



## OPEN ACCESS

## EDITED BY

Joon Young Kwak,  
Korea Institute of Science and Technology  
(KIST), Republic of Korea

## REVIEWED BY

Beomseok Kang,  
Georgia Institute of Technology, United States  
Tuo Shi,  
Zhejiang Lab, China

## \*CORRESPONDENCE

Rachmad Vidya Wicaksana Putra  
✉ rachmad.putra@tuwien.ac.at

## SPECIALTY SECTION

This article was submitted to  
Neuromorphic Engineering,  
a section of the journal  
Frontiers in Neuroscience

RECEIVED 05 February 2023

ACCEPTED 24 March 2023

PUBLISHED 12 April 2023

## CITATION

Putra RVW, Hanif MA and Shafique M (2023)  
RescueSNN: enabling reliable executions on  
spiking neural network accelerators under  
permanent faults. *Front. Neurosci.* 17:1159440.  
doi: 10.3389/fnins.2023.1159440

## COPYRIGHT

© 2023 Putra, Hanif and Shafique. This is an  
open-access article distributed under the terms  
of the [Creative Commons Attribution License  
\(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or reproduction  
in other forums is permitted, provided the  
original author(s) and the copyright owner(s)  
are credited and that the original publication in  
this journal is cited, in accordance with  
accepted academic practice. No use,  
distribution or reproduction is permitted which  
does not comply with these terms.

# RescueSNN: enabling reliable executions on spiking neural network accelerators under permanent faults

Rachmad Vidya Wicaksana Putra<sup>1\*</sup>, Muhammad Abdullah Hanif<sup>2</sup>  
and Muhammad Shafique<sup>2</sup>

<sup>1</sup>Embedded Computing Systems, Institute of Computer Engineering, Technische Universität Wien (TU Wien), Vienna, Austria, <sup>2</sup>eBrain Lab, Division of Engineering, New York University Abu Dhabi (NYUAD), Abu Dhabi, United Arab Emirates

To maximize the performance and energy efficiency of Spiking Neural Network (SNN) processing on resource-constrained embedded systems, specialized hardware accelerators/chips are employed. However, these SNN chips may suffer from permanent faults which can affect the functionality of weight memory and neuron behavior, thereby causing potentially significant accuracy degradation and system malfunctioning. Such permanent faults may come from manufacturing defects during the fabrication process, and/or from device/transistor damages (e.g., due to wear out) during the run-time operation. However, the impact of permanent faults in SNN chips and the respective mitigation techniques have not been thoroughly investigated yet. Toward this, we propose RescueSNN, a novel methodology to mitigate permanent faults in the compute engine of SNN chips *without requiring additional* retraining, thereby significantly cutting down the design time and retraining costs, while maintaining the throughput and quality. The key ideas of our RescueSNN methodology are (1) analyzing the characteristics of SNN under permanent faults; (2) leveraging this analysis to improve the SNN fault-tolerance through effective fault-aware mapping (FAM); and (3) devising lightweight hardware enhancements to support FAM. Our FAM technique leverages the fault map of SNN compute engine for (i) minimizing weight corruption when mapping weight bits on the faulty memory cells, and (ii) selectively employing faulty neurons that do not cause significant accuracy degradation to maintain accuracy and throughput, while considering the SNN operations and processing dataflow. The experimental results show that our RescueSNN improves accuracy by up to 80% while maintaining the throughput reduction below 25% in high fault rate (e.g., 0.5 of the potential fault locations), as compared to running SNNs on the faulty chip without mitigation. In this manner, the embedded systems that employ RescueSNN-enhanced chips can efficiently ensure reliable executions against permanent faults during their operational lifetime.

## KEYWORDS

spiking neural networks, accelerators, fault tolerance, manufacturing defects, reliability, resilience, permanent faults

## 1. Introduction

In recent years, SNNs have shown a potential for achieving high accuracy with ultra-low power/energy consumption due to their sparse spike-based operations (Putra and Shafique, 2020). Moreover, SNNs can perform unsupervised learning with unlabeled data using spike-timing-dependent plasticity (STDP), which is highly desired for real-world applications (e.g., autonomous agents like UAVs and robotics), especially due to two reasons: these systems are typically subjected to unforeseen scenarios (Putra and Shafique, 2021b, 2022a); and gathering unlabeled data is cheaper than labeled ones (Rathi et al., 2019). An SNN architecture supporting unsupervised learning is shown in Figure 1A. To maximize the performance and energy efficiency of SNN processing, specialized SNN accelerators/chips are employed (Painkras et al., 2013; Akopyan et al., 2015; Davies et al., 2018; Frenkel et al., 2019). However, these SNN chips may suffer from permanent faults, which can occur during: (1) *chip fabrication process* due to manufacturing defects, as fabricating an SNN chip with millions-to-billions of nano-scale transistors with 100% correct functionality is difficult, and even worsen due to the aggressive technology scaling (Hanif et al., 2018, 2021; Zhang et al., 2018); and (2) *run time operation* due to device/transistor wear out and damages, that are caused by Hot Carrier Injection (HCI), Bias Temperature Instability (BTI), electromigration, or Time Dependent Dielectric Breakdown (TDDB) (Radetzki et al., 2013; Werner et al., 2016; Hanif et al., 2018, 2021; Baloch et al., 2019; Mercier et al., 2020).

Permanent faults can affect the functionality of the compute engine of SNN accelerators/chips, including the local weight memory/registers and neurons, by corrupting the weight values and neuron behavior (i.e., membrane potential dynamics and spike generation), thereby degrading the accuracy, as shown in Figure 2. For instance, permanent faults can change the weight values through stuck-at 0 and 1, as shown in Figure 1B. Simply stopping the executions on faulty chips at run time will lead to a short operational lifetime, while discarding the faulty chips at design time will lead to low yield and increase the per-unit cost of the non-faulty chip. Therefore, *alternate low-cost solutions for mitigating permanent faults in the SNN compute engine<sup>1</sup> are required*. These solutions will prolong the operational lifetime of SNN chips. Moreover, such solutions also increase the applicability of wafer-scale chips for SNNs where embracing permanent faults is important to maintain the yield.

**Targeted Problem:** *How can we efficiently mitigate permanent faults in the SNN compute engine (i.e., the local weight registers and neurons) on the accuracy, thereby improving the SNN fault tolerance and maintaining the throughput.* The efficient solution to this problem will enable reliable SNN executions on faulty chips without the need for retraining for energy-constrained embedded systems, such as IoT-Edge devices and autonomous agents.

<sup>1</sup> For conciseness, we use "SNN compute engine" or "compute engine" interchangeably to represent the compute engine of an SNN accelerator/chip.

### 1.1. State-of-the-art and their limitations

Besides discarding the faulty chips, the standard VLSI fault tolerance techniques like Dual Modular Redundancy (DMR) (Vadlamani et al., 2010), Triple Modular Redundancy (TMR) (Lyons and Vanderkulk, 1962), and Error Correction Code (ECC) (Sze, 2000), may be used for mitigating permanent faults. However, they require extra (redundant) hardware and/or executions which incur huge area and energy overheads. State-of-the-art works have studied different design aspects for understanding faults in SNNs and devising mitigation techniques, as follows.

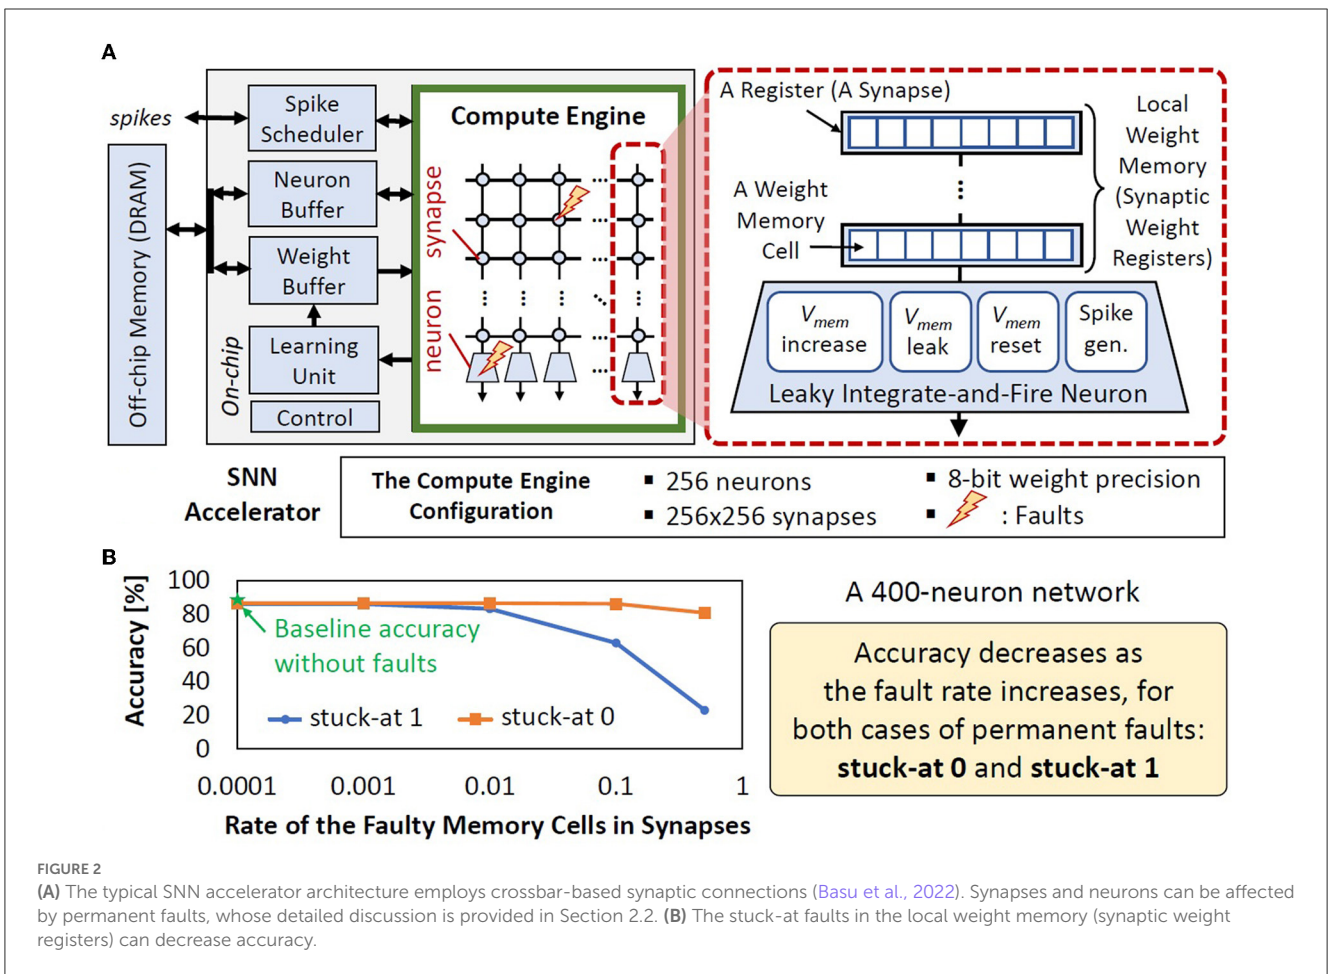
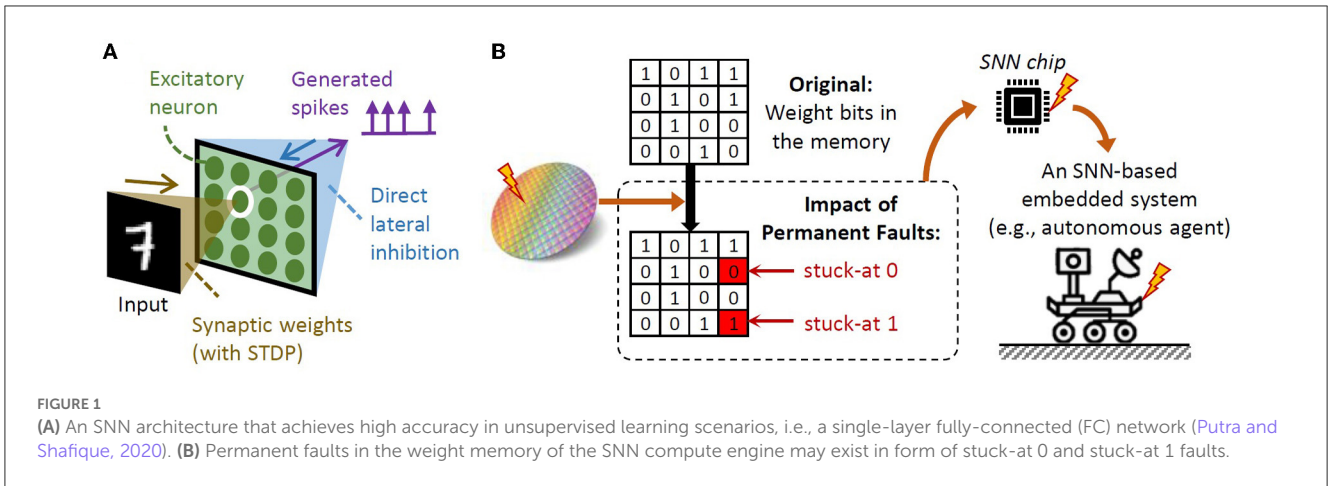
- **SNN fault modeling:** Possible faults that can affect an SNN have been identified in Vatajelu et al. (2019). In the analog domain, fault modeling for analog neuron hardware and its fault tolerance strategies have been investigated in El-Sayed et al. (2020) and Spyrou et al. (2021), which are out of the scope of this work as we target SNN implementation in the digital domain.
- **SNN fault tolerance:** Previous works studied the impact of faults on SNN weights without considering the underlying hardware architectures and processing dataflows (Schuman et al., 2020; Rastogi et al., 2021) and each discussing a specific fault, like bit-flip or synapse removal. Recent works devised mitigation techniques for faults in the weight memories of an SNN hardware (Putra et al., 2021a,b, 2022a), while work of Putra et al. (2022b) aimed at addressing transient faults.

