

SwitchX: Gmin-Gmax Switching for Energy-efficient and Robust Implementation of Binarized Neural Networks on ReRAM Xbars

ABHIROOP BHATTACHARJEE and PRIYADARSHINI PANDA, Yale University

Memristive crossbars can efficiently implement Binarized Neural Networks (BNNs) wherein the weights are stored in high-resistance states (HRS) and low-resistance states (LRS) of the synapses. We propose SwitchX mapping of BNN weights onto ReRAM crossbars such that the impact of crossbar non-idealities, that lead to degradation in computational accuracy, are minimized. Essentially, SwitchX maps the binary weights in such a manner that a crossbar instance comprises of more HRS than LRS synapses. We find BNNs mapped onto crossbars with SwitchX to exhibit better robustness against adversarial attacks than the standard crossbarmapped BNNs, the baseline. Finally, we combine SwitchX with state-aware training (that further increases the feasibility of HRS states during weight mapping) to boost the robustness of a BNN on hardware. We find that this approach yields stronger defense against adversarial attacks than adversarial training, a state-ofthe-art software defense. We perform experiments on a VGG16 BNN with benchmark datasets (CIFAR-10, CIFAR-100 and TinyImagenet) and use Fast Gradient Sign Method ($\epsilon = 0.05$ to 0.3) and Projected Gradient Descent ($\epsilon = \frac{2}{255}$ to $\frac{32}{255}$, $\alpha = \frac{2}{255}$) adversarial attacks. We show that SwitchX combined with state-aware training can yield upto \sim 35% improvements in clean accuracy and \sim 6–16% in adversarial accuracies against conventional BNNs. Furthermore, an important by-product of SwitchX mapping is increased crossbar power savings, owing to an increased proportion of HRS synapses, which is furthered with state-aware training. We obtain upto ~21-22% savings in crossbar power consumption for state-aware trained BNN mapped via SwitchX on 16×16 and 32×32 crossbars using the CIFAR-10 and CIFAR-100 datasets.

CCS Concepts: • **Hardware** → *Safety critical systems*;

Additional Key Words and Phrases: Binarized neural network, ReRAM crossbar, non-idealities, switchedmapping, adversarial robustness

ACM Reference format:

Abhiroop Bhattacharjee and Priyadarshini Panda. 2023. *SwitchX*: Gmin-Gmax Switching for Energy-efficient and Robust Implementation of Binarized Neural Networks on ReRAM Xbars. *ACM Trans. Des. Autom. Electron. Syst.* 28, 4, Article 60 (May 2023), 21 pages.

https://doi.org/10.1145/3576195

© 2023 Association for Computing Machinery.

1084-4309/2023/05-ART60 \$15.00

https://doi.org/10.1145/3576195

This work was supported in part by C-BRIC, a JUMP center sponsored by DARPA and SRC, Google Research Scholar Award, the National Science Foundation (Grant1947826), TII (Abu Dhabi), the DARPA **AI Exploration** (**AIE**) program, and the DoE MMICC center SEA-CROGS (AwardDE-SC0023198).

Authors' addresses: A. Bhattacharjee and P. Panda, Department of Electrical Engineering, Yale University 10 Hillhouse Avenue, New Haven, CT-06511; emails: abhiroop.bhattacharjee@yale.edu, priya.panda@yale.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

1 INTRODUCTION

Memristive crossbars have received significant focus for their ability to realize *Deep Neural Networks* (DNNs) by efficiently performing *Multiply-and-Accumulate* (MAC) operations using analog dot-products [44, 48, 54]. These systems have been realized using a wide range of emerging *Non-Voltalile-Memory* (NVM) devices such as, *Resistive RAM* (ReRAM), *Phase Change Memory* (PCM), *Ferroelectric FET* (FeFET), and Spintronic devices [9, 14, 35, 46]. These devices exhibit high on-chip storage density, non-volatility, low leakage, and low-voltage operation and thus, enable compact and energy-efficient implementation of DNNs [1, 7, 23].

In the recent years, *Binarized Neural Networks* (BNNs) have emerged as efficient and reasonably accurate low-precision models to implement DNNs [21]. BNNs consist of binary synaptic weights and activations, *viz*. {-1,+1}. The binarization enables lower computational complexity and power consumption for MAC operations. BNNs are crossbar-friendly in the sense that they can be implemented on crossbars in the manner shown in Figure 1 (see Normal Mapping). Here, the binarized weights are programmed as conductances of the synaptic devices, such as ReRAMs, at the cross-points. A weight value of "+1" corresponds to a *Low Resistance State* (LRS) while that of "-1" corresponds to a *High Resistance State* (HRS). During inference, the activations of BNNs are fed into each row *i* of the crossbar as analog voltages (generated using Digital-to-Analog Converters (DACs)). As per Ohm's Law, the voltages interact with the device conductances at the cross-points *G_{ij}* and produce a current. Consequently, by Kirchhoff's current law, the net output current sensed at each column *j* is the sum of currents through each device, i.e.,:

$$I_j = \sum_i G_{ij} * V_i.$$
⁽¹⁾

The above dot-product operation is depicted in Figure 2 (Left).

It is a known fact that memristive crossbar arrays possess non-idealities such as, interconnect parasitics, process variations in the synaptic devices, driver and sensing resistances, and so on. [6, 23]. These non-idealities manifest as imprecise dot-product currents causing accuracy degradation when DNNs are mapped onto crossbars. Many previous works [6, 9, 23, 32, 33] have used frameworks to capture the impact of circuit noise or non-idealities in crossbars and proposed noiseaware retraining of DNNs to mitigate accuracy losses (see Table 1). However, [6, 9, 23, 32, 33] do not study the impact of crossbar non-idealities on the robustness of neural networks against adversarial attacks. Adversarial attacks are structured, yet, small perturbations on the input, that fool a DNN causing high confidence misclassification. This vulnerability severely limits the deployment and potential safe-use of DNNs for real-world applications [5, 39]. Recently there have been investigative works such as [3, 42, 43], wherein the authors have shown that crossbar non-idealities while degrading performance, can improve the adversarial attack resilience of hardware-mapped DNNs in comparison to baseline software DNNs without additional efficiency-driven hardware optimizations (see Table 1). Furthermore, the authors in [16, 38, 39, 45] show that hardware optimization techniques, such as quantization, model compression, and pruning can be leveraged to improve the adversarial robustness of DNNs in addition to energy-efficiency.

In this work, we introduce *SwitchX* based mapping of BNN weights onto ReRAM crossbars (see Figure 1 for *SwitchX* mapping), whereby we map the binarized weights in a manner such that the number of HRS synapses always dominate every crossbar instance. We show that *SwitchX* mapping interfere with hardware non-idealities to yield benefits in terms of adversarial robustness when compared with normally-mapped BNNs (the baseline). This is achieved by boosting the feasibility of HRS states in crossbars to tackle crossbar non-idealities and improve the natural and adversarial performance (robustness) of BNNs for their secure deployment on edge-devices.

SwitchX



Fig. 1. Pictorial representation of Normal Mapping and *SwitchX* approaches showing mapping of a 3×3 binary weight matrix on 3×3 crossbar with non-idealities. In the case of Normal Mapping, "+1"s are mapped to LRS states, and "-1"s are mapped to HRS states to obtain G_{ideal} . In the case of *SwitchX*, the mapping depends on the mean of the weight matrix. If mean < 0, then the process is similar to Normal Mapping. If mean > 0, then "+1"s are mapped to HRS and "-1"s to LRS. This ensures that HRS states are always in majority in a crossbar array.



Fig. 2. (Left) A 2×2 ideal ReRAM crossbar with weights programmed as synaptic conductances (G_{ij} s); (Right) A 2×2 non-ideal crossbar with the resistive and device-level non-idealities marked. These non-idealities lead to imprecise dot-product currents and that manifests as accuracy degradation when neural networks are evaluated on crossbars.

It has been shown in earlier works that the weights of a BNN encoded as resistance states in crossbars contribute significantly to the power dissipated by the crossbars [8, 18, 27]. An important by-product of *SwitchX* mapping is the reduction in the overall power expended by the crossbar instances during BNN inference, owing to an increased proportion of HRS synapses.

