

# Software-Level Accuracy Using Stochastic Computing With Charge-Trap-Flash Based Weight Matrix

Varun Bhatt\*, Shalini Shrivastava†, Tanmay Chavan† and Udayan Ganguly†

\* Department of Computing Science, University of Alberta, Edmonton, Canada  
Email: vbhatt@ualberta.ca

† Department of Electrical Engineering, Indian Institute of Technology Bombay, Mumbai, India

**Abstract**—The in-memory computing paradigm with emerging memory devices has been recently shown to be a promising way to accelerate deep learning. Resistive processing unit (RPU) has been proposed to enable the vector-vector outer product in a crossbar array using a stochastic train of identical pulses to enable one-shot weight update, promising intense speed-up in matrix multiplication operations, which form the bulk of training neural networks. However, the performance of the system suffers if the device does not satisfy the condition of linear conductance change over around 1,000 conductance levels. This is a challenge for nanoscale memories. Recently, Charge Trap Flash (CTF) memory was shown to have a large number of levels before saturation, but variable non-linearity. In this paper, we explore the trade-off between the range of conductance change and linearity. We show, through simulations, that at an optimum choice of the range, our system performs nearly as well as the models trained using exact floating point operations, with less than 1% reduction in the performance. Our system reaches an accuracy of 97.9% on MNIST dataset, 89.1% and 70.5% accuracy on CIFAR-10 and CIFAR-100 datasets (using pre-extracted features). We also show its use in reinforcement learning, where it is used for value function approximation in Q-Learning, and learns to complete an episode the mountain car control problem in around 146 steps. Benchmarked to state-of-the-art, the CTF based RPU shows best in class performance to enable software equivalent performance.

**Index Terms**—Deep Learning, Neuromorphic Hardware, Crossbar Array, Flash

## I. INTRODUCTION

Deep Learning [1] has become the core driving force of artificial intelligence (AI). Applications such as image recognition, playing games, self-driving cars, and AI assistants are all made possible with the help of deep learning. At the core of deep learning lies artificial neural networks (ANNs) [2]. ANNs are trained using large sets of data to approximate a function that explains the given data. Training is done using backpropagation [3], in which the weights of the neural network are updated based on gradient descent update rule.

The majority of the operations in training ANNs are matrix multiplications. Graphics processing units (GPUs) and Tensor processing units (TPUs) are specialized digital hardware designed to speed up this matrix multiplication. With faster computation cores, the bottleneck is currently in memory systems and data transfer [4]. Moreover, training ANNs for

a typical real-world application requires hundreds of years of GPU time [5], leading to high energy costs.

In-memory computing [6] is an emerging paradigm, where data transfer is minimized by storing data and performing computation at the same place. Crossbar arrays with non-volatile memory have been shown to use lower energy, while also reaping the benefits of in-memory computation. Unfortunately, most of the devices struggle with precision and hence, the resulting performance of the system is not on par with their digital counterparts.

Gokmen and Vlasov [7] proposed a hypothetical resistive processing unit (RPU) that can be used to accelerate ANNs while being more energy-efficient than GPUs and having a negligible loss in accuracy. A crossbar architecture with a stochastic weight update rule allowed matrix multiplication in  $\mathcal{O}(1)$  time. Linearity in weight update of the cross-point device and a high number of conductance levels were shown to be necessary to ensure good accuracy.

Various approaches with nanoscale emerging memories like PCM [8] and RRAM [9] have shown insufficient linearity to enable RPU as the sole memory. Recently, traditional charge trap flash memory has shown promising linearity [10], [11]. However, their performance in the RPU framework has not been explored.

In this paper, we present a charge trap flash device that can act as a cross-point device in the RPU framework. We experimentally show a high number of conductance levels and approximately linear updates by choosing appropriate pulse width and voltage for weight update. Through simulations, we show that it indeed leads to a good accuracy when tested on MNIST, CIFAR-10 and CIFAR-100 datasets. In addition to supervised learning problems, we also successfully train a reinforcement learning agent on the Mountain Car environment.

## II. RELATED WORK

Matrix-vector multiplication and vector-vector outer product form the bulk of operations while training neural network. RPU [7] speeds up this computation using stochastic multiplication and hypothetical devices with linear weight updates.

Electronic synapses that have been proposed, such as nanoscale memristive synapses, may not have the gradual learn-

ing required for RPU. Phase-change memory (PCM) based synapse has gradual positive conductance change, but abrupt negative conductance change, which requires novel synapse circuit design with enhanced controller complexity as well as a dual precision approach. Successful methods supplement weight storage in low precision but compact PCM with high precision but area inefficient CMOS based memory to achieve high performance [6], [12]–[14].

With resistive random-access memories (RRAMs), multiple devices are required to obtain sufficiently gradual weight change to enable software equivalent learning [15], [16]. Additionally, RRAM (HfO<sub>2</sub>/PCMO/NbO<sub>2</sub>) and PCM based memory has additional process complexity / cost to be integrated into CMOS [17].

Floating-gate has been explored as an analog memory for neural networks extensively [18]. However, horizontal floating-gate flash memory has been replaced by vertical charge trap flash memory with storage in silicon nitride traps for advanced technology nodes [19].

In contrast with memristor, a silicon-oxide-nitride-oxide (SONOS) based charge trap flash memory has significantly gradual conductance change with conductance saturation after 100 pulses [10]. This may be compared to 20 pulses for PCM [8], or 20 pulses for PCMO based RRAM [9]. Maximum conductance change was between 5-20% of the range of conductance and noise was around 5%-10% of the range of conductance. A dual precision approach in which one flash cell has a 1x factor and another has an 8x factor to define the weight was required to obtain software level accuracy on MNIST. The weight updates also required varying pulse voltage and time, which would incur additional circuit costs.