In summary, the state-of-the-art works still focus on permanent fault modeling and injection only on the weight memories of SNN hardware. Hence, *the impact of permanent faults in the SNN compute engine (i.e., synapses and neurons) and the respective fault mitigation techniques, especially with a focus on avoiding retraining costs, are still unexplored*. To study the challenges of mitigating permanent faults, we perform an experimental case study in Section 1.2.

### 1.2. Case study and research challenges

In this case study, we consider an SNN accelerator architecture in Figure 2A. We assume all neurons are not faulty, and inject permanent faults (i.e., stuck-at 0 or 1) on the weight registers with random distribution and different rates of faulty memory cells to see the significance of faulty registers on accuracy. Details on the experimental setup are discussed in Section 4. From the experimental results in Figure 2B, we make the following key observations.

- Classification accuracy decreases as the rate of faulty memory cells increases for both stuck-at 0 and stuck-at 1 scenarios, thereby showing the negative impact of permanent faults in the synapses.
- In the stuck-at 0 case, the stored weight value will either stay the same or decrease from the original value. In the case of decreased weight value, the corresponding neuron



will require more stimulus (input spikes) to increase its membrane potential and reach the threshold potential for generating a spike, which represents recognition of a specific class. However, in an SNN model, multiple neurons may be responsible to recognize the same class. Therefore, if the neuron with faulty weight bits cannot recognize the input class, then other neurons may recognize it. As consequence, the accuracy degradation caused by stuck-at 0 in memory cells is relatively small and negligible in some cases.

- In the stuck-at 1 case, the stored weight value will either stay the same or increase from the original value. In the case of increased weight value, the corresponding neuron will require less stimulus (input spikes) to increase its membrane potential and reach the threshold potential for generating a spike, which represents recognition of a specific class. Therefore, this neuron may become more active to generate more spikes for any input classes, which leads to more misclassification. As consequence, the accuracy degradation caused by stuck-at 1 in

memory cells is more significant/noticeable than the stuck-at 0 case.

- Combinations of fault types and fault rates lead to different accuracy, which represents different fault patterns in real-world chips, rendering it unpredictable at design time.

Based on these observations, we outline the following research challenges to devise an efficient solution for the targeted problem.

- *The mitigation technique should not employ retraining*, as retraining needs huge compute/memory costs, processing time, and a training dataset that may not be available in certain cases due to the restriction policies. Moreover, retraining is not a scalable approach considering the large number of fabricated chips, as it needs to consider a unique fault map from each chip thereby retraining per chip. Note, the fault map information can be obtained through the standard wafer/chip test procedure after fabrication, hence this test does not introduce new cost and only incurs a typical cost for chip test (Xu et al., 2020; Fan et al., 2022).
- *The mitigation should have minimal performance/energy overhead at run time* as compared to that of the baseline design without fault mitigation technique, thereby making it applicable for energy-constrained embedded systems.
- *The technique should not avoid the use of faulty SNN components (i.e., synapses and neurons)*, as it means omitting the entire computations in the respective columns of the SNN compute engine, which leads to throughput reduction.
- *SNNs require a specialized permanent fault mitigation technique* as compared to other neural network computation models (e.g., deep neural networks), since SNNs have different operations and dataflows.

### 1.3. Our novel contributions

To address the above challenges, we propose RescueSNN, a novel methodology that enables reliable executions on SNN accelerators under permanent faults. To the best of our knowledge, this work is the first effort that mitigates permanent faults in the SNN accelerators/chips. Following are the key steps of the RescueSNN methodology (the overview is shown in Figure 3).

- **Analyzing the SNN fault tolerance** to understand the impact of faulty components (i.e., synapses and neurons) on accuracy considering the given fault rates.
- **Employing the fault-aware mapping (FAM) techniques** to safely map SNN weights and neuron operations on faulty compute engine, thereby maintaining accuracy and throughput. Our FAM techniques leverage the fault map of the compute engine to perform the following key mechanisms.

1. Mapping the significant weight bits on the non-faulty memory cells of the synapses (weight registers) to minimally pollute/change the weight values.

2. Selectively employing faulty neurons that do not cause significant accuracy degradation at inference, based on the behavior of their membrane potential dynamics and spike generation.

- **Devising simple hardware enhancements** to enable efficient FAM techniques. Our enhancements shuffle the weight bits from the synapses by employing simple combinational logic units, such as multiplexers, so that these weight bits can be used for SNN computations.

**Key Results:** We evaluate our RescueSNN methodology using Python-based simulations (Hazan et al., 2018) on a multi-GPU machine. The experimental results show that our RescueSNN improves the SNN accuracy without retraining by up to 80% and 47% on the MNIST and the Fashion MNIST respectively, from the SNN processing without fault mitigation.

## 2. Backgrounds

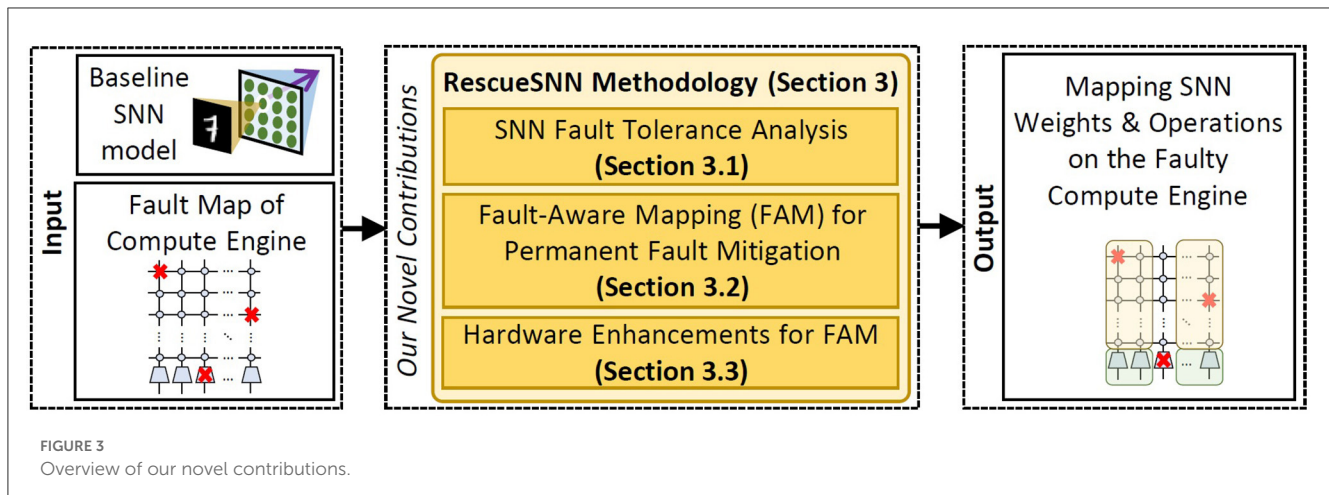
### 2.1. Spiking neural networks

**Overview:** SNNs are the brain-inspired computational models that employ action potentials (spikes) to encode the information (Maass, 1997; Putra and Shafique, 2022b). In an SNN model, neurons and synapses are connected in a specific architecture (Mozafari et al., 2019; Tavanaei et al., 2019). In this work, we consider a fully-connected (FC)-based network architecture as shown in Figure 1A, since it has demonstrated a high accuracy when employing bio-plausible learning rules, e.g., the spike-timing-dependent plasticity (STDP). This network architecture connects each data of the input (e.g., a pixel of an image) to all excitatory neurons for generating spikes. Each spike is used for inhibiting the excitatory neurons except the one from which the spike is generated. Such bio-plausible learning rules offer efficient learning mechanisms as they perform local learning in each synapse by leveraging spike events without any global loss function, thereby enabling unsupervised and online learning capabilities, which are especially beneficial for autonomous agents (Pfeiffer and Pfeil, 2018; Putra and Shafique, 2022a). In this work, we employ the pair-based weight-dependent STDP and bound each weight value ( $wgh$ ) within the defined range, i.e.,  $wgh = [0, 1]$ , because this approach has been widely used by previous works (Diehl and Cook, 2015; Srinivasan et al., 2017; Hazan et al., 2019). Here, the pair-wise weight-dependent STDP learning rule is used, as it defines the maximum allowed weights, which is suitable for fixed-point format (see Eq. 1).

$$\Delta wgh = \begin{cases} -\eta_{pre} \cdot x_{post} \cdot wgh^\mu & \text{on presynaptic spike} \\ \eta_{post} \cdot x_{pre} \cdot (wgh_m - wgh)^\mu & \text{on postsynaptic spike} \end{cases} \quad (1)$$

$\Delta wgh$  denotes the weight update,  $\eta_{pre}$  and  $\eta_{post}$  denote the learning rate for pre- and post-synaptic spike, while  $x_{pre}$  and  $x_{post}$  denote the pre- and post-synaptic traces, respectively.  $wgh_m$  denotes the maximum allowed weight,  $wgh$  denotes the current weight, and  $\mu$  denotes the weight dependence factor.