Contributions: In summary, the key contributions of this work are as follows:

- (1) We comprehensively analyse how *SwitchX* mapping of binarized weights onto ReRAM crossbars manifests as an increase in the robustness and adversarial stability of the mapped networks. Note, although the results in this work have been presented typically for ReRAM crossbars, the *SwitchX* approach and its benefits are not limited to ReRAM crossbars and can be extended to crossbars with other NVM devices such as, PCM or FeFET.
- (2) We carry out experiments on a state-of-the-art neural network architecture (VGG16) [49] using benchmark datasets (CIFAR-10, CIFAR-100 [28], and TinyImagenet). We propose a

	Crossbar non-idealities included				
Work	New weight-mapping strategy	Crossbar energy- efficiency	Non-ideality aware performance (accuracy)	Non-ideality driven adversarial robustness	Noise-aware training or fine-tuning
[6, 9, 23, 32, 33]	×	×	\checkmark	×	\checkmark
[3, 42, 43]	×	×	\checkmark	\checkmark	×
[18, 19]	\checkmark	\checkmark	×	×	×
[12]	×	×	\checkmark	\checkmark	\checkmark
NEAT [4]	×	\checkmark	\checkmark	\checkmark	\checkmark
Our work (SwitchX)	\checkmark	\checkmark	\checkmark	\checkmark	×

Table 1. A Comparisor	Table to Show the Contributions	of Previous	Works and Ou	SwitchX Approach to
	Mapping on Non-ide	al Crossbar	S	

novel graphical approach by plotting *robustness maps* to evaluate the adversarial robustness of BNNs on crossbars and perform a fair comparison between *SwitchX* and Normal Mapping.

- (3) We find that *SwitchX* approach when combined with state-aware training (explained in Section 5) increases the feasibility of higher HRS mapping. This significantly boosts the robustness (both clean and adversarial accuracies) of the mapped BNNs. Further, *SwitchX* combined with state-aware training outperforms the case when *SwitchX* is combined with Adversarial training, a state-of-the-art software defense against adversarial attacks.
- (4) We also carry out ablation studies and ascertain that the *SwitchX* approach can unleash robustness to BNNs mapped onto crossbars with various specifications and dimensions ranging from 16×16 , 32×32 , 64×64 to 128×128 .
- (5) Apart from robustness benefits, we also find that the above approach can lead to significant power savings (~21-22% for 16 × 16 and 32 × 32 crossbars using CIFAR-10 and CIFAR-100 datasets) with respect to normally mapped BNNs on crossbars, thereby bringing in crossbar energy-efficiency. Note, as we see later (in Section 6.3) that the crossbar energy-efficiency via *SwitchX* mapping decreases on increasing the crossbar size, we thus report the overall crossbar power savings of *SwitchX* BNNs only for 16 × 16 and 32 × 32 crossbars.

2 RELATED WORKS

BNNs have been widely explored as a hardware-friendly method to implement neural networks on crossbar-based architectures. Particularly, early works such as [11, 36] have proposed optimized hardware architectures using RRAM crossbars to facilitate BNN inference with higher energyefficiency, throughput and parallelism, and have been shown to be resilient against device-level variations. Another work [20] has further proposed RRAM noise-aware retraining of a BNN model to further improve noise-tolerance of crossbar-mapped BNNs on error-prone crossbar platforms. Later works such as [51] have proposed a dual-activation and dual-synapse based implementation of BNNs on RRAM crossbar-arrays (referred to as Complementary Resistive Cell (CRC) arrays). This helps perform device noise-resilient XNOR operations on the RRAM crossbars and furthers the energy-efficiency and throughput during BNN inference. Furthermore, other works focusing on energy-efficient implementations of BNN on crossbars include [57], which implements an ADC-less RRAM accelerator to carry out BNN inference and outperforms standard state-of-theart crossbar-based inference accelerators such as ISAAC [47], Pipelayer [50], and FloatPIM [22]. However, in [57], RRAM device-noise models are not included during BNN inference and hence, hardware-realistic BNN accuracies are not reported. Recent works such as [37, 52] include mature RRAM noise models in addition to device-level variations such as stuck-at-fault defects and have

proposed architectural modifications as well as noise-mitigation strategies for BNNs to achieve similar accuracy on hardware with state-of-the-art multi-precision counterparts. However, majority of the above works focus on energy-efficiency and throughput as the key goals to optimize for BNN inference on crossbar-arrays. Moreover, none of the above works have explored the area of adversarial robustness for BNNs on crossbars in presence of hardware non-idealities.

Our work SwitchX explores non-ideality aware adversarial robustness of BNNs in presence of resistive crossbar non-idealities in addition to device-level variations. It should be noted that although there has been a complementary work, called *Non-linearity aware Training* (NEAT) [4] for non-ideality aware adversarial robustness of DNNs on crossbars, it involves additional training cost of fine-tuning a pretrained software DNN model before deploying onto non-ideal crossbars. In addition, NEAT focuses on robustness of multi-precision DNN models but does not take into account the hardware overheads of implementation of multibit weights (or conductances) with an ensemble of memristive devices (using bit-slicing) on crossbars. To this end, SwitchX uses binary weights that are simple to program on crossbars with a single synaptic NVM device and hence, are more crossbar-friendly. Furthermore, there have been previous works such as, [18, 19] (see Table 1), introducing techniques that also increase the proportion of high resistance ReRAM states in crossbars leading to better crossbar energy-efficiencies. But, none of them have focused on adversarial robustness in presence of resistive crossbar non-idealities. It should also be noted that SwitchX, unlike [4, 6, 9, 12, 23, 32, 33] in Table 1, is not noise-aware retraining or fine-tuning of neural networks but rather is a simple weight-mapping strategy that helps mitigate the impact of crossbar noise.

3 BACKGROUND

3.1 Memristive Crossbars and their Non-idealities

Memristive crossbar arrays have been employed to implement *Matrix-Vector-Multiplications* (**MVMs**) in neural networks in an analog manner. Traditionally, a crossbar (see Figure 2 (Left)), consists of a 2D array of synaptic NVM devices interfaced with DACs, *Analog-to-Digital Converters* (**ADCs**), and a write circuit to program the synapses. The synaptic devices at the crosspoints are programmed to a particular value of conductance (G_{ij}). The MVM operations during inference are performed by converting the digital inputs to a neural network into analog voltages on the *Read Wordlines* (**RWLs**) using DACs, and sensing the output current flowing through the **bitlines** (**BLs**) using the ADCs [23]. For an ideal crossbar, if V_{in} is the input voltage vector, *Iout_{ideal}* is the output current vector and G_{ideal} is the conductance matrix, then:

$$Iout_{ideal} = V_{in} * G_{ideal}.$$
(2)

Non-idealities and equivalent conductance matrix: The analog nature of the computation leads to various non-idealities resulting in errors in the computation of MVMs. These include various linear resistive non-idealities in the crossbars. Figure 2 (Right) shows the equivalent circuit for a crossbar array accounting for various peripheral and parasitic non-idealities (*Rdriver*, *Rwire_row*, *Rwire_col*, and *Rsense*) modeled as parasitic resistances. The cumulative effect of all the non-idealities results in the deviation of the output current from its ideal value (i.e., Equation (2)), resulting in an *Iout*_{non-ideal} vector. The relative deviation of *Iout*_{non-ideal} from its ideal value is measured by *non-ideality factor* (NF) [6] as

$$NF = (Iout_{ideal} - Iout_{non-ideal})/Iout_{ideal}.$$
(3)

Thus, NF is a direct measure of crossbar non-idealities, i.e., increased non-idealities induce a greater value of NF, affecting the accuracy of the neural network mapped onto them [6, 23, 24].

In this work, we take a trained BNN model with binary weights ($\{-1, +1\}$) and map onto ReRAM crossbars using only two conductance states- G_{MIN} (HRS) and G_{MAX} (LRS) (details on mapping of negative and positive weights of the BNN along with circuit implementation have been presented in Section 4.1). From Figure 1, once the binarized weights of a BNN are mapped onto crossbars to obtain G_{ideal} , we integrate the resistive interconnect non-idealities of the crossbar and synaptic device variations to convert G_{ideal} into $G_{non-ideal}$ (see Section 4.2). This completes the mapping of the weights of the BNN onto a non-ideal crossbar. Thus, we have

$$Iout_{non-ideal} = V_{in} * G_{non-ideal}.$$
(4)

Unless otherwise stated, we carry out experiments on ReRAM crossbars with $R_{MIN} = 20k\Omega$ and device ON/OFF ratio of 10 [17]. The resistive are non-idealities as follows: $Rdriver = 1k\Omega$, $Rwire_row = 5\Omega$, $Rwire_col = 10\Omega$ and $Rsense = 1k\Omega$ [6, 13, 23, 55]. The device-to-device variations in ReRAM conductances are modeled using a Gaussian distribution around the nominal device conductances with $\sigma/\mu = 10\%$ [10, 55]. For our experiments, by performing SPICE simulations with the ReRAM device model [25], we identified that the binary analog voltages input to the ReRAM crossbars for BNN inference can be +0.1/-0.1 V to maintain linear ReRAM I-V characteristics [35]. In this work, we have assumed that the correct mapping of the BNN weights to the memristive conductances (HRS or LRS) without stuck-at-fault defects is ensured post-training [56]. We use a hardware evaluation framework in Pytorch [2, 23] to map the BNNs onto non-ideal crossbars and investigate the cumulative impact of the resistive and device-level non-idealities on networks mapped with *SwitchX* technique (explained in Section 4.1) on the robustness of neural networks.

