

Improving DNN Fault Tolerance using Weight Pruning and Differential Crossbar Mapping for ReRAM-based Edge AI

Geng Yuan^{*1}, Zhiheng Liao^{*2}, Xiaolong Ma¹, Yuxuan Cai¹, Zhenglun Kong¹, Xuan Shen¹, Jingyan Fu², Zhengang Li¹, Chengming Zhang³, Hongwu Peng⁴, Ning Liu¹, Ao Ren⁵, Jinhui Wang⁶, Yanzhi Wang¹

¹Northeastern University, ²North Dakota State University, ³Washington State University, ⁴University of Connecticut,

⁵Chongqing University, ⁶University of South Alabama

¹{yuan.geng, ma.xiaol, cai.yuxu, kong.zhe, shen.xu, li.zhen, yanzhi.wang}@northeastern.edu,

²{zhiheng.liao, jingyan.fu}@ndsu.edu, ³chengming.zhang@wsu.edu, ⁴hongwu.peng@uconn.edu,

⁵ren.ao@cqu.edu.cn, ⁶jwang@southalabama.edu

Abstract—Recent research demonstrated the promise of using resistive random access memory (ReRAM) as an emerging technology to perform inherently parallel analog domain in-situ matrix-vector multiplication—the intensive and key computation in deep neural networks (DNNs). However, hardware failure, such as stuck-at-fault defects, is one of the main concerns that impedes the ReRAM devices to be a feasible solution for real implementations. The existing solutions to address this issue usually require an optimization to be conducted for each individual device, which is impractical for mass-produced products (e.g., IoT devices). In this paper, we rethink the value of weight pruning in ReRAM-based DNN design from the perspective of model fault tolerance. And a differential mapping scheme is proposed to improve the fault tolerance under a high stuck-on fault rate. Our method can tolerate almost an order of magnitude higher failure rate than the traditional two-column method in representative DNN tasks. More importantly, our method does not require extra hardware cost compared to the traditional two-column mapping scheme. The improvement is universal and does not require the optimization process for each individual device.

I. INTRODUCTION

DNNs have achieved significant progress in many fields such as computer vision [1], natural language processing [2], and medical diagnosis [3]. It becomes a fundamental building block of many artificial intelligence (AI) applications, such as facilitating edge devices in the Internet of Things (IoT) system. However, to pursue better performance (i.e., accuracy), the ever-growing size and computation load make the DNN difficulty to be deployed on edge devices in IoT system, since edge devices usually have limited resources and require ultra-low power consumption.

ReRAM enabled in-memory computing, as one of the emerging technologies, has shown advantages such as near-zero leakage-power, high active-power efficiency, non-volatility, and concise circuit structure. Accordingly, it is especially suitable for edge devices in IoT. ReRAM cells can be used to form a crossbar structure to conduct in-situ dot products, such as the convolution computations in DNNs. However, one key challenge that prevents the ReRAM from

becoming a practical technology and used in real hardware implementations is the reliability issue. For example, the stuck-at-fault defect will lead the ReRAM cell's conductance to be fixed to a high or low conductance state, regardless of its programmed, as known as the stuck-on fault or stuck-off fault. For a DNN implementation, when mapping the well-trained DNN model to the ReRAM crossbars, the stuck-at-fault defects result in a mismatch between target weights and the actual values mapped onto the crossbar, and eventually lead to an undesirable accuracy drop.

Several methods are proposed to address the stuck-at-fault defects, such as permuting the order of the crossbar rows and columns [4], retraining the network by considering the defect locations [5], [6]. These methods require individual optimization process for each ReRAM crossbar, and may also introduce complex control circuits, resulting in additional hardware overhead. Although these methods are effective in mitigating the accuracy drop caused by the stuck-at-fault defects, even [7] can restore 99% of the accuracy drop, but for mass-produced IoT products, applying optimization for each individual product (IoT device) will bring a huge time cost, and is not realistic. Thus, it is desirable to have a universal approach to improve the DNN fault tolerance.

DNN weight pruning technique is another active research area that has been intensively studied. It can effectively reduce the number of DNN weights as well as the computations, achieving high-efficiency DNN inference. Especially, the unstructured pruning allows the weights at arbitrary locations to be pruned [8], which achieves relative low accuracy degradation under a high pruning ratio.