We consider the Leaky Integrate-and-Fire (LIF) neuron model, as it has low computational complexity with high bio-plausible behavior (Izhikevich, 2004). Meanwhile, the synapse is modeled by weight value ( $wgh$ ) which represents the strength of the synaptic connection between the corresponding neurons. Furthermore, an SNN model typically employs a specific spike coding scheme to encode/decode data information into/from spikes. In this work, we consider the rate coding scheme which employs the frequency of spikes to proportionally represent the data, i.e., a higher data value is represented by a higher number of spikes. The reason is that, the rate coding scheme achieves high accuracy using bio-plausible STDP-based learning rules, which perform efficient learning mechanisms under unsupervised settings, thereby enabling energy-efficient and smart computing systems (Diehl and Cook, 2015; Rathi et al., 2019; Putra and Shafique, 2020).

**SNN Accelerator Architecture:** We consider the typical SNN accelerator and compute engine architectures shown in Figures 2A, 4 respectively, which are adapted from the design in Frenkel et al. (2019). We focus on the compute engine, as it is responsible for generating spikes which determine the accuracy. The compute engine has a synapse crossbar, and each synapse has a register that stores a weight value. We use 8-bit weight precision as it has a good trade-off for accuracy-memory (Putra and Shafique, 2020, 2021a). Each bit is stored in a weight memory cell. To optimize the chip area, each synapse adds its weight with an accumulated value from the previous synapse in the same column, and stores the cumulative value in a 32-bit register. In this manner, only a single connection is required between a neuron and the connected synapses. Furthermore, each column of the compute engine has a single LIF neuron. If a spike comes to a neuron, the membrane potential ( $V_{mem}$ ) increases by the respective weight ( $wgh$ ), otherwise, the  $V_{mem}$  decreases/leaks. If the  $V_{mem}$  reaches the threshold potential ( $V_{th}$ ), a spike is generated, and then  $V_{mem}$  goes back to the reset potential ( $V_{reset}$ ). Hence, a LIF neuron has four main operations: (1)  $V_{mem}$  increase, (2)  $V_{mem}$  leak, (3)  $V_{mem}$  reset, and (4) spike generation, as shown in Figure 2A. Neuron operations (1)–(3) define the membrane potential dynamics.

## 2.2. Permanent fault modeling

### 2.2.1. Overview

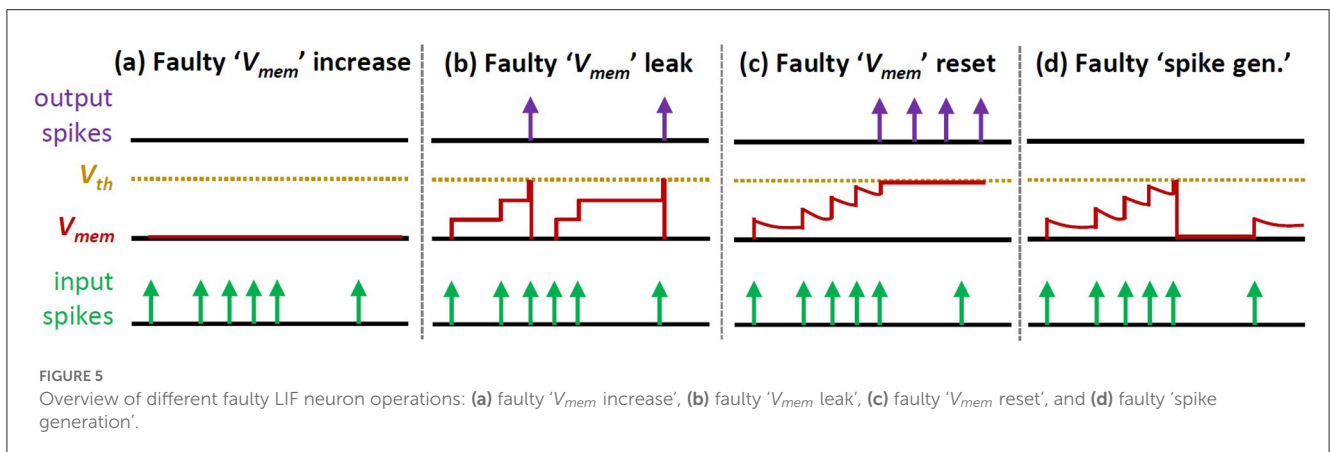
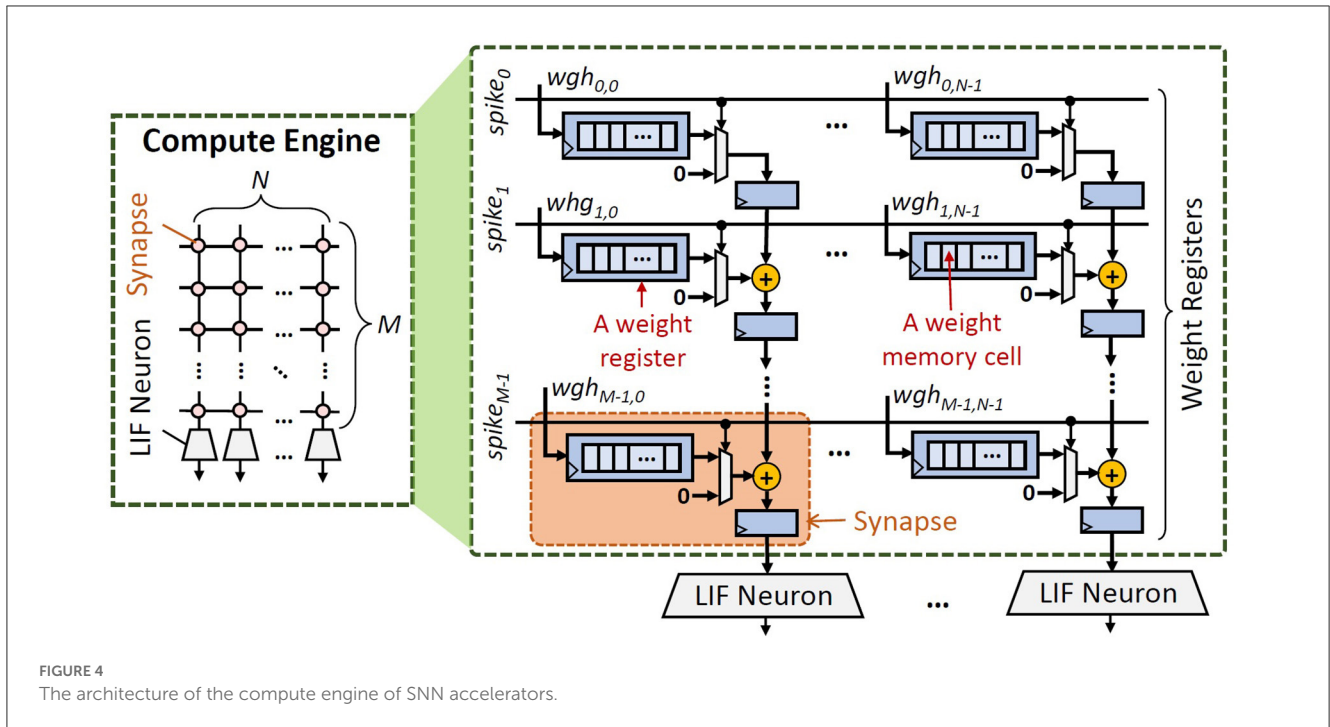
An SNN compute engine consists of two main components, i.e., synapses and neurons, which have different hardware circuitry. Therefore, we need to define a fault model for each component to achieve fast design space exploration.

1. **Synapses:** Each synapse hardware uses a register to store a weight value. Therefore, each permanent fault in a synapse can affect a single weight bit in the form of either a stuck-at 0 or a stuck-at 1 fault.
2. **Neurons:** Each neuron hardware depends on the neuron model to facilitate its operations. Therefore, permanent faults can manifest in different forms depending upon the type of operation being executed on the (digital) neuron hardware, as discussed in the following (see an overview in Figure 5).

- *Faults in the ‘ $V_{mem}$  increase’ operation* make the neuron unable to increase its membrane potential. As a result, this neuron cannot generate any spikes (i.e., a dormant neuron).
- *Faults in the ‘ $V_{mem}$  leak’ operation* make the neuron unable to decrease its membrane potential. Hence, this neuron acts like the Integrate-and-Fire (IF) neuron model.
- *Faults in the ‘ $V_{mem}$  reset’ operation* make the neuron unable to reset its membrane potential. As a result, this neuron will continuously generate spikes.
- *Faults in the ‘spike generation’* make the neuron unable to generate spikes (i.e., dormant neuron).

### 2.2.2. Fault generation and distribution

Previous studies have shown that permanent faults occur in random locations of a chip, hence leading to a certain fault map (Stanisavljević et al., 2011; Radetzki et al., 2013; Werner et al., 2016; Zhang et al., 2018; Mercier et al., 2020). Following are the key steps to generate and distribute permanent faults on the SNN compute engine, as shown in Figure 6.



1. We consider a weight memory cell and a neuron operation as the potential fault locations.
2. We generate permanent faults based on the given fault rate and distribute them randomly across the potential fault locations. The fault rate represents the ratio between the total number of faulty weight memory cells and neuron operations to the total number of potential fault locations (i.e., the total number of weight memory cells and neuron operations).
3. If a fault occurs in a local weight memory cell, then we randomly select the type of stuck-at fault (i.e., either stuck-at 0 or stuck-at 1). Meanwhile, if a fault occurs in a neuron operation, then we randomly select the type of permanent fault operation.

### 3. RescueSNN methodology

The overview of our RescueSNN methodology is shown in Figure 7, and its key steps are discussed in the following subsections.

