Association for Computing Machinery, 2019.
- [15] W. Zhang, X. Peng, H. Wu, B. Gao, H. He, Y. Zhang, S. Yu, and H. Qian, "Design guidelines of rram based neural-processing-unit: A joint device-circuit-algorithm analysis," in *Proceedings of the 56th Annual Design Automation Conference 2019, DAC '19*, (New York, NY, USA), Association for Computing Machinery, 2019.
- [16] J. Yue, Y. Liu, F. Su, S. Li, Z. Yuan, Z. Wang, W. Sun, X. Li, and H. Yang, "Aeris: Area/energy-efficient 1t2r reram based processing-in-memory neural network system-on-a-chip," in *Proceedings of the 24th Asia and South Pacific Design Automation Conference, ASPDAC '19*, (New York, NY, USA), p. 146–151, Association for Computing Machinery, 2019.
- [17] M. Hu, J. P. Strachan, Z. Li, E. M. Grafals, N. Davila, C. Graves, S. Lam, N. Ge, J. J. Yang, and R. S. Williams, "Dot-product engine for neuromorphic computing: Programming 1t1m crossbar to accelerate matrix-vector multiplication," in *2016 53rd ACM/EDAC/IEEE Design Automation Conference (DAC)*, pp. 1–6, IEEE, 2016.
- [18] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [19] A. Krizhevsky, G. Hinton, *et al.*, "Learning multiple layers of features from tiny images," 2009.
- [20] 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, Ieee, 2009.