Although the weight pruning techniques are used for compressing the DNN model to save hardware and speed up the inference, it is still considered as a trade-off between model accuracy and hardware costs. However, the essence of weight pruning is to zero out the redundant weights in the DNNs. This provides a unique possibility to improve the tolerance of DNNs to the stuck-off fault because the pruned weight is possible to fall exactly on the ReRAM crossbar stuck-off fault location, avoiding the mismatch between target weight and the

*These Authors contributed equally.

actual value represented by the ReRAM cell.

In this paper, we rethink the value of weight pruning in ReRAM-based DNN design from the perspective of model fault tolerance. We argue that weight pruning is no longer just a technique that sacrifices accuracy in exchange for hardware reduction. Unstructured weight pruning can also be used as a technique to improve model fault tolerance, meanwhile resulting in improved DNN accuracy on ReRAM-based design.

Our experimental results demonstrate that lower than a certain pruning ratio, the ability of model fault tolerance increases as the pruning ratio goes up. After exceeding a certain pruning ratio, the model fault tolerance decreases again. The pruning can only improve the model fault tolerance to the stuck-off fault, but as reported in [9], the failure rate of stuck-on fault can be $5.2\times$ higher than the stuck-off fault and become the major factor for the degradation of accuracy. Thus, we propose a differential mapping scheme to mitigate the impact of stuck-on fault and improve the overall accuracy without additional hardware costs.

The contributions of our paper can be summarized as:

- We propose to utilize unstructured pruning to improve fault tolerance of the ReRAM-based DNN in the edge AI. The improvement is statistically achieved for all models, which avoids conducting optimizations for each individual device.
- We explore the relation between model sparsity and its robustness and propose a hierarchical progressive pruning method, which can find the best-suited pruning ratio for DNNs fault tolerance.
- We propose to use a differential mapping scheme to convert the stuck-off fault tolerance achieved by pruning to stuck-on fault tolerance, which can effectively improve the overall fault tolerance when the stuck-on fault is the major fault defect.
- We validate the effectiveness of proposed methods on potential tasks for IoT applications, including image classification and object detection.

II. BACKGROUND

A. ReRAM Crossbar and ReRAM-based DNN Accelerator

Resistive RAM (ReRAM), as an emerging memory device, has many promising features, including non-volatility, nearly zero leakage power, high integration density, and high scalability [10], [11]. Every memory cell comprises a ReRAM element to store one or more data bits. In addition to data storage, its memory array structure, as known as the ReRAM-based crossbar, is capable of performing in-situ dot product between an input vector (i.e., input feature maps) and stored numbers (i.e., weights).

For DNNs, the hardware accelerators have been intensively studied in recent years [12]–[20]. Specifically, specialized hardware acceleration using ReRAM-based crossbars has been studied in recent years. Previous work, such as PRIME [21] and ISAAC [22], leverages in-situ computation to avoid the

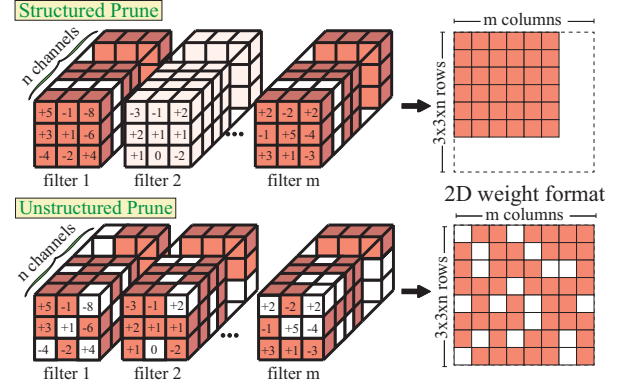


Fig. 1. Example of structured pruning and unstructured pruning.

tremendous cost of data movement and efficiently compute multiply-accumulate operations, the most intensive computation in DNNs. Recent work, TIMELY [23] saves the energy cost of data movements and D/A and A/D domain conversion by enhancing analog data locality to keep computations in analog domain.