#### 3.1. SNN fault tolerance analysis under permanent faults

SNN fault tolerance analysis is important to understand how a given SNN model will behave considering a specific operating condition (e.g., a combination of certain fault rate, type of stuck-at fault, architecture of the compute

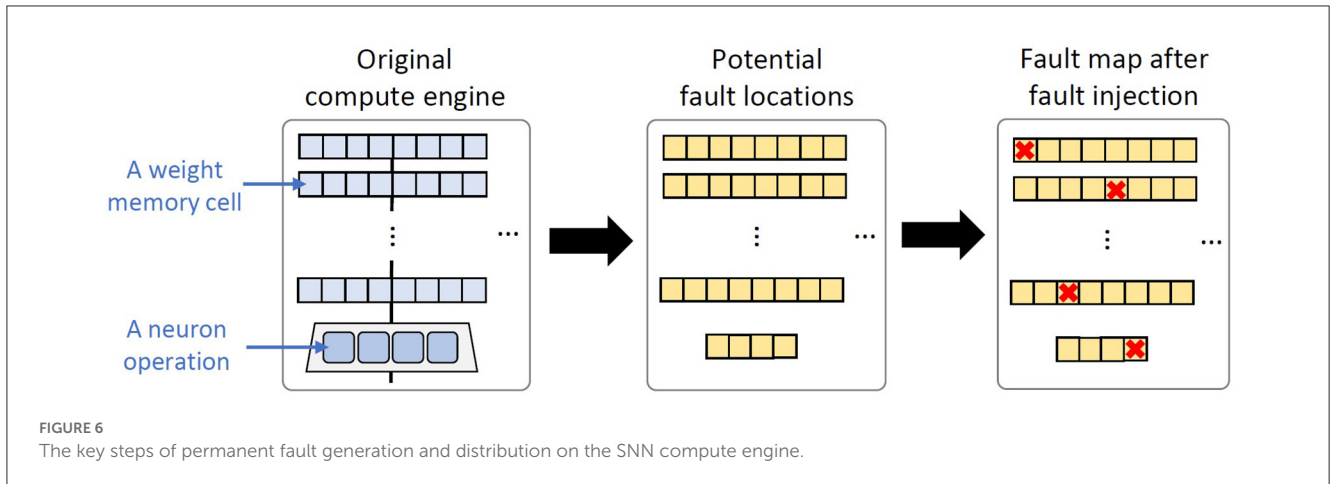


FIGURE 6 The key steps of permanent fault generation and distribution on the SNN compute engine.

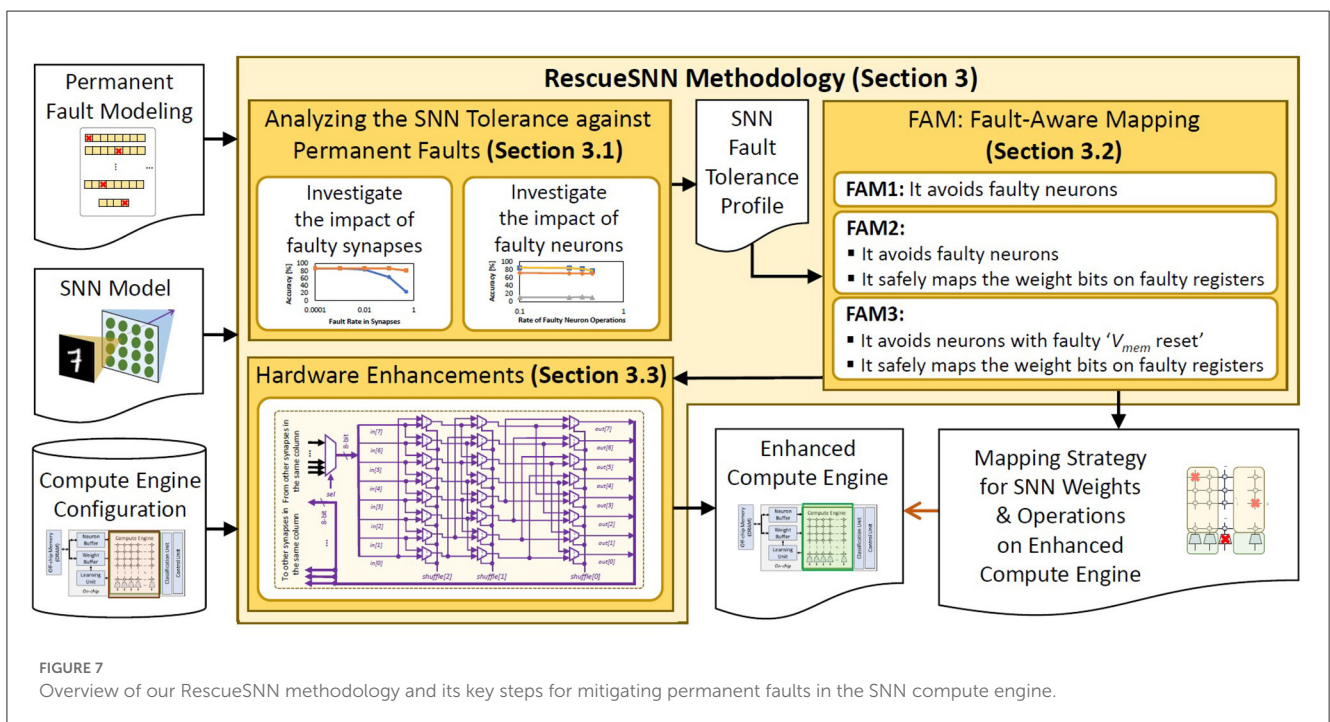


FIGURE 7 Overview of our RescueSNN methodology and its key steps for mitigating permanent faults in the SNN compute engine.

engine, etc.). This analysis provides information which can be leveraged for devising an efficient fault mitigation technique. Therefore, our RescueSNN methodology investigates the interaction between the faulty components (i.e., synapses and neurons) and the obtained accuracy. To do this, we perform the following experimental case studies, while considering a 400-neuron network with fully-connected architecture like in Figure 1A.

- We study the accuracy under faulty weight registers, by injecting a specific stuck-at fault (i.e., either stuck-at 0 or stuck-at 1) in the weight registers, while considering fault-free neurons. Experimental results are shown in Figure 2B. We observe that both stuck-at 0 and stuck-at 1 faults can degrade accuracy. Therefore, the mitigation technique should address both stuck-at faults.

- We study the accuracy under faulty neuron operations, by injecting faults on the neuron operations, while considering fault-free weight registers. Experimental results are shown in Figure 8, from which we make the following observations.

1. Faulty 'spike generation', ' $V_{mem}$  increase', and ' $V_{mem}$  leak' operations have tolerable accuracy, since their faulty behavior does not dominate the spiking activity, and/or the function of the corresponding faulty neurons for classification may be substituted by other neurons that recognize the same class. Therefore, these neurons can still be used for SNN processing.
2. Faulty ' $V_{mem}$  reset' operations cause significant accuracy degradation, as these operations make the corresponding neurons dominate classification. Therefore, these neurons should not be used for SNN processing.

Note, complex SNN models with multiple layers and different computational architectures (e.g., convolutional and fully-connected) may have different observation results as compared to results in Figure 8. However, previous work has observed similar trends to our study, i.e., neurons with faulty ' $V_{mem}$  reset' operations continuously generate spikes (so-called *saturated neurons*) and cause the most significant accuracy degradation than other types of faulty neuron operations (Spyrou et al., 2021). It also identified that, saturated neurons affect classification accuracy at any layer of SNN models, as these faulty neurons always dominate the classification activity which results in a significant accuracy degradation. This finding is consistent with the insights provided by our study. However, it is still challenging to achieve high accuracy when employing STDP-based learning on complex SNN models (Rathi et al., 2023), thereby hindering their applicability for diverse applications, such as systems with online training and unsupervised learning requirements (e.g., autonomous mobile agents). Therefore, in this work, we consider the FC-based SNNs shown in Figure 1A to enable multiple advantages, such as high accuracy, unsupervised learning capabilities, and efficient online training.

### 3.2. The proposed fault-aware mapping

Permanent faults in SNN chips can be identified at the design time and at the run time. The post-fabrication test can be employed to find a set of locations of faults (due to manufacturing defects) in the SNN compute engine (i.e., fault map) at the design time (Zhang et al., 2018; Putra et al., 2021a). Meanwhile, the online test strategy like the Built-In Self-Test (BIST) techniques can be employed to obtain the fault map (due to device wear out or physical damages) at the run time (Baloch et al., 2019; Wang et al., 2019; Mercier et al., 2022). Our RescueSNN methodology leverages this fault map to safely map the SNN weights and operations on the compute engine, thereby minimizing the negative impact of permanent faults. To do this, the RescueSNN employs Fault-Aware Mapping (FAM) techniques that mitigate the faults in synapses and neurons through the following key mechanisms.

1. **FAM for Synapses:** The significant weight bits are placed on the non-faulty weight memory cells, while the insignificant bits are placed on the faulty ones, by performing a simple bit-shuffling technique. The significance of weight bits can be identified by experiments that observe the accuracy after modifying a specific bit (Putra et al., 2021a). In general, previous studies have observed that the significance of a weight bit is proportional to its bit location. For instance, in 8-bit fixed-point precision, bit-7 has the highest significance than other bits. Furthermore, a synapse may have a single faulty bit or multiple faulty bits. Therefore, we propose a mapping strategy that can address both cases using the following steps (see an overview in Figure 9).

- We identify the faulty weight bits (e.g., through the post-fabrication testing) to obtain information of the fault map and fault rate in each synapse hardware.
- We identify the maximum fault rate in each synapse hardware for safely storing a weight. In this work, we

consider a maximum of 2 faulty bits from an 8-bit weight, based on the fault rates that offer tolerable accuracy from analysis in Section 3.1.