Recently, a similar charge trap flash device has been programmed by quantum tunneling to show extremely gradual programming of 1,000-10,000 levels, which gives a 10-100x improvement over literature [10]. The maximum conductance change per spike is controlled to <1% of the range while the noise is 0.1% of the range. However, linearity is not available in the entire range, which is essential for RPU applications. An important question is whether, by reducing the range of conductance, a smaller but more linear range can be found, which would enable software equivalent RPU, despite experimentally measured noise.

### III. BACKGROUND

#### A. Artificial Neural Networks

Artificial Neural Networks work based on the principle of multi-layered perceptron [20, Chapter 6]. Each layer of neurons performs a weighted linear combination of its inputs, applies a non-linear function, and passes the output to the next layer. Mathematically, given an input vector  $\mathbf{x}$  and a weight matrix  $\mathbf{W}$ , a fully connected layer  $i$  outputs

$$\mathbf{a}^{(i)} = \phi^{(i)}(\mathbf{W}^{(i)}\mathbf{x}^{(i)}), \quad (1)$$

where  $\phi^{(i)}$  is some non-linear function called the activation. This operation is repeated for all layers, giving the output  $\hat{\mathbf{y}} = f(\mathbf{x}, \mathbf{W})$ .

In machine learning, neural networks are used to approximate the function between the input data and a target. Gradient descent is used to minimize a loss function ( $\mathcal{L}(\hat{\mathbf{y}}, \mathbf{y})$ ) between the output of the neural network ( $\hat{\mathbf{y}}$ ) and the true target ( $\mathbf{y}$ ). The gradients are calculated efficiently using backpropagation [3].

Backpropagation uses chain rule to propagate the gradients to the lower layers, given the gradients of the higher layers. Let  $\mathbf{z}^{(i)} = \mathbf{W}^{(i)}\mathbf{x}^{(i)}$  and  $\delta^{(i)} = \frac{\partial \mathcal{L}}{\partial \mathbf{z}^{(i)}}$ . Then,

$$\delta^{(i-1)} = \mathbf{W}^{(i)T} \delta^{(i)} \odot \phi^{(i-1)'}(\mathbf{z}^{(i-1)}) \quad (2)$$

$$\frac{\partial \mathcal{L}}{\partial \mathbf{W}^{(i)}} = \delta^{(i)} \mathbf{x}^{(i)T} \quad (3)$$

where  $\phi^{(i-1)'}$  are the gradients of the activation functions and  $\odot$  is the Hadamard (element-wise) product. Equations 1, 2, and 3, along with the gradient descent update, form the core of training a neural network.

#### B. Resistive Processing Unit

Resistive processing units (RPUs) [7] attempt to speed up the computation of the matrix-vector multiplication (Equations 1, 2) and vector-vector outer product (Equation 3). For efficient hardware implementation, devices are arranged in a crossbar architecture with device conductance at each cross point representing a weight.

First, Ohm's law, combined with Kirchhoff's current law, is used to enable multiply-accumulate operation naturally in hardware. During forward pass (Equation 1), passing voltage proportional to  $\mathbf{x}^{(i)}$  to the rows makes the current at the columns equal to the output of the layer  $\mathbf{W}^{(i)}\mathbf{x}^{(i)}$ . Similarly, during backward pass (Equation 2), passing voltage proportional to  $\delta^{(i)}$  to the columns makes the current at the rows equal to  $\mathbf{W}^{(i)T}\delta^{(i)}$ , which is required for back-propagating the gradient.

Second, weight update by a simple stochastic AND operation is performed directly on non-volatile memory elements. The outer product (Equation 3) is calculated using stochastic multiplication. Two pulse trains, with probability of high voltage proportional to  $\mathbf{x}^{(i)}$ ,  $\delta^{(i)}$  respectively, are generated and passed through rows and columns respectively. The voltage levels are set such that the resistive device updates its weight by  $\Delta w$  when the pulses coincide, and there is no change when the pulses don't coincide. Since the expected number of coincidences is proportional to  $\mathbf{x}^{(i)}\delta^{(i)}$ , the total weight update is proportional to the gradient in expectation. Figure 1 shows an example of pulse trains and the resulting update.

The crossbar architecture and the stochastic weight update makes RPU more energy and area efficient compared to high precision digital multiplication blocks [7].

### IV. FLASH SYNAPSE

#### A. Experimental Device

We use a CTF capacitor (Figure 2), which is fabricated as described by Sandhya et al. [21]. The device is fabricated on an n-Si substrate with 4 nm thermal SiO<sub>2</sub> as a tunnel oxide, 6

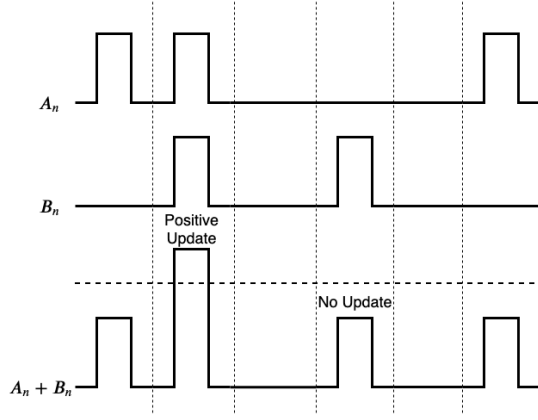


Fig. 1. Analog multiplication using Stochastic pulse trains in RPU: Analog numbers are represented by a stochastic pulse train of identical pulses where the probability of high voltage in trains  $A_n$  and  $B_n$  is proportional to  $x^{(i)}$ ,  $\delta^{(i)}$  respectively. Updates occur at the coincidences, i.e.,  $\text{AND}(A_n, B_n)$ , which have a probability proportional to  $x^{(i)}\delta^{(i)}$ . Polarity is reversed for negative updates.

nm LPCVD  $\text{Si}_3\text{N}_4$  as charge trap layer (CTL), 12 nm MOCVD  $\text{Al}_2\text{O}_3$  as blocking oxide, and n+ polysilicon on 12" substrate by Applied Materials cluster tool. Aluminum is used as a back contact. A self-aligned B implant and anneal is done to provide a source for minority carriers for fast programming as shown in Figure 2a.