B. Hardware Failure

Similar to the other nano-scale devices, hardware failure that greatly degrade the accuracy of neural networks perplexes the application of ReRAM [4], [5]. Some techniques have been proposed to model and detect the faults in ReRAM crossbar array including fault model with testing scheme [24], a march algorithm to cover defect [25], and analyzing impacts of stuck-at-fault defects on accuracy of a sparse coding network [26]. Various schemes to tolerate faults in ReRAM crossbar array have been proposed in some prior works. In hardware level, isolating faulty ReRAM is proposed by switching off access transistor [27]. However, for the large ReRAM crossbar array implemented in complicated neural networks, tremendous routing and area overheads are inevitable, utilizing controlling individual transistor access. In [5], redundant columns of ReRAM crossbars are utilized as a substitution for the defect ReRAM columns. But this method not only introduces nontrivial area overhead, but also increases design complexity of peripheral circuit. Unlike hardware level schemes, fault-aware network retraining with weight mapping/remapping are commonly composed in software level optimizations [4]–[6]. They retrain the neural network by forcing the weights to be the same as the fault value on corresponding crossbar locations. However, the retraining and remapping will introduce tremendous computing and interaction costs.

C. DNN Weight Pruning

DNN Weight pruning [8], [28]–[31], as one of the DNN model compression techniques, has also been investigated to reduce the weight storage and improve performance for ReRAM-based acceleration designs [32], [33]. The weight pruning consists of two main categories: 1) structured pruning and 2) unstructured pruning, as shown in Figure 1.

1) *Structured pruning*: **Structured pruning** prunes the entire channel(s)/filter(s) of DNN weights [34], [35]. After removing the pruned weights, the size of weight matrix can be reduced in a structured manner and a regular shape of the weight matrix can be maintained. This makes structured pruning hardware friendly. By applying structured pruning to ReRAM-based DNN accelerator design, the number of ReRAM crossbars and the corresponding peripheral circuits can be reduced and hence an effective reduction of hardware cost in power and area. However, due to the coarse pruning granularity of structured pruning, the pruned DNN models usually suffer from considerable accuracy loss.

2) *Unstructured pruning*: **Unstructured pruning** adopts a fine-grained pruning strategy, which prunes the weights at arbitrary locations [36], [37]. This ensures higher flexibility to find the optimized pruned location, which is beneficial to preserve a higher model accuracy. However, in order to accelerate the computation by leveraging unstructured sparsity (i.e., zeros in an irregular shape), it is inevitable to introduce additional indices to locate the non-zero weights during the computation. It is generally considered non-hardware friendly. Thus, hardware designs, such as ReRAM-based accelerator, cannot take advantage of unstructured pruning to achieve hardware reduction and computation acceleration [31].

III. THE IMPACT OF DEFECT ON DNN

Previous studies have shown the promising features and performances of ReRAM-based crossbars to be used in DNN accelerations. And compared to the conventional CMOS-based designs, ReRAM-based design achieves a great advantage in terms of power and area and is especially suitable for IoT devices. However, the defect (e.g., stuck-at-fault defect) caused by the hardware imperfection during the manufacturing process or aging with time is still one of the key challenges that prevent the ReRAM to become a practical technology in computing and memory devices.

For example, Figure 2 shows the accuracy degradation caused by stuck-at-fault defects on ResNet18 using CIFAR10 dataset. The defect model reported in [9] is used. Note that the value of DNN weights can be either positive or negative, but the weights stored in ReRAM crossbars are represented by the conductance value of ReRAM cells, which can only be positive values. Thus, there are two main mapping methods to make sure the results can be calculated correctly. The general way is to decompose each weight into a positive portion and a negative portion, then map them to two separate ReRAM crossbar columns. Another way is to add an offset to the original weights when mapping them to the crossbars. This will shift all negative weights to the positive range. By doing so, with a certain cost of offset circuitry, half of the crossbars can be saved compared to the two-column mapping method. As the Figure 2 shows, both mapping methods suffer severe accuracy loss compared to the ideal case, where the offset method shows much worse tolerance to the defect.

The current solutions to handle the stuck-at-fault defects (e.g., defect-aware retraining, crossbar permutation, utilizing

redundant crossbars) require an optimization process for each individual device (chip). Obviously, the cost of retraining and remapping for each device is too high, and it is impractical for the mass-produced IoT products. In addition, as IoT devices are used, new defects may occur, which may significantly mitigate the impact of the optimization methods.