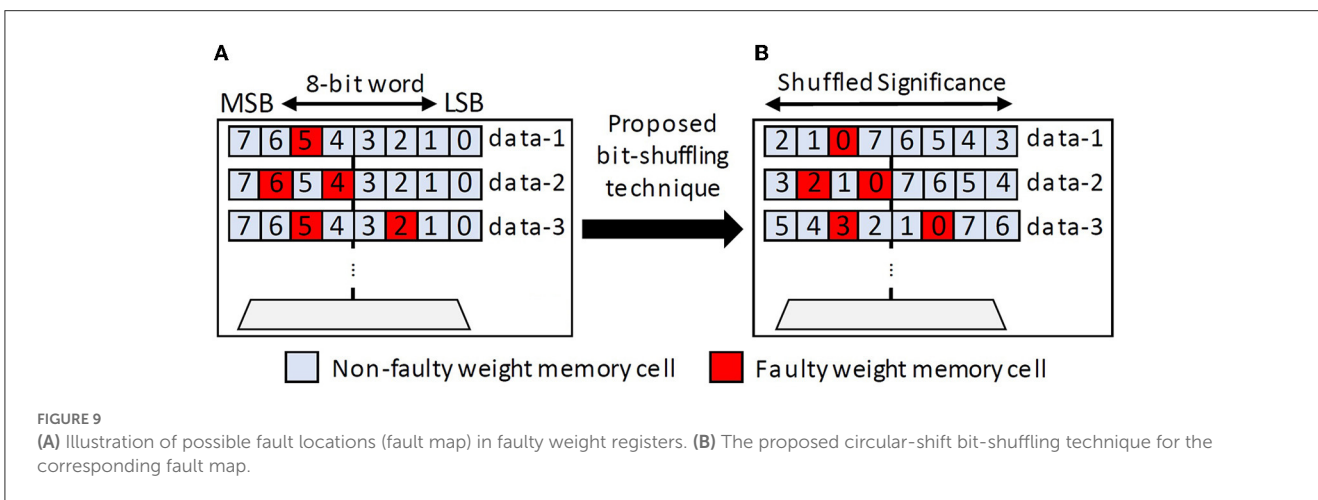
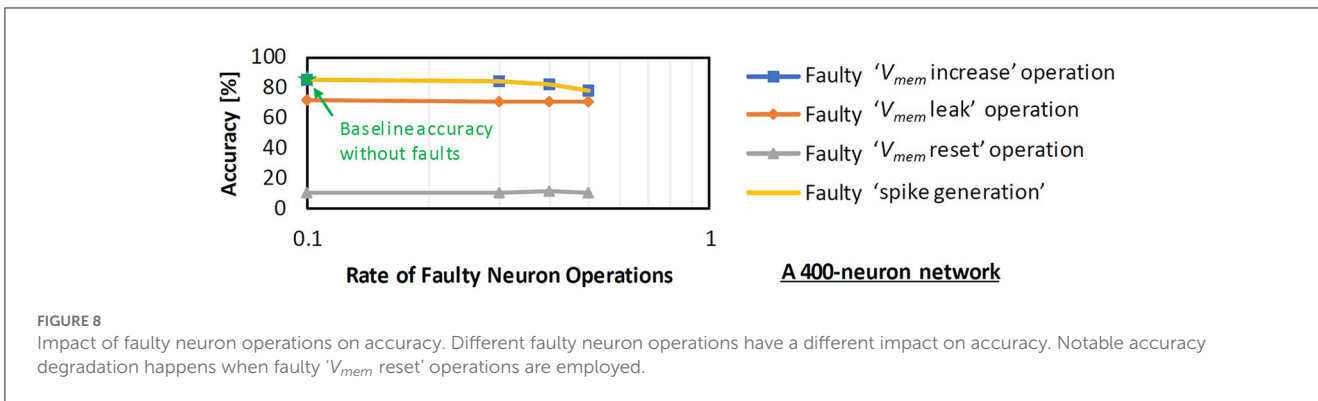
- We identify the segment in each synapse with the highest number of subsequent non-faulty memory cells. This information is leveraged for maximizing the possibility of storing the significant bits in the non-faulty cells. Hence, we also examine the corner case (i.e., the right-most and left-most cells) as possible subsequent non-faulty memory cells; see the third row of Figure 9B with data-3.
  - We perform a circular-shift technique for each data word to efficiently implement bit-shuffling.
2. **FAM for Neurons:** The use of neurons should be avoided if they have faulty ' $V_{mem}$  reset' operations, as these faulty operations cause significant accuracy degradation. Meanwhile, neurons with other types of faults can still be used for SNN processing, as their faulty behavior does not dominate the spiking activity. Different SNN operations that aim at recognizing the same input class are mapped on both the faulty and fault-free neurons for maintaining throughput, while compensating the loss from the faulty ' $V_{mem}$  increase', ' $V_{mem}$  leak', and 'spike generation' operations.

Afterward, we leverage these mechanisms for devising three mapping strategies, as the variants of our FAM technique (i.e., FAM1, FAM2, and FAM3), which provide trade-offs between accuracy and mapping complexity, as discussed in the following.

- **FAM1:** It avoids mapping the SNN weights and operations on the columns of compute engine that have faulty neurons, as shown in Figure 10a, as faulty neurons can reduce the accuracy more than faulty registers, especially in the case of faulty ' $V_{mem}$  reset'. However, FAM1 does not mitigate the negative impact of faults in the registers, hence the accuracy improvement is sub-optimal. The benefit of FAM1 is due to its simple mechanism which enables a low-complexity control mechanism.
- **FAM2:** It maps the SNN weights and operations on the columns of compute engine that have fault-free neurons (just like FAM1) and employs a bit-shuffling technique to map the significant weight bits on the non-faulty memory cells, as shown in Figure 10b. Therefore, FAM2 can improve the SNN fault tolerance at the cost of a more complex control mechanism than FAM1.
- **FAM3:** It selectively maps the SNN weights and operations on the columns of compute engine that do not have faulty ' $V_{mem}$  reset' operations, as well as maps the significant weight bits on the non-faulty memory cells using a bit-shuffling technique, as shown in Figure 10c. Therefore, FAM3 can enhance the SNN fault tolerance as compared to FAM1 at the cost of a more complex control mechanism, and can improve the throughput as compared to FAM1 and FAM2.

Information regarding how to map the SNN weights and operations on the compute engine is provided through software program (e.g., firmware), thereby enabling the applicability and flexibility of the proposed FAM technique (e.g., FAM1, FAM2, or





FAM3) for different possible fault maps on the compute engine. The meta data of this information is stored in the on-chip buffer, which can be accessed for operations in the compute engine.

### 3.3. Our hardware enhancements for FAM

Our FAM2 and FAM3 strategies may make the weight bits stored in a shuffled form. Therefore, an additional mechanism is required for converting these weight bits into the original order, so that they can be used for SNN executions. Toward this, we propose *lightweight hardware enhancements to support the re-shuffling mechanism to undo the data transformation, i.e., through a simple 8-bit barrel shifter*. The key idea is to re-shuffle the order of output wires from each synapse into the original order, so that the corresponding weights can be used directly for neuron operations. To optimize the overheads (e.g., area), we share a *hardware enhancement block* (HEB) with all synapses in the same column of compute engine, and different synapses will access the enhancement block at a different time; see Figure 11. In this manner, the number of HEBs is equal to the number of columns in the SNN compute engine. Furthermore, to control the functionality of HEBs, we employ an *enhancement control unit* (ECU). This ECU stores the bit-shifting information and uses it for controlling the barrel shifter in HEBs. For each column of the compute engine, the ECU employs (1) a dedicated selector signal *sel* to determine

which weight should be processed in the HEB at a time, and (2) a set of registers that stores bit-shifting information *shuffle[2:0]* for all weights in the same column.

## 4. Evaluation methodology

For evaluating the RescueSNN methodology, we employ the experimental setup shown in Figure 12. We use the fully-connected network shown in Figure 1A with a different number of neurons, to evaluate the generality of our RescueSNN methodology. For conciseness, we represent a network with *i*-number of neurons as *N<sub>i</sub>*. We use the MNIST and Fashion MNIST as the workloads, and adopt the same test conditions as used widely by the SNN community (Diehl and Cook, 2015). For comparison, we consider the SNN without fault mitigation as the baseline.

**Fault Generation and Injection:** Permanent faults are generated based on the fault modeling in Section 2.2 of the revised manuscript. To do this, we first generate binary values (i.e., 0 and 1) based on the given fault rate while considering the potential fault locations (shown in Figure 13). Here, “0” represents a non-faulty memory cell in synapses or a non-faulty operation in neurons; while “1” represents a faulty memory cell in synapses or a faulty operation in neurons. These binary values are then randomly distributed into an array that represents the potential fault locations, so that each value corresponds to a specific weight memory cell or a specific neuron operation. Figure 13 shows the

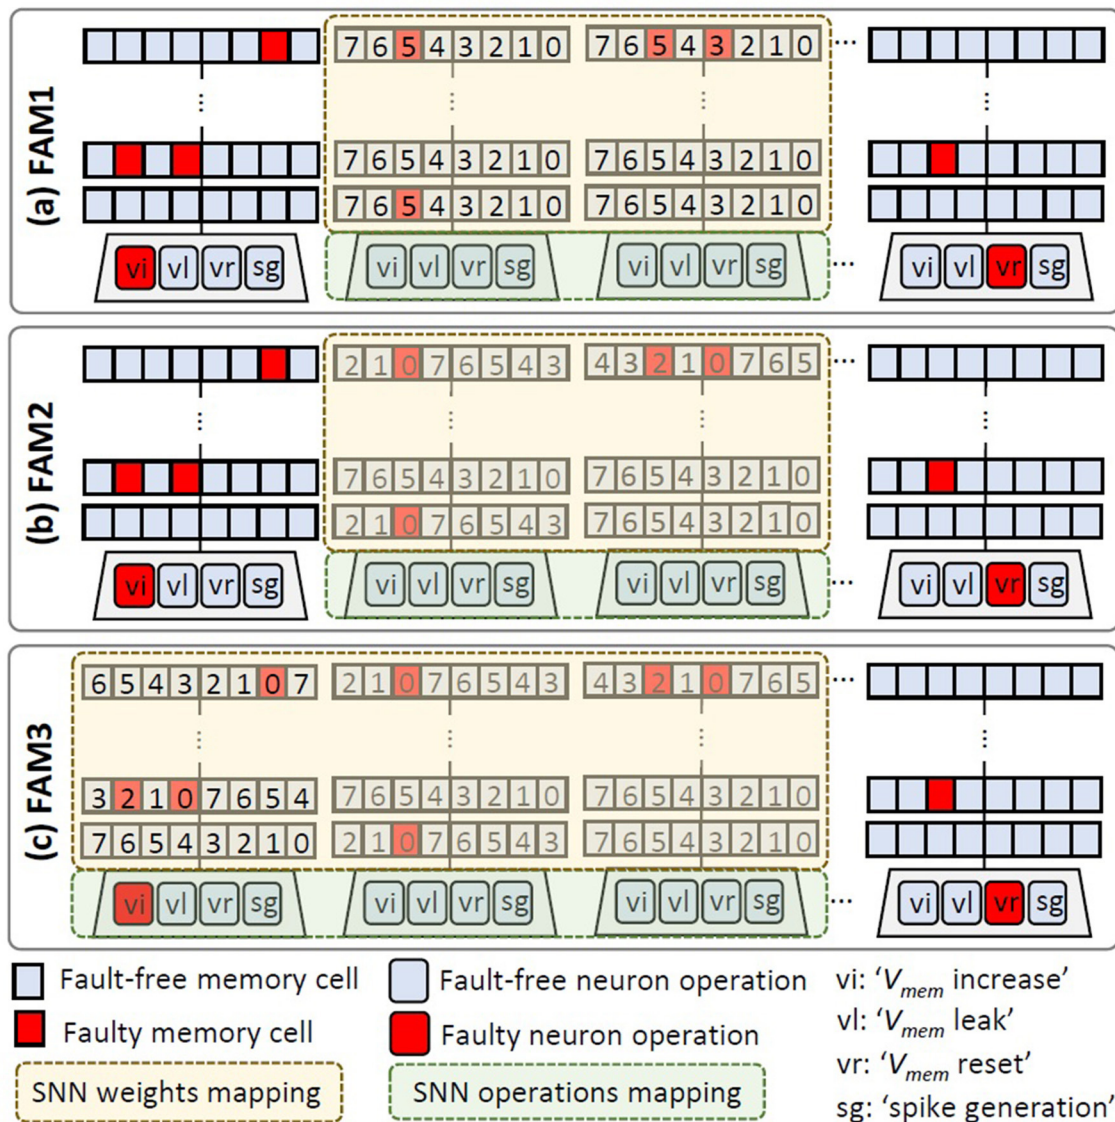


FIGURE 10  
Our FAM strategies: (a) FAM1, (b) FAM2, and (c) FAM3.

