

PANTHER: A Programmable Architecture for Neural Network Training Harnessing Energy-efficient ReRAM

Aayush Ankit, Izzat El Hajj, Sai Rahul Chalamalasetti, Sapan Agarwal, Matthew Marinella, Martin Foltin, John Paul Strachan, Dejan Milojicic, Wen-mei Hwu, and Kaushik Roy.

Abstract—The wide adoption of deep neural networks has been accompanied by ever-increasing energy and performance demands due to the expensive nature of training them. Numerous special-purpose architectures have been proposed to accelerate training: both digital and hybrid digital-analog using resistive RAM (ReRAM) crossbars. ReRAM-based accelerators have demonstrated the effectiveness of ReRAM crossbars at performing matrix-vector multiplication operations that are prevalent in training. However, they still suffer from inefficiency due to the use of serial reads and writes for performing the weight gradient and update step.

A few works have demonstrated the possibility of performing outer products in crossbars, which can be used to realize the weight gradient and update step without the use of serial reads and writes. However, these works have been limited to low precision operations which are not sufficient for typical training workloads. Moreover, they have been confined to a limited set of training algorithms for fully-connected layers only.

To address these limitations, we propose a bit-slicing technique for enhancing the precision of ReRAM-based outer products, which is substantially different from bit-slicing for matrix-vector multiplication only. We incorporate this technique into a crossbar architecture with three variants catered to different training algorithms. To evaluate our design on different types of layers in neural networks (fully-connected, convolutional, etc.) and training algorithms, we develop PANTHER, an ISA-programmable training accelerator with compiler support. Our design can also be integrated into other accelerators in the literature to enhance their efficiency. Our evaluation shows that PANTHER achieves up to $8.02\times$, $54.21\times$, and $103\times$ energy reductions as well as $7.16\times$, $4.02\times$, and $16\times$ execution time reductions compared to digital accelerators, ReRAM-based accelerators, and GPUs, respectively.

Index Terms—Accelerators, resistive random-access memory (ReRam), neural networks, training

1 INTRODUCTION

Deep Neural Networks (DNNs) have seen wide adoption due to their success in many domains such as image processing, speech recognition, and natural language processing. However, DNN training requires substantial amount of computation and energy which has led to the emergence of numerous special-purpose accelerators [1]. These accelerators have been built using various circuit technologies, including digital CMOS logic [2], [3] as well as hybrid digital-analog logic based on ReRAM crossbars [4], [5].

ReRAM crossbars are circuits composed of non-volatile elements that can perform Matrix-Vector Multiplication (MVM) in the analog domain with low latency and energy consumption. Since MVM operations dominate the performance of DNN inference and training, various inference [4], [5], [6] and training [7], [8] accelerators have been built using these crossbars. However, while inference algorithms do not modify matrices during execution, training algorithms modify them during the weight gradient and update step (weight gradient computation followed by the weight up-

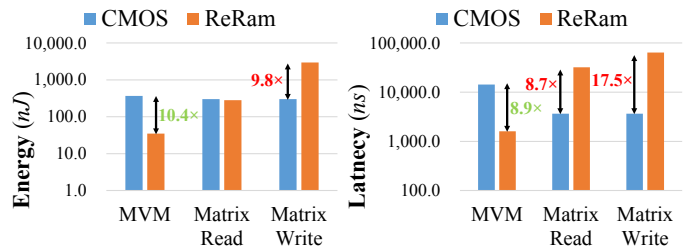


Fig. 1. Comparing CMOS and ReRAM Primitives

date). For this reason, training accelerators [7], [8] require frequent reads and write to crossbar cells to realize weight gradient and update operations. These reads and writes to ReRAM crossbars are performed one row at a time (like a typical memory array), and are referred to as *serial reads and writes* in this paper.

Figure 1 compares the energy and latency of CMOS and ReRAM technologies for various primitive operations. As shown, MVM consumes $\approx 10.4\times$ less energy and has $\approx 8.9\times$ lower latency with ReRAM over CMOS (at same area) for a 32 nm technology node. However, reading and writing the entire matrix consumes much higher energy and latency with ReRAM. Particularly, ReRAM writing energy and latency are an order of magnitude higher due to the cost of the program-verify approach which requires tens of pulses [9]. Therefore, the use of serial reads and writes during training takes away the overall benefits gained from using ReRAM for acceleration.

To overcome this issue, recent demonstrations [10], [11]

- A. Ankit, and K. Roy are with the Department of Electrical and Computer Engineering, Purdue University.
- I. E. Hajj is with the Department of Computer Science, American University of Beirut.
- S. R. Chalamalasetti, M. Foltin, J. P. Strachan, and D. Milojicic are with the Hewlett Packard Labs.
- S. Agarwal, and M. Marinella are with the Sandia National Labs.
- W. Hwu is with the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign.

Correspondance email: aankit@purdue.edu

have shown that Outer Product Accumulate (OPA) operations can be performed in crossbars to realize the weight gradient and update operations without the use of serial reads and writes. The OPA operation is performed by applying two input vectors at the rows and the columns of a crossbar simultaneously, to update each cell depending on the inputs at the corresponding row and column. However, these demonstrations are limited to low-precision inputs/outputs (2-4 bits) and weights (2-5 bits) which is not sufficient for the typical training workloads [12], [13]. Moreover, they are confined to Stochastic Gradient Descent (SGD) with batch size of one for fully-connected layers only.

To address these limitations, we propose a bit-slicing technique for achieving higher precision OPA operations by slicing the bits of the output matrix weights across multiple crossbars. While bit-slicing has previously been done for MVM operations [4], bit-slicing matrices to also support OPA operations is substantially different. For MVM, the rows and the crossbar cells are inputs and the columns are outputs, whereas for OPA, the rows and the columns are both inputs and the outputs are the crossbar cells themselves. Moreover, bit-slicing OPA presents additional constraints for the distribution of bits across the slices. First, weights are constant during MVM, but they change during OPA, which necessitates support for overflow within each slice and accounting for saturation. Second, MVM favors fewer bits per slice to reduce analog-to-digital Converter (ADC) precision requirements [4], but we show that OPA favors more bits per slice. Third, MVM favors homogeneous slicing of bits (equal number of bits per slice), but we show that OPA favors heterogeneous slicing.

We incorporate our proposed technique for enhancing OPA precision into a crossbar architecture that performs both MVM and OPA operations at high precision. We present three variants of the crossbar architecture that are catered to different training algorithms: SGD, mini-batch SGD, and mini-batch SGD with large batches. Using this crossbar architecture, we build PANTHER, a Programmable Architecture for Neural Network Training Harnessing Energy-efficient ReRAM. We use PANTHER to evaluate our design on different layer types (fully-connected, convolutional, etc.) and training algorithms. Our design can also be integrated into existing training accelerators in the literature to enhance their efficiency. Our evaluation shows that PANTHER achieves up to $8.02\times$, $54.21\times$, and $2,358\times$ energy reductions as well as $7.16\times$, $4.02\times$, and $119\times$ execution time reductions compared to digital accelerators, ReRAM-based accelerators, and GPUs, respectively.

We make the following contributions:

- A bit-slicing technique for implementing high-precision OPA operations using ReRAM crossbars (Section 3)
- A crossbar-based architecture, that embodies this bit-slicing technique, with three variants for different training algorithms (Section 4)
- An ISA-programmable accelerator with compiler support to evaluate different types of layers in neural networks and training algorithms (Section 5)

We begin with a background on the use of ReRAM crossbars for DNN training (Section 2).

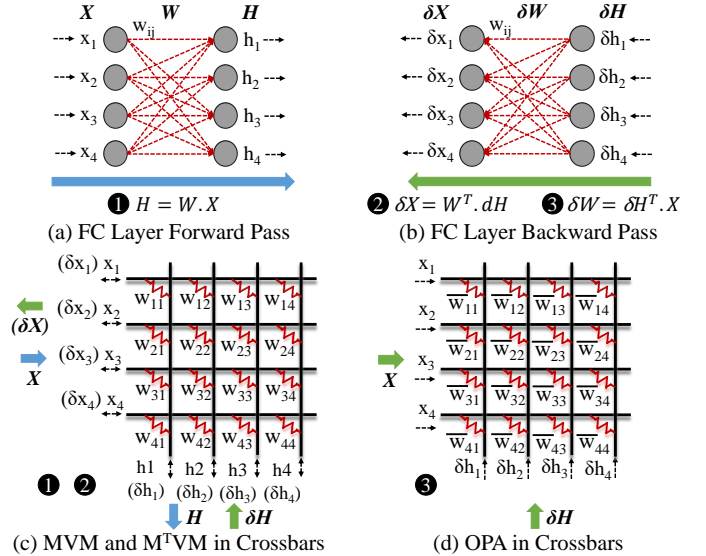


Fig. 2. FC Layer Matrix Operations in Crossbars

2 BACKGROUND

2.1 Deep Neural Network Training

Typical DNN training comprises of iterative updates to a model's weights in order to optimize the loss based on an objective function. Equations 1–4 show the steps involved in DNN training based on the Stochastic Gradient Descent (SGD) algorithm [14]. Equation 1 constitutes the forward pass which processes an input example to compute the activations at each layer. Equation 2 computes the output error and its gradient based on a loss function using the activations of the final layer. Equation 3 constitutes the backward pass which propagates the output error to compute the errors at each layer. Finally, equation 4 computes the weight updates to minimize the error.

$$H^{(l+1)} = W^{(l)} X^{(l)}, X^{(l+1)} = \sigma(H^{(l+1)}) \quad (1)$$

$$E = Loss(X^{(L)}, y), \delta H^{(L)} = \nabla E \odot \sigma'(X^{(L)}) \quad (2)$$

$$\delta H^{(l)} = [(W^{(l+1)})^T \delta H^{(l+1)}] \odot \sigma'(X^{(l)}) \quad (3)$$

$$\frac{\partial E}{\partial W^l} \text{ (or } \delta W^l) = X^{(l)} (\delta H^{(l+1)})^T, W^l = W^l - \eta * \frac{\partial E}{\partial W^l} \quad (4)$$

2.2 Using Crossbars for Training

The most computationally intensive DNN layers that are typical targets for acceleration are the *fully-connected* layers and the *convolutional* layers. We use fully-connected layers as an example to show how ReRAM crossbars can be used to accelerate DNN training workloads.

2.2.1 Overview of Fully Connected (FC) Layers

Figures 2(a) and (b) illustrate the operations involved during training in a FC layer. The training involves three types of matrix operations: ① *activation*, ② *layer gradients*, and ③ *weight gradients*. Activation corresponds to an MVM operation with the weight matrix (W), as shown in Equation 1. Layer gradients correspond to an MVM operation with the transpose of the weight matrix (hereon denoted as M^TVM), as shown in Equation 3. Weight gradients correspond to an outer product operation, the result of which is accumulated to the weight matrix based on the learning rate (η), as shown

in Equation 4. Therefore weight gradients and updates together can be viewed as an Outer Product Accumulate (OPA) operation on the weight matrix.