3.2 Adversarial Attacks and Defense

Neural networks are vulnerable to adversarial attacks in which the model gets fooled by applying precisely calculated small perturbations on the input [39]. Goodfellow et al. [15] proposed *Fast Gradient Sign Method* (FGSM) to generate adversarial attacks (X_{adv}) by linearization of the loss function (L) of the trained models with respect to the input (X) as shown in Equation (5).

$$X_{adv} = X + \epsilon \times sign(\nabla_x L(\theta, X, y_{true})).$$
⁽⁵⁾

Here, y_{true} is the true class label for the input X; θ denotes the model parameters (weights, biases, and so on). The quantity $\Delta = \epsilon \times sign(\nabla_x L(\theta, X, y_{true}))$ is the net perturbation which is controlled by ϵ . It is noteworthy that gradient propagation is a crucial step in unleashing an adversarial attack. Furthermore, the contribution of a gradient to Δ would vary for different layers of the network depending upon the activations [39]. In addition to FGSM-based attacks, multi-step variants of FGSM, such as **Projected Gradient Descent** (**PGD**) [34] have also been proposed that cast stronger attacks. The PGD attack, shown in Equation (6), is an iterative attack over *n* steps. In each step *i*, perturbations of strength α are added to X_{adv}^{i-1} . Note, that X_{adv}^0 is created by adding random noise to the clean input *X*. Additionally, for each step, X_{adv}^i is projected on a *Norm ball* [34], of radius ϵ . In other words, we ensure that the maximum pixel difference between the clean and adversarial inputs is ϵ .

$$X_{adv} = \sum_{i=1}^{n} X_{adv}^{i-1} + \alpha sign(\nabla_{x} \mathcal{L}(\theta, X, y_{true})).$$
(6)

To build resilience against small adversarial perturbations, defense mechanisms such as gradient masking or obfuscation [40] have been proposed. Such methods construct a model devoid of useful gradients, thereby making it difficult to create an adversarial attack. Furthermore, adversarial training [15, 29, 30, 34] is the state-of-the-art and strongest-known software defense against

adversarial attacks. Here, the training dataset is augmented with adversarial examples so that the network learns to predict them correctly. The authors in [3] have also shown that hardware non-idealities can intrinsically lead to defense via gradient obfuscation against adversarial perturbations, thereby making a neural network on hardware adversarially robust than baseline software models.

In this work, we explore how *SwitchX* mapping of BNNs on non-ideal crossbars can yield adversarial robustness when compared with Normal Mapping of BNN weights. Furthermore, we show that when *SwitchX* is combined with state-aware training of BNN (discussed in Section 5), it intrinsically unleashes even stronger defense on hardware than adversarial training.

4 METHODOLOGY

4.1 SwitchX Mapping

For the baseline or Normal Mapping of BNNs onto crossbars, "+1" weights are mapped to G_{MAX} or LRS and "'-1" to G_{MIN} or HRS as shown in Figure 1 (see Normal Mapping). In this work, we perform switched-mapping of the weights of a BNN onto crossbars, termed as *SwitchX*. As shown in Figure 1 for *SwitchX*, we compute the mean of the binarized weight matrix of the BNN to be mapped onto a crossbar instance. If the mean is positive-valued (mean > 0), we switch the mapping with "+1" values to HRS states and "-1" values to LRS states. If the mean is negative-valued or zero (mean ≤ 0), we do not perform any switching and conduct Normal Mapping. This approach ensures that the proportion of HRS states in a crossbar array is always higher when mapping BNN weights. As noted in Figure 1, the ideal conductance matrix G_{ideal} obtained from the switched mapping is again converted into $G_{non-ideal}$ taking the crossbar non-idealities into account.

Circuit implementation of SwitchX for BNN inference: Figure 3 shows the manner in which BNN inference is carried out accurately on an m × m crossbar via SwitchX approach. Standard techniques involve a dual-crossbar approach to map neural networks onto crossbars with negative weights. However, in this work, we follow a single crossbar-based approach similar to [26, 53] for mapping both positive and negative BNN weights as ReRAM conductances, thereby reducing hardware overheads. The binary digital inputs of a BNN $(A_1 - A_m)$ are converted into analog voltages $(V_{i1}-V_{im})$ using DACs (level-shifters producing +0.1/-0.1 V) which are input to the crossbar. The software model weights of a BNN are centered around zero, i.e., $\{-1, +1\}$. In the crossbars, the binary weights of the BNN ($\{-1, +1\}$) are mapped as $\{G_{MIN}, G_{MAX}\}$ synaptic conductances that are not symmetrically centered around zero. Hence, to conduct accurate BNN inference, we need to perform some transformations in hardware to ensure that the trained software weights after being mapped to conductance values still remain symmetrically centered around zero, i.e., $\left\{\frac{-(G_{MAX}-G_{MIN})}{2}, \frac{+(G_{MAX}-G_{MIN})}{2}\right\}$. To this end, an additional column of ReRAM devices with conductances of value $G_0 = \frac{(\tilde{G}_{MAX} + G_{MIN})}{2}$ is added (highlighted in blue) and $R_0 = R$ is set to facilitate the transformation. Note, this is viable for the case when the ReRAM devices under consideration can be programmed to a level between G_{MIN} and G_{MAX} having a conductance equal to $\frac{(G_{MAX}+G_{MIN})}{2}$. In case, the ReRAM devices are bimodal, we realize $G_0 = (G_{MAX} + G_{MIN})$ using a parallel combination of two devices- one in HRS state and the other in LRS state, and set $R_0 = \frac{R}{2}$. The resulting crossbar currents are sensed by the ADC unit consisting of transconductance amplifiers (generating voltages V_{o1} - V_{om}) followed by comparators (with $V_{REF} = 0$) to produce binarized activations. However, for weights mapped onto crossbars via SwitchX transformation, wherein HRS and LRS states are interchanged for the case when mean > 0, a corresponding inverse transformation is also necessary in the end to ensure the correct activations $(X_1 - X_m)$ at the output. This is done by the digital module consisting of inverters and multiplexers (highlighted in violet). The SEL signal input to the multiplexers is set to "0" for mean > 0 scenario to conduct switched-mapping and



Fig. 3. Circuit for *SwitchX* implementation during BNN inference shown for an m × m memristive crossbar (highlighted in green). Here, $V_{i1} - V_{im}$ are analog voltages input to the crossbar and $X_1 - X_m$ are the output binary activations.

"1" for mean ≤ 0 scenario to conduct Normal Mapping to produce the correct output activations $(X_1 - X_m)$. Effectively, during *SwitchX* implementation, based on the value of mean, the analytical operations are governed by Equations (7) and (8).

For mean ≤ 0 ,

$$\sum_{i} a_{i} * w_{ij} = \sum_{i} V_{i} * G_{ij} - \sum_{i} V_{i} * G_{0}.$$
(7)

Here, the subtraction operation is needed to ensure that the trained binarized weights after being mapped to conductance values in the crossbars remain symmetrically centered around zero, i.e., $\left\{\frac{-(G_{MAX}-G_{MIN})}{2}, \frac{+(G_{MAX}-G_{MIN})}{2}\right\}$. Note, the above equation is also true for normally mapping a BNN irrespective of the value of mean [26].

For mean > 0,

$$\sum_{i} a_i * w_{ij} = -\left[\sum_{i} V_i * G_{ij} - \sum_{i} V_i * G_0\right].$$
(8)

Basically, for switched-mapping in Equation (8), the additional inverse operation with respect to Equation (8) is emulated using the digital inverse transformation module (highlighted in violet in Figure 3).

Trends for non-ideality factor (NF): Figure 4 shows the variation in NF for different crossbar sizes for a given binary weight matrix with a higher proportion of "+1" values (mean > 0) for



Fig. 4. Box Plot of error bars to show variations in NF for various crossbar sizes for Normal and *SwitchX* mapping.