potential locations/components that can be affected by permanent faults to cause faulty memory cells as well as faulty ' $V_{mem}$  increase', ' $V_{mem}$  leak', ' $V_{mem}$  reset', and 'spike generation' operations. For each fault in the weight memory cells (synapses), we randomly determine the type of fault, i.e., either stuck-at 0 or stuck-at 1. In stuck-at 0 case, value 0 is injected to the corresponding memory cell; while in stuck-at 1 case, value 1 is injected. Meanwhile, each fault in neurons corresponds to either faulty ' $V_{mem}$  increase', ' $V_{mem}$  leak', ' $V_{mem}$  reset', or 'spike generation' operation (as shown in Figure 5). Each faulty behavior in the corresponding neuron is realized through different approaches as described in the following.

- **Faulty ' $V_{mem}$  increase' operation:** It is mainly caused by faulty addition in the ' $V_{mem}$  increase' part, hence  $V_{mem}$  is not increased despite there are incoming spikes.

- **Faulty ' $V_{mem}$  leak' operation:** It is mainly caused by faulty subtraction in ' $V_{mem}$  leak' part, hence  $V_{mem}$  is not decreased despite there are no incoming spikes.
- **Faulty ' $V_{mem}$  reset' operation:** It is mainly caused by faulty comparison in ' $V_{mem}$  reset' part, hence the spike generator is activated to continuously generate spikes.
- **Faulty 'spike generation':** It is mainly caused by faulty multiplexing in 'spike generation' part, hence the spike generator is always deactivated and no output spikes are produced.

**Accuracy Evaluation:** We use the Python-based simulations (Hazan et al., 2018), which run on Nvidia RTX 2080 Ti GPUs, while considering SNN accelerator architecture shown in Figures 2A, 4.

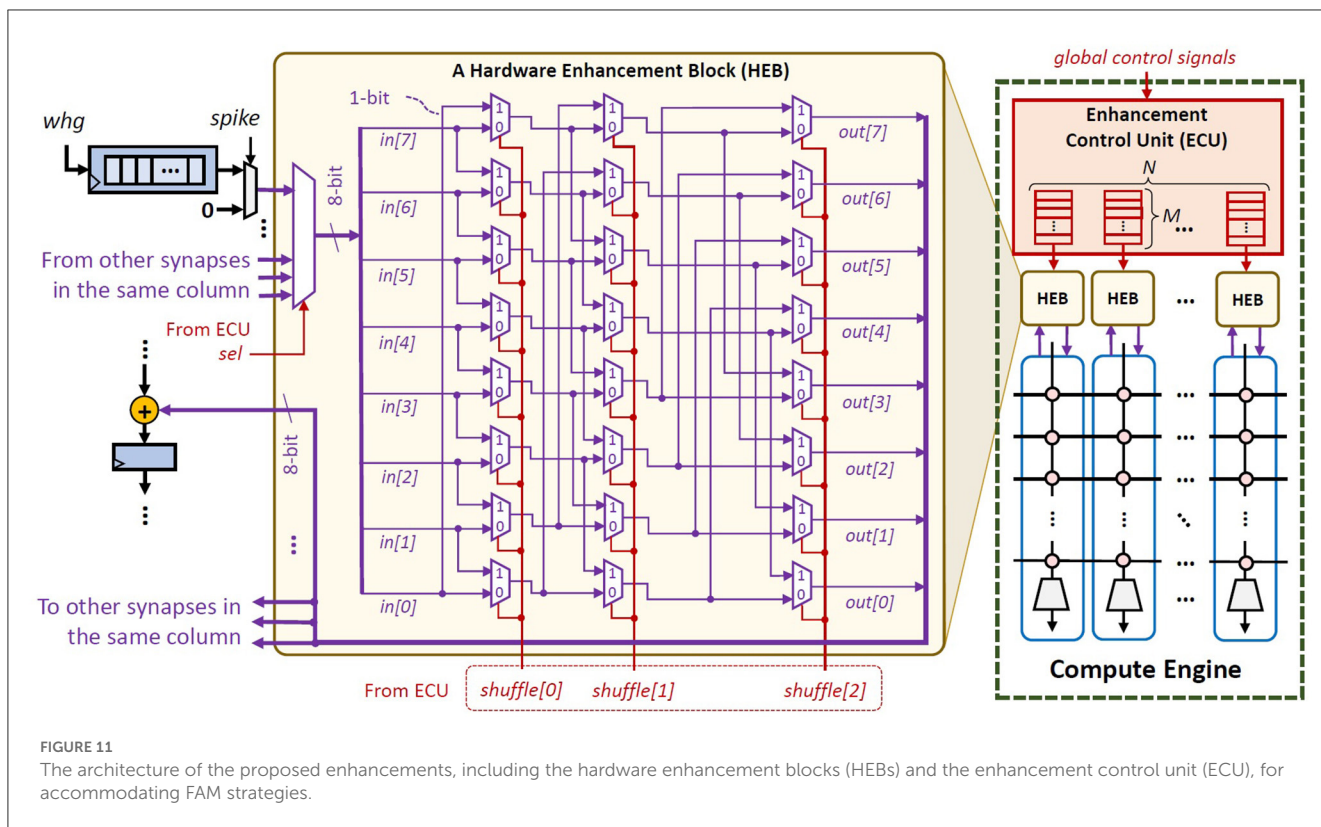


FIGURE 11 The architecture of the proposed enhancements, including the hardware enhancement blocks (HEBs) and the enhancement control unit (ECU), for accommodating FAM strategies.

**Hardware Evaluation:** We evaluate the area, energy consumption, and throughput of both the original compute engine (without enhancements) and the enhanced compute engine using our RescueSNN methodology. To do this, we design RTL codes for both (original and enhanced) compute engines, then synthesize them using the Cadence Genus tool considering a 65nm CMOS technology to obtain their area, power consumption, and timing (i.e., a clock cycle latency for SNN processing on the compute engine). Afterward, we calculate the required number of cycles and computation latency for processing an input sample (i.e., *latency-per-sample*), considering the timing from synthesis and the mapping strategy on active synapses and neurons in the compute engine. Then, we estimate the throughput by computing the number of samples that can be processed within 1 s of SNN inference based on the information of latency-per-sample. Furthermore, we also estimate the energy consumption by leveraging the information of power consumption from synthesis and latency-per-sample for SNN inference. The estimation of throughput and energy consumption is also performed using the Python-based simulation framework (Hazan et al., 2018).

## 5. Results and discussion

We evaluate different design aspects including accuracy, throughput, energy consumption, and area as discussed in the following.

### 5.1. Maintaining accuracy

Figure 14 presents the experimental results for the accuracy of different fault mitigation techniques, i.e., the baseline and our FAM-based strategies including FAM1, FAM2, and FAM3. We observe that the baseline suffers from a significant accuracy degradation as shown by ❶, because it does not mitigate faults in synapses and neurons, thereby leading to unreliable SNN executions. The significant accuracy degradation is mainly due to the fault model for faulty ' $V_{mem}$  reset' operation that makes the corresponding neuron generate spikes continuously once its membrane potential  $V_{mem}$  reaches the threshold potential  $V_{th}$ , thereby dominating the classification activity and leading to high misclassification. We also observe that FAM1 significantly improves the SNN fault tolerance as compared to the baseline, because FAM1 avoids the use of faulty neurons, especially for faulty ' $V_{mem}$  reset' operations, as shown by ❷. Our FAM2 improves the SNN fault tolerance even more as compared to FAM1, since FAM2 also mitigates faults in the weight registers in addition to avoiding the use of faulty neurons, as shown by ❸. Meanwhile, our FAM3 also significantly improves the SNN fault tolerance from baseline and FAM1, and obtains comparable accuracy to FAM2, since FAM3 mitigates faults in weight registers and selectively uses faulty neurons. It achieves up to 80% accuracy improvement compared to the baseline on the MNIST dataset, as shown by ❹. We also observe that the same reasons are also applicable to different workloads, thereby leading the accuracy profiles for Fashion MNIST to have similar trends to the accuracy profiles for MNIST. These results show that *our FAM strategies (FAM1, FAM2, and FAM3) are effective for mitigating permanent*



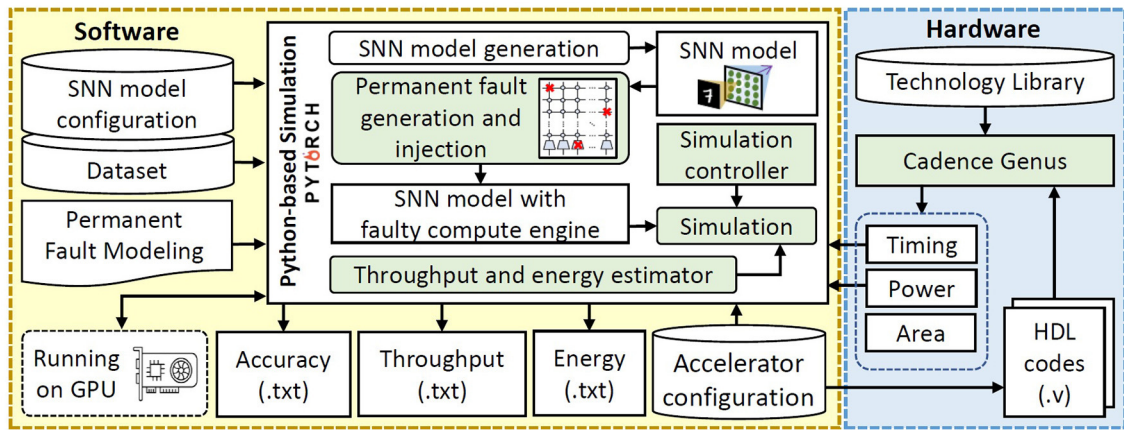


FIGURE 12 Overview of the experimental setup and tools flow.