Question 1: Is there a way that can improve the tolerance of stuck-at-fault defects uniformly instead of handling each IoT device individually?

Weight pruning technique, which creates zeros weights in DNN models, might be a potential solution for this question. This feature seems to be naturally beneficial to improve the stuck-off tolerance of DNN models, because when mapping the weights on to the ReRAM crossbars, some location of stuck-off faults may overlap with the weights that have been already pruned to zeros during the training (pruning) process on software. The unstructured pruning is usually considered non-hardware friendly [31] and is not adopted in hardware designs. However, for our case, the unstructured pruning would be a better choice than structured pruning, because unstructured pruning can create more zeros while maintaining a higher accuracy.

Question 2: Can the unstructured pruning be used in hardware design to improve tolerance of stuck-at-fault? If true, what is the relation between the pruning ratio and the model's fault tolerance?

However, as reported in [9], the failure rate of stuck-on fault is as $5.2\times$ higher than stuck-off fault. This will make the stuck-on fault become the most severe issue and hinder the overall tolerance improvement achieved by pruning (that only affects stuck-off fault).

Question 3: Is there an efficient way to improve the overall fault tolerance when stuck-on fault is the major issue?

In the following paper, we will answer the above three questions.

IV. IMPROVING DNN FAULT TOLERANCE USING WEIGHT PRUNING

Stuck-at-faults is one of the main factors that causes the mismatch between software trained model and actual model mapped onto the hardware (e.g., ReRAM crossbar). This mismatch usually leads to undesirable performance degradation (i.e., accuracy drop) on hardware implementations. To answer




Fig. 2. Impact of stuck-at faults on accuracy using ResNet18 on CIFAR10.

the Question 1 and Question 2, we investigate using the weight pruning technique to improve the DNN tolerance to the stuck-at fault without conducting optimizations for each individual device.

A. Problem Formulation

We define the occurrence of mismatch between a trained weight W and its actual value represented by the ReRAM crossbar W_{act} as:

$$v(w, w_{act}) = \begin{cases} 0 & w_{act} = w, \\ 1 & w_{act} \neq w. \end{cases} \quad (1)$$

Given the failure rate of stuck-off fault and stuck-on fault of a ReRAM crossbar as P_{sa0} and P_{sa1} , then the expectation of the occurrence of a DNN weight mapped on a ReRAM crossbar that is different to the trained value can be represented as:

$$E = 0 \cdot p(v = 0) + 1 \cdot p(v = 1) = P_{sa0} + P_{sa1}. \quad (2)$$

However, when incorporating the weight pruning, the locations of the stuck-off fault may overlap with the pruned weights (0s), then the mismatch is avoided. Considering the pruning ratio of a DNN $R_p = \frac{\#_of_zero_weights}{total_ \#_of_weights}$, then the expectation of the occurrence of weight mismatch becomes:

$$E' = P_{sa0} \cdot (1 - R_p) + P_{sa1}. \quad (3)$$

Obviously, the E' is smaller than E , which indicates that the weight pruning can statistically reduce the weight mismatch caused by stuck-off fault and can potentially mitigate the accuracy degradation.

B. Exploring the Effects of Weight Pruning on DNN Stuck-at Fault Tolerance

As we discussed above, the weight pruning seems an effective technique can be used to reduce the weight mismatch caused by stuck-off fault. Theoretically, a higher prune ratio leads to a higher probability that the stuck-off fault location overlaps with the pruned weight. However, for real DNN models, the accuracy degradation caused by the stuck-off fault is not decreasing monotonically as the prune ratio increasing.

The Figure 3 shows the results of accuracy drop caused by stuck-off fault under different unstructured pruning ratio on CIFAR10 dataset using ResNet18 models. As shown in Figure 3 (a), we evaluate the models with the prune ratio from 0 (i.e., not prune) to 0.9 (i.e., prune 90 % weights). The original accuracy represents the model accuracy after pruning without incorporating stuck-off fault. We can observe that even 90% of weights are pruned, the model can still retain a high accuracy (similar to the unpruned model) due to the high model redundancy on the CIFAR10 dataset and high flexibility of unstructured pruning. We can also observe that as the pruning ratio increases, the accuracy drop caused by the stuck-off fault decreases and reaches its minimum value at the pruning ratio of 0.6, then the accuracy drop goes up again.