Fig. 5. Hardware evaluation platform for mapping the BNN parameters onto non-ideal crossbars following *SwitchX* technique.

a fully-connected layer of a neural network and a given set of input voltages drawn randomly from a uniform distribution. It is observed that NF increases with increasing crossbar size (16×16 to 64×64). Further for a given crossbar size, *SwitchX* results in a decrease in the value of NF with respect to the case of Normal Mapping. This occurs because *SwitchX* increases the effective resistance of a crossbar array, thereby minimizing the effect of interconnect parasitic resistances and device-level variations, by increasing the proportion of HRS states. *Thus, BNNs mapped onto crossbars via SwitchX approach would suffer less interference from non-idealities on MVM operations which would thereby lead to lesser accuracy degradation.*

4.2 Hardware Evaluation Framework

As discussed in Section 3.1, we use a framework in Pytorch, based on RxNN [23], to map trained BNNs onto non-ideal memristive crossbars via *SwitchX* technique and investigate the cumulative impact of the circuit and device-level non-idealities on their adversarial robustness. Figure 5 illustrates the overall simulation framework that is used for non-ideal crossbar evaluation using *SwitchX* technique. The entire platform, being based on Python, enables better integration between the software model and the simulation framework. In the platform, a Python wrapper is built that unrolls each and every convolution operation in the software BNN into MAC operations. The matrices obtained are then zero-padded (in case the size of the weight matrix is not an exact multiples of the crossbar size) and partitioned into multiple crossbar instances consisting of the binarized weights. The next stage (functional modeling) of the platform converts the binarized weights W to suitable conductance states G (HRS or LRS as marked in green and blue colors,

Synapse characteristics		Non-idealities		
Parameter	Value	Parameter	Value	
R _{MIN}	20 kΩ	Rdriver, Rsense	$1 \text{ k}\Omega, 1 \text{ k}\Omega$	
R _{MAX}	200 kΩ	Rwire_row, Rwire_col	5 Ω, 10 Ω	
ON/OFF ratio	10	Synaptic variation	10%	

Table 2. Table Showing Default Crossbar and NVM Device Specificationsfor Non-ideality Aware Performance Evaluation of BNNs

Table 3.	Baseline Software M	odels using V	GG16 BNN	with CIFAR-10,	CIFAR-100 and	TinyImagenet
		Datasets Si	ubjected to	FGSM Attack		

Dataset	Clean Accuracy (%)	Adversarial accuracies (%)					
		$\epsilon = 0.05$	<i>ϵ</i> = 0.1	<i>ε</i> = 0.15	$\epsilon = 0.2$	$\epsilon = 0.25$	<i>ε</i> = 0.3
CIFAR-10	88.92	44.63	39.81	35.87	32.7	30.58	28.17
CIFAR-100	55.04	15.97	13.29	11.51	10.01	8.82	7.98
TinyImagenet	47.4	6.28	5.47	4.8	4.25	3.93	3.7
CIFAR-10 with state-aware training ($ \delta = 1e - 3$)	87.55	46.61	40.63	35.64	31.82	28.86	26.38

respectively) using *SwitchX* approach based on proportions of "+1"s and "-1"s in the crossbar instances. Thereafter, the circuit-level parasitic non-idealities are integrated via circuit laws (*Kirchoff's laws*) and linear algebraic operations written in Python [2]. This integration is similar to the numerical operations adopted in the *RxNN* framework, which is a recent work that accurately models circuit non-idealities while evaluating DNNs on crossbars during inference. *RxNN* has been shown to closely match the results obtained using SPICE models for interconnect parasitics and achieves significantly high speed up during DNN inference on non-ideal crossbars. Furthermore, the ReRAM device variations are included with gaussian profiling. The default specifications of the NVM devices as well as the values of non-idealities have been listed in Table 2, which are used for our experiments. The non-ideal synaptic conductances *G'* are then integrated into the original Pytorch-based BNN model to conduct inference. This framework, thus, enables us to analyze the impact of intrinsic crossbar non-idealities on mapping BNNs with *SwitchX*.

5 EXPERIMENTS

We conduct experiments on a VGG16 BNN architecture with benchmark datasets- CIFAR-10, CIFAR-100 and TinyImagenet. The CIFAR-10 and CIFAR-100 datasets consist of 50,000 training and 10,000 test RGB images of size 32×32 belonging to 10 and 100 classes, respectively. The Tiny-Imagenet dataset is a larger and complex dataset consisting of RGB images of size 64×64 . It is a subset of the Imagenet dataset having 100,000 training images and 10,000 test images from 200 different classes. After training the BNNs on software, we launch FGSM or PGD attacks by adding adversarial perturbations to the clean test inputs and record the adversarial accuracies in each case. This forms our baseline software models (first two rows in Table 3).

State-aware training: In standard scenarios, the binarization of weights during softwaretraining occurs with respect to a threshold (δ) value of 0.0. That is, a weight value greater (or lesser) than 0.0 is quantized as "+1" (or "-1"), respectively. State-aware training is a method to make the distribution of binarized weights across a channel more non-uniform [18]. In this approach, the threshold (δ) is a hyperparameter assuming a positive value when there are more negative weights in a channel and a negative value otherwise. The weights are now quantized according to the new threshold value. For the VGG16 BNN, we found that a value $|\delta| = 1e - 3$ for CIFAR-10 and TinyImagenet datasets and $|\delta| = 5e - 4$ for CIFAR-100 dataset were optimal to retain the performance of the BNNs upon state-aware training on software and hence, selected for the upcoming experiments. Subsequently, *SwitchX*-mapping on such networks increases the proportion of HRS states in the crossbar instances (implying a greater reduction in crossbar non-ideality factor). State-aware training on VGG16 BNN with CIFAR-10 dataset yields a baseline software model as shown in Table 3 (fourth row). Data for VGG16/CIFAR-100 and VGG16/TinyImagenet with state-aware training have not been shown for brevity. We extensively analyze the benefits of this approach in the upcoming sections.

Modes of adversarial attack: For the adversarial attacks (FGSM/PGD) on the crossbar-mapped models of the BNNs, we consider two modes:

- (1) *Software-inputs-on-hardware (SH) mode*: The adversarial perturbations for each attack are created using the software-based baseline's loss function and then added to the clean input that yields the adversarial input. The generated adversaries are then fed to the crossbarmapped BNN. This is a kind of black-box adversarial attack.
- (2) Hardware-inputs-on-hardware (HH) mode: The adversarial inputs are generated for each attack using the loss from the crossbar-based hardware models. This is a kind of white-box adversarial attack. It is evident that HH perturbations will incorporate the intrinsic hardware non-idealities and thus will cast stronger attacks than SH.

For all FGSM attacks in this work, we have $\epsilon = \{0.05, 0.1, 0.15, 0.2, 0.25, 0.3\}$. Also, we consider PGD attacks iterated over 7 steps with $\alpha = 2/255$ and $\epsilon = \{2/255, 4/255, 8/255, 16/255, 32/255\}$. Note, for all upcoming analyses which include the impact of interconnect parasitic non-idealities, to have a feasible runtime for the simulations during BNN inference that are out by first mapping weights on crossbars via *SwitchX* method and then integrating with non-idealities and variations (see Figure 5), we have shown results on crossbar sizes of 16×16 or 32×32 .

6 RESULTS AND DISCUSSION

6.1 Reduction in Adversarial Noise Sensitivity

Earlier works [31, 39] have identified a metric termed as *Adversarial Noise Sensitivity* (ANS) to quantify the sensitivity of each layer of a neural network to adversarial perturbations. ANS for a layer *l* in a neural network, subjected to an adversarial attack, is defined in terms of an error ratio as follows:

$$ANS_{l} = \frac{||A_{adv}^{l} - A^{l}||_{2}}{||A^{l}||_{2}},$$
(9)

where, A^l and A^l_{adv} are respectively the clean and adversarial activation values of the layer *l*. ANS is a simple metric to evaluate how each layer contributes to the net adversarial perturbation during the gradient propagation.

In Figure 6 we show results for a VGG16 BNN (trained using state-aware training) mapped via *SwitchX*, and a standard VGG16 BNN mapped normally onto 32×32 crossbars ($R_{MIN} = 20k\Omega$ and ON/OFF ratio = 10), both adversarially perturbed using HH mode of FGSM attack. We plot the ANS values for all the convolutional layers and the final fully-connected layer (marked as FC). For the case of *SwitchX*-mapped BNN we observe lower ANS values implying a reduced error amplification effect [31] and hence, lesser sensitivity and better stability to the induced adversarial perturbations. However, the authors in [39] also agree that a lower value of ANS across the layers of a neural network would not necessarily imply improved adversarial robustness under all circumstances. Hence, in Section 6.2, we propose a novel graphical approach by plotting "*robustness maps*" to evaluate the adversarial robustness of BNNs upon *SwitchX* mapping onto non-ideal crossbars.



Fig. 6. ANS values plotted across different layers of VGG16 BNN trained on CIFAR-10 for Normal Mapping and *SwitchX* mapping on 32×32 crossbars. The *SwitchX* BNN is trained via state-aware training. Here, layers 3, 6, 11, and 16 denote the Pooling layers and FC marks the final fully-connected layer. Results are shown for HH mode of FGSM attack with $\epsilon = 0.1, 0.2, 0.3$.