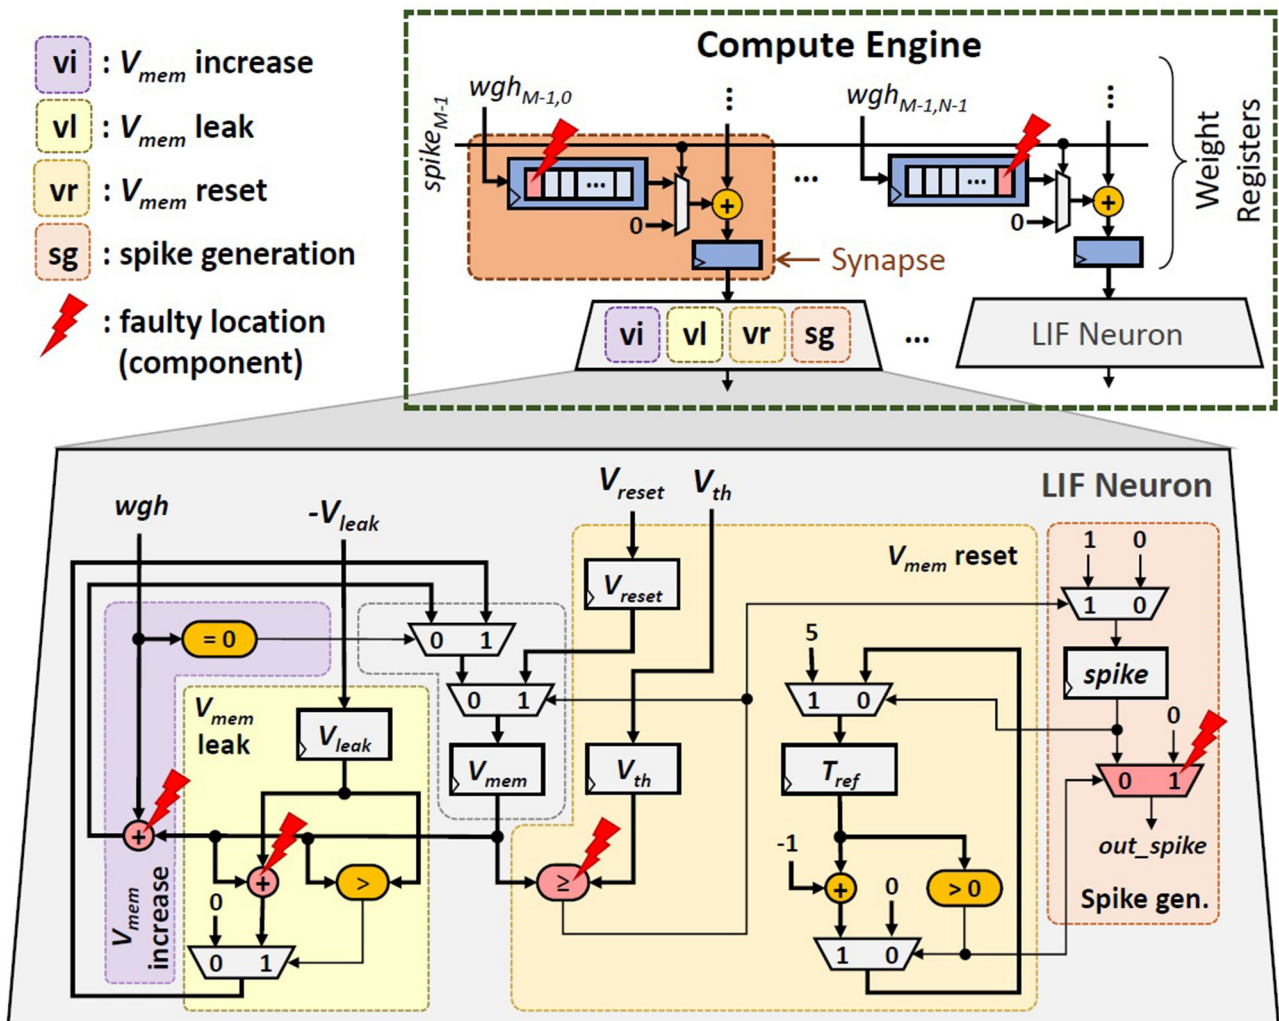


FIGURE 13 The potential locations/components that can be affected by permanent faults to cause faulty memory cells as well as faulty ' $V_{mem}$  increase', ' $V_{mem}$  leak', ' $V_{mem}$  reset', and 'spike generation' operations.



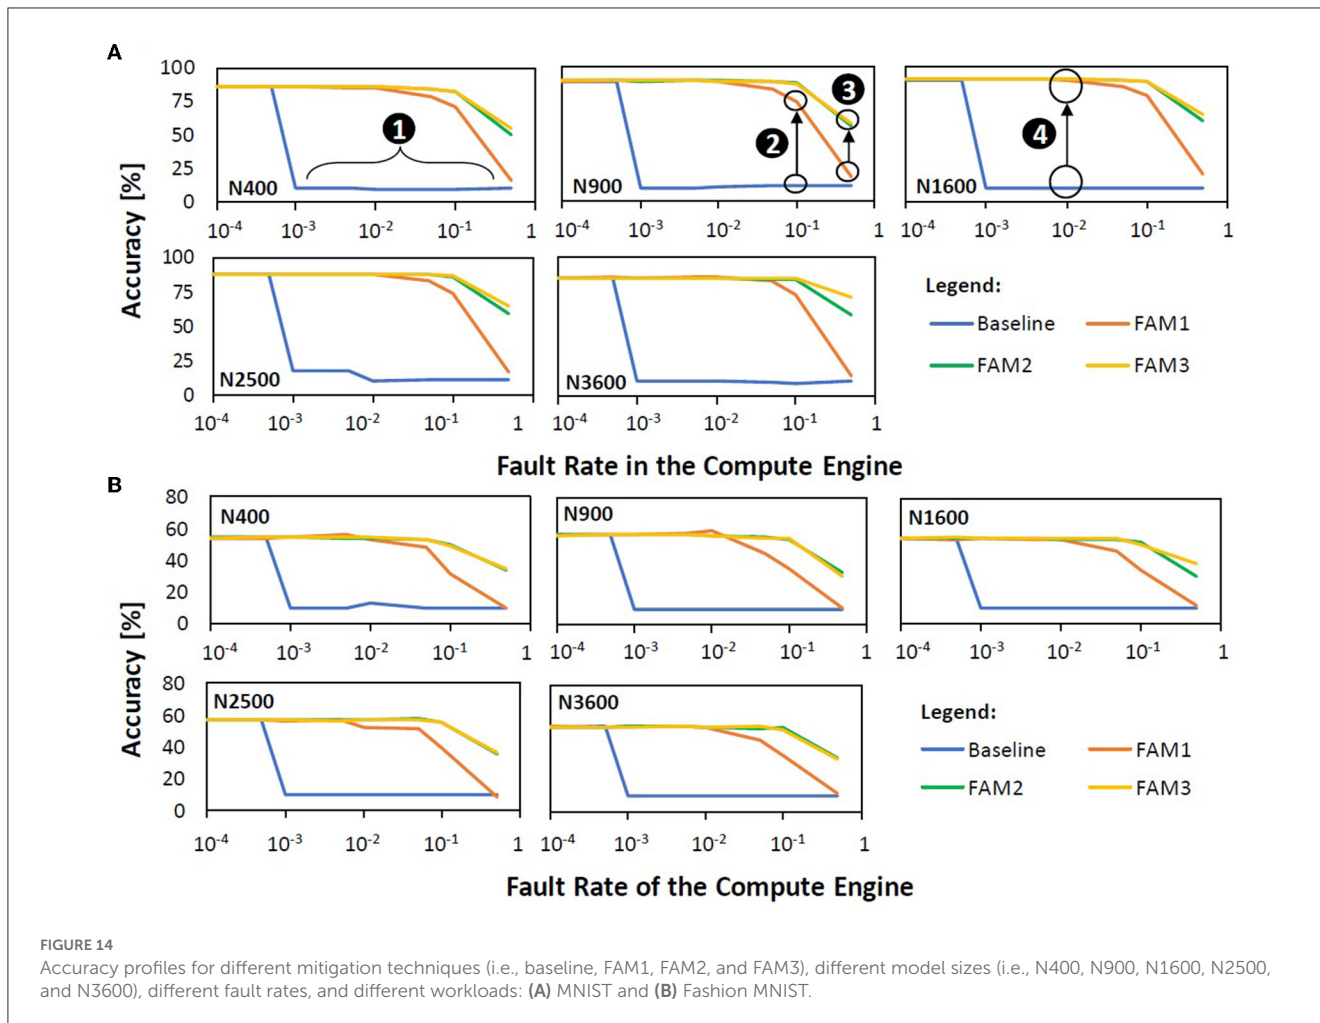


FIGURE 14 Accuracy profiles for different mitigation techniques (i.e., baseline, FAM1, FAM2, and FAM3), different model sizes (i.e., N400, N900, N1600, N2500, and N3600), different fault rates, and different workloads: (A) MNIST and (B) Fashion MNIST.

faults in the compute engine without retraining, across different model sizes, fault rates, and workloads.