2.2.2 Activation and Layer Gradients in Crossbars

Figure 2(c) shows how a ReRAM crossbar can be used to compute activation and layer gradients. The weights of the matrix (W) are stored in the crossbar cells as the conductance state [15]. The MVM operation is realized by applying the input vector (X) as voltages on the rows of crossbar. Subsequently, the output vector (H) is obtained as currents from the columns. The M^TVM operation is realized by applying the input vector (δH) as voltages on the columns of the crossbar. Subsequently, the output vector (δX) is obtained as currents from the rows.

Both MVM and M^TVM operations execute $O(n^2)$ multiply-and-accumulate operations in one computational step in the analog domain (n is the crossbar size). Therefore, ReRAM crossbars can be leveraged to design highly efficient primitives for activation and layer gradient computations. For this reason, they have been extensively considered for DNN inference [4], [5], [6] and training [7], [8] accelerators.

2.2.3 Weight Gradients and Updates in Crossbars

Figure 2(d) shows how a ReRAM crossbar can be used to compute weight gradients. The OPA operation can be realized by applying the inputs (X and δH) as voltages on the crossbar’s rows and columns, respectively. The change ($\bar{w}_{ij} - w_{ij}$) in the value stored at a cross-point (i, j) is equal to the product of the voltage on row i and column j (details in Section 3). Therefore, the outer product operation in the crossbar is naturally fused with the weight matrix accumulate operation.

The OPA operation executes $O(n^2)$ multiply-and-accumulate operations in one computational step in the analog domain. It avoids serial reads and writes to ReRAM crossbar cells, which is important because reads and writes have orders of magnitude higher cost (energy and latency) than in-crossbar computations (MVM, M^TVM , OPA). Therefore, ReRAM crossbars can be leveraged to design highly efficient primitives for weight gradient computation and weight update.

The aforementioned technique has been demonstrated with low-precision inputs/outputs (2-4 bits) and weights (2-5 bits) on the SGD training algorithm for FC layers only [10], [11]. In this paper, we enhance the technique with architecture support to increase its precision and cater to a multiple training algorithms and different layer types.

3 ENHANCING RERAM-BASED OPA PRECISION

DNN workloads require 16 to 32 bits of precision for training [12], [13]. However, input digital-to-analog converters (DACs), crossbar cells, and output ADCs cannot support such levels of precision due to technology limitations and/or energy considerations. For this reason, accelerators that use ReRAM crossbars for MVM/ M^TVM operations typically achieve the required precision with bit-slicing [4], [6], [16], where matrix bits are sliced across the cells of multiple crossbars, input bits are streamed at the crossbar

rows/columns, and shift-and-add logic is used to combine the output bits at each column/row across crossbars (slices).

Bit-slicing matrices to also support OPA operations is different because both the rows and columns are simultaneously applied as inputs and the outputs are the crossbar cells themselves. Moreover, bit-slicing for OPA operations presents additional constraints for the choice of bit distribution across slices. This section describes our technique for bit-slicing the OPA operation (Section 3.1), and discusses the constraints it adds to the choice of bit distribution and how we address them (Sections 3.2 to 3.4).

3.1 Bit Slicing the OPA Operation

Figure 3(a) illustrates how the OPA operation is performed when 2-bit inputs are applied at the rows and the columns. The digital row input is encoded in the time-domain using pulse-width modulation. The digital column input is encoded in the amplitude-domain using pulse-amplitude modulation. Both pulse-width and pulse-amplitude modulations can be implemented using DACs. The weight change in a cell depends on the duration and the amplitude of the pulses applied on the corresponding row and column respectively, thereby realizing an analog OPA operation [10], [11].

To perform an OPA operation with 16-bit inputs, naively increasing the DAC resolution is infeasible because DAC power consumption grows rapidly with resolution (N) as:

$$P_{\text{DAC}} = \beta(2^N / N + 1)V^2 f_{\text{clk}} [17] \quad (5)$$

Instead, we propose an architectural scheme to realize a 16-bit OPA operation by bit-streaming the row input bits, bit-slicing the column input bits, and bit-slicing the matrix weights across multiple crossbars.

Figure 3(b) illustrates how we stream row input bits, m bits at a time over $16/m$ cycles. Meanwhile column input bits are left-shifted by m -bits every cycle. Since the number of cycles decrease linearly with m while the cycle duration increases exponentially with m due to pulse-width modulation of row input, we choose $m = 1$ to minimize total latency. Using $m = 1$ also means that the row DACs are just inverters, thereby having low power consumption.

Figure 3(c) shows how we slice column input bits across crossbars. Only one weight W_{ij} is shown for clarity. In each cycle, the left-shifted column input is divided into chunks of p bits ($p = 2$ in this example) and each chunk is applied to the corresponding crossbar.

Figure 3(d) illustrates the steps for a 16-bit \times 16-bit OPA operation at one crosspoint in the crossbar, resulting in a 32-bit output value for each matrix weight. It puts together the bit-streaming of the row input vector b and bit-slicing of the column input vector a with $p = 4$. Each dot represents a partial product ($a_n \cdot b_n$), and the color corresponds to a specific weight slice (crossbar). Thus, the net accumulation to a slice is the result of all partial products of the specific color. The updated weight after a time step T_n can be expressed as:

$$W_{\text{updated}} = W_{\text{old}} + \sum_{n=0}^n (a \ll n) * b_n \quad (6)$$

Crossbars store data in unsigned form. To enable positive and negative weight updates (δW), we represent inputs in

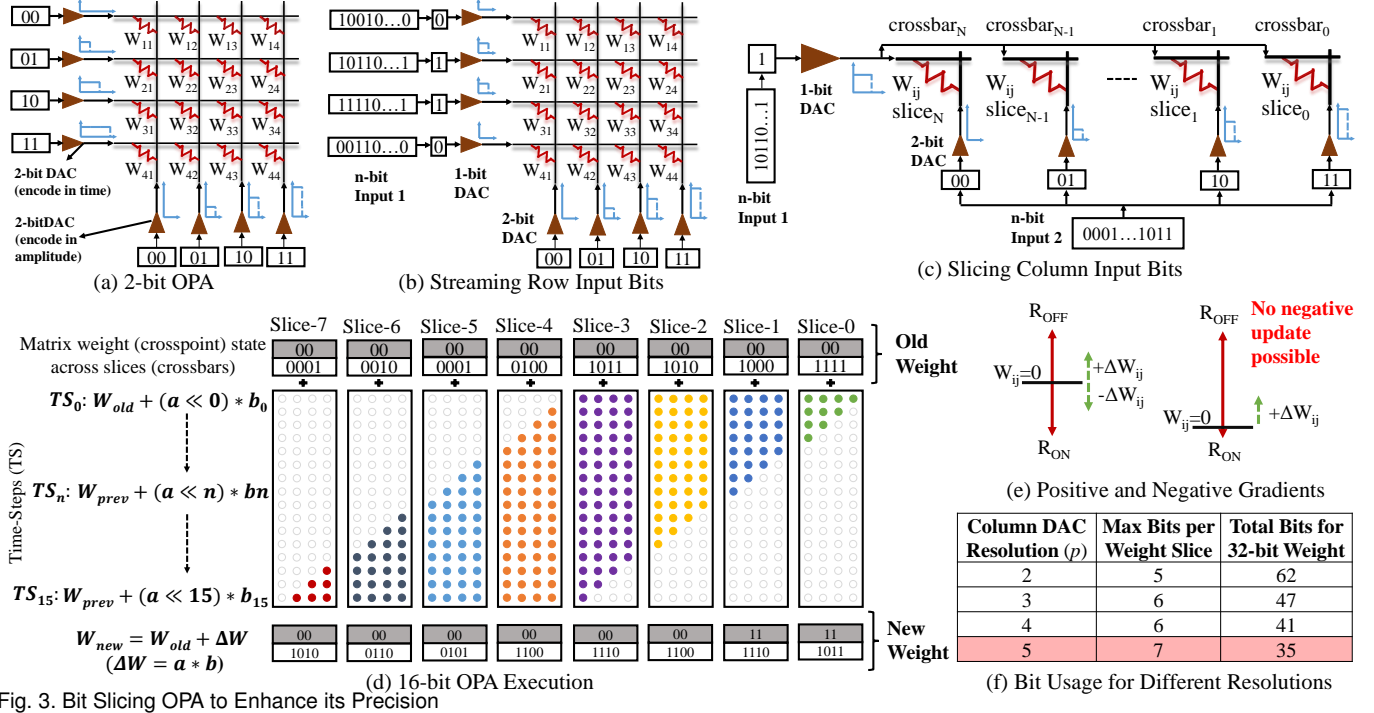


Fig. 3. Bit Slicing OPA to Enhance its Precision

the signed magnitude representation. To enable a symmetric representation of positive and negative weight updates, we bias each device such that, a zero weight (W_{ij}) is represented by the memory state $(R_{ON} + R_{OFF})/2$, as shown in Figure 3(e). Hence, the signed magnitude computation and biased data representation enable both positive and negative updates to weights. This is important as both polarities of updates are equally important in DNN training. Such a biased-representation can be implemented by adding an extra column per crossbar (128 rows, 128 columns) with minimal area/energy cost [18].

3.2 Bits to Handle Overflow

For MVM/ M^T VM, the matrix weights are inputs to the operation and they do not change. In contrast, for OPA, the matrix weights are accumulated with the values resulting from the outer product. As a result, the weight slice stored in a crossbar cell may overflow, either from multiple accumulations within one OPA or over multiple OPAs. We handle this overflow by provisioning weight slices with additional bits to store the carry (shaded bits shown in Figure 3(d)).

Propagating carry bits to other slices would require serial reads and writes which incur high overhead. For this reason, we do not propagate the carry bits immediately. Instead, they are kept in the slice and participate in future MVM/ M^T VM and OPA operations on the crossbar.

The carry bits cannot be kept in the weight slice indefinitely because eventually the weight slice may get saturated i.e. crossbar cell at maximum/minimum state for positive/negative update. Saturation is detrimental for *trainability* (desirable loss reduction during training) because it freezes training progress due to the absence of weight change. For this reason, we employ a periodic *Carry Resolution Step* (CRS) which executes infrequently to perform carry propagation using serial reads and writes. We evaluate the impact of the number of bits provisioned per slice and the CRS frequency on saturation and accuracy in Section 7.1.

3.3 Number of Slices vs. Bits Per Slice

When slicing matrix bits across multiple crossbars, there is a tradeoff between the number of slices and the number of bits per cell in each slice. MVM operations favor using more slices and fewer bits per slice. The reason is that energy increases linearly with the number of crossbars, and non-linearly with the precision of a crossbar due to the increase in ADC precision required to support it. Therefore, using more slices with fewer bits each is better for energy consumption.

In contrast, OPA favors having fewer slices with more bits per slice. The reason is that OPA introduces carry bits to each slice and having more slices with fewer bits each increases the overhead from the carry bits. For example, Figure 3(f) shows that with 2 bits per slice, 62 total bits are required to represent the 32-bit weight while capturing the carry bits adequately.

To strike a balance, we choose $p = 4$, since $p > 4$ requires a device precision that exceeds ReRAM technology limits [15]. A 4-bit DAC resolution is feasible because DAC power does not increase rapidly at low resolution (Equation 5). By choosing $p = 4$, our MVM/ M^T VM operations consume more energy than other ReRAM-based accelerators. However, our more energy efficient OPA operations compensate because they avoid the need for expensive serial reads and writes.

3.4 Heterogeneous Weight Slicing

MVM operations favor homogeneous bit-slicing. Increasing the precision of a slice while decreasing the precision of another is always an unfavorable tradeoff because energy increases nonlinearly with the precision of a crossbar. In contrast, for OPA operations where crossbar values change, provisioning more bits for slices that experience more weight updates helps reduce the frequency of saturation, thereby ensuring trainability at low CRS frequency.

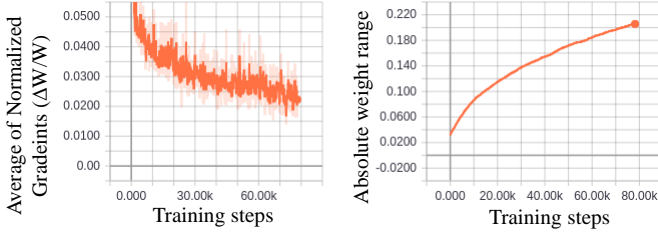


Fig. 4. Weight Gradients across Training Steps

Heterogeneous weight slicing provisions more bits for matrix slices that change more frequently. The frequency of change is impacted by two factors: OPA asymmetry and the small weight gradient range in DNNs. OPA asymmetry is illustrated in Figure 3(d) where the central slices receive more partial products (dots) than the edge slices, which motivates increasing precision for the central slices. Small weight gradient range is shown in Figure 4 where weight updates form a very small fraction (2% – 5%) of the overall weight range for $\geq 95\%$ of training steps, which motivates increasing precision of the lower slices. We evaluate the impact of heterogeneous weight slicing on energy and accuracy in Section 7.2.

4 MATRIX COMPUTATION UNIT (MCU)

The techniques described in Section 3 are incorporated into a Matrix Computation Unit (MCU) for DNN training accelerators. This section first describes the MCU’s organization (Section 4.1). It then describes the three variants of the MCU optimized for SGD (Section 4.2), mini-batch SGD (Section 4.3), and mini-batch SGD with large batches (Section 4.4).

4.1 MCU Organization

Figure 5 illustrates the organization of the MCU. Performing an MVM operation with the MCU is illustrated by the red arrow. Digital inputs stored in the *XBarIn* registers are fed to the crossbar rows through the *Input Driver*. The output currents from the crossbar columns are then converted to digital values using *ADC* and stored in the *XBarOut* registers.

Performing a M^TVM operation in the MCU is illustrated by the purple arrow in Figure 5. The key difference compared to the MVM operation is the addition of multiplexers to supply inputs to crossbar columns instead of rows and to read outputs from crossbar rows instead of columns.

MVM and M^TVM operations require 16 to 32 bits of precision for training [2]. We use 16-bit fixed-point representation for input/output data and 32-bit fixed-point representation for weight data which ensures sufficient precision [12].

Performing an OPA operation in the MCU is illustrated by the blue arrow in Figure 5. Digital inputs stored in the *XBarIn* registers are fed to the crossbar rows through the *Input Driver*. Digital inputs stored in the *XBarOut* registers are fed to the crossbar columns through the *Input Driver*. The effect of this operation is that the outer product of the input vectors is accumulated to the matrix stored in the crossbar. To support positive and negative inputs, the input drivers

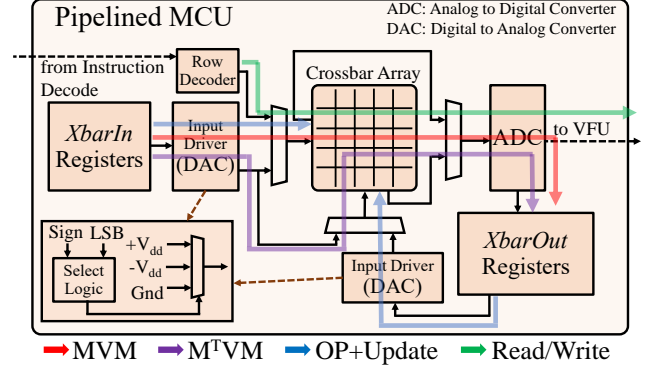


Fig. 5. Matrix Computation Unit

TABLE 1
Dataflow for SGD

Time Step	MCU0 (Layer1)	MCU1 (Layer2)	MCU2 (Layer3)
0	MVM (a_0) (a_1)		
1		MVM (a_1) (a_2)	
2			MVM (a_2) (a_3)
3			M^TVM (δh_3) (δh_2)
4		M^TVM (δh_2), (δh_1)	OP (a_2 , δh_3) (∇W_3)
5	OP (a_0 , δh_1) (∇W_1)	OP (a_1 , δh_2) (∇W_2)	

in Figure 5 use the sign bit (MSB) to drive the crossbar rows and columns with positive or negative voltages.

4.2 Variant #1 for SGD Acceleration

SGD-based training performs example-wise gradient descent. First, an input example performs a forward pass (MVM) to generate activations - H^1 . Next, the error computed with respect to the activation of the output layer is back propagated (M^TVM) to compute the layer gradients - δX^1 . Finally, the activations and layer gradients are used to update (OPA) the weight matrix - W^1 , before the next input example is supplied.

Table 1 illustrates the logical execution of matrix operations in three MCUs for a three-layer DNN with an input example a_0 . Each time step shows the operations executed on each MCU and their inputs/outputs. For example, at time step 0, MCU0 performs an MVM operation on input a_0 to compute the output a_1 . The illustration assumes that each layer maps on one MCU and does not show the interleaved nonlinear operations for clarity. For a layer size larger than one MCU capacity (128×128 matrix), the layer is partitioned across multiple MCUs (see Section 5.3).

Variant #1 of the MCU uses a single crossbar to perform all three matrix operations: MVM, M^TVM , and OPA. This variant is suitable for SGD because, as shown in Table 1, the three matrix operations are data dependent and will never execute concurrently. However, this variant creates structural hazards for mini-batch SGD as described in Section 4.3.

4.3 Variant #2 for Mini-Batch SGD Acceleration

Mini-batch SGD performs batch-wise gradient descent. Like SGD, each input performs MVM, M^TVM , and OPA to compute activations, layer gradients, and weight gradients/updates, respectively. However, the weight update is only reflected at the end of a batch to be used by the inputs of the next batch.

Table 2 illustrates the logical execution of matrix operations for a batch of five inputs, where $a_n m$ refers to the m^{th}

TABLE 2
Dataflow for Mini-Batch SGD

Time Step	MCU0 (Layer1)	MCU1 (Layer2)	MCU2 (Layer3)
0	MVM (a_{00}) (a_{01})		
1	MVM (a_{10}) (a_{11})	MVM (a_{01}) (a_{02})	
2	MVM (a_{20}) (a_{21})	MVM (a_{11}) (a_{12})	MVM (a_{02}) (a_{03})
3	MVM (a_{30}) (a_{31})	MVM (a_{21}) (a_{22})	MVM (a_{12}) (a_{13})
4	MVM (a_{40}) (a_{41})	MVM (a_{31}) (a_{32})	M ^T VM (δh_{03}) (δh_{02})
5		M ^T VM (δh_{02}), (δh_{01})	M ^T VM (δh_{13}) (δh_{12})
6		MVM (a_{41}) (a_{42})	MVM (a_{32}) (a_{33})
7		M ^T VM (δh_{12}), (δh_{11})	M ^T VM (δh_{23}) (δh_{22})
8		M ^T VM (δh_{22}), (δh_{21})	MVM (a_{42}) (a_{43})
9		M ^T VM (δh_{32}), (δh_{31})	M ^T VM (δh_{43}) (δh_{42})
10		M ^T VM (δh_{42}), (δh_{41})	
9-12	OP (a_{n0} , δh_{n1}) (∇W_{n1})	OP (a_{n1} , δh_{n2}) (∇W_{n2})	OP (a_{n2} , δh_{n3}) (∇W_{n3})

Iterate for $n=1$ to 4

activation of n^{th} input. MVM operations can be executed for multiple input examples concurrently in a pipelined fashion (MVM (a_{10}) (a_{11}), MVM (a_{01}) (a_{02}) in Table 2). Additionally, the MVM and M^TVM operations for different inputs in the batch can also execute in parallel during the same timestep, provided that there is no structural hazard on the MCU. The desire to eliminate such structural hazards motivates Variant #2.

Variant #2 of the MCU eliminates structural hazards in mini-batch SGD by storing two copies of the matrix on different crossbars, enabling the MCU to perform MVM and M^TVM in parallel. This replication improves the *energy-delay product* for a batch. With $< 2\times$ increase in area, we improve the batch latency by $O(L)$, where L is the number of layers. The ISA instruction for performing MVM/M^TVM (Section 5.2) is designed to enable the compiler (Section 5.3) to schedule these two operations in parallel on the same MCU.

The OPA operations are executed at the end of the mini-batch (steps 9-12 in Table 2) to reflect the weight updates for the entire batch. These OPA operations require that the vectors involved are saved until then. Variant #2 saves these vectors in shared memory. However, if the batches are large, this approach puts too much stress on the shared memory which motivates Variant #3 (Section 4.4).

4.4 Variant #3 for Mini-Batch SGD with Large Batches

For mini-batch SGD with very large batch sizes, saving the vectors in shared memory requires large shared memory size which degrades storage density. Variant #3 alleviates the pressure shared memory size by maintaining three copies of each crossbar. The first two copies enable performing MVM and M^TVM in parallel, similar to Variant #2. The third copy is used to perform the OPA operation eagerly, as soon as its vector operands are available, without changing the matrices being used by the MVM and M^TVM operations.

Performing OPA eagerly avoids saving vectors until the end, reducing the pressure on the shared memory. However, using a third crossbar for OPA requires serial reads and writes to commit the weight updates to the first and the second crossbars for MVM and M^TVM in the next batch. Section 7.6 discusses the impact of these design choices.

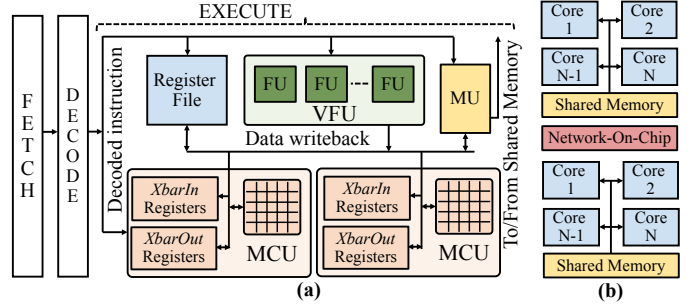


Fig. 6. Architecture Overview

5 PROGRAMMABLE ACCELERATOR

The MCU described in Section 4 can be integrated with prior ReRAM-based training accelerators [7], [8] to improve their efficiency. We develop a programmable training accelerator named PANTHER to evaluate our design by extending the PUMA ReRAM-based inference accelerator [6]. This section describes PANTHER’s organization (Section 5.1), ISA considerations (Section 5.2), compiler support (Section 5.3), and an example of how to implement convolutional layers (Section 5.4).

5.1 Accelerator Organization

PANTHER is a spatial architecture organized in three tiers: nodes, tiles, and cores. A node consists of multiple tiles connected via an on-chip network, and a tile consists of multiple cores connected to a shared memory, as illustrated in Figure 6(b). A core consists of multiple MCUs for executing matrix operations, a digital CMOS-based vector functional unit (VFU) for executing arithmetic operations and non-linear functions, a register file, and a load/store memory unit. A core also features an instruction execution pipeline making the accelerator ISA-programmable. To support DNNs whose model storage exceeds a node’s total MCU capacity, multiple nodes can be connected via an interconnect. This organization is similar to PUMA’s [6] and is not a contribution of this paper. The key distinction from PUMA is the MCU which supports M^TVM and OPA operations, not just MVM operations, as described in Section 4.

5.2 ISA Considerations

The PUMA [6] ISA includes *mvm* instructions executed by crossbars, arithmetic/logic/nonlinear instructions executed by the VFU, load/store instructions to access shared memory, send/receive instructions to communicate with other tiles, and control flow instructions. We extend the PUMA ISA to also include a *mcu* instruction for executing all three matrix operations (MVM, M^TVM, OPA) on the MCU.

The *mcu* instruction takes six 3-bit masks, where each mask corresponds to one of the MCUs on the core (up to six). The three bits in the mask correspond to the three supported matrix operations (MVM, M^TVM, OPA). If multiple bits are set, then the instruction executes the operations concurrently. For example, if mask 0 is set to ‘110’ and mask 1 is set to ‘011’, then MCU 0 will execute MVM and M^TVM simultaneously and MCU 1 will execute M^TVM and OPA simultaneously. The incorporation of all three operations into a single instruction is important for being able to execute

them concurrently in order to leverage the parallelism in batch-wise training on Variant #2 (Section 4.3) and Variant #3 (Section 4.4). Furthermore, having separate masks for each MCU within a core helps leverage the parallelism across matrix operations [6]. The mask is generated by fusing different MCU operations as discussed in Section 5.3. The *mcu* instruction does not take source and destination operands since these are implied to by *XBarIn* and *XBarOut*.

The semantic of the OPA operation is that it takes effect at the end of the execution when a special *halt* instruction is invoked. This semantic allows the same code to work for any of the three MCU variants, making the choice of variant a microarchitectural consideration and the ISA agnostic to it. The implementation of the OPA semantic on each of the variants is as follows. Consider the case when all three bits of an MCU’s mask are set. In Variant #1, MVM and M^TVM will be serialized on the same crossbar, while the operands of OPA will be saved to shared memory then applied to that crossbar when *halt* is invoked. In Variant #2, MVM and M^TVM will be executed in parallel on the two crossbar copies, while the operands of OPA will be treated like in Variant #1. In Variant #3, MVM and M^TVM will be executed in parallel on the first two crossbar copies, while the operands of OPA will be applied to the third crossbar. The values of the third crossbar will then be copied to the first two crossbars when *halt* is invoked.

5.3 Compiler Support

The PUMA [6] compiler provides a high-level programming interface in C++ that allows programmers to express models in terms of generic matrix and vector operations. The compiler is implemented as a runtime library that builds a computational graph when the code is executed then compiles the graph to PUMA ISA code. The compiler partitions matrices into sub-matrices and maps these sub-matrices to different MCUs, cores, and tiles. It then maps the operations in the graph to different MCUs, cores, and tiles accordingly, inserting communication operations where necessary. The compiler then linearizes the graph, creating an instruction sequence for each core. It performs register allocation for each sequence, spilling registers to shared memory if necessary. Finally, it generates ISA code for each core, collectively comprising a kernel that runs on the accelerator.

We make the following extensions to the PUMA compiler to support PANTHER. We extend the application programming interface (API) to allow programmers to define training matrices that support MVM, M^TVM , and OPA operations. We extend the intermediate representation to represent these matrices and include them in the partitioning. We also add an analysis and transformation pass for identifying MCU operations in the graph that can be fused and fusing them. This pass fuses MCU operations that do not have data dependences between them and that use different MCUs on the same core or use the same MCU but are different types of operations (MVM, M^TVM , OPA). The fusing process is iterative because every time operations are fused, new dependences are introduced to the graph. Finally, we extend the code generator to support the new *mcu* ISA instruction.

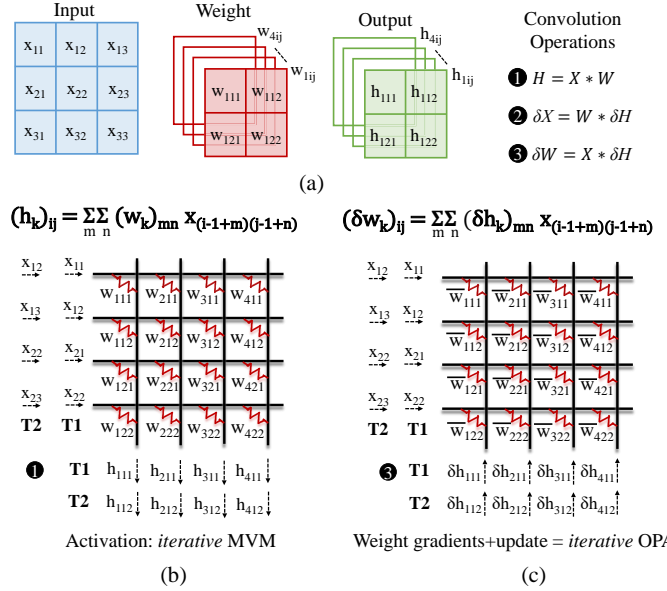


Fig. 7. Convolutional Layer Matrix Operations in Crossbars

Note that since the model weights are not updated until the *halt* instruction at the end, the scope of a kernel is a single batch. Multiple batches are executed by invoking the kernel multiple times on different input data.

5.4 Implementing Convolutional Layers

ReRAM-based OPA has one-to-one correspondence to the weight gradient/update operation for FC layers (discussed in Section 2.2.3). By integrating this technique into a programmable accelerator with compiler support, we enable the mapping of more complex layers on top of it such as convolutional layers. This section describes how convolutional layers can be implemented in our accelerator.

Figure 7(a) shows a typical convolution layer and the associated operations during training. Like with FC layers, convolutional layers performs three types of matrix operations: 1 *activation*, 2 *layer gradients*, and 3 *weight gradients*. Unlike FC layers, these operations are all convolutions (*).

5.4.1 Activation and Layer Gradients

Figure 7(b) shows how the convolution operation for activation is implemented in the crossbar on top of the MVM primitive. This approach is similar to that used in existing accelerators [8]. The crossbar stores the *convolution kernel* in the form of linearized filters (w_k), where each column corresponds to the weights associated with a specific output channel (h_k). The convolution operation to compute activations is implemented as an *iterative* MVM operation. An iteration is represented as a time step (T1/T2) in Figure 7(b), and corresponds to a specific (i,j) pair. A block of input features (X) is applied to the crossbar’s rows as *convolution data* in each iteration. In a similar manner, the convolution operation for layer gradients (not shown in the figure) is realized using iterative M^TVM . The next layer’s errors (δH) are used as the convolution data and flipped filters (vertically and horizontally) are used as the convolution kernel.

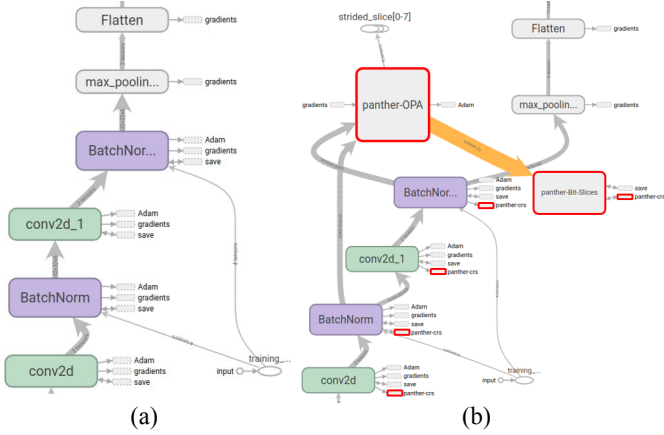


Fig. 8. Computational graph obtained using TensorBoard for (a) example model (b) example model with PANTHER OPA

5.4.2 Weight Gradients

Figure 7(c), shows our proposed technique for implementing the weight gradients convolution operation and weight update in the crossbar on top of the OPA primitive. The weight gradient computation uses input features (X) as the convolution data and output feature’s errors (δH) as the convolution kernel. Each iteration is represented as a time step (T1, T2) in Figure 7(c), and corresponds to a specific (i,j) pair. On every iteration, the output feature’s errors are applied on the columns, in a depth major order. Simultaneously, by applying the portion of input features that generate the corresponding activations (H) on the rows, a *partial convolution* is obtained between X and δH . Striding across the output feature’s errors and input features for n^2 time steps, where n is size of one output feature map, realizes the full convolution operation. Convolutions for different output feature maps are performed in parallel across the crossbar’s columns, using the same weight data layout as used in MVM and M^T MV operations. To the best of our knowledge, our work is the first to formulate the weight gradients convolution operation in terms of outer products.

5.4.3 Comparison with Other Accelerators

Existing ReRAM-based training accelerators such as PipeLayer [8] do not compute the weight gradient convolutions using outer products, but rather, they compute them using MVM operations. This requires writing the convolution kernel (δH) on the crossbar because the convolution operation here uses *non-stationary* data (δH) as the convolution kernel. The drawback of this approach is that the latency and energy consumption of the serial reads and writes is very high, taking away from the overall efficiency provided by ReRAM-based MVMs.

6 METHODOLOGY

6.1 Architecture Simulator

We extend the PUMA [6] simulator to model the MCU unit and its associated instructions. The PUMA simulator is a detailed cycle-level architecture simulator that runs applications compiled by the compiler, in order to evaluate the execution of benchmarks. The simulator models all the necessary events that occur in an execution cycle, including

TABLE 3
Summary of platforms

Parameter	PANTHER (1 node)	Base _{digital} (1 node)	2080-Ti (1 card)
SIMD lanes	108 M	108 M	4352
Technology	CMOS-ReRam (32 nm)	CMOS (32 nm)	CMOS (12 nm)
Frequency	1 GHz	1 GHz	1.5 GHz
Area	117 mm ²	578 mm ²	750 mm ²
TDP	105 W	839 W	250 W
On-Chip Memory	72.4 MB	72.4 MB	29.5 MB

TABLE 4
Details of workloads

Layer	C	M	H/W	R/S	E/F	Wt (MB)	In (MB)	Ops/B
CNN-Vgg16								
Conv1	3	64	32	3	32	0.003	0.006	368.640
Conv2	32	64	32	3	16	0.035	0.063	92.160
Conv3	64	128	16	3	16	0.141	0.031	209.455
Conv4	128	128	16	3	8	0.281	0.063	52.364
Conv5	128	256	8	3	8	0.563	0.016	62.270
Conv6	256	256	8	3	8	1.125	0.031	62.270
Conv7	256	256	8	3	4	1.125	0.031	15.568
Conv8	256	512	4	3	4	2.250	0.008	15.945
Conv9	512	512	4	3	4	4.500	0.016	15.945
Conv10	512	512	4	3	2	4.500	0.016	3.986
Conv11	512	512	2	3	2	4.500	0.004	3.997
Conv12	512	512	2	3	2	4.500	0.004	3.997
Conv13	512	512	2	3	1	4.500	0.004	0.999
Dense14	512	4096	-	-	-	4.000	0.001	1.000
Dense15	4096	4096	-	-	-	32.000	0.008	1.000
Dense16	4096	100	-	-	-	0.781	0.008	0.990
MLP-L4								
Dense1	1024	256	-	-	-	0.500	0.002	0.996
Dense2	256	512	-	-	-	0.250	0.000	0.998
Dense3	512	512	-	-	-	0.500	0.001	0.998
Dense4	512	10	-	-	-	0.010	0.001	0.909

compute, memory and NoC transactions. To estimate power and timing of the CMOS digital logic components, their RTL implementations are synthesized to the IBM 32nm SOI technology library, and evaluated using the Synopsys Design Compiler. For the on-chip SRAM memories, the power and timing estimates are obtained from Cacti 6.0. Subsequently, the power and timing of each component are incorporated in the cycle-level simulator in order to estimate the energy consumption.

MCU Modelling. Since the MCU is built with analog components and cannot be synthesized with publicly available libraries, we adopted the models from past works [4], [10] and ADC survey [19]. We use the ReRam crossbar array and sample-and-hold circuit models in ISAAC [4]. We used capacitive DACs and Successive Approximation Register (SAR) ADCs. The DAC area and power are estimated using the equations described in Saberi et al. [20]. The ADCs for different precisions namely 8-12 bits operating at a sampling frequency of 1GHz are obtained from the ADC survey [19]. The ADC optimization technique in Newton [21] is incorporated to avoid unnecessary ADC conversions.

6.2 Functional Simulator

We implement a functional simulator using TensorFlow that models PANTHER’s bit-sliced OPA technique. This simulator enables performing design space exploration (for accuracy) on large-scale DNNs to explore the bounds on heterogeneous weight slicing and CRS frequency for trainability. Here, a layer’s weights are represented as a multi-dimensional tensor of shape $S \times M \times N$, where S corresponds to a weight slice (discussed in Figure 3 (d)), and M and N correspond to the weight matrix’s dimensions respectively. Each weight slice can have a unique bit-precision, to model heterogeneous configurations (Section 3.4). The weight values beyond the range permissible by the bit-

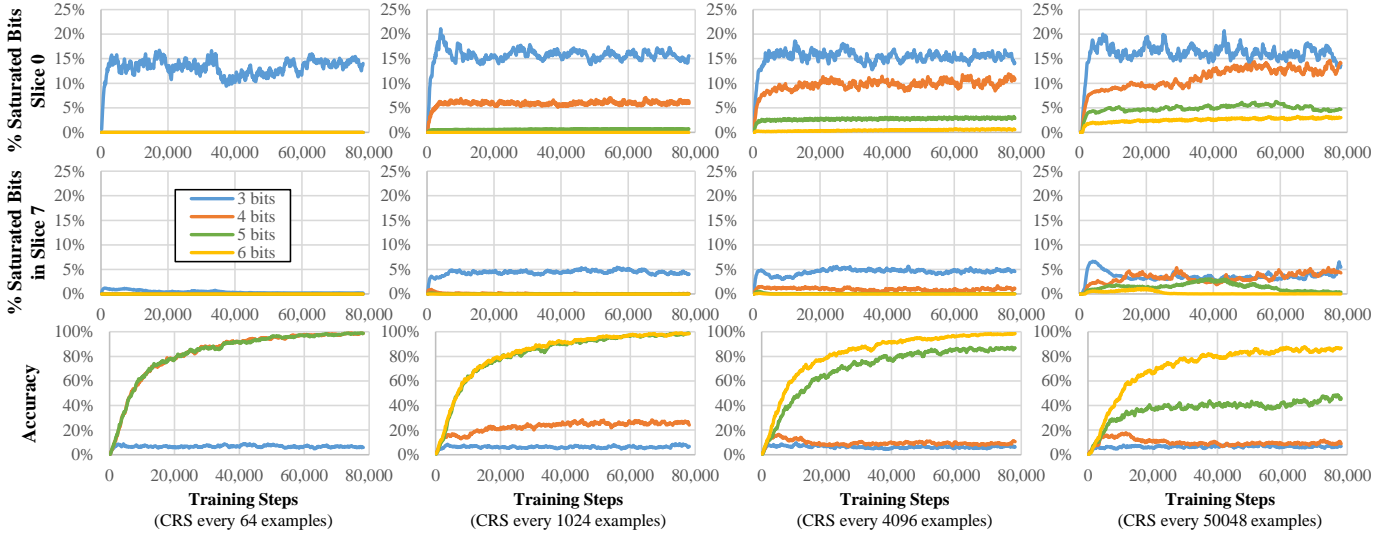


Fig. 9. Impact of Slice Bits and CRS Frequency on Accuracy

precision are clipped to model a slice’s saturation. Subsequently, the weight update operation in native TensorFlow is modified to quantize and bit-slice the computed weight gradients and then update the previous copy of weights (already quantized and bit-sliced). Figures 8 (a) and (b) show the computational graphs for an example neural network model, and the example model augmented with PANTHER OPA operation (shown in red) respectively.

6.3 Baselines

We evaluate PANTHER against three weight-stationary ASIC baselines: $Base_{digital}$, $Base_{mvm}$, and $Base_{opa/mvm}$, as well as one NVIDIA GPU platform - Turing RTX 2080-Ti (2080-Ti).

$Base_{digital}$ uses a digital version of the MCU where weights are stored in an SRAM array within the core and matrix operations are performed with a digital VFU. $Base_{digital}$ is an adaptation of the digital baseline used in PUMA [6]. As shown in the PUMA work, this digital baseline is an optimistic estimate of the Google TPU [3]. It is optimistic because it uses weight-stationary MVM computations similar to TPU, but assumes that the entire model is mapped using on-chip SRAM, thereby avoiding the off-chip memory access costs in TPU. Therefore, our comparisons with $Base_{digital}$ also serve as a lower-bound on PANTHER’s improvements compared to TPU. The objective of comparing with $Base_{digital}$ is to demonstrate the benefit of ReRAM-based computing over pure digital approaches.

$Base_{mvm}$ uses ReRAM for MVM and M^TVM , and a digital VFU for OPA with serial reads/writes to the crossbar. $Base_{opa/mvm}$ is a replication of PipeLayer’s [8] approach described in Section 5.4.3 and only applies to convolutional layers. It uses ReRAM for MVM and M^TVM , and realizes OPA with ReRAM MVMs and serial reads/writes. The objective of comparing with $Base_{mvm}$ and $Base_{opa/mvm}$ is to demonstrate the benefit of ReRAM-based OPA operations.

Configurations. $Base_{mvm}$ and $Base_{opa/mvm}$ use 32-bit weights sliced across 16 slices with 2 bits each, which is optimal since crossbars only do MVM/ M^TVM . PANTHER uses heterogeneous weight slicing with 32-bit weights represented using 39 bits sliced across 8 slices distributed from MSB to LSB like so: 44466555 (unless otherwise specified).

For this reason, PANTHER consumes 17.5% higher energy for MVM/ M^TVM than $Base_{mvm}$ and $Base_{opa/mvm}$ due to higher ADC precision. We also use a CRS frequency of 1024 steps (unless otherwise specified) which achieves similar accuracy as the software implementation. For all three ASIC baselines and PANTHER, the hierarchical organization uses 138 tiles per node, with 8 cores per tile and 2 MCUs per core. Table 3 summarizes the platforms. Note that both $Base_{mvm}$ and $Base_{opa/mvm}$ have same platform parameters as PANTHER.

6.4 Workloads

We use a 4-layered MLP model and Vgg-16 CNN model on SVHN and CIFAR-100 datasets, respectively. Table 4 details the layer details of the two models and their computational intensity (operations to byte ratio). The individual layers of the chosen MLP and CNN models span a wide range of computational intensity observed across the spectrum of neural network workloads. Thus, our workloads are well representative of the large variety of layer types found in neural network models such as fully-connected, 2 D-convolution, point-wise convolution, etc.

Similar to other ReRAM training accelerators [7], [8], we use fixed-point arithmetic which has been shown to be successful for training large DNNs [13]. We use the CIFAR-100 dataset for CNN which is comparable to the ImageNet dataset in terms of training difficulty [22], [23]. However, ImageNet’s large image sizes make it difficult to run the training flow without actual hardware (CIFAR-100 requires 2 days and ImageNet requires 1 month on the simulator).

7 EVALUATION

7.1 Impact of Slice Bits and CRS Frequency on Accuracy

Figure 9 shows the impact of the number of bits used per slice (uniform weight slicing) and CRS frequency for the CNN benchmark. We analyze the percentage of saturated cells per slice for a lower order and higher order slice, and their implications on CNN’s Top-5 training accuracy.

Using 3 bits per slice shows significantly higher percentage of saturated cells for the lower order slice (Slice

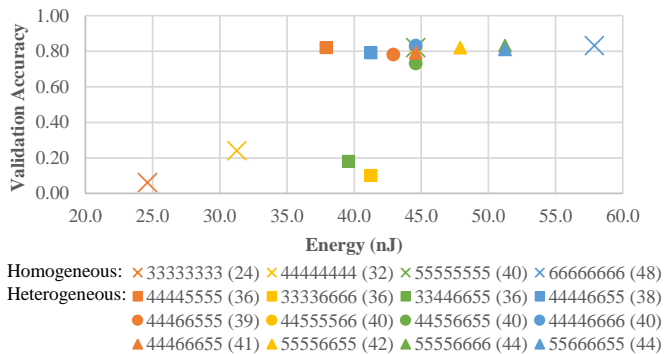


Fig. 10. Heterogeneous Weight-Slicing

0) than other configurations. Further, increasing the CRS frequency does not reduce the saturation fraction of Slice 0 at 3-bits. Consequently, the training accuracy with 3-bits slices remains very low throughout the training steps.

Using 4 bits per slice performs well at high CRS frequency (CRS every 64 steps), but does not scale well at lower CRS frequencies. A high CRS frequency is undesirable due to the high cost of serial reads and writes incurred during carry propagation between discrete slices.

Slices with 5-bits and 6-bits are robust to repeated weight updates as they exhibit lower saturation for both lower order and higher order slices even at low CRS frequencies (every 1024 or 4096 steps). Note that the cost of a CRS operation at low frequency (every 1024 steps) has negligible impact on overall energy and performance ($\leq 4.8\%$).

Figure 9 also motivates heterogeneous weight slicing because it shows that the higher order slice has significantly lower saturation in general than the lower order slice.

7.2 Impact of Heterogeneous Weight Slicing

Figure 10 shows the accuracy and energy of sixteen slicing configurations. Generally speaking, increasing the total number of bits improves accuracy by reducing saturation, but it also increases energy because it requires higher precision ADCs for MVM and M^TVM . The graph shows that heterogeneous weight slicing enables favourable accuracy-energy tradeoffs, enabling lower energy at comparable accuracy or better accuracy at comparable energy. Provisioning ≥ 4 bits for the four higher order slices (4 – 7) and ≥ 5 bits for the four lower order slices (0 – 3) ensures desirable accuracy. Any configuration using 3 bit slices (irrespective of total bits) leads to significant accuracy degradation. Note that the configuration used in the rest of the evaluation (44466555) is not a Pareto-optimal one, so our energy numbers in the rest of the evaluation are underestimated.

7.3 Variant #1 SGD Energy Comparison

Figure 11 compares the layer-wise energy consumption of PANTHER’s Variant #1 to that of all three baselines for SGD.

Base_{digital}. Compared to *Base_{digital}*, we achieve $7.01\times-8.02\times$ reduction in energy. This advantage is due to the energy efficiency of computing MVM, M^TVM , and OPA in ReRAM.

Base_{mvm}. Compared to *Base_{mvm}*, we achieve $31.03\times-54.21\times$ reductions in energy for FC layers (Layers 1-4 in

MLP and 14-16 in CNN) and $1.47\times-31.56\times$ for convolution layers (Layers 1-13), with the later (smaller) convolution layers showing larger reductions. Recall that *Base_{mvm}* uses serial reads and writes to perform the OPA operation with digital logic. While the large convolutional layers can amortize these reads and writes, the FC layers and small convolutional layers do not have enough work to do so which is why they suffer relatively. In contrast, PANTHER avoids these reads and writes by performing OPA in the crossbar (11.37 nJ).

Base_{opamvm}. *Base_{opamvm}* behaves similarly to *Base_{mvm}*. Recall that both baselines perform serial reads and writes to crossbars for OPA, but *Base_{mvm}* uses CMOS VFUs while *Base_{opamvm}* uses ReRAM MVMs. Since ReRAM MVMs and CMOS OPAs have comparable energy consumption (35.10 nJ and 37.28 nJ respectively), the overall energy of the two baselines is similar.

7.4 Variant #2 Mini-Batch SGD Energy

Figure 12 compares the layer-wise energy consumption of Variant #2 of PANTHER to that of all three baselines for Mini-Batch SGD with batch size 64. Compared to SGD results (Figure 11), the key difference is that having multiple batches before weight updates amortizes the cost of serial reads and writes in *Base_{mvm}* and *Base_{opamvm}* (smaller blue bar). Our energy improvements therefore come mainly from reducing OPA energy. Energy is reduced by $1.61\times-2.16\times$ for fully connected layers for *Base_{mvm}* and *Base_{opamvm}*. It is reduced by $1.18\times-1.63\times$ and $1.22\times-2.45\times$ for convolutional layers for *Base_{mvm}* and *Base_{opamvm}*, respectively.

For very large batch sizes such as 1,024 (not shown in the figure), ReRAM writes can be completely amortized by *Base_{mvm}* and *Base_{opamvm}*. In this case, PANTHER reduces energy by $\simeq 1.18\times$ compared to *Base_{mvm}* and *Base_{opamvm}* due to reducing OPA energy. However, batch sizes preferred by ML practitioners for DNN training (32, 64) are typically smaller than what is required to amortize the ReRAM memory access costs because large batch sizes have adverse effects on DNN generalization [24].

7.5 Variant #2 Execution Time

Figure 13 compares the layer-wise execution time of Variant #2 to all three baselines for different batch sizes.

Base_{digital}. Compared to *Base_{digital}*, we have consistently lower execution time due to faster MVM, M^TVM , and OPA operations in ReRAM.

Base_{mvm}. For MLPs with small batch sizes, *Base_{mvm}* significantly suffers because the ReRAM write latency is not amortized. However, for larger batch sizes and for CNNs, the ReRAM write latency is amortized. Nevertheless, we still outperform *Base_{mvm}* across all batch sizes because of lower latency ReRAM OPA. In fact, our advantage grows with batch size because OPA consumes a larger percentage of the total time for larger batches since the forward and backward passes benefit from pipeline parallelism whereas OPA operations are serialized at the end.

Base_{opamvm}. *Base_{opamvm}* behaves similarly to *Base_{mvm}* for convolutional layers.

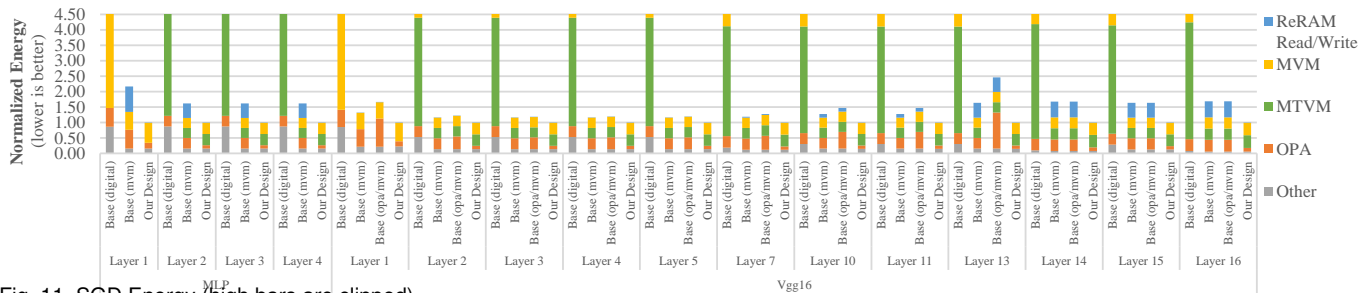


Fig. 11. SGD Energy (high bars are clipped)

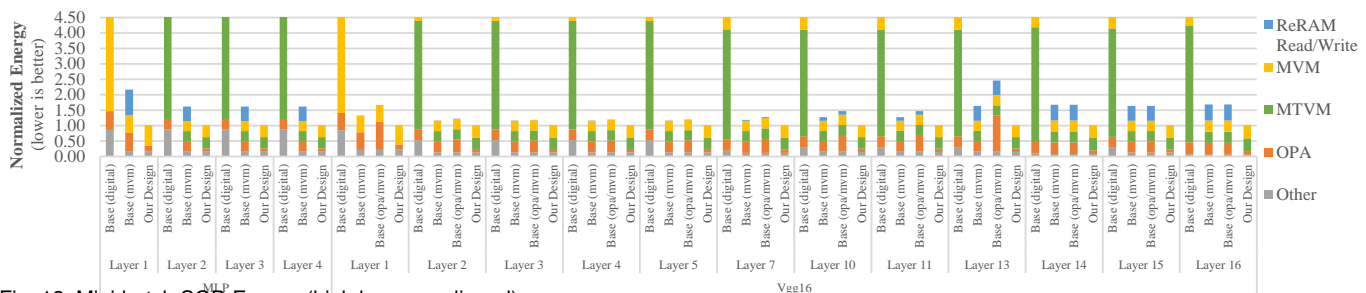


Fig. 12. Mini-batch SGD Energy (high bars are clipped)

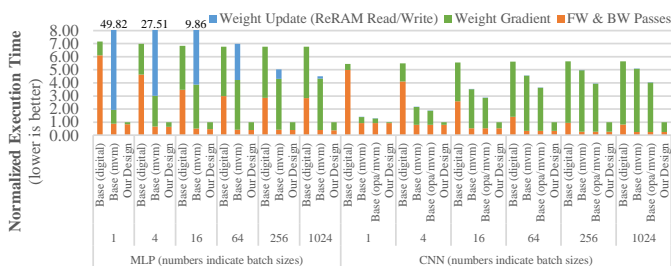


Fig. 13. Execution Time

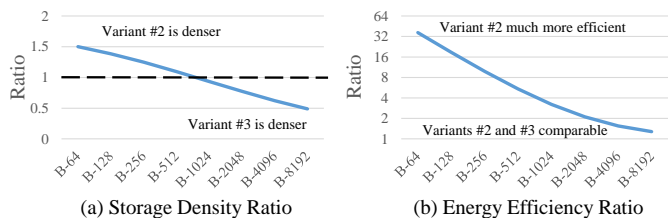


Fig. 14. Variant #2 vs. Variant #3

7.6 Comparing Variants #2 and #3

Increasing the batch size for mini-batch SGD increases Variant #2's shared memory requirements for storing all activations and layer gradients in the batch, degrading its storage density. Variant #3 uses a third crossbar for eagerly computing and storing weight gradients, thereby keeping shared memory requirements low at the expense of higher energy to commit the updates to the other crossbars at the end. Figure 14 shows that Variant #2 has better storage density and energy efficiency for small batch sizes, while Variant #3 has better storage density for very large batch sizes at comparable energy efficiency.

7.7 Comparison with GPUs

Figure 15 compares the energy consumption and execution time of Variant #2 with a 2080-Ti GPU for SGD (batch size 1) and Mini-Batch SGD (batch sizes 64 and 1k). Our design

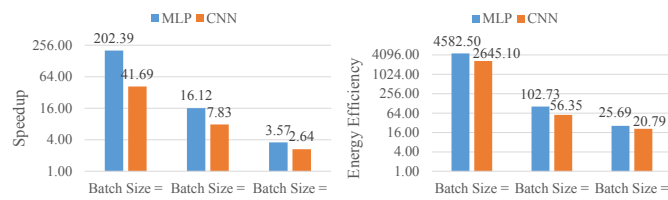


Fig. 15. PANTHER's speedup and energy-efficiency compared to GPU

significantly reduces energy consumption and execution time due to the use of energy-efficient and highly parallel ReRAM-based matrix operations.

GPUs rely on data reuse to hide memory access latency. For this reason, their relative performance is worse for MLP compared compared to CNN, and for smaller batch sizes compared to larger ones. Our design enables efficient training for a wide spectrum of batch sizes (small to large). Training based on small batch sizes is common in emerging applications such as lifelong learning [25] and online reinforcement learning [26], where training does not rely on any earlier collected dataset.

7.8 Sensitivity to ReRAM endurance

ReRAM devices have finite switching (1 to 0, 0 to 1) endurance of 10^9 conservative writes [27], [28], which limits their applicability towards on-chip memories for typical workloads. However, the small magnitude of typical weight updates make ReRAM feasible for DNN training. Considering a 5% average conductance change per batch, the lifetime of a chip will be ≈ 6 years (assuming 50% reduction from failed training flows), for 1,000 trainings per year where each training is comprised of 100 epochs, 64 batch-size and 1M training examples (typical parameters in state-of-the-art image recognition benchmarks [29]). While weight slicing makes lower order slices more prone to degradation arising from limited endurance, adding redundancy at lower order slices and higher endurance from technology improvements (currently shown in spintronics [30]) can make the chip more robust.

8 RELATED WORK

Various ReRAM-based training accelerators [7], [8] have been proposed, but they rely on expensive serial reads and writes to accomplish weight updates. We avoid these reads and writes by leveraging the in-crossbar OPA operations [10], [11], and extending their precision for practical trainability. Our crossbar architecture can be used to enhance existing accelerators.

ReRAM-based accelerators have also been proposed for DNN inference [4], [5], [6], [16], graph processing [31], scientific computing [32], and general purpose data parallel applications [33]. Our work focuses on DNN training.

Analog [34], [35] and DRAM-based [36], [37], [38] accelerators have been proposed as alternatives to digital-CMOS accelerators. Our work uses ReRAM as an alternative.