6.2 Robustness Analysis

Earlier approaches to quantify adversarial robustness of neural networks have been based on *Adversarial Loss* (AL) metric which is the difference between the clean accuracy and the adversarial accuracy for a given value of ϵ [2, 38, 39]. Note, clean accuracy is the natural accuracy of a neural network when not under attack. A reduction in the value of AL is said to improve the adversarial robustness of the network. However, AL is not always a suitable metric for the assessment of robustness since reduction in AL does not convey whether the cause is a decrease in clean accuracy or an increase in adversarial accuracy or both. For instance, while evaluating a network with CIFAR-10 dataset, we find the clean and the adversarial accuracies (for a particular ϵ) to be 10% each. Then the value of AL is zero implying that the network is exceptionally robust. However, this is absurd since a 10% accuracy is random for CIFAR-10 dataset implying that the network is arbitrarily predicting and is not trained.

In this work, we evaluate robustness of the mapped BNNs on crossbars graphically as shown in Figures 7 and 8. For a specific mode of attack (SH or HH) and a given crossbar size, we plot Δ *Clean Accuracy*, the difference between clean accuracy of the mapped network in question and the corresponding clean accuracy of the software baseline, on the x-axis. Δ Adversarial Accuracy (for a particular ϵ value) which is the difference between the adversarial accuracy of the mapped network in question and the corresponding adversarial accuracy of the software baseline is plotted on the y-axis. We term this plot as a *"robustness map"*. The value of \triangle *Clean Accuracy* is negative since BNNs when mapped on hardware suffer accuracy loss owing to non-idealities. The region bounded by the line y = -x and the y-axis denotes that the absolute increase in the adversarial accuracy is higher than the absolute degradation in the clean accuracy. If a point lies in this region closer to the y-axis, it implies greater adversarial accuracy with lower clean accuracy loss and hence, greater robustness. Therefore, this is our favorable region. As we move farther away from the y-axis in the favorable region, the robustness reduces as the neural network suffers from high loss of clean accuracy, and this has been shown by the variation in color gradient from dark to lightbrown (dark shade implying higher robustness). Likewise, the region bounded by the line y = xand the y-axis is where the mapped-network is highly vulnerable to adversarial attacks and hence, the unfavorable region. Here, as we move farther away from the y-axis in the unfavorable region, the robustness reduces and the degree of unfavorability increases. This has been shown by the variation in color gradient from dark to light-yellow (darker shade implying higher robustness).



Fig. 7. (a) and (b) Robustness maps for VGG16 BNN using CIFAR-100 dataset for SH and HH modes of FGSM attack respectively; (c) and (d) Robustness maps for VGG16 based BNN using CIFAR-100 dataset for SH and HH modes of PGD attack respectively.



Fig. 8. (a) and (b) Robustness maps for VGG16 based BNN using CIFAR-10 dataset for SH and HH modes of FGSM attack respectively considering the case of *SwitchX* combined with state-aware training; (c) and (d) Robustness maps for VGG16 based BNN using CIFAR-10 dataset for SH and HH modes of FGSM attack respectively considering the case of *SwitchX* combined with adversarial training.

This approach to assess robustness of a network is comprehensive and accurate since, it takes into account the cumulative impact of both clean accuracy and adversarial accuracy (which is a strong function of the clean accuracy). The closer a point is to the region marked with dark-brown color, the better the robustness of the neural network. Note, in Figures 7 and 8, the circular points correspond to mappings on 16×16 crossbars while the triangular points correspond to mappings on 32×32 crossbars. All the details pertaining to the baseline software models have been listed in Table 3. Furthermore, the results in Figures 7 and 8 are for BNNs mapped onto crossbars having ReRAM device ON/OFF ratio of 10 with $R_{MIN} = 20k\Omega$.

Figure 7 shows the robustness maps for BNNs based on VGG16 network with CIFAR-100 dataset for both SH and HH modes of attack. Figure 7(a) and (b) pertain to FGSM attack with ϵ varying from 0.05 to 0.3 with step size of 0.05. We find that *SwitchX* imparts greater clean accuracy (~3% for 32 × 32 crossbar) as well as better adversarial accuracies on hardware for both modes of attack, with the points corresponding to *SwitchX* situated closer to the dark-brown portion favorable region than the corresponding points for Normal Mapping. This is a consequence of the reduction in non-ideality factor in case of *SwitchX* with respect to Normal Mapping as discussed in Section 3.1. Note, the points for 32 × 32 crossbars are situated farther from the favorable region than the corresponding points for 16 × 16 crossbars. This is owing to greater non-idealities that exist in case of a 32 × 32 crossbar than a 16 × 16 crossbar (Figure 4). We further observe that the points for *SwitchX BNN* for a given crossbar size are more closely packed than the corresponding points for Normal BNN. This implies that even on increasing the perturbation strength (ϵ), lesser adversarial loss is observed in case of *SwitchX BNN*.



Fig. 9. Robustness maps for VGG16 BNN with- (a) CIFAR-100 dataset for HH mode of FGSM attack mapped onto 32 × 32 crossbars with $R_{MIN}/R_{MAX} = 20k\Omega/200k\Omega$ (circle), $R_{MIN}/R_{MAX} = 50k\Omega/250k\Omega$ (diamond) and $R_{MIN}/R_{MAX} = 50k\Omega/500k\Omega$ (triangle); (b) TinyImagenet dataset for HH mode of FGSM attack mapped onto 32 × 32 crossbars (triangle) and 16 × 16 crossbars (circle) with $R_{MIN}/R_{MAX} = 20k\Omega/200k\Omega$.

Figure 7(c) and (d) also present similar results but for a PGD attack with ϵ varying from 2/255 to 32/255 with step size of 2/255. In this case, the robustness is very high for SH mode of attack as compared to HH mode of attack, with points corresponding to 16 × 16 crossbars situated inside the favorable region. Similar to the case of FGSM attack, *SwitchX* outperforms a Normal BNN in terms of robustness for both modes of attack. Here, points for different ϵ values, given a style of mapping and crossbar size, are more closely packed than the corresponding points of FGSM attack. This implies that hardware non-idealities interfere more with PGD attacks than FGSM attacks resulting in lesser accuracy degradation.

Effect of varying R_{MIN} of ReRAM devices: Previous works such as [3, 6] have shown that the impact of resistive crossbar non-idealities (or NF) decreases upon increasing R_{MIN} of memristive devices. To this end, we plot robustness maps in Figure 9(a) to compare Normal and *SwitchX* BNNs upon increasing R_{MIN} from 20 $k\Omega$ to 50 $k\Omega$. We find that for both Normal and *SwitchX* BNNs, data points corresponding to $R_{MIN} = 50k\Omega$ (triangle) are closer to the favorable region than for $R_{MIN} = 20k\Omega$ (circle) for the same ON/OFF ratio of 10, signifying reduced impact of non-idealities on the crossbar-mapped network models. Further, for $R_{MIN} = 50k\Omega$, *SwitchX* BNNs show performance improvement in terms of robustness over Normal BNNs (~1% improvement in clean accuracy and adversarial accuracies for HH based FGSM attack on 32×32). Also, in Figure 9(a), we increase R_{MIN} from 20 $k\Omega$ to 50 $k\Omega$ and reduce the device ON/OFF ratio to 5 (diamond-shaped points). We find that even at a lower ON/OFF ratio of 5, *SwitchX* BNN outperforms the corresponding Normal BNN in terms of robustness is a stronger function of R_{MIN} than R_{MAX} . This is because for both *SwitchX* and Normal BNNs, improvement in robustness is greater on traversing from the circular to the diamond-shaped points than from the diamond-shaped points to the triangular points.

Efficacy of *SwitchX* combined with state-aware training: Figure 8 shows the robustness maps for BNNs based on VGG16 network with CIFAR-10 dataset for both SH and HH modes of FGSM attack. Here, we find that there is a significant benefit in terms of improvement in clean accuracy (~10%) and adversarial accuracies (~2–5%) due to *SwitchX* for a 32×32 crossbar with

respect to Normal Mapping. Similar to the case for CIFAR-100 dataset, we find that *SwitchX* outperforms a Normal BNN in terms of robustness for both modes of attack. We now analyze cases when *SwitchX* is combined with state-aware training as well as adversarial training. The results for FGSM attack are summarized:

6.2.1 With Adversarial Training. From Figure 8(c) and (d), we find that SwitchX combined with adversarial training significantly boosts the robustness of the mapped BNN both in terms of clean as well as adversarial accuracy improvement. For 16 × 16 crossbar, the points (in red) for different ϵ lie in the close vicinity of the boundary of the favorable region while for 32 × 32 crossbar, the rise in clean accuracy with respect to Normal BNN is very high (~20%). Furthermore, given a style of mapping and crossbar size, the points for adversarial training are more closely packed than the corresponding points of standalone SwitchX BNN or Normal BNN, implying lesser accuracy losses on increasing perturbation strength of the attack.