### B. Working as Synapse

The program/erase operation is based on FN tunneling. When a positive pulse is applied to the gate, electrons from the channel tunnel through the 4 nm tunnel oxide to be trapped in the CTL, i.e., programming (Figure 2b). To erase, a negative pulse is applied to the gate. Electrons are ejected from the CTL by tunneling through the tunnel oxide (Figure 2c).

Programming and erasing results in a threshold voltage shift ( $\Delta v_T$ ). The threshold voltage ( $v_T$ ) is translated to drain current ( $i_d$ ), which indicates the synaptic conductance ( $g$ ) as follows:

$$i_d = \beta_1(v_{GS} - v_T)v_{DS} \quad (4)$$

$$i_d = g \cdot v_{DS} \quad (5)$$

$$\Delta g \propto -\beta_2 \cdot \Delta v_T \quad (6)$$

where  $\beta_1$  and  $\beta_2$  are proportionality constants [22]. Erasing ( $\Delta v_T < 0$ ) results in potentiation ( $\Delta g > 0$ ), while programming ( $\Delta v_T > 0$ ) results in depression ( $\Delta g < 0$ ). Henceforth, we use  $v_T$  and  $g$  interchangeably since they are simply the scaled version of each other. An approximately linear and gradual change of conductance with the pulse number can be designed by pulse-width modulation [11].

### C. Experimental Data

1) *Curve Fitting Device Updates*: We experimentally calculate the pulse amplitude and pulse width that gives an approximately linear weight change. Figure 3a shows the experimental data of  $v_T$  vs pulse number for LTD (using a pulse of +12.5V and 0.85ms width) and LTP (using a pulse of




Fig. 2. (a) Charge Trap Flash (CTF) schematic. Energy band diagram showing charge transport by quantum tunneling to charge/discharge silicon nitride atomic defects during: (b) Programming and (c) Erasing




Fig. 3. (a) Experimental data (dots) and its curve fitting (lines) using the equation  $v_T(n) = x_1(n)^{x_2} + x_3$ . (b) Mean  $\Delta v_T$  vs  $v_T$  shows that  $v_T$  shift is non-uniform. (c) Repeated measurements (6 times) of (a) is used to estimate the noise as a fraction of mean  $\Delta v_T$  vs  $v_T$ . The experimental  $\sigma/\mu$  is 30-40% for LTP and  $\sim 10\%$  for LTD.

-12.5V and 15ms width). The scatter points are the observed data and the solid lines are the corresponding curve fits.

The curves were fit using the equation  $v_T(n) = x_1(n)^{x_2} + x_3$  to minimize the mean squared error, with  $x_1, x_2, x_3$  being the curve fit variables. The equation for  $\Delta v_T$  was then found by setting  $\Delta v_T(n) = v_T(n+1) - v_T(n)$  to get

$$\Delta v_T(v_T) = x_1 x_2 \left( \frac{v_T - x_3}{x_1} \right)^{\frac{x_2 - 1}{x_2}} \quad (7)$$

We define  $\Delta_0^+(g)$  as the positive change in  $v_T$  when  $v_T = g$  (using LTD data) and  $\Delta_0^-(g)$  as the negative change in  $v_T$  when  $v_T = g$  (using LTP data). Figure 3b shows the variation of  $\Delta v_T$  with  $v_T$ . The results of the curve fit gave  $x_1, x_2, x_3 = 9.55 \times 10^{-4}, 7.19 \times 10^{-1}, -3.22 \times 10^{-1}$  respectively for LTD and  $x_1, x_2, x_3 = -2.38 \times 10^{-3}, 5.80 \times 10^{-1}, -1.12 \times 10^{-1}$  for LTP respectively, which implies that

$$\Delta_0^+(g) = 4.50(g + 0.32)^{-0.39} \times 10^{-5} \quad (8)$$

$$\Delta_0^-(g) = -1.74(-g - 0.11)^{-0.72} \times 10^{-5} \quad (9)$$

2) *Characterization of Device Noise*: To find the noise in the updates, LTP and LTD experiments were repeated six times on the same device to characterize the variation within a device. For each experiment, a curve was fit and the corresponding  $\Delta v_T$  was found. Then, for each  $v_T$ , the standard deviation ( $\sigma$ ) of the evaluation of all six  $\Delta v_T$  was found. Figure 3c shows the standard deviation as a percentage of mean vs  $v_T$  for LTD and LTP. This standard deviation is a measure of variation over time within a flash device - interpreted as noise. To simplify the simulations,  $\sigma$  was set to a high constant for all  $v_T$  in our experiments.

#### D. CTF in RPU array

1) *Simulating Device Updates*: The conductances of a CTF device are always positive, but the weights can be negative. Thus, two devices are required to represent both positive and negative weights. Mathematically, the weight

$$w = k(g_1 - g_2) \quad (10)$$

The scaling constant  $k$  is used to control the range of device conductance. In hardware, 2 CTF devices are arranged as shown in Figure 4a. Applying voltages to the gates of the devices generates currents at the drain and source respectively. These currents are added to implement Equation 10.

$\Delta w$  is not constant since  $\Delta g$  is a function of the current device conductances, and whether the update is positive or negative. The update is also noisy. Accommodating all these modifications, the positive and negative updates are given by

$$\Delta^+ g = \Delta_0^+(g) + N \quad (11)$$

$$\Delta^- g = \Delta_0^-(g) + N \quad (12)$$

where  $N \sim \mathcal{N}(0, \sigma)$  is the noise.

2) *Controlling Linearity and Noise*: Since the range of  $w$  only depends on the dataset and the step size,  $k$  controls the range of  $v_T$  used, and hence, the noise, linearity, and the number of levels available. For example, Gokmen and Vlasov [7] showed that the required range of  $w$  was  $(-0.3, 0.3)$ , when training on MNIST dataset. Based on Equation 10, a conductance range of  $\frac{0.3}{k}$  on each device is sufficient to




Fig. 4. CTF in RPU array: Combination of two CTF used for representing  $g_1$  (blue) and  $g_2$  (grey). (a) The voltage applied at the gates generates currents at the source and drain respectively which are subtracted to produce weights capable of assuming positive and negative values. (b) Weight Update: Stochastic pulse trains  $A_n$  and  $B_n$  are applied to Gate and S/D shorted configuration respectively to produce an  $\text{AND}(A_n, B_n)$  operation based voltage summation at the CTF device. (c) Crossbar architecture with bit line (BL) and word line (WL). (d) Each unit cell in the crossbar is the combination of two CTF.




Fig. 5. Effect of  $k$  on various device properties: (a) the required range of  $g$  and linearity, (b) maximum standard deviation as a fraction of mean  $\Delta v_T$  and the number of levels available. As  $k$  increases, the required range of  $g$  is reduced. With appropriate centering, the range can be restricted to regions of high linearity and low noise. However, a smaller range also reduces the number of levels in the range.




Fig. 6. Two possible updates cycles where polarities of  $X_n$  and  $D_n$  can be set independently. Since  $\Delta w \propto \Delta g_1 - \Delta g_2$ , the updates resulting from these cycles are of the required polarity and magnitude. Any combination of these cycles (for example, alternating between them every iteration) is also valid.

represent this range. Hence, a higher  $k$  implies a lower required range of  $g$ , which can be observed in Figure 5a.

Constraining  $g$  to a lower range improves linearity (Figure 5a). It also allows us to stay in the region with low noise, leading to lower maximum standard deviation as a fraction of mean  $\Delta v_T$  (Figure 5b). But as a trade-off, the number of levels available before it goes out of the range of  $g$  is reduced (Figure 5b). In Section V, we show the effect of this trade-off on the performance of the system. In addition to the range, the center-point of the conductance range is optimized by trial and error to improve linearity.

3) *Circuit Design Considerations*: Performing an addition or subtraction of pulse trains is easier from a hardware perspective than an AND operation [7]. To perform a positive update, two positive polarity pulse trains can be added such that a positive voltage pulse results at the coincidences. The polarities can be reversed to perform a negative update. Since  $x^{(i)}$  and  $\delta^{(i)}$  are applied to the two ends of the crossbar, the polarity of the pulse trains must depend independently on the corresponding  $x^{(i)}$  or  $\delta^{(i)}$  and not the product  $x^{(i)}\delta^{(i)}$ . The input  $x^{(i)}$  can be assumed to be positive since inputs are generally normalized between 0 and 1 and the common non-linear activations functions used in a neural network like sigmoid or ReLU only output positive values.

Two possible update cycles with these constraints and the corresponding pulse polarities are shown in Figure 6. We always use the positive cycle in our experiments.

Weight update in hardware for CTF devices is done by

applying the voltage at the gate with respect to Source-Drain connected to the ground (Figure 4b).

As proposed by Gokmen and Vlasov [7], non-linear activation functions and their gradient can be implemented using an external circuitry. For the special case of ReLU activation, this external circuitry can be simplified. ReLU simply passes forward the positive inputs and blocks the negative inputs. The gradient is hence, 1 for positive inputs and 0 for negative inputs.

Figure 4c shows the crossbar architecture with 4 word lines (WL) for applying voltages and 2 bit lines (BL) to read the currents. Figure 4d shows a single unit cell in the crossbar with 2 CTF devices. Algorithm 1 describes the steps for calculating the weight update while simulating a CTF device.

---

**Algorithm 1:** Update calculation in CTF device simulation.

---

**Input:** Gradients ( $\delta^{(i)}$ ); Inputs ( $x^{(i)}$ ); Length of pulse trains ( $PL$ ); Input scaling constant ( $C$ ); Weight update functions  $\Delta_0^+$ ,  $\Delta_0^-$ ; Device conductances  $G_1^{(i)}$  and  $G_2^{(i)}$ ; Noise ( $\sigma$ ).

**Output:** Updated values of the device conductances of the layer  $G_1^{(i)}$ ,  $G_2^{(i)}$ .

```

1 for each cross-point do
2   Let  $g_1, g_2$  be the device conductances
3   Find  $x^{(i)}, \delta^{(i)}$  corresponding to the cross point
4   Sample  $X_1, \dots, X_{PL} \sim \text{Bernoulli}(|Cx^{(i)}|)$ 
5   Sample  $D_1, \dots, D_{PL} \sim \text{Bernoulli}(|C\delta^{(i)}|)$ 
6   Set the polarity of all  $D_n$  equal to the sign of  $\delta^{(i)}$ 
7   for each coincidence in  $X_n \wedge D_n$  do
8     Sample noise  $N \sim \mathcal{N}(0, \sigma)$ 
9     if  $\delta^{(i)} < 0$  then
10       $\Delta^+ g_1 \leftarrow \Delta_0^+(g_1) + N$ 
11       $g_1 \leftarrow g_1 + \Delta^+ g_1$ 
12    else
13       $\Delta^+ g_2 \leftarrow \Delta_0^+(g_2) + N$ 
14       $g_2 \leftarrow g_2 + \Delta^+ g_2$ 
15    end
16  end
17 end

```

---

## V. EXPERIMENTS AND RESULTS

To test the performance of a neural network with flash synapse as the cross point device, we performed three experiments. We trained neural networks for supervised classification of digits in the MNIST dataset [23], images in the CIFAR dataset [24], and for reinforcement learning in the Mountain Car environment [25]. All neural network operations were performed by simulating CTF devices as described in section IV-D. As a baseline in all the experiments, we performed the neural network training using exact floating point operations.

Table I shows the list of hyperparameters used in the experiments. A combination of manual tuning and grid search was used to find these hyperparameters. Hyperparameters

TABLE I  
HYPERPARAMETERS.

Hyperparameter	Value
Update step size ( $\alpha$ )	MNIST: 0.01 CIFAR: 0.1 Mountain Car: 0.00625
Initial weights ( $w_0$ )	Kaiming uniform [26]
Weight scaling factor ( $k$ )	$600\alpha$
Initial device conductance ( $g_{1,0}, g_{2,0}$ )	$-0.2 \pm \frac{w_0}{2k}$
Pulse train length $PL$	10
Input scaling factor $C$	$\sqrt{\frac{\alpha}{PL \cdot \Delta^+ g(c) \cdot k}}$

related to the CTF device and RPU were kept constant for all the experiments.

### A. MNIST

MNIST dataset consists of 60,000 training and 10,000 test images of 10 handwritten digits, each of size 28x28 pixels.

A fully connected neural network with 2 hidden layers consisting of 256 and 128 neurons respectively, was used for classification. The neural network was trained for 10 epochs. Two sets of experiments were performed, with noise standard deviation ( $\sigma$ ) being 10% of the mean in one and 100% of the mean in the other.

Figure 7a shows the learning curves with 10% noise and 100% noise respectively, compared with that of the baseline. The curves are averaged over the 10 runs and one standard error is shaded. The final accuracies with the flash device are  $98.07 \pm 0.05\%$  and  $97.91 \pm 0.06\%$  with 10% noise and 100% noise respectively. The final accuracy of the baseline is  $98.05 \pm 0.07\%$ .

#### 1) Effect of Weight Scaling Factor ( $k$ ) on Performance:

As described in Section IV-D, changing  $k$  leads to a trade-off between linearity, noise, and the number of pulses available. To study its effect on the performance, we adjust  $k$  and measure the test and train accuracies.

Figure 7b shows the variation of train and test accuracies for different values of  $k$  at a noise level of 10%. The highest train accuracy of 99.84% was obtained for  $k = 6$ , with the corresponding test accuracy being 98.15%.

Higher values of  $k$  used a lower range of device conductances, which reduced the precision of the system since  $\Delta_0^+(g)$  and  $\Delta_0^-(g)$  are unchanged. Lower values of  $k$  used a larger range of device conductances. Since the conductance change became more non-linear on either extreme, the performance declined.

2) Noise Analysis: To study the effect of noise on the performance, we run the MNIST experiments with noise level varying from 0% to 500% and find the test accuracy after 3 epochs.

Figure 7c shows the accuracy as a function of noise, averaged over 4 runs. The accuracy is  $97.5 \pm 0.07\%$  without noise,  $97.3 \pm 0.14\%$  with 100% noise, and drops to  $93.4 \pm 0.18\%$  at 500% noise.

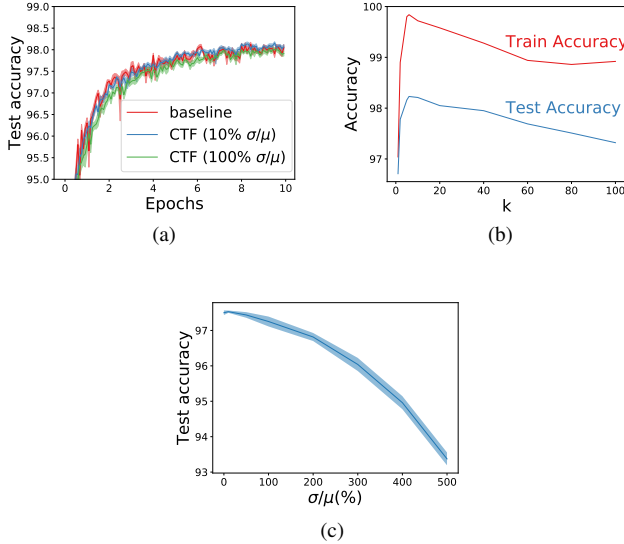


Fig. 7. MNIST experiments: (a) Test accuracy as a function of the number of epochs on MNIST dataset (averaged over 10 runs, one standard error shaded). The difference in accuracy between floating point baseline and flash synapse RPU is negligible with 10% noise and around 0.1% with 100% noise. (b) Train and test accuracy as a function of  $k$  on the MNIST dataset. Accuracies are low for very low and very high values of  $k$ , with  $k = 6$  being the best value. (c) Test accuracy as a function of update noise on MNIST dataset after 3 epochs (averaged over 4 runs, one standard error shaded). The accuracy drops by less than 0.3% with 100% noise and by around 4% with 500% noise.