Many accelerators use digital CMOS technology for accelerating DNNs, including those that mainly target inference [1] or also target training [39]. Our work uses hybrid digital-analog computation based on ReRAM crossbars, not just CMOS.

Recent works have explored training DNNs with reduced precisions in floating-point arithmetic domain such as bfloat16 [40], float8 [41] as well as fixed-point arithmetic domain [13], [42]. While floating-point arithmetic is not amenable to ReRAM-based hardware (without modifications), the reductions in fixed-point precision can be exploited in PANTHER by reducing the MCU width (number of slices) to improve training energy and time.

ReRAM technology suffers from imprecise writes due to non-idealities (noise and non-linearity) and manufacturability issues (stuck-at-faults and process variations). However, the iterative nature of DNN training and careful re-training helps recover the accuracy loss from non-idealities [43], faults [44], and variations [45]. Re-training is a fine-tuning process (typically 1 epoch) with insignificant cost compared to training.

9 CONCLUSION

We propose a bit-slicing technique for enhancing the precision of ReRAM-based OPA operations to achieve sufficient precision for DNN training. We incorporate our technique into a crossbar architecture that performs high-precision MVM and OPA operations, and present three variants catered to different training algorithms: SGD, mini-batch SGD, and mini-batch SGD with large batches. Finally, to evaluate our design on different layer types and training algorithms, we develop PANTHER, an ISA-programmable training accelerator with compiler support. Our evaluation shows that PANTHER achieves up to $8.02\times$, $54.21\times$, and $103\times$ energy reductions as well as $7.16\times$, $4.02\times$, and $16\times$ execution time reductions compared to digital accelerators, ReRAM-based accelerators, and GPUs, respectively. The proposed accelerator explores the feasibility of ReRAM technology for DNN training by mitigating their serial read and write limitations, and can pave the way for efficient design of future machine learning systems.

ACKNOWLEDGEMENT

This work was supported by the Center for Brain-inspired Computing (C-BRIC), one of six centers in JUMP, a DARPA

sponsored Semiconductor Research Corporation (SRC) program; and Hewlett Packard Labs. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

REFERENCES

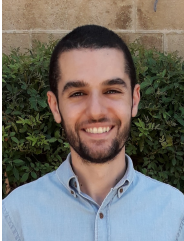
- [1] Vivienne Sze, Yu-Hsin Chen, Tien-Ju Yang, and Joel Emer. Efficient processing of deep neural networks: A tutorial and survey. *arXiv preprint arXiv:1703.09039*, 2017.
- [2] Yunji Chen, Tao Luo, Shaoli Liu, Shijin Zhang, Liqiang He, Jia Wang, Ling Li, Tianshi Chen, Zhiwei Xu, Ninghui Sun, et al. DaDianNao: A machine-learning supercomputer. In *Proceedings of the 47th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 609–622. IEEE Computer Society, 2014.
- [3] Norman P. Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, Al Borchers, Rick Boyle, Pierre-luc Cantin, Clifford Chao, Chris Clark, Jeremy Coriell, Mike Daley, Matt Dau, Jeffrey Dean, Ben Gelb, Tara Vazir Ghaemmaghami, Rajendra Gottipati, William Gulland, Robert Hagmann, C. Richard Ho, Doug Hogberg, John Hu, Robert Hundt, Dan Hurt, Julian Ibarz, Aaron Jaffey, Alek Jaworski, Alexander Kaplan, Harshit Khaitan, Daniel Killebrew, Andy Koch, Naveen Kumar, Steve Lacy, James Laudon, James Law, Diemthu Le, Chris Leary, Zhuyuan Liu, Kyle Lucke, Alan Lundin, Gordon MacKean, Adriana Maggiore, Maire Mahony, Kieran Miller, Rahul Nagarajan, Ravi Narayanaswami, Ray Ni, Kathy Nix, Thomas Norrie, Mark Omernick, Narayana Penukonda, Andy Phelps, Jonathan Ross, Matt Ross, Amir Salek, Emad Samadiani, Chris Severn, Gregory Sizikov, Matthew Snelham, Jed Souter, Dan Steinberg, Andy Swing, Mercedes Tan, Gregory Thorson, Bo Tian, Horia Toma, Erick Tuttle, Vijay Vasudevan, Richard Walter, Walter Wang, Eric Wilcox, and Doe Hyun Yoon. In-datacenter performance analysis of a tensor processing unit. In *Proceedings of the 44th Annual International Symposium on Computer Architecture*, ISCA'17, pages 1–12, New York, NY, USA, 2017. ACM.
- [4] Ali Shafiee, Anirban Nag, Naveen Muralimanohar, Rajeev Balasubramanian, John Paul Strachan, Miao Hu, R Stanley Williams, and Vivek 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, pages 14–26. IEEE Press, 2016.
- [5] Ping Chi, Shuangchen Li, Cong Xu, Tao Zhang, Jishen Zhao, Yongpan Liu, Yu Wang, and Yuan 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, pages 27–39, Piscataway, NJ, USA, 2016. IEEE Press.
- [6] Aayush Ankit, Izzat El Hajj, Sai Rahul Chalamalasetti, Geoffrey Ndu, Martin Foltin, R Stanley Williams, Paolo Faraboschi, Wenmei W Hwu, John Paul Strachan, Kaushik Roy, and Milojevic Dejan. Puma: A programmable ultra-efficient memristor-based accelerator for machine learning inference. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 715–731. ACM, 2019.
- [7] Ming Cheng, Lixue Xia, Zhenhua Zhu, Yi Cai, Yuan Xie, Yu Wang, and Huazhong Yang. Time: A training-in-memory architecture for memristor-based deep neural networks. In *Proceedings of the 54th Annual Design Automation Conference 2017*, page 26. ACM, 2017.
- [8] Linghao Song, Xuehai Qian, Hai Li, and Yiran Chen. Pipelayer: A pipelined ReRAM-based accelerator for deep learning. In *High Performance Computer Architecture (HPCA), 2017 IEEE International Symposium on*, pages 541–552. IEEE, 2017.
- [9] Emmanuelle J Merced-Grafals, Noraica Dávila, Ning Ge, R Stanley Williams, and John Paul Strachan. Repeatable, accurate, and high speed multi-level programming of memristor 1t1r arrays for

- power efficient analog computing applications. *Nanotechnology*, 27(36):365202, 2016.
- [10] Matthew J Marinella, Sapan Agarwal, Alexander Hsia, Isaac Richter, Robin Jacobs-Gedrim, John Niroula, Steven J Plimpton, Engin Ipek, and Conrad D James. Multiscale co-design analysis of energy, latency, area, and accuracy of a ReRAM analog neural training accelerator. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 8(1):86–101, 2018.
- [11] Pritish Narayanan, Alessandro Fumarola, Lucas L Sanches, Kohji Hosokawa, SC Lewis, Robert M Shelby, and Geoffrey W Burr. Toward on-chip acceleration of the backpropagation algorithm using nonvolatile memory. *IBM Journal of Research and Development*, 61(4/5):11–1, 2017.
- [12] Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaev, Ganesh Venkatesh, et al. Mixed precision training. *arXiv preprint arXiv:1710.03740*, 2017.
- [13] Shuang Wu, Guoqi Li, Feng Chen, and Luping Shi. Training and inference with integers in deep neural networks. *arXiv preprint arXiv:1802.04680*, 2018.
- [14] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning internal representations by error propagation. Technical report, California Univ San Diego La Jolla Inst for Cognitive Science, 1985.
- [15] Miao Hu, Catherine Graves, Can Li, Yunning Li, Ning Ge, Eric Montgomery, Noraica Davila, Hao Jiang, R Stanley Williams, J Joshua Yang, Qiangfei Xia, and John Paul Strachan. Memristor-based analog computation and neural network classification with a dot product engine. *Advanced Materials*, 2018.
- [16] Mahdi Nazm Bojnordi and Engin Ipek. Memristive Boltzmann machine: A hardware accelerator for combinatorial optimization and deep learning. In *High Performance Computer Architecture (HPCA), 2016 IEEE International Symposium on*, pages 1–13. IEEE, 2016.
- [17] Yulhwa Kim, Hyungjun Kim, Daehyun Ahn, and Jae-Joon Kim. Input-splitting of large neural networks for power-efficient accelerator with resistive crossbar memory array. In *Proceedings of the International Symposium on Low Power Electronics and Design*, page 41. ACM, 2018.
- [18] Son Ngoc Truong and Kyeong-Sik Min. New memristor-based crossbar array architecture with 50-% area reduction and 48-% power saving for matrix-vector multiplication of analog neuromorphic computing. *Journal of semiconductor technology and science*, 14(3):356–363, 2014.
- [19] Boris Murmann. ADC performance survey 1997-2011. <http://www.stanford.edu/~murmann/adcsurvey.html>, 2011.
- [20] Mehdi Saberi, Reza Lotfi, Khalil Mafinezhad, and Wouter 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*, 58(8):1736–1748, 2011.
- [21] Anirban Nag, Rajeev Balasubramonian, Vivek Srikumar, Ross Walker, Ali Shafiee, John Paul Strachan, and Naveen Murali-manohar. Newton: Gravitating towards the physical limits of crossbar acceleration. *IEEE Micro*, 38(5):41–49, 2018.
- [22] Yonatan Geifman. cifar-vgg. <https://github.com/geifmany/cifar-vgg/blob/master/README.md>, 2018.
- [23] BVLC. caffe. <https://github.com/BVLC/caffe/wiki/Models-accuracy-on-ImageNet-2012-val>, 2017.
- [24] Dominic Masters and Carlo Luschi. Revisiting small batch training for deep neural networks. *arXiv preprint arXiv:1804.07612*, 2018.
- [25] Zhiyuan Chen and Bing Liu. Lifelong machine learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 10(3):1–145, 2016.
- [26] Bradley C Stadie, Sergey Levine, and Pieter Abbeel. Incentivizing exploration in reinforcement learning with deep predictive models. *arXiv preprint arXiv:1507.00814*, 2015.
- [27] Zhiqiang Wei, Y Kanzawa, K Arita, Y Katoh, K Kawai, S Muraoka, S Mitani, S Fujii, K Katayama, M Iijima, et al. Highly reliable taor reram and direct evidence of redox reaction mechanism. In *2008 IEEE International Electron Devices Meeting*, pages 1–4. IEEE, 2008.
- [28] J Joshua Yang, M-X Zhang, John Paul Strachan, Feng Miao, Matthew D Pickett, Ronald D Kelley, G Medeiros-Ribeiro, and R Stanley Williams. High switching endurance in tao x memristive devices. *Applied Physics Letters*, 97(23):232102, 2010.
- [29] Wei Yang. pytorch-classification. <https://github.com/bearpaw/pytorch-classification/blob/master/TRAINING.md>, 2017.
- [30] Xuanyao Fong, Yusung Kim, Rangharajan Venkatesan, Sri Harsha Choday, Anand Raghunathan, and Kaushik Roy. Spin-transfer torque memories: Devices, circuits, and systems. *Proceedings of the IEEE*, 104(7):1449–1488, 2016.
- [31] Linghao Song, Youwei Zhuo, Xuehai Qian, Hai Li, and Yiran Chen. GraphR: Accelerating graph processing using ReRAM. In *High Performance Computer Architecture (HPCA), 2018 IEEE International Symposium on*, pages 531–543. IEEE, 2018.
- [32] Ben Feinberg, Uday Kumar Reddy Vengalam, Nathan Whitehair, Shibo Wang, and Engin Ipek. Enabling scientific computing on memristive accelerators. In *2018 ACM/IEEE 45th Annual International Symposium on Computer Architecture, ISCA'18*, pages 367–382. IEEE, 2018.
- [33] Daichi Fujiki, Scott Mahlke, and Reetuparna Das. In-memory data parallel processor. In *Proceedings of the Twenty-Third International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 1–14. ACM, 2018.
- [34] Robert LiKamWa, Yunhui Hou, Julian Gao, Mia Polansky, and Lin Zhong. RedEye: analog ConvNet image sensor architecture for continuous mobile vision. In *Proceedings of the 43rd International Symposium on Computer Architecture, ISCA'16*, pages 255–266. IEEE Press, 2016.
- [35] Prakalp Srivastava, Mingu Kang, Sujan K Gonugondla, Sungmin Lim, Jungwook Choi, Vikram Adve, Nam Sung Kim, and Naresh Shanbhag. PROMISE: an end-to-end design of a programmable mixed-signal accelerator for machine-learning algorithms. In *Proceedings of the 45th Annual International Symposium on Computer Architecture, ISCA'18*, pages 43–56. IEEE Press, 2018.
- [36] Mingyu Gao, Jing Pu, Xuan Yang, Mark Horowitz, and Christos Kozyrakis. Tetris: Scalable and efficient neural network acceleration with 3d memory. In *Proceedings of the Twenty-Second International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 751–764. ACM, 2017.
- [37] Duckhwan Kim, Jaeha Kung, Sek Chai, Sudhakar Yalamanchili, and Saibal Mukhopadhyay. Neurocube: A programmable digital neuromorphic architecture with high-density 3d memory. In *Proceedings of the ACM/IEEE 43rd Annual International Symposium on Computer Architecture, ISCA'16*, pages 380–392. IEEE, 2016.
- [38] Shuangchen Li, Dimin Niu, Krishna T Malladi, Hongzhong Zheng, Bob Brennan, and Yuan Xie. Drisa: A DRAM-based reconfigurable in-situ accelerator. In *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 288–301. ACM, 2017.
- [39] Swagath Venkataramani, Ashish Ranjan, Subarno Banerjee, Dipankar Das, Sasikanth Avancha, Ashok Jagannathan, Ajaya Durg, Dheemanth Nagaraj, Bharat Kaul, Pradeep Dubey, and Anand Raghunathan. ScaleDeep: a scalable compute architecture for learning and evaluating deep networks. In *Proceedings of the 44th Annual International Symposium on Computer Architecture, ISCA'17*, pages 13–26, New York, NY, USA, 2017. ACM.
- [40] Dhiraj Kalamkar, Dheevatsa Mudigere, Naveen Mellempudi, Dipankar Das, Kunal Banerjee, Sasikanth Avancha, Dharma Teja Vooturi, Nataraj Jammalamadaka, Jianyu Huang, Hector Yuen, et al. A study of bfloat16 for deep learning training. *arXiv preprint arXiv:1905.12322*, 2019.
- [41] Naigang Wang, Jungwook Choi, Daniel Brand, Chia-Yu Chen, and Kailash Gopalakrishnan. Training deep neural networks with 8-bit floating point numbers. In *Advances in neural information processing systems*, pages 7675–7684, 2018.
- [42] Yukuan Yang, Shuang Wu, Lei Deng, Tianyi Yan, Yuan Xie, and Guoqi Li. Training high-performance and large-scale deep neural networks with full 8-bit integers. *arXiv preprint arXiv:1909.02384*, 2019.
- [43] Sapan Agarwal, Steven J Plimpton, David R Hughart, Alexander H Hsia, Isaac Richter, Jonathan A Cox, Conrad D James, and Matthew J Marinella. Resistive memory device requirements for a neural algorithm accelerator. In *2016 International Joint Conference on Neural Networks (IJCNN)*, pages 929–938. IEEE, 2016.
- [44] Chenchen Liu, Miao Hu, John Paul Strachan, and Hai Helen Li. Rescuing memristor-based neuromorphic design with high defects. In *Proceedings of the 54th Annual Design Automation Conference 2017*, page 87. ACM, 2017.
- [45] Lerong Chen, Jiawen Li, Yiran Chen, Qiuping Deng, Jiyuan Shen, Xiaoyao Liang, and Li Jiang. Accelerator-friendly neural-network training: Learning variations and defects in RRAM crossbar. In *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2017*. IEEE, mar 2017.