6.2.2 With State-aware Training. As discussed in Section 5, state-aware mapping combined with *SwitchX* can lead to an increase in the proportion of HRS states in the crossbar instances, and thus a reduction in the non-ideality factor of the crossbars. From Figure 8(a) and (b), we find that this approach significantly boosts the robustness of the mapped BNN both in terms of clean as well as adversarial accuracy improvements. For both 16×16 and 32×32 crossbars, the points (in red) for different ϵ lie in the close vicinity of the boundary of the favorable region. We find that for 32×32 crossbar, the rise in clean and adversarial accuracies is so large that it becomes comparable to the case of standalone *SwitchX BNN* mapped on a smaller 16×16 crossbar. Overall, we find the rise in clean and adversarial accuracies is ~35% and ~6-16% greater than Normal BNN, respectively. In this case also, the points for different ϵ values, given a style of mapping and crossbar size, are more closely packed than the corresponding points of standalone SwitchX BNN or Normal BNN, implying lesser accuracy losses on increasing perturbation strength of the attack. Interestingly, we find that this approach emerges as a stronger defense against adversarial attacks than SwitchX combined with adversarial training (a state-of-the-art software defense), the defensive action being more pronounced for larger crossbar sizes (32×32 in Figure 8(a) and (b)).

Furthermore, we observe similar results in Figure 9(b) using the large and complex TinyImagenet dataset with VGG16 BNN. For crossbar sizes of 16 × 16 and 32 × 32, we find *SwitchX* BNNs to outperform the corresponding Normal BNNs in terms of both clean and adversarial accuracies. There is ~5.5% improvement in clean accuracies and ~1–2% improvement in FGSM adversarial accuracies for HH mode of attack. With state-aware training and *SwitchX* mapping combined, the improvements on 16 × 16 and 32 × 32 crossbars shoot upto ~21.2% for clean accuracies and ~4.5–8.3% for adversarial accuracies.

Effect of varying device ON/OFF ratio in crossbars: The results in [6] show that on increasing the ReRAM device ON/OFF ratio (increasing the value of HRS at a constant value of LRS) in crossbars, the non-ideality factor decreases. This should translate to robustness benefits for our *SwitchX* approach that in itself increases the feasibility of HRS states in crossbars. In Figure 10, we map BNNs onto 32×32 crossbars having device ON/OFF ratios of 10 and 100 respectively. We then compare the robustness of a standalone *SwitchX* BNN and a *SwitchX* BNN combined with state-aware training for a VGG16 network with CIFAR-10 dataset for both cases. We find that mapping BNNs on crossbars having higher ON/OFF ratios boosts the robustness of the mapped BNNs against both SH and HH based adversarial attacks (~2–4% for standalone *SwitchX*). Further, there are ~6% and ~2% improvements in clean accuracy respectively for standalone *SwitchX* and *SwitchX* combined with state-aware training.



Fig. 10. Robustness maps for VGG16 BNN with CIFAR-10 dataset for SH (Left) and HH (Right) modes of FGSM attack respectively mapped onto 32×32 crossbars with device ON/OFF ratios of 10 (red) and 100 (blue) for both standalone *SwitchX* BNN and *SwitchX* BNN + state-aware training.



Fig. 11. Robustness maps for VGG16 BNN with CIFAR-10 dataset for SH (Left) and HH (Right) modes of FGSM attack respectively mapped onto 64×64 (circle) and 128×128 (triangle) crossbars for both standalone *SwitchX* BNN and *SwitchX* BNN + state-aware training.

Effect of synaptic device variations on larger crossbars with greater R_{MIN} :

Here, we increase R_{MIN} of the NVM devices to $200k\Omega$ (at ON/OFF ratio of 10) so that the impact of interconnect parasitic non-idealities can be minimized, thereby making simulations on larger crossbar sizes with a modest runtime feasible. In this scenario, the synaptic device variations become predominant and determine the robustness of the crossbar-mapped BNN models. For the robustness maps in Figure 11, we assume the NVM device variations to be Gaussian with $\sigma/\mu = 20\%$ and the BNN mappings are carried out on 64×64 and 128×128 crossbars. The results are consistent with that in Figure 8(a) and (b), indicating that the *SwitchX* method is applicable to BNN mapped onto larger crossbar sizes.

6.3 Impact of SwitchX on Crossbar Power Consumption

Owing to low leakage of the memristive devices, the V.I power consumption in the crossbars constitutes a significantly large portion of the overall power expended in low-precision BNN



Fig. 12. (a) Plot showing power savings in 16×16 crossbar for standalone *SwitchX* as well as *SwitchX* combined with state-aware training with respect to the baseline of Normal BNN for VGG16 BNN; (b) Plot showing layer-wise crossbar power consumption for the different convolutional layers in the VGG16 BNN using CIFAR-10 and CIFAR-100 datasets in case of mapping on 16×16 crossbars.



Fig. 13. Plot showing power savings on various crossbar sizes with *SwitchX* mapping.

inference on ReRAM crossbars [8, 18, 27]. Thus, in this section, we show an interesting by-product of *SwitchX* mapping, whereby we can also reduce the crossbar power consumption compared to normally mapped BNNs (shown in Figure 12). Here, the simulations for estimating the *V.I* power consumed by the crossbars during BNN inference has been performed using ReRAM devices with $R_{MIN} = 20k\Omega$ and ON/OFF ratio of 10 using the NeuroSim tool [41], and the analog voltages input to the crossbar are +0.1/-0.1 V (obtained via SPICE simulations as specified in Section 3.1).

Trends for energy-efficiency in crossbars: Figure 13 shows a plot of the power savings observed when randomly generated binary weight matrices with a higher proportion of "+1" values are mapped via *SwitchX* onto crossbars of sizes ranging from 16×16 , to 128×128 . Here, input voltages to the crossbars are drawn from a uniform distribution. The case of "90% HRS states" implies that 10% of the values in the BNN weight matrix were "-1", while "75% HRS states" implies that 25% of the values were "-1", i.e., the weight matrix is less non-uniformly distributed than the former. Similarly, "60% HRS states" implies that 40% of the values in the BNN weight matrix were "-1". We find that the power savings (~7–34% on 64×64 crossbars) increase when the distribution of BNN weights becomes more non-uniform (from "60% HRS states" to "90% HRS states") for different crossbar sizes. A more non-uniform distribution would imply greater proportion of HRS states on *SwitchX* mapping, thereby translating into greater crossbar power savings owing to lower dot-product currents. Note, state-aware training in BNNs increases the non-uniformity in the distribution of HRS-LRS synapses in crossbars when mapped using *SwitchX*.

From Figure 12(a) (data shown for a 16 \times 16 crossbar), we find that standalone *SwitchX* approach leads to ~9% power-savings on an average with respect to Normal BNN for a VGG16 network with

Work	Type of non-ideality	Energy-efficiency	Adversarial Robustness (%)
REBNN [51]	N.A.	4.26×	N.A.
Lattice [57]	N.A.	7.68×	N.A.
[12]	Device noise	N.A.	~6% \uparrow (Black-box)
SwitchX	Resistive and Device noises	1.21-1.22×*	~6–16% ↑ (White-box)

Table 4. Table Showing Comparison with Previous Works of BNNs on Crossbars

The improvements in energy-efficiency and adversarial accuracy (wherever applicable) have been quoted as is from the previous works. Here, N.A. stands for not applicable.

*Crossbar energy-efficiency.

CIFAR-100 dataset. While for CIFAR-10 dataset, there is ~8% power saving for standalone *SwitchX*, it increases to ~22% when *SwitchX* is combined with state-aware training. Furthermore, on carrying out similar experiments on 32×32 crossbars, we obtain ~21% power savings on combining *SwitchX* with state-aware training (data not shown for brevity). For CIFAR-100 dataset, we obtain ~19% power savings on 16 × 16 crossbars on combining *SwitchX* with state-aware training. Figure 12(b) shows the layer-wise normalized average power consumption by the network with CIFAR-100 and CIFAR-10 dataset. We find that overall for each convolutional layer of the network, the power consumed by the *SwitchX* BNN is lesser than the Normal BNN. Furthermore, this reduction becomes even more significant between "conv7" and "conv8" layers when *SwitchX* is combined with state-aware training (for CIFAR-10 dataset). This result is in accordance with Figure 13 which shows that the power savings on crossbar arrays increase significantly when the HRS-LRS state-distribution becomes more non-uniform.