The drop in accuracy with higher noise is expected since the actual updates can be very different from the updates required for gradient descent. But, as shown in section IV-C, 100% noise is well above those found experimentally in the flash device. Hence, it acts as a lower bound on the obtainable accuracy.

## B. CIFAR

CIFAR dataset consists of 50,000 training and 10,000 test images of real world objects. Each image is colored and 32x32 pixels in size. CIFAR-10 consists of 10 classes of images, while CIFAR-100 consists of 100 classes of images.

Since convolutional neural networks (CNNs) are generally used for classification on these datasets, we follow the methodology used by Ambrogio et al. [14] to compare our device with the baseline. A pre-trained CNN, specifically, ResNet-50 [27] pre-trained on the ImageNet [28] dataset, is used for feature extraction. The CIFAR images were resized, normalized and passed through the pre-trained network. The activations of the last hidden layer were considered as features.

Once the features were extracted, a neural network with no hidden layers was trained to classify the images based on the features. The neural network was trained for 10 epochs.

Figure 8a shows the learning curves with 10% noise and 100% noise respectively, for CIFAR-10 dataset. The final accuracies with the flash device are  $89.21 \pm 0.09\%$  and  $89.07 \pm 0.09\%$  respectively. The final accuracy of the baseline is  $89.6 \pm 0.07\%$ .

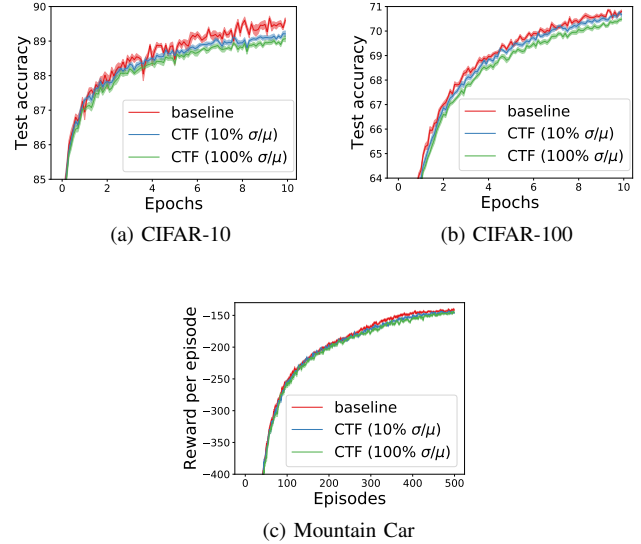


Fig. 8. CIFAR and Mountain Car experiments: (a) Test accuracy as a function of the number of epochs on CIFAR-10 dataset (averaged over 10 runs, one standard error shaded). The difference in accuracy between floating point baseline and flash synapse RPU is around 0.4% with 10% noise and around 0.5% with 100% noise. (b) Test accuracy as a function of the number of epochs on CIFAR-100 dataset (averaged over 10 runs, one standard error shaded). The difference in accuracy between floating point baseline and flash synapse RPU is around 0.1% with 10% noise and around 0.3% with 100% noise. (c) Reward per episode as a function of the number of episodes on the Mountain Car environment (averaged over 100 runs, the standard error is less than the line width). The difference in reward between floating point baseline and flash synapse RPU is around 4% with 10% noise and around 3% with 100% noise.

Figure 8b shows the same for the CIFAR-100 dataset. The final accuracies with the flash device are  $70.67 \pm 0.07\%$  and  $70.48 \pm 0.07\%$  with 10% and 100% noise respectively. The final accuracy of the baseline is  $70.8 \pm 0.06\%$ .

## C. Mountain Car

Mountain Car is a control problem in which the agent should drive a car to the top of the mountain. The agent observes its current horizontal position and its velocity. It can move forward, move backward or do nothing. Since the agent can't accelerate enough to reach the peak by just moving forward, it needs to move back and forth to build enough momentum before being able to reach the peak [25]. The agent gets a reward of -1 at every time step until it reaches the goal, and hence, it needs to reach the goal as quickly as possible.

We used tile coding [30, pg. 217] to extract features from the observations and used a neural network with no hidden layers on top of it to predict the state-action values (Q-values) for each action. Mathematically,  $\hat{q}(s, a; \mathbf{W})$  provided an approximation of  $Q(s, a)$ , for each state  $s$  and action  $a$ . The weights were updated using Q-learning [31] update:

$$\mathbf{W} \leftarrow \mathbf{W} + \alpha(R + \gamma \max_{a'} \hat{q}(S', a'; \mathbf{W}) - \hat{q}(S, A; \mathbf{W})) \nabla_{\mathbf{W}} \hat{q}(S, A; \mathbf{W}) \quad (13)$$

TABLE II  
COMPARISON OF OUR WORK WITH THE PREVIOUS WORKS ON MNIST DATASET.

Authors	Precision	Programming	Devices per Weight	MNIST Accuracy	Applications Demonstrated
Ambrogio et al. [14]	Dual Precision: High precision, volatile DRAM + Low precision non-volatile PCM	Analog pulse V and time	2 PCM + DRAM	<b>97.95%</b>	Supervised Learning - MNIST, CIFAR-10, CIFAR-100
Nandakumar et al. [29]	Dual Precision: High precision, volatile CMOS + Low precision, non-volatile PCM	Analog pulse V and time	2 PCM + SRAM	97.40%	Supervised Learning - MNIST
Agarwal et al. [10]	<b>Single precision</b>	Analog pulse V and time	<b>2 SONOS flash</b>	97.6%	Supervised Learning - File Types, MNIST
Agarwal et al. [10]	Dual Precision: High & Low precision CTF by relative weight	Analog pulse V and time	4 SONOS flash	<b>98%</b>	Supervised Learning - File Types, MNIST
Nandakumar et al. [8]	<b>Single Precision</b>	<b>Stochastic Identical Pulse Train</b>	2 PCM	83%	Supervised Learning - MNIST
Babu et al. [9]	<b>Single Precision</b>	<b>Stochastic Identical Pulse Train</b>	<b>2 PCMO</b>	88.1%	Supervised Learning - MNIST
This work	<b>Single Precision</b>	<b>Stochastic Identical Pulse Train</b>	<b>2 CTF</b>	<b>97.9%</b>	<b>Supervised Learning - MNIST, CIFAR-10, CIFAR-100; Reinforcement Learning - Mountain Car</b>

where  $S$  is the current state,  $A$  is the action chosen,  $R$  is the reward obtained,  $S'$  is the next state,  $\alpha$  is the step size, and  $\gamma$  is the discount factor. The gradient calculation and weight update in Equation 13 was performed by simulating the flash device.

Action selection was done using epsilon-greedy strategy with  $\epsilon = 0.1$ . Hash-based tile coding software by Sutton [32] was used for feature extraction, with 8 equally sized tiles per dimension and 16 tilings.

The agent was trained for 500 episodes, with each episode being terminated either on reaching the goal or after 1000 steps. The experiment was repeated 100 times and the total reward obtained from each episode was recorded.

Figure 8c shows the total reward per episode as a function of the number of episodes with 10% noise and 100% noise respectively. The floating point baseline obtains a reward of  $-143 \pm 1.6$  (which implies that it takes around 143 steps to complete an episode). With the flash device, the reward is  $-147 \pm 1.8$  with 10% noise and  $-146 \pm 2$  with 100% noise.

## VI. DISCUSSIONS

We show that the CTF device works as a replacement for floating point update in various applications. In all the experiments, the performance of our device was close to that of the floating point baseline. It was also fairly robust to the experimentally measured noise of 10-40% in updates which is crucial for analog computing.

Classification on MNIST dataset showed that a multi-layer neural network can be trained using the CTF device. Classification on CIFAR-100 dataset showed that even in the

regime of a large number of classes and relatively low data, the performance is on par with the floating point updates. Training an agent on Mountain Car environment showed that our method is not just restricted to the supervised learning setting, but can also be used in other settings that use neural networks.

Table II shows that comparison of various current approaches. Among various approaches for in-memory computing, precision enhancement of low precision but compact nanoscale memory like Phase Change Memory (PCM) with high precision but area inefficient CMOS memory enables high performance on MNIST dataset [14], [29]. Further, single precision approaches with RPU based stochastic identical pulse based weight update show degraded performance of 83% for PCM [8] and of 88% for PCMO based RRAM [9] on MNIST dataset. Agarwal et al. [10] have shown a single precision approach based on SONOS based Flash memory with analog pulse control with voltage and time to record a performance of 97.6% on MNIST. This technology is based on NOR flash memory like programming scheme using high current/power technique of channel hot electrons (CHE). Enhancing precision by a dual precision technique with more flash devices per weight and control circuit to enable a periodic carry improves MNIST performance to 98%.

In comparison, our flash memory is programmed with the low current/power/energy FN tunneling technique. Stochastic pulse train based RPU is demonstrated, eschewing the need for variable pulses with analog voltage levels and pulse time controls. The low rate of conductance change, high linearity produces a peak performance of 97.9% - which is robust to

experimentally measured noise levels. Further, our method produces excellent performance on various ANN applications like classification on CIFAR-10, CIFAR-100 datasets, and reinforcement learning on Mountain Car environment - demonstrating excellent generalization.

## VII. CONCLUSIONS

In this paper, we proposed a charge trap flash device in an RPU architecture to accelerate deep neural networks while maintaining software-level accuracy. The resistive processing unit speeds up vector-matrix and vector-vector multiplication operations, which are ubiquitously used in the backpropagation algorithm to train deep neural networks. We engineered the magnitude and the width of the pulse used to update the weights using the flash device. The updates were shown to be linear, gradual and symmetric, which is necessary for good performance.

We then simulated the device to train neural networks on MNIST, CIFAR-10 and CIFAR-100 datasets. In each case, the accuracy of the system was close to the floating point baseline, showing excellent generalization. The system was also robust to noise in weight updates, with less than 1% drop in accuracy when the simulated noise was 10x the experimentally observed value. We also demonstrated the generality of the method by applying it to a reinforcement learning method on the Mountain Car environment. The performance of our system matched the software baseline in this experiment too. Such implementation is benchmarked against the state of the art demonstrations to show best-in-class performance - indicating a promising hardware option for in-memory computing.

## ACKNOWLEDGMENTS

The work was supported in part by SRC and NNETRA Project (DST & MeitY, Govt. of India).

## REFERENCES

- [1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, p. 436, 2015.
- [2] F. Rosenblatt, "The perceptron: A probabilistic model for information storage and organization in the brain," *Psychological Review*, vol. 65, no. 6, p. 386, 1958.
- [3] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *Nature*, vol. 323, no. 6088, p. 533, 1986.
- [4] M. Bauer, H. Cook, and B. Khailany, "CudaDMA: Optimizing GPU memory bandwidth via warp specialization," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. ACM, 2011, p. 12.
- [5] M. Moravčík, M. Schmid, N. Burch, V. Lisý, D. Morrill, N. Bard *et al.*, "DeepStack: Expert-level artificial intelligence in heads-up no-limit poker," *Science*, vol. 356, no. 6337, pp. 508–513, 2017.
- [6] M. Le Gallo, A. Sebastian, R. Mathis, M. Manica, H. Giefers, T. Tuma *et al.*, "Mixed-precision in-memory computing," *Nature Electronics*, vol. 1, no. 4, p. 246, 2018.
- [7] T. Gokmen and Y. Vlasov, "Acceleration of deep neural network training with resistive cross-point devices: Design considerations," *Frontiers in Neuroscience*, vol. 10, p. 333, 2016.
- [8] S. Nandakumar, M. Le Gallo, I. Boybat, B. Rajendran, A. Sebastian, and E. Eleftheriou, "A phase-change memory model for neuromorphic computing," *Journal of Applied Physics*, vol. 124, no. 15, 2018.
- [9] A. V. Babu, S. Lashkare, U. Ganguly, and B. Rajendran, "Stochastic learning in deep neural networks based on nanoscale PCMO device characteristics," *Neurocomputing*, vol. 321, pp. 227–236, 2018.
- [10] S. Agarwal, D. Garland, J. Niroula, R. B. Jacobs-Gedrim, A. Hsia, M. S. Van Heukelom *et al.*, "Using floating-gate memory to train ideal accuracy neural networks," *IEEE Journal on Exploratory Solid-State Computational Devices and Circuits*, vol. 5, no. 1, pp. 52–57, 2019.
- [11] S. Shrivastava, T. Chavan, and U. Ganguly, "Ultra-low energy charge trap flash based synapse enabled by parasitic leakage mitigation," *arXiv preprint arXiv:1902.09417*, 2019.
- [12] M. Suri, O. Bichler, D. Querlioz, O. Cueto, L. Perniola, V. Sousa *et al.*, "Phase change memory as synapse for ultra-dense neuromorphic systems: Application to complex visual pattern extraction," in *International Electron Devices Meeting*. IEEE, 2011, pp. 4–4.
- [13] O. Bichler, M. Suri, D. Querlioz, D. Vuillaume, B. DeSalvo, and C. Gamrat, "Visual pattern extraction using energy-efficient "2-PCM synapse" neuromorphic architecture," *IEEE Transactions on Electron Devices*, vol. 59, no. 8, pp. 2206–2214, 2012.
- [14] S. Ambrogio, P. Narayanan, H. Tsai, R. M. Shelby, I. Boybat, C. di Nolfo *et al.*, "Equivalent-accuracy accelerated neural-network training using analogue memory," *Nature*, vol. 558, no. 7708, p. 60, 2018.
- [15] I. Boybat, M. Le Gallo, S. Nandakumar, T. Moraitis, T. Parnell, T. Tuma *et al.*, "Neuromorphic computing with multi-memristive synapses," *Nature Communications*, vol. 9, no. 1, p. 2514, 2018.
- [16] A. Shukla, S. Prasad, S. Lashkare, and U. Ganguly, "A case for multiple and parallel RRAMs as synaptic model for training SNNs," in *International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2018, pp. 1–8.
- [17] V. Sze, Y.-H. Chen, T.-J. Yang, and J. S. Emer, "Efficient processing of deep neural networks: A tutorial and survey," *Proceedings of the IEEE*, vol. 105, no. 12, pp. 2295–2329, 2017.
- [18] O. Fujita and Y. Amemiya, "A floating-gate analog memory device for neural networks," *IEEE transactions on electron devices*, vol. 40, no. 11, pp. 2029–2035, 1993.
- [19] D. Kang, W. Jeong, C. Kim, D. Kim, Y. Cho, K. Kang *et al.*, "256 Gb 3 b/Cell V-nand flash memory with 48 stacked WL layers," *IEEE Journal of Solid-State Circuits*, vol. 52, no. 1, pp. 210–217, 2017.
- [20] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.
- [21] C. Sandhya, U. Ganguly, N. Chattar, C. Olsen, S. M. Seutter, L. Date *et al.*, "Effect of SiN on performance and reliability of charge trap flash (CTF) under Fowler–Nordheim tunneling program/erase operation," *IEEE Electron Device Letters*, vol. 30, no. 2, pp. 171–173, 2008.
- [22] Y. Taur and T. H. Ning, *Fundamentals of Modern VLSI Devices*. Cambridge University Press, 2013.
- [23] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [24] A. Krizhevsky, "Learning multiple layers of features from tiny images," Citeseer, Tech. Rep., 2009.
- [25] A. W. Moore, "Efficient memory-based learning for robot control," 1990.
- [26] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 1026–1034.
- [27] —, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [28] J. Deng, W. Dong, R. Socher, L. Li, K. Li, and F. Li, "ImageNet: A large-scale hierarchical image database," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2009.
- [29] S. R. Nandakumar, M. L. Gallo, I. Boybat, B. Rajendran, A. Sebastian, and E. Eleftheriou, "Mixed-precision architecture based on computational memory for training deep neural networks," in *IEEE International Symposium on Circuits and Systems, ISCAS*, 2018, pp. 1–5.
- [30] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. MIT press, 2018.
- [31] C. J. C. H. Watkins and P. Dayan, "Q-learning," *Machine Learning*, vol. 8, no. 3-4, pp. 279–292, 1992.
- [32] R. S. Sutton, "Tile coding software – reference manual," <http://incompleteideas.net/tiles/tiles3.html>, 2017, accessed: 2019-07-13.