The reason is that, under a relative low pruning ratio, a higher weight pruning ratio can create more zero weights in

the DNN model, which is possible to let more stuck-off fault locations overlap with those pruned weights (0s) and hence a less accuracy drop caused by the mismatch between trained weights and their actual mapped values. However, when the pruning ratio is relatively high, few weights are left in the model, and each of them becomes very important. Thus, the robustness of the model is reduced, and the faults that hit the non-zero weights will lead to more significant damage to the accuracy. A similar trend can also be observed in Figure 3 (b), which shows the comparison of accuracy drop under different stuck-off fault rates. We can find that weight pruning is more effective in improving fault tolerance under a higher fault rate.

Figure 4 shows similar comparisons on ImageNet dataset using ResNet18. Since the image classification task on ImageNet dataset is much complicated than on CIFAR10 dataset, the model is more sensitive to the stuck-at faults. We can observe that by using weight pruning, the accuracy drop caused by stuck-off fault can be mitigated by up to 2% (Top-1 error) for the case under a 0.3 pruning ratio with a 0.024 fault rate compared with 0 pruning ratio case.




Fig. 3. Results on CIFAR10 dataset using ResNet18 under different pruning ratios: (a) Original model accuracy and accuracy drop caused by the stuck-off fault (0.016 fault rate), and (b) comparison of accuracy drop under different rate of stuck-off fault.




Fig. 4. Results on ImageNet dataset using ResNet18 under different pruning ratios: (a) Original model accuracy and accuracy drop caused by the stuck-off fault (0.016 fault rate), and (b) comparison of accuracy drop under different rate of stuck-off fault.

C. A Hierarchical Progressive Pruning Method

In order to mitigate the accuracy drop caused by stuck-at faults without losing network robustness, we propose a hierarchical progressive pruning method, which can find the best-suited pruning ratio for DNNs. Our method is derived by considering two key observations: 1) as the pruning ratio increases, the accuracy drop caused by stuck-at-fault defects will first decrease until reaching a minimum point, then

increase again; 2) Larger DNN layers (i.e., usually the later layers in a neural network) have higher redundancy and can be pruned with higher pruning ratio without compromising accuracy.

The overall flow of our hierarchical progressive pruning method is shown in Algorithm 1. We first divide a given DNN model into several blocks according to the layer sizes. Thus, each block contains several layers that have similar layer size (e.g., layers with 64 channels). At the beginning, all the blocks are in the active model block list B . We apply a small starting pruning ratio p and get the pruned model M . Note that, for this step, any pruning algorithm can be used to get the pruned model. And we use unstructured pruning in our method since, in this paper, we focus on creating more zeros rather than reducing hardware cost. Then, we add stuck-at-fault defects on the model M and evaluate the model accuracy with fault defect, and repeat many times to get the average accuracy. If the accuracy performs better (or worse but within a threshold) then, we increase the pruning ratio and run test the accuracy again. Otherwise, we remove the first block (which has the smallest block size) from B and test the accuracy with the same pruning ratio again. In this way, we can gradually increase the pruning ratio while narrow down the selected blocks. As a result, we can obtain a pruned model with different pruning ratios in block-wise while delivering the best fault tolerance.

D. Differential Mapping Scheme

By incorporating the weight pruning, the accuracy drop caused by stuck-off fault can be mitigated. However, the performance of hardware implementation needs to consider both the stuck-off fault and stuck-on fault. Moreover, according to [9], the failure rate of stuck-on fault can be $5.2\times$ higher than the failure rate of stuck-off fault. In this case, the stuck-on fault becomes the major issue leading to the accuracy drop rather than the stuck-off fault. To solve this problem, we propose a differential ReRAM crossbar mapping scheme to improve the network tolerance to stuck-on fault.

The traditional mapping scheme decomposes the weight into a positive magnitude portion w^+ and a negative magnitude portion w^- , and uses two ReRAM cells to represent positive and negative portions separately.

Different from the traditional mapping scheme that decomposes the weight into a positive magnitude portion w^+ and a negative magnitude portion w^- and sums them to reconstruct the original weight value during the computation, our proposed differential mapping scheme represents the weight value by using the difference between two ReRAM cells.