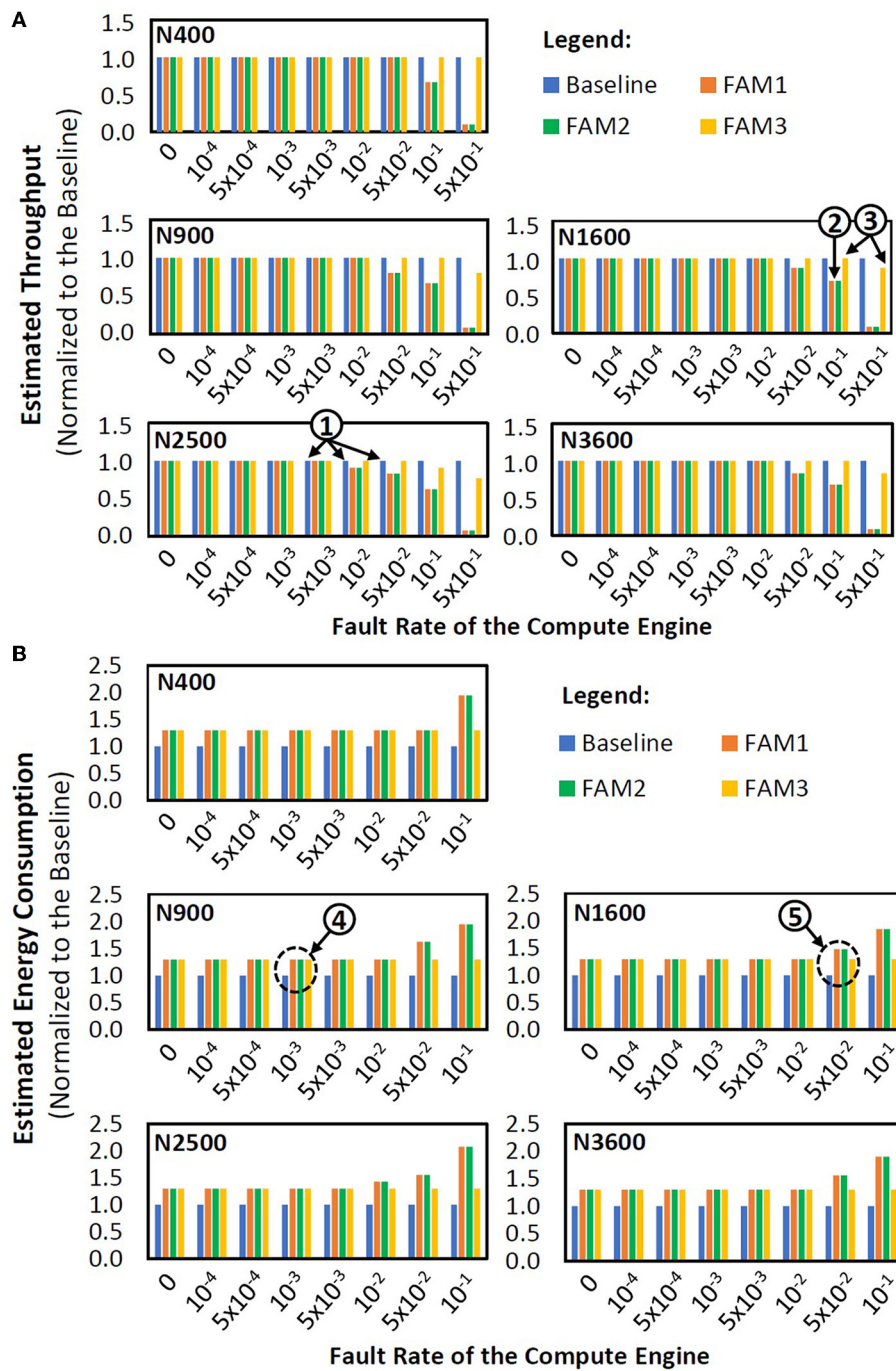
maintaining the throughput across different model sizes, fault rates, and workloads.

### 5.2. Maintaining throughput

Figure 15A presents the experimental results on the throughput of different mitigation techniques, i.e., the baseline and our FAM strategies (FAM1, FAM2, and FAM3). We observe that the baseline has the highest throughput across different model sizes and fault rates, as it uses all synapses and neurons for performing SNN executions, as shown by ❶. Meanwhile, FAM1 and FAM2 may suffer from throughput reduction because they avoid the use of faulty neurons, thereby omitting the corresponding columns of the SNN compute engine. For instance, FAM1 and FAM2 may suffer from 30% throughput reduction for N1600 with 0.1 fault rate, as shown by ❷. Meanwhile, our FAM3 can maintain the throughput close to the baseline (e.g., keeping the throughput reduction below 25% in a 0.5 fault rate), thereby improving the throughput significantly as compared to FAM1 and FAM2. The reason is that, FAM3 omits the columns of compute engine only if the corresponding neurons have faulty 'V<sub>mem</sub> reset' operations. For instance, FAM3 has less than 15% throughput reduction for N1600, as indicated by ❸. These results show that our FAM3 is effective for

### 5.3. Energy consumption and area overheads

Figure 15B shows the experimental results on the energy consumption of different mitigation techniques, i.e., the baseline and our FAM strategies (FAM1, FAM2, and FAM3). We observe that different techniques have comparable energy for small fault rates, as shown by label-❹. The reason is that small fault rates have a low probability of faulty neurons, hence the resource utilization for different techniques is similar. For large fault rates, FAM1 and FAM2 have higher energy consumption than the baseline and FAM3, as shown by label-❺. The reason is that large fault rates have a high probability of faulty neurons, hence the resource utilization for different techniques is different, i.e., FAM1 and FAM2 avoid the use of faulty neurons, thereby incurring higher compute latency and energy consumption. The baseline and FAM3 have comparable energy since FAM3 employs simple hardware enhancements: (1) multiplexing operations in each HEB which are shared for all synapses in the same column of compute engine, and (2) registers



**FIGURE 15** (A) Throughput across different mitigation techniques, different model sizes, and different fault rates for both MNIST and Fashion MNIST, as they have a similar number of SNN weights and operations. (B) Energy consumption across different mitigation operations techniques, different model sizes, and different fault rates for both MNIST and Fashion MNIST, as they have a similar number of SNN weights and operations.

accesses in ECU, thereby minimizing the energy consumption overhead for FAM3 (i.e., within 30%). For area footprint, the original compute engine consumes around 6.27 mm<sup>2</sup> of area, while the one with proposed enhancements consumes around 8.56 mm<sup>2</sup> of area. Therefore, the proposed enhancements incur about 36.5% of area overhead, which encompasses about 36.2% of ECU and about 0.3% of HEBs. The area of ECU dominates the total area

of enhancements since it mainly employs a set of 3-bit registers (i.e., 256x256 registers), which incurs a larger area as compared to HEBs (i.e., 256x25 multiplexers). These results show that *our FAM3 achieves minimum overheads in terms of energy consumption and area across different model sizes, fault rates, and workloads.*

In summary, the above discussions show that *our RescueSNN methodology can effectively mitigate permanent faults in the*

*SNN chips without retraining.* Since our RescueSNN addresses permanent faults during both the design time and the run time, it increases the yield of SNN chips, as well as enables efficient and reliable SNN executions during their operational lifetime. Furthermore, our RescueSNN also avoids carbon emission as it does not need any retraining, thereby offering an environment-friendly solution (Strubell et al., 2019, 2020).

## 5.4. Further discussion

In general, we observe that a faulty ' $V_{mem}$  reset' operation can cause significant accuracy degradation as it deteriorates the neuron from the expected behavior. The reason is that, the generated (faulty) spikes will affect how the SNN model understands the information, since an SNN model typically employs a certain spike coding scheme, i.e., rate coding in this work. Therefore, a neuron with faulty ' $V_{mem}$  reset' operation will generate a high number of spikes and dominate the classification activity, thereby leading to high misclassification and significant accuracy degradation. We also observe that, a higher number of spikes generated by faulty ' $V_{mem}$  reset' operation also indicates that the SNN model performs more frequent neuron operations that correspond to spike generation. This condition leads to higher power/energy consumption for SNN processing, which has been observed and studied in previous works (Krithivasan et al., 2019; Park et al., 2020; Putra and Shafique, 2023).

**Comparisons with Retraining Technique:** In a standard chip fabrication process, manufactured chips are evaluated in a wafer/chip test procedure (i.e., wafer acceptance test and chip probing test). This test procedure aims at evaluating the quality of each chip, including any faults in the chip (Xu et al., 2020; Fan et al., 2022). In this step, the permanent faults and the corresponding fault map information from manufacturing defects are identified. Therefore, this step does not introduce new cost, and only requires a typical cost for a standard wafer/chip test procedure (Xu et al., 2020; Fan et al., 2022). In the retraining technique, the fault map information is then incorporated in the retraining process considering how the weights and neuron operations are mapped on the SNN compute engine, i.e., so-called *fault-aware training* (FAT). In this manner, the SNN model is expected to adapt to the presence of faults, hence maintaining high accuracy. This indicates that, the retraining technique requires (1) fault map information from the chip test procedure, and (2) a full training dataset, which may be unavailable due to restriction policy. *Furthermore, each chip has a unique fault map which requires its own retraining process, thereby incurring huge time and energy costs.* Otherwise, the retraining technique will not be effective. Meanwhile, our proposed FAM technique in RescueSNN methodology leverages the fault map information to safely map the weights and neuron operations on the SNN compute engine. It ensures that the SNN processing is not negatively affected by permanent faults, thereby maintaining high accuracy. Although each chip has a unique fault map which requires a specific mapping, the cost for devising the mapping strategy is significantly lower than the cost of retraining. Furthermore, our FAM technique does not require any training dataset, hence it is highly applicable to a wide range of SNN applications.

**Benefits and Limitations of Pruning:** Neurons in the fully-connected (FC)-based SNN architecture shown in Figure 1A can be pruned while keeping the accuracy close to that of the original network, considering that a high rate of faulty ' $V_{mem}$  increase' operations does not significantly degrade accuracy. The benefits of pruning in FC-based architecture have been demonstrated in previous work (Rathi et al., 2019), including reduction of memory footprint and energy consumption. The pruning technique is suitable if we rely on offline training, i.e., an SNN model is trained offline with the training dataset, and the knowledge learnt from the training phase is kept unchanged during inference at run time. However, the pruning technique is not suitable if we consider SNN-based systems that need to update their knowledge regularly at run time to adapt to different operational environments (i.e., so-called *dynamic environments*) such as autonomous mobile agents, e.g., unmanned ground vehicles (UGVs). The reason is that, SNN-based systems may encounter new input features in different environments and the offline-trained knowledge may not be representative for recognizing the corresponding classes, thereby leading to low accuracy at run time and requiring online training to update their knowledge (Putra and Shafique, 2021b, 2022a). Therefore, SNN models with unpruned neurons and unsupervised learning capabilities are beneficial for learning and recognizing new features in (unlabeled) data samples from the operational environments during online training. In summary, users can select which SNN model to employ depending on the design requirements. An alternative is employing the FC-based SNNs shown in Figure 1A with/without pruning since they can enable multiple benefits, such as high accuracy when employing STDP-based learning under unsupervised settings, and efficient online training capabilities.

## 6. Conclusion

We propose RescueSNN, a novel methodology for mitigating permanent faults in SNN chips. RescueSNN leverages the fault map of compute engine to perform fault-aware mapping for SNN weights and operations, and employs efficient hardware enhancements for the proposed mapping technique. The results show that RescueSNN improves the SNN fault tolerance without retraining. As a result, faulty SNN chips can be rescued and used for reliable SNN processing during their operational lifetime.

## Data availability statement

Publicly available datasets were used in this study. These datasets can be found at the following sites: <http://yann.lecun.com/exdb/mnist/> (MNIST), <http://fashion-mnist.s3-website.eu-central-1.amazonaws.com/> (Fashion MNIST).

## Author contributions

RP, MH, and MS contributed to the concept development of the framework and methods. RP implemented the framework and wrote the first draft of the manuscript under the guidance of MS. RP and MH organized, analyzed, and discussed the experiments.

MS supervised the research work and gave valuable scientific and technical feedbacks at each step. All authors contributed to manuscript revision, read, and approved the submitted version.

## Funding

This work was supported in parts by the Center for Artificial Intelligence and Robotics (CAIR), funded by Tamkeen under the NYUAD Research Institute Award CG010, and the Center for Cyber Security (CCS), funded by Tamkeen under the NYUAD Research Institute Award G1104. This work was also partially supported by the project “eDLAuto: An Automated Framework for Energy-Efficient Embedded Deep Learning in Autonomous Systems”, funded by the NYUAD Research Enhancement Fund (REF). The authors acknowledge TU Wien Bibliothek for the publication fee support through its Open Access Funding Program.

## References

- Akopyan, F., