

---

# Exploring Bit-Slice Sparsity in Deep Neural Networks for Efficient ReRAM-Based Deployment

---

Jingyang Zhang<sup>1</sup>, Huanrui Yang<sup>1</sup>, Fan Chen<sup>1</sup>, Yitu Wang<sup>2</sup>, Hai Li<sup>1</sup>

<sup>1</sup>Department of Electrical and Computer Engineering, Duke University

<sup>2</sup>School of Microelectronics, Fudan University

{jingyang.zhang, huanrui.yang, fan.chen, hai.li}@duke.edu, ytwang16@fudan.edu.cn

## Abstract

Emerging resistive random-access memory (ReRAM) has recently been intensively investigated to accelerate the processing of deep neural networks (DNNs). Due to the *in-situ* computation capability, analog ReRAM crossbars yield significant throughput improvement and energy reduction compared to traditional digital methods. However, the power hungry analog-to-digital converters (ADCs) prevent the practical deployment of ReRAM-based DNN accelerators on end devices with limited chip area and power budget. We observe that due to the limited bit-density of ReRAM cells, DNN weights are bit sliced and correspondingly stored on multiple ReRAM bitlines. The accumulated current on bitlines resulted by weights directly dictates the overhead of ADCs. As such, bitwise weight sparsity rather than the sparsity of the full weight, is desirable for efficient ReRAM deployment. In this work, we propose *bit-slice*  $\ell_1$ , the first algorithm to induce bit-slice sparsity during the training of dynamic fixed-point DNNs. Experiment results show that our approach achieves  $2\times$  sparsity improvement compared to previous algorithms. The resulting sparsity allows the ADC resolution to be reduced to 1-bit of the most significant bit-slice and down to 3-bit for the others bits, which significantly speeds up processing and reduces power and area overhead.

## 1 Introduction

Although the promising performance of Deep neural network (DNN) models have been demonstrated in various real-world tasks [1, 2], the intensive computing and memory requirements of DNN processing make its deployment extremely difficult, especially on end devices with limited resources and rigid power budget [3, 4]. The challenges of efficient deployment of large DNN models have motivated researches on model compression, including pruning [3, 4] and quantization [5, 6]. Coupled with algorithm development, customized CMOS DNN accelerators are extensively investigated to take full advantage from model compression algorithms. For example, ESE [7] is optimized to achieve high computation efficiency on element-wise sparse DNNs, while DNPU [8] supports low-precision, dynamic fixed-point operations. However, these digital approaches typically require that most of the network weights to be stored off-chip, resulting in a large performance penalty for memory access.

In the meantime, the emerging resistive random-access memory (ReRAM) provides a novel mixed-signal design paradigm. In general, DNN weights are encoded as the ReRAM cell conductance, while the core computing pattern in DNN processing, i.e., massive matrix-vector multiplications, can be executed *in-situ* in one computing-in-memory (CIM) cycle without moving data back and forth in the memory hierarchy. Indeed, prior ReRAM-based accelerators have demonstrated two orders of magnitude advantages in energy, performance and chip footprint, over their digital counterparts [9, 10]. However, the conversion between digital and analog domains, especially the analog-to-digital converters (ADCs), limit the effectiveness of these CIM designs to a certain extent because they normally account for  $> 60\%$  power and  $> 30\%$  area overhead [9].

We observe that each operand (i.e. weight) is bit-sliced across multiple ReRAM bitlines (located in the same row) due to the limited cell bit density. The accumulated currents on bitlines dictate ADC’s bit-resolution, which in turn determines the size and power consumption of the ADCs, as the ADC overhead generally increases exponentially with its resolution [9]. Higher sparsity in each bit-slice is desired to reduce the accumulated currents. Based on this observation, we propose *bit-slice*  $\ell_1$ , a novel sparsity regularization that applies  $\ell_1$  penalization to all bit-slices of each fixed-point weight elements during training to induce bit-slice sparsity. Unlike previous weight-grade sparsity methods, such finer-grained sparsity distribution achieves balanced sparsity when mapped to practical ReRAM crossbars, resulting in efficient deployment and significantly reduced ADC overhead. Existing work explored bit-partition [11] and dynamic bit-level fusion/decomposition [12] in efficient DNN accelerator designs, but none of these works considered the sparsity within each bit-slice. Therefore, our work on bit-slice sparsity provides new opportunities to effectively exploit sparsity in sparse accelerators, as initially demonstrated in [13].

We apply *bit-slice*  $\ell_1$  regularization to the training process of a 8-bit dynamic fixed-point DNN to show its effectiveness. We assume each slice contains two bits because 2 bits/cell is the most common multi-level cell type in current ReRAM technology. Please note that as technology advances, our approach can be easily extended to support more bits per slice. Experiment results shows that *bit-slice*  $\ell_1$  achieves  $2\times$  sparsity improvement compared to previous full-number pruning algorithms. The resulting sparsity can significantly speeds up processing and reduces power and area overhead in ReRAM deployment. To the best of our knowledge, this is the first algorithm specifically designed to train DNN models that are friendly for the bit sliced deployment on ReRAM crossbars.

## 2 Proposed method

In this section, we first describe the procedure for training a DNN with dynamic fixed-point quantization, which fits the requirement for ReRAM deployment. Then we introduce our bit-slice  $\ell_1$  regularizer, which aims at providing bit-slice sparsity for the ReRAM instead of the sparsity on the full-number weights. Finally, we present the whole training routine of our method as we apply the proposed regularizer to the dynamic fixed-point training process. The training routine and the proposed regularizer is demonstrated in Figure 1.

### 2.1 Dynamic fixed-point quantization

As observed by Polino et al. [6], the weight of different layers may have various dynamic ranges. Keeping the dynamic range of each layer is important for maintaining the performance of the model, especially after low-precision quantization [6]. Therefore for each layer, we need to first compute its dynamic range and scale the weight to the range of  $[0, 1]$  before applying quantization. Since state-of-the-art ReRAM based accelerators often map negative and positive weight elements to separated crossbars [10], here we ignore the sign of weight elements and only focus on quantizing their absolute values. The dynamic range of a layer  $W_l$  is defined as:

$$S(W_l) = \lceil \log_2(\max_{w_l^i \in W_l} (|w_l^i|)) \rceil, \quad (1)$$

where  $i$  is the index of each weight element in layer  $l$ .

Then we apply uniform quantization to the scaled weight. Considering the scaling factor  $2^{-S(W_l)}$  applied on the weight elements, the quantization step size of a  $n$ -bit quantization would be  $Q_{step} = 2^{S(W_l)-n}$ , and, the weight element  $w_l^i$  will be quantized to:

$$B(w_l^i) = \lfloor \frac{w_l^i}{Q_{step}} \rfloor. \quad (2)$$

This mapping will guarantee all  $B(w_l^i)$  are within the range of  $[0, 2^n - 1]$ . The  $n$ -bit binary representation of  $B(w_l^i)$  will be stored in the ReRAM for computation. The dynamic range of the original weight can be recovered as  $Q(w_l^i) = B(w_l^i) \cdot Q_{step}$ , which can be easily implemented with a shifting operation after each layer’s computation on the ReRAM crossbar. The quantization precision is set to 8 bits in this work, which is efficient in hardware deployment without significant accuracy loss.

### 2.2 Bit-slice $\ell_1$

After quantization, the quantized weight  $B(w_l^i)$  can be represented in binary form as  $B(w_l^i) = \sum_{j=0}^7 b_j \cdot 2^j$ . Then for ReRAM mapping,  $B(w_l^i)$  will be sliced into four 2-bit slices, i.e.  $\{b_7, b_6\}$ ,




Figure 1: Illustration of the training routine and our bit-slice  $\ell_1$ .

Table 1: Results on MNIST

Method	Accuracy	Ratio of non-zero wights				Average
		$\hat{B}^3$	$\hat{B}^2$	$\hat{B}^1$	$\hat{B}^0$	
Pruned	97.99%	1.08%	5.87%	8.42%	17.42%	8.20±5.94%
$\ell_1$	97.99%	1.19%	5.21%	7.01%	11.36%	6.19±3.65%
$B\ell_1$	97.67%	<b>0.84%</b>	<b>4.02%</b>	<b>4.27%</b>	<b>9.58%</b>	<b>4.68±3.14%</b>

$\{b_5, b_4\}$ ,  $\{b_3, b_2\}$ , and  $\{b_1, b_0\}$  where  $b_7$  is the MSB and  $b_0$  is the LSB, and be mapped onto 4 separated crossbars. Here we propose the *bit-slice*  $\ell_1$  regularizer, which can apply  $\ell_1$  regularization to all the bit-slices simultaneously in order to reach a sparse mapping on all ReRAM crossbars. Formally, the bit-slice process can be represented as  $B(w_l^i) = \sum_{k=0}^3 \hat{B}_l^{i,k} \cdot 2^{2k}$ , where  $\hat{B}_l^{i,k}$  will be an integer within the range of  $[0, 3]$ . The bit-sliced  $\ell_1$  of weight  $W_l$  is therefore defined as:

$$B\ell_1(W_l) := \sum_{i,k} \hat{B}_l^{i,k}. \quad (3)$$

Note that the  $B\ell_1$  regularizer takes the full weight  $W_l$  as input for training. This property enables the regularizer to smoothly fit into the training routine of a dynamic fixed-point DNN.

### 2.3 Training routine

The proposed  $B\ell_1$  regularizer enables us to achieve bit-slice sparsity by training from scratch. Yet we find it would be more efficient in reaching higher sparsity by starting from a pretrained, element-wise sparse model, such as a model trained with the  $\ell_1$  regularizer.

We follow the training procedure proposed in [5] to train the dynamic fixed-point network. Specifically, we keep full precision weights during the training. As shown in Figure 1, for each step, we first quantize  $w_l^i$  to  $B(w_l^i)$  as described in Equation (2), then the  $w_l^i$  is replaced with the recovered quantized weight  $Q(w_l^i)$ . We use  $Q(w_l^i)$  to do the forward pass, compute the cross entropy loss  $\mathcal{L}_{CE}$  and the penalty of the  $B\ell_1$  regularizer. The gradient is then accumulated to  $Q(w_l^i)$  with full precision, which will be used as the new  $w_l^i$  for the next step. The update rule for each training step can be formally formulated as (index  $i$  is omitted here for clarity):

$$q^{(t)} = Q(w_l^{(t)}), \quad w_l^{(t+1)} = q^{(t)} - lr \times (\nabla_q \mathcal{L}_{CE}(q^{(t)}) + \alpha \nabla_q B\ell_1(q^{(t)})) \quad (4)$$

## 3 Experiment results

We test the proposed bit-slice  $\ell_1$  on the MNIST benchmark [14] with a toy model consisting of two linear layers, and on the CIFAR-10 dataset [15] with VGG-11 [1] and ResNet-20 [16]. All the models are implemented and trained in the deep learning framework PyTorch<sup>1</sup>.

Table 1 and Table 2 summarize the performance of the proposed bit-slice  $\ell_1$  on MNIST and CIFAR-10, respectively. Note that in these two tables and the rest of this section,  $\hat{B}^3$ ,  $\hat{B}^2$ ,  $\hat{B}^1$ , and  $\hat{B}^0$  represent the 4 slices of the bit-slice weights, from the most significant to the least significant respectively. The sparsity is computed across the whole model. We take normal  $\ell_1$  regularization as a baseline,

<sup>1</sup>Codes for the experiments available at [https://github.com/zjysteven/bitsslice\\_sparsity](https://github.com/zjysteven/bitsslice_sparsity)

Table 2: Results on CIFAR-10

Model	Method	Accuracy	Ratio of non-zero wights				Average
			$\hat{B}^3$	$\hat{B}^2$	$\hat{B}^1$	$\hat{B}^0$	
VGG-11	Pruned	88.93%	0.86%	28.30%	34.14%	33.39%	24.17±13.65%
	$\ell_1$	<b>89.39%</b>	0.39%	9.37%	18.43%	22.19%	12.59±8.45%
	$B\ell_1$	89.33%	<b>0.21%</b>	<b>3.57%</b>	<b>7.09%</b>	<b>10.71%</b>	<b>5.40±3.92%</b>
ResNet-20	Pruned	89.22%	1.10%	8.07%	21.92%	43.96%	18.76±16.36%
	$\ell_1$	<b>90.62%</b>	0.44%	4.71%	14.37%	33.16%	13.17±12.60%
	$B\ell_1$	89.66%	<b>0.31%</b>	<b>3.34%</b>	<b>11.99%</b>	<b>31.39%</b>	<b>11.76±12.12%</b>




Figure 2: Bit-slice sparsity of VGG-11 on CIFAR-10 during training.

which is applied to the full weight. Without regularization,  $\hat{B}^3$  can be pruned to about 1% non-zero weights, while  $\hat{B}^2$ ,  $\hat{B}^1$ , and especially  $\hat{B}^0$  still have a large amount of non-zero elements, resulting in significant imbalance between the sparsity of bit-slice weights. The element-wise sparsity induced by normal  $\ell_1$  regularization is able to improve bit-slice sparsity; while our bit-slice  $\ell_1$  achieves higher sparsity in all test cases. Bit-slice  $\ell_1$  also mitigates the unbalanced sparsity as shown by the lower standard variance compared to normal  $\ell_1$ 's results. These results support that the proposed bit-slice  $\ell_1$  fits better for ReRAM, to which bit-slice sparsity is of great importance.

Figure 2 compares the percentage of non-zero bit-slice elements during the training with original  $\ell_1$  and the proposed bit-slice  $\ell_1$ . It can be clearly observed that bit-slice  $\ell_1$  reduces the number of non-zero bit-slices faster than normal  $\ell_1$  regularization from the very beginning, which again proves that bit-slice  $\ell_1$  is a better option for regularizing bit-slice weights.

In simulation, we map the achieved 8-bit weights with bit-slice sparsity onto 4 groups of  $128 \times 128$  ReRAM crossbars (XBs), with each group storing 2 bits of the 8-bit weights.  $XB_{3,2,1,0}$  store the 2-bit slices from the MSB to the LSB respectively. According to the sparsity level achieved, we can apply ADCs with different resolutions to the 4 groups of XBs. As illustrated in Table 3, ISAAC [9] needs to use 8-bit ADCs to store the weights without bit-slice sparsity even after ADC optimization. However, with the bit-slice sparsity achieved by bit-slice  $\ell_1$ , we can use 1-bit and 3-bit ADCs instead.

Table 3: ADC Overhead Saving with Bit-Slice Sparsity

	w/o Bit-Slice Sparsity		w/ Bit-Slice Sparsity		
	Resolution	Resolution	Energy Saving	Speedup	Area Saving
$XB_3$	8 bit	1 bit	28.4×	8×	2×
$XB_{2,1,0}$	8 bit	3 bit	14.2×	2.67×	2×

According to [17], the power of ADC is approximately proportional to  $2^N/(N + 1)$  and the sensing time of ADC is directly proportional to  $N$ . Here,  $N$  denotes the resolution of ADC. Therefore, with bit-slice sparsity, the 1-bit ADC of  $XB_0$  can achieve  $28.4\times$  energy saving and  $8\times$  sensing time speedup. Meanwhile, the 3-bit ADC can achieve  $14.2\times$  energy saving and  $2.67\times$  sensing time speedup. From the area perspective, the area of a 6-bit ADC is approximately the half of an 8-bit ADC but the area varies little when the resolution is lower than 6. Thus, with bit-slice sparsity, the ADC can achieve  $2\times$  area saving.

## 4 Conclusion

In conclusion, we proposed bit-slice  $\ell_1$  regularizer, the first algorithm specifically designed to train DNN models that are friendly for the bit sliced deployment on ReRAM crossbars. The proposed method can induce higher and more balanced sparsity levels among bit-slices of DNN weights comparing to traditional element-wise sparsity inducing training methods. The achieved bit-slice sparsity will enable significant reduction on the ADC energy and area overhead, and will further improve inference speed for ReRAM deployment.

## Acknowledgement

This work was supported in part by NSF CNS-1822085 and NSF CSR-1717885.

## References

- [1] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [2] K. Makantasis, K. Karantzas, A. Doulamis, and N. Doulamis. Deep supervised learning for hyperspectral data classification through convolutional neural networks. In *2015 IEEE IGARSS*. IEEE, 2015.
- [3] Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. In *Advances in neural information processing systems*, pages 1135–1143, 2015.
- [4] Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks. In *Advances in neural information processing systems*, pages 2074–2082, 2016.
- [5] P. Gysel. Ristretto: Hardware-oriented approximation of convolutional neural networks. *arXiv preprint arXiv:1605.06402*, 2016.
- [6] A. Polino, R. Pascanu, and D. Alistarh. Model compression via distillation and quantization. *arXiv preprint arXiv:1802.05668*, 2018.
- [7] S. Han, J. Kang, H. Mao, Y. Hu, X. Li, Y. Li, D. Xie, H. Luo, S. Yao, Yu Wang, H. Yang, and W. J. Dally. ESE: Efficient speech recognition engine with sparse lstm on fpga. In *Proceedings of FPGA '17*, 2017.
- [8] D. Shin, J. Lee, J. Lee, and H. Yoo. 14.2 dnpu: An 8.1tops/w reconfigurable cnn-rnn processor for general-purpose deep neural networks. In *2017 IEEE ISSCC*, 2017.
- [9] 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)*, 2016.
- [10] L. Song, X. Qian, H. Li, and Y. Chen. Pipelayer: A pipelined rram-based accelerator for deep learning. In *2017 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, 2017.
- [11] S. Ghodrati, H. Sharma, S. Kinzer, A. Yazdanbakhsh, K. Samadi, N. Kim, D. Burger, and H. Esmaeilzadeh. Mixed-Signal Charge-Domain Acceleration of Deep Neural networks through Interleaved Bit-Partitioned Arithmetic. *arXiv e-prints*, 2019.
- [12] H. Sharma, J. Park, N. Suda, L. Lai, B. Chau, V. Chandra, and H. Esmaeilzadeh. Bit fusion: Bit-level dynamically composable architecture for accelerating deep neural networks. In *Proceedings of ISCA*, 2018.
- [13] T. Yang, H. Cheng, C. Yang, I. Tseng, H. Hu, H. Chang, and H. Li. Sparse rram engine: Joint exploration of activation and weight sparsity in compressed neural networks. In *Proceedings of ISCA*, 2019.
- [14] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [15] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- [16] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *IEEE CVPR*, 2016.
- [17] M. Saberi, R. Lotfi, K. Mafinezhad, and W. A. Serdijn. Analysis of power consumption and linearity in capacitive digital-to-analog converters used in successive approximation adcs. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 2011.