Given a weight value w scaled to the range $[-1,1]$, then it will be mapped as:

$$w_a = \begin{cases} 1 & w \geq 0, \\ 1 - |w| & w < 0, \end{cases} \quad w_b = \begin{cases} 1 - w & w > 0, \\ 1 & w \leq 0, \end{cases}$$

where the w_a and w_b are the two ReRAM cells (with normalized range to $[0,1]$) used to represent the weight w . During the computations, the original weight value is obtained by using $w_a - w_b$.

Algorithm 1: Hierarchical Progressive Pruning.

```

1 Initialization:
2  $M_{init}$  : well-trained model,
3  $th$  : threshold of acceptable accuracy drop,
4 Model blocks  $B = [b_1, b_2, b_3, \dots, b_n]$ ,
5  $R_{range} = [r_1, r_2, r_3, \dots, r_m]$  candidate pruning ratios
   in ascending order,
6 Current model  $M_{cur} = M_{init}$ ,
7 Best model  $M_{best} = M_{init}$  ;

8 for  $p$  in  $R_{range}$  do
9   Prune blocks  $B$  in  $M_{cur}$  with prune ratio  $p$ ;
10  Get pruned model  $M$ ;
11  for num of test runs do
12    Add stuck-at-fault defects on model  $M$ ;
13    Evaluate model accuracy;
14  end
15  Get average model accuracy  $Acc(M)$ ;
16  if  $Acc(M_{cur}) - Acc(M) < th$  then
17    if  $Acc(M) > Acc(M_{cur})$  then
18       $M_{best} = M$ ;
19    end
20     $M_{cur} = M$ ;
21  else
22     $B.pop\_front()$  // remove first block in  $B$ ;
23  end
24  if  $B$  is empty then
25    break;
26  end
27 end
28 return  $M_{best}$  // model with highest accuracy

```

Our differential mapping scheme ensures a more number of 1s to be mapped on the ReRAM crossbars. There is always at least one 1 to be mapped on one of the two ReRAM cells. And both of the ReRAM cell will be 1 when weight value is zero. In this way, the weight pruning will create more 1s, which will increase the opportunity that the stuck-on faults overlap with those 1s and mitigate the accuracy drop caused by the stuck-on fault. It can improve the performance significantly, since the failure rate of stuck-on fault is higher than stuck-off fault and will lead to more significant accuracy drop, although the differential mapping scheme only improve the tolerance to the stuck-on fault.

The Figure 5 shows the hardware implementation of our proposed differential mapping scheme. Compared with traditional hardware implementation [38] that weight values are calculated by two ReRAM columns and an OpAmp-based subtractor, the differential mapping method still utilizes similar arithmetic circuit without increasing hardware cost [39]. When performing differential mapping, the w_a and w_b values are mapped to the ReRAM cells in the G_a and G_b columns, respectively. When reading the weight value, the current sum of the G_a and G_b columns is subtracted by the arithmetic

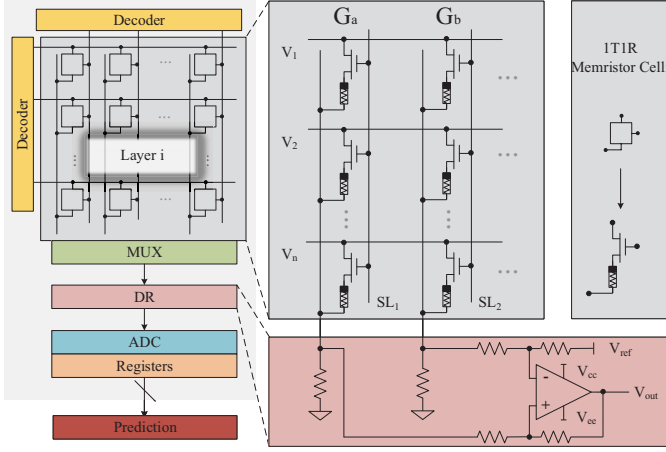


Fig. 5. Hardware implementation of proposed differential mapping scheme. circuit to obtain an effective activation value.