Aayush Ankit received the B.Tech. (2015) from Indian Institute of Technology (BHU), Varanasi. He is pursuing PhD degree in Electrical and Computer Engineering at Purdue University and has been a research assistant to Prof. Kaushik Roy since 2015. His research interests lie in hardware and software design for efficient machine learning. During his PhD, he has done internships as – ML Architect at HPE Labs, Palo Alto (2017); CPU Designer at Intel Corporation, Hillsboro (2017); and GPU Architect at Samsung

ACL, San Jose (2019).



Izzat El Hajj is an Assistant Professor in the Department of Computer Science at the American University of Beirut. His research interests are in software support for emerging processor and memory technologies. Izzat received his M.S. (2014) and Ph.D. (2019) in Electrical and Computer Engineering at the University of Illinois at Urbana-Champaign and his B.E. (2011) in Electrical and Computer Engineering at the American University of Beirut.



Sai Rahul Chalamalasetti is a Sr. Research Engineer in the Systems Architecture Lab of the Hewlett Packard Labs in Palo Alto, CA. Before that he was a hardware engineer in the Silicon Design Lab of the HP Server division in Houston, TX. He is involved in designing In-Memory Machine Learning Accelerators, Next Generation Memory Interface (Gen-Z) and Application Accelerations using FPGAs. He is a Senior Member of the IEEE and act as a reviewer and program committee member for research conferences.

He received the M.S degree and Ph.D. degree from the University of Massachusetts Lowell in 2009 and 2012, respectively, both in computer engineering. He has authored or co-authored over twenty five publications in peer-reviewed journals and conference proceedings.



Sapan Agarwal is a Senior Member of Technical Staff at Sandia National Laboratories. He is leading projects to develop analog in-memory computing accelerators for high performance computing, database accelerators, and neural network training and inference. He is also leading the development of simulation tools to model the impact of device level radiation events on a running algorithm. He created a CrossSim, a crossbar simulator for analog neural network accelerators that allows for the co-design of novel

devices and algorithms by modeling the impact of device properties on neural algorithms. He has invented several new semiconductor device concepts including new low power transistors, artificial analog synapses, and wide band gap LEDs. He works with the IEEE International Roadmap for Devices and Systems to map out future computing architectures and devices.



Matthew Marinella received the Ph.D. degree in electrical engineering from Arizona State University, Tempe, AZ, USA, in 2008, under the supervision of D. Schroder. He is currently a Principal Member of the Technical Staff with Sandia National Labs, Livermore, CA, USA. He is a Principal Investigator for Sandia's Nonvolatile Memory Program and numerous neuromorphic and low-power computing projects. He chairs the Emerging Memory Devices Section for the IRDS Roadmap Beyond CMOS Chapter. He serves on

various technical program committees, including the IEEE International Conference on Rebooting Computing.



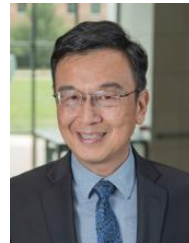
John Paul Strachan is a research team leader at Hewlett Packard Labs in Palo Alto, CA. His team develops computational accelerators from emerging device technology with interests in applications from machine learning, network security, and optimization. He received his undergraduate and PhD from MIT and Stanford University, respectively. He has over 50 patents and has co-authored or authored over 80 peer-reviewed papers. He has previously worked on nanomagnetic devices for memory for which he

was awarded the Falicov Award from the American Vacuum Society, and has developed sensing systems for precision agriculture in a company which he co-founded. He serves in many professional societies including IEEE IEDM ExComm, the Nanotechnology Council ExComm, past program chair and steering member of the International Conference on Rebooting Computing and related communities.



Dejan Milojcic is a distinguished technologist and director at Hewlett Packard Labs, Palo Alto, CA [1998-]. Previously, he worked in the OSF Research Institute, Cambridge, MA [1994-1998] and Institute "Mihajlo Pupin", Belgrade, Serbia [1983-1991]. His areas of expertise include system software and distributed systems. He received his PhD from University of Kaiserslautern, Germany (1993); and MSc/BSc from Belgrade University, Serbia (1983/86). Dejan was a managing director of the Open Cirrus

Cloud Computing testbed (2007-2011). Dejan has over 200 papers, 2 books and 52 granted patents. Dejan is an IEEE Fellow (2010), ACM Distinguished Engineer (2008), and HKN and USENIX member. He was president of IEEE Computer Society (2014) and IEEE Presidential candidate (2019). He has been on many conference program committees and journal editorial boards.



Wen-mei Hwu joined NVIDIA in February 2020 as Senior Distinguished Research Scientist, after spending 32 years at the University of Illinois at Urbana-Champaign, where he was a Professor, Sanders-AMD Endowed Chair, Acting Department Head and Chief Scientist of the Parallel Computing Institute. Hwu and his Illinois team developed the superblock compiler scheduling and optimization framework that has been adopted by virtually all modern vendor and open-source compilers today. For his research

contributions, Hwu received the ACM SigArch Maurice Wilkes Award, the ACM Grace Murray Hopper Award, the IEEE Computer Society Charles Babbage Award, the ISCA Influential Paper Award, the MICRO Test-of-Time Award, the IEEE Computer Society B. R. Rau Award, the CGO Test-of-Time Award and the Distinguished Alumni Award in CS of the University of California, Berkeley. He has also won numerous best paper awards for major conferences. He is a Fellow of IEEE and ACM.



Kaushik Roy received the B.Tech. degree from IIT Kharagpur, India and the Ph.D. degree from University of Illinois at Urbana-Champaign, USA in 1990. He was with Texas Instruments, Dallas, TX where he was involved in field-programmable gate array architecture development and low-power circuit design. He joined the Electrical and Computer Engineering Faculty, Purdue University, West Lafayette, IN in 1993, where he is currently an Edward G. Tiedemann Jr. Distinguished Professor. He has authored over 600

papers in refereed journals and conferences, holds 15 patents, graduated 60 Ph.D. students, and is a coauthor of two books. Dr. Roy has received the U.S. National Science Foundation Career Development Award, IBM Faculty Partnership Award, the ATT/Lucent Foundation Award, the 2005 SRC Technical Excellence Award, the SRC Inventors Award, the Purdue College of Engineering Research Excellence Award, the Humboldt Research Award, the IEEE Circuits and Systems Society Technical Achievement Award, the Distinguished Alumnus Award from IIT Kharagpur, and the Fulbright-Nehru Distinguished Chair. He has also won numerous best paper awards for major conferences, and has been on the Editorial board for top journals.