6.4 Comparison with Previous Works

In Section 2, we have specified that there are several previous works on neural network implementations on RRAM crossbar-arrays, such as [51, 57], where the key goal is to obtain high energyefficiency and not adversarial robustness. Although recently, [12] proposed a RRAM device noiseaware training methodology to generate BNN models that help improve natural as well as adversarial robustness on crossbars with device variations, it does not take into account the impact of circuit-level resistive non-idealities such as crossbar-interconnect parasitics. Also, the noise-aware training methodology in [12] does not translate to any benefits in terms of energy-efficiency for BNNs on crossbars. However, our SwitchX mapping strategy has an interesting by-product of improving crossbar energy-efficiency in addition to providing better adversarial robustness (see Section 6.3). In Table 4, we quantitatively compare our results with some of the prior works on BNNs implemented on crossbars, including [12], to put our work in context. Since the scope of this work is primarily to improve the adversarial robustness of BNNs on non-ideal crossbars, we do not achieve the best of energy-efficiency. However, we obtain significantly better robustness against stronger white-box adversarial attacks compared to [12] even upon the consideration of resistive crossbar non-idealities, which have a larger disruptive effect on the performance of neural networks.

7 CONCLUSION

In this work, we propose *SwitchX*-mapping of binary weights onto crossbars in such a manner that a crossbar array always comprises of a greater proportion of HRS than LRS. We perform a comprehensive analysis on crossbar arrays comprising of interconnect and device-level non-idealities, whereby *SwitchX* reduces the overall impact of non-idealities during inference. This effect manifests as reduced accuracy losses and increased robustness against adversarial attacks for BNNs mapped onto crossbars. We also combine *SwitchX* with state-aware training that further

increases the feasibility of HRS states during weight mapping to boost the adversarial robustness of BNNs on crossbars. With an increase in the feasibility of HRS synapses on crossbars, *SwitchX* also helps improve the crossbar energy-efficiencies, in addition to robustness.

REFERENCES

- [1] Aayush Ankit, Izzat El Hajj, Sai Rahul Chalamalasetti, Georey Ndu, Martin Foltin, R. Stanley Williams, Paolo Faraboschi, Wen-mei W. Hwu, John Paul Strachan, Kaushik Roy, et al. 2019. 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. 715–731.
- [2] Abhiroop Bhattacharjee and Priyadarshini Panda. 2020. Rethinking non-idealities in memristive crossbars for adversarial robustness in neural networks. arXiv preprint arXiv:2008.11298 (2020).
- [3] Abhiroop Bhattacharjee, Abhishek Moitra, and Priyadarshini Panda. 2021. Efficiency-driven hardware optimization for adversarially robust neural networks. In 2021 Design, Automation & Test in Europe Conference & Exhibition (DATE). IEEE, 884–889.
- [4] Abhiroop Bhattacharjee, Lakshya Bhatnagar, Youngeun Kim, and Priyadarshini Panda. 2021. NEAT: Non-linearity aware training for accurate, energy-efficient and robust implementation of neural networks on 1T-1R crossbars. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 41, 8 (2021), 2625–2637.
- [5] Nicholas Carlini, Anish Athalye, Nicolas Papernot, Wieland Brendel, Jonas Rauber, Dimitris Tsipras, Ian Goodfellow, Aleksander Madry, and Alexey Kurakin. 2019. On evaluating adversarial robustness. arXiv preprint arXiv:1902.06705 (2019).
- [6] Indranil Chakraborty, Mustafa Fayez Ali, Dong Eun Kim, Aayush Ankit, and Kaushik Roy. 2020. GENIEx: A generalized approach to emulating non-ideality in memristive xbars using neural networks. In 2020 57th ACM/IEEE Design Automation Conference (DAC). IEEE, 1–6.
- [7] I. Chakraborty, A. Jaiswal, A. K. Saha, S. K. Gupta, and K. Roy. 2020. Pathways to efficient neuromorphic computing with non-volatile memory technologies. *Applied Physics Reviews* 7, 2 (2020), 021308.
- [8] Meng-Fan Chang, Albert Lee, Pin-Cheng Chen, Chrong Jung Lin, Ya-Chin King, Shyh-Shyuan Sheu, and Tzu-Kun Ku. 2015. Challenges and circuit techniques for energy-efficient on-chip nonvolatile memory using memristive devices. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems* 5, 2 (2015), 183–193.
- [9] Wei-Hao Chen, Win-San Khwa, Jun-Yi Li, Wei-Yu Lin, Huan-Ting Lin, Yongpan Liu, Yu Wang, Huaqiang Wu, Huazhong Yang, and Meng-Fan Chang. 2017. Circuit design for beyond von Neumann applications using emerging memory: From nonvolatile logics to neuromorphic computing. In 2017 18th International Symposium on Quality Electronic Design (ISQED). IEEE, 23–28.
- [10] Pai-Yu Chen, Binbin Lin, I-Ting Wang, Tuo-Hung Hou, Jieping Ye, Sarma Vrudhula, Jae-sun Seo, Yu Cao, and Shimeng Yu. 2015. Mitigating effects of non-ideal synaptic device characteristics for on-chip learning. In Proceedings of the 2015 IEEE/ACM International Conference on Computer-Aided Design. IEEE, 194–199.
- [11] Wei-Hao Chen, Kai-Xiang Li, Wei-Yu Lin, Kuo-Hsiang Hsu, Pin-Yi Li, Cheng-Han Yang, Cheng-Xin Xue, En-Yu Yang, Yen-Kai Chen, Yun-Sheng Chang, et al. 2018. A 65nm 1Mb nonvolatile computing-in-memory ReRAM macro with sub-16ns multiply-and-accumulate for binary DNN AI edge processors. In *Proceedings of the 2018 IEEE International Solid-State Circuits Conference*. IEEE, 494–496.
- [12] Sai Kiran Cherupally, Jian Meng, Adnan Siraj Rakin, Shihui Yin, Injune Yeo, Shimeng Yu, Deliang Fan, and Jae-Sun Seo. 2022. Improving the accuracy and robustness of RRAM-based in-memory computing against RRAM hardware noise and adversarial attacks. *Semiconductor Science and Technology* 37, 3 (2022), 034001.
- [13] Yexin Deng, Peng Huang, Bing Chen, Xiaolin Yang, Bin Gao, Juncheng Wang, Lang Zeng, Gang Du, Jinfeng Kang, and Xiaoyan Liu. 2012. RRAM crossbar array with cell selection device: A device and circuit interaction study. *IEEE Transactions on Electron Devices* 60, 2 (2012), 719–726.
- [14] Deliang Fan, Yong Shim, Anand Raghunathan, and Kaushik Roy. 2015. STT-SNN: A spin-transfer-torque based softlimiting non-linear neuron for low-power artificial neural networks. *IEEE Transactions on Nanotechnology* 14, 6 (2015), 1013–1023.
- [15] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and harnessing adversarial examples. arXiv preprint arXiv:1412.6572 (2014).
- [16] Shupeng Gui, Haotao Wang, Haichuan Yang, Chen Yu, Zhangyang Wang, and Ji Liu. 2019. Model compression with adversarial robustness: A unified optimization framework. In *Proceedings of the Advances in Neural Information Processing Systems*. H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.), Vol. 32, Curran Associates, Inc. Retrieved from https://proceedings.neurips.cc/paper/2019/file/2ca65f58e35d9ad45bf7f3ae5cfd08f1-Paper. pdf.

- [17] Basma Hajri, Hassen Aziza, Mohammad M. Mansour, and Ali Chehab. 2019. RRAM device models: A comparative analysis with experimental validation. *IEEE Access* 7 (2019), 168963–168980.
- [18] Yintao He, Ying Wang, Xiandong Zhao, Huawei Li, and Xiaowei Li. 2020. Towards state-aware computation in ReRAM neural networks. In 2020 57th ACM/IEEE Design Automation Conference (DAC). IEEE, 1–6.
- [19] Yintao He, Ying Wang, Huawei Li, and Xiaowei Li. 2021. Saving energy of RRAM-based neural accelerator through state-aware computing. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* (2021), 1–1. DOI:https://doi.org/10.1109/TCAD.2021.3103147
- [20] Tifenn Hirtzlin, Marc Bocquet, J.-O. Klein, Etienne Nowak, Elisa Vianello, J.-M. Portal, and Damien Querlioz. 2019. Outstanding bit error tolerance of resistive ram-based binarized neural networks. In Proceedings of the 2019 IEEE International Conference on Artificial Intelligence Circuits and Systems. IEEE, 288–292.
- [21] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. 2016. Binarized neural networks. Advances in Neural Information Processing Systems 29 (2016).
- [22] Mohsen Imani, Saransh Gupta, Yeseong Kim, and Tajana Rosing. 2019. Floatpim: In-memory acceleration of deep neural network training with high precision. In Proceedings of the 2019 ACM/IEEE 46th Annual International Symposium on Computer Architecture. IEEE, 802–815.
- [23] Shubham Jain, Abhronil Sengupta, Kaushik Roy, and Anand Raghunathan. 2020. RxNN: A framework for evaluating deep neural networks on resistive crossbars. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 40, 2 (2020), 326–338.
- [24] Shubham Jain and Anand Raghunathan. 2019. CxDNN: Hardware-software compensation methods for deep neural networks on resistive crossbar systems. ACM Transactions on Embedded Computing Systems (TECS) 18, 6 (2019), 1–23.
- [25] Jiang et al. 2014. Stanford university resistive-switching random access memory (RRAM) verilog-A model. NanoHUB (2014). DOI:https://doi.org/doi:/10.4231/D37H1DN48
- [26] Yong Kim, Won Hee Jeong, Son Bao Tran, Hyo CheonWoo, Jihun Kim, Cheol Seong Hwang, Kyeong-Sik Min, and Byung Joon Choi. 2019. Memristor crossbar array for binarized neural networks. AIP Advances 9, 4 (2019), 045131.
- [27] Yulhwa Kim, Hyungjun Kim, and Jae-Joon Kim. 2018. Neural network-hardware co-design for scalable RRAM-based BNN accelerators. arXiv preprint arXiv:1811.02187 (2018).
- [28] Alex Krizhevsky, Geoffrey Hinton, et al. 2009. Learning multiple layers of features from tiny images. (2009).
- [29] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. 2016. Adversarial machine learning at scale. arXiv preprint arXiv:1611.01236 (2016).
- [30] Hyeungill Lee, Sungyeob Han, and Jungwoo Lee. 2017. Generative adversarial trainer: Defense to adversarial perturbations with GAN. arXiv preprint arXiv:1705.03387 (2017).
- [31] Ji Lin, Chuang Gan, and Song Han. 2019. Defensive quantization: When efficiency meets robustness. arXiv preprint arXiv:1904.08444 (2019).
- [32] Beiye Liu, Hai Li, Yiran Chen, Xin Li, Tingwen Huang, Qing Wu, and Mark Barnell. 2014. Reduction and IR-drop compensations techniques for reliable neuromorphic computing systems. In 2014 IEEE/ACM International Conference on Computer-Aided Design (ICCAD). IEEE, 63–70.
- [33] Beiye Liu, Hai Li, Yiran Chen, Xin Li, Qing Wu, and Tingwen Huang. 2015. Vortex: Variation-aware training for memristor x-bar. In Proceedings of the 52nd Annual Design Automation Conference. 1–6.
- [34] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2017. Towards deep learning models resistant to adversarial attacks. arXiv preprint arXiv:1706.06083 (2017).
- [35] Adnan Mehonic, Dovydas Joksas, Wing H. Ng, Mark Buckwell, and Anthony J. Kenyon. 2019. Simulation of inference accuracy using realistic RRAM devices. *Frontiers in Neuroscience* 13 (2019), 593.
- [36] Leibin Ni, Zichuan Liu, Hao Yu, and Rajiv V. Joshi. 2017. An energy-efficient digital ReRAM-crossbar-based CNN with bitwise parallelism. *IEEE Journal on Exploratory Solid-state Computational Devices and Circuits* 3 (2017), 37–46.
- [37] Chen Nie, Zongwu Wang, Qidong Tang, Chenyang Lv, Li Jiang, and Zhezhi He. 2022. Cross-layer designs against nonideal effects in ReRAM-based processing-in-memory system. In Proceedings of the 2022 23rd International Symposium on Quality Electronic Design. IEEE, 1–6.
- [38] Priyadarshini Panda, Indranil Chakraborty, and Kaushik Roy. 2019. Discretization based solutions for secure machine learning against adversarial attacks. *IEEE Access* 7 (2019), 70157–70168.
- [39] Priyadarshini Panda. 2020. QUANOS: Adversarial noise sensitivity driven hybrid quantization of neural networks. In Proceedings of the ACM/IEEE International Symposium on Low Power Electronics and Design. 187–192.
- [40] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z. Berkay Celik, and Ananthram Swami. 2017. Practical black-box attacks against machine learning. In Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security. 506–519.
- [41] Xiaochen Peng, Shanshi Huang, Yandong Luo, Xiaoyu Sun, and Shimeng Yu. 2019. DNN+ NeuroSim: An end-to-end benchmarking framework for compute-in-memory accelerators with versatile device technologies. In Proceedings of the 2019 IEEE International Electron Devices Meeting (IEDM). IEEE, 32–5.

- [42] Deboleena Roy, Indranil Chakraborty, Timur Ibrayev, and Kaushik Roy. 2020. Robustness hidden in plain sight: Can analog computing defend against adversarial attacks? arXiv: 2008.1201 (2020).
- [43] Deboleena Roy, Indranil Chakraborty, Timur Ibrayev, and Kaushik Roy. 2021. On the intrinsic robustness of NVM crossbars against adversarial attacks. In Proceedings of the 2021 58th ACM/IEEE Design Automation Conference (DAC). IEEE, 565–570.
- [44] Catherine D. Schuman, Thomas E. Potok, Robert M. Patton, J. Douglas Birdwell, Mark E. Dean, Garrett S. Rose, and James S. Plank. 2017. A survey of neuromorphic computing and neural networks in hardware. arXiv preprint arXiv:1705.06963 (2017).
- [45] Vikash Sehwag, Shiqi Wang, Prateek Mittal, and Suman Jana. 2020. Hydra: Pruning adversarially robust neural networks. Advances in Neural Information Processing Systems 33 (2020), 19655–19666.
- [46] Abhronil Sengupta, Yong Shim, and Kaushik Roy. 2016. Proposal for an all-spin artificial neural network: Emulating neural and synaptic functionalities through domain wall motion in ferromagnets. *IEEE Transactions on Biomedical Circuits and Systems* 10, 6 (2016), 1152–1160.
- [47] Ali Shafiee, Anirban Nag, Naveen Muralimanohar, Rajeev Balasubramonian, John Paul Strachan, Miao Hu, R. Stanley Williams, and Vivek Srikumar. 2016. ISAAC: A convolutional neural network accelerator with in-situ analog arithmetic in crossbars. ACM SIGARCH Computer Architecture News 44, 3 (2016), 14–26.
- [48] Mrigank Sharad, Georgios Panagopoulos, and Kaushik Roy. 2012. Spin neuron for ultra low power computational hardware. In 70th Device Research Conference. IEEE, 221–222.
- [49] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014).
- [50] Linghao Song, Xuehai Qian, Hai Li, and Yiran Chen. 2017. Pipelayer: A pipelined reram-based accelerator for deep learning. In Proceedings of the 2017 IEEE International Symposium on High Performance Computer Architecture (HPCA). IEEE, 541–552.
- [51] Linghao Song, You Wu, Xuehai Qian, Hai Li, and Yiran Chen. 2019. Rebnn: In-situ acceleration of binarized neural networks in reram using complementary resistive cell. *CCF Transactions on High Performance Computing* 1, 3 (2019), 196–208.
- [52] Qidong Tang, Zhezhi He, Fangxin Liu, Zongwu Wang, Yiyuan Zhou, Yinghuan Zhang, and Li Jiang. 2022. HAWIS: Hardware-aware automated width search for accurate, energy-efficient and robust binary neural network on ReRAM dot-product engine. In *Proceedings of the 2022 27th Asia and South Pacific Design Automation Conference (ASP-DAC)*. IEEE, 226–231.
- [53] Son Ngoc Truong and Kyeong-Sik Min. 2014. 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 (2014), 356–363.
- [54] H.-S. PhilipWong, Heng-Yuan Lee, Shimeng Yu, Yu-Sheng Chen, Yi Wu, Pang-Shiu Chen, Byoungil Lee, Frederick T. Chen, and Ming-Jinn Tsai. 2012. Metal–Oxide RRAM. Proc. IEEE 100, 6 (2012), 1951–1970.
- [55] Lixue Xia, Peng Gu, Boxun Li, Tianqi Tang, Xiling Yin, Wenqin Huangfu, Shimeng Yu, Yu Cao, Yu Wang, and Huazhong Yang. 2016. Technological exploration of RRAM crossbar array for matrix-vector multiplication. *Journal of Computer Science and Technology* 31, 1 (2016), 3–19.
- [56] Baogang Zhang, Necati Uysal, Deliang Fan, and Rickard Ewetz. 2020. Handling stuck-at-fault defects using matrix transformation for robust inference of DNNs. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 39, 10 (2020), 2448–2460. DOI: https://doi.org/10.1109/TCAD.2019.2944582
- [57] Qilin Zheng, Zongwei Wang, Zishun Feng, Bonan Yan, Yimao Cai, Ru Huang, Yiran Chen, Chia-Lin Yang, and Hai Helen Li. 2020. Lattice: An ADC/DAC-less ReRAM-based processing-in-memory architecture for accelerating deep convolution neural networks. In *Proceedings of the 2020 57th ACM/IEEE Design Automation Conference*. IEEE, 1–6.

Received 25 April 2022; revised 27 September 2022; accepted 30 November 2022