V. EVALUATION

A. Experimental Setup

We evaluate our methods on CIFAR10 dataset and ImageNet dataset using ResNet18 for image classification task. We also validate our methods on YOLO-v4 for object detection task using MS COCO dataset. The ReRAM failure model is adopted from [9] with the ratio of stuck-off and stuck-on fault is 1:5.2. The ADMM-based pruning algorithm [40] is used for weight pruning during the hierarchical progressive pruning process. All model training, pruning, and accuracy evaluations are conducted on a GPU server using PyTorch API. Each accuracy result is obtained by averaging the results of 100 runs.

B. Comparison Results

We first evaluate the model fault tolerance via model accuracy on image classification task. As shown in Figure 6, we compare the model fault tolerance results optimized by our methods (including unstructured pruning and differential crossbar mapping) to the original unpruned model with traditional two-column mapping (as discussed in Section III). We do not compare to the offset mapping scheme since the fault tolerance of the offset scheme is worse than the traditional two-column scheme. On CIFAR10 dataset, the original model accuracy is 94.1% without introduce the stuck-at faults. Under failure rate of 0.001, our method only has 0.2% accuracy drop on average, where the accuracy drop of the traditional mapping is 1.2%. It can be observed that a severe accuracy drop occurs to the traditional mapping under failure rate of 0.005, where our optimizations can preserve a high accuracy under a failure rate of 0.01. For the ImageNet, since the classification task on ImageNet is harder than CIFAR10, the network is more sensitive to the stuck-at faults. As we can see, our optimizations clearly provides a better fault tolerance than the traditional two-column mapping.

We also validate our methods on object detection task. Figure 7 shows the comparison results on MS COCO dataset using YOLOv4 for object detection task.




Fig. 6. Comparison results on CIFAR10 dataset and ImageNet using ResNet18 under different failure rate.




Fig. 7. Comparison results of YOLOv4 on MS COCO dataset under different failure rate.

All the results shows that our method can tolerate almost an order of magnitude higher failure rate than the traditional two-column method, which demonstrates the effectiveness of our optimizations.

VI. CONCLUSION

In this paper, we rethink the value of weight pruning in ReRAM-based DNN design from the perspective of model fault tolerance. We propose a hierarchical progressive pruning method, which can find the best-suited pruning ratio for DNNs to improve the model fault tolerance. And a differential mapping scheme is proposed to improve the fault tolerance under a high stuck-on fault rate.

ACKNOWLEDGEMENT

This work is funded by the National Science Foundation Awards CCF-1937500 and CNS-1909172.

REFERENCES

- [1] I. Goodfellow *et al.*, *Deep learning*. MIT press, 2016.
- [2] J. Devlin *et al.*, “Bert: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of NAACL-HLT*, 2019.
- [3] C. Gongye *et al.*, “New passive and active attacks on deep neural networks in medical applications,” in *ICCAD*, 2020.
- [4] L. Chen *et al.*, “Accelerator-friendly neural-network training: Learning variations and defects in rram crossbar,” in *DATE*. IEEE, 2017.
- [5] C. Liu *et al.*, “Rescuing memristor-based neuromorphic design with high defects,” in *DAC*. IEEE, 2017.
- [6] L. Xia *et al.*, “Fault-tolerant training with on-line fault detection for rram-based neural computing systems,” in *DAC*, 2017.
- [7] B. Zhang *et al.*, “Handling stuck-at-fault defects using matrix transformation for robust inference of dnn,” *IEEE T-CAD*, 2020.
- [8] A. Ren *et al.*, “Admm-nn: An algorithm-hardware co-design framework of dnn using alternating direction method of multipliers,” 2018.
- [9] C. Chen *et al.*, “Rram defect modeling and failure analysis based on march test and a novel squeeze-search scheme,” *IEEE Transactions on Computers*, 2015.
- [10] L. Chua, “Memristor-the missing circuit element,” *IEEE Transactions on circuit theory*, vol. 18, no. 5, pp. 507–519, 1971.

- [11] G. Yuan *et al.*, "Memristor crossbar-based ultra-efficient next-generation baseband processors," in *IEEE 60th International Midwest Symposium on Circuits and Systems (MWSCAS)*, 2017.
- [12] C. Ding *et al.*, "Circnn: accelerating and compressing deep neural networks using block-circulant weight matrices," in *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture*, 2017.
- [13] H. Fang *et al.*, "Multivariate time series classification using spiking neural networks," *CoRR*, 2020.
- [14] Y. Wang *et al.*, "Towards ultra-high performance and energy efficiency of deep learning systems: an algorithm-hardware co-optimization framework," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018.
- [15] H. Fang *et al.*, "A general framework to map neural networks onto neuromorphic processor," in *ISQED*, 2019.
- [16] X. Ma *et al.*, "An area and energy efficient design of domain-wall memory-based deep convolutional neural networks using stochastic computing," in *19th ISQED*, 2018.
- [17] H. Li *et al.*, "Real-time mobile acceleration of dnns: From computer vision to medical applications," in *Proceedings of the 26th Asia and South Pacific Design Automation Conference*, 2021.
- [18] H. Fang *et al.*, "Scalable noc-based neuromorphic hardware learning and inference," in *IJCNN*, 2018.
- [19] H. Fang, A. Shrestha *et al.*, "Exploiting neuron and synapse filter dynamics in spatial temporal learning of deep spiking neural network," in *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, 2020.
- [20] H. Fang *et al.*, "Encoding, model, and architecture: Systematic optimization for spiking neural network in fpgas," in *2020 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*. IEEE, 2020.
- [21] P. Chi *et al.*, "Prime: A novel processing-in-memory architecture for neural network computation in rram-based main memory," in *ISCA*, 2016.
- [22] A. Shafiee *et al.*, "Isaac: A convolutional neural network accelerator with in-situ analog arithmetic in crossbars," in *ISCA*, 2016.
- [23] W. Li *et al.*, "Timely: Pushing data movements and interfaces in pim accelerators towards local and in time domain," in *ISCA*, 2020.
- [24] S. Kannan *et al.*, "Modeling, detection, and diagnosis of faults in multilevel memristor memories," *IEEE T-CAD*, 2015.
- [25] C.-Y. Chen *et al.*, "Rram defect modeling and failure analysis based on march test and a novel squeeze-search scheme," *IEEE Transactions on Computers*.
- [26] P. Sheridan and W. D. Lu, "Defect considerations for robust sparse coding using memristor arrays," in *IEEE/ACM NANOARCH*. IEEE, 2015.
- [27] H. Manem *et al.*, "Design considerations for variation tolerant multilevel cmos/nano memristor memory," in *Great lakes symposium on VLSI*, 2010.
- [28] L. Deng *et al.*, "Model compression and hardware acceleration for neural networks: A comprehensive survey," *Proceedings of the IEEE*, 2020.
- [29] X. Ma *et al.*, "Resnet can be pruned 60x: Introducing network purification and unused path removal (p-rm) after weight pruning," in *2019 IEEE/ACM International Symposium on Nanoscale Architectures (NANOARCH)*, 2019.
- [30] Z. Li *et al.*, "6.7 ms on mobile with over 78% imagenet accuracy: Unified network pruning and architecture search for beyond real-time mobile acceleration," *arXiv preprint arXiv:2012.00596*, 2020.
- [31] X. Ma *et al.*, "Non-structured dnn weight pruning—is it beneficial in any platform?" *arXiv preprint arXiv:1907.02124*, 2019.
- [32] G. Yuan *et al.*, "An ultra-efficient memristor-based dnn framework with structured weight pruning and quantization using admm," in *IEEE/ACM ISLPED*, 2019.
- [33] X. Ma *et al.*, "Tiny but accurate: A pruned, quantized and optimized memristor crossbar framework for ultra efficient dnn implementation," *2020 25th ASP-DAC*, 2020.
- [34] W. Wen *et al.*, "Learning structured sparsity in deep neural networks," in *NeurIPS*, 2016.
- [35] Y. He *et al.*, "Channel pruning for accelerating very deep neural networks," in *ICCV*. IEEE, 2017.
- [36] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding," *arXiv preprint arXiv:1510.00149*, 2015.
- [37] J. Frankle and M. Carbin, "The lottery ticket hypothesis: Finding sparse, trainable neural networks," *ICLR*, 2018.
- [38] O. Krestinskaya *et al.*, "Memristive gan in analog," *Nature Scientific Reports*, 2020.
- [39] O. Krestinskaya *et al.*, "Neuromemristive circuits for edge computing: A review," *IEEE TNNLS*, 2020.
- [40] T. Zhang *et al.*, "A systematic dnn weight pruning framework using alternating direction method of multipliers," in *ECCV*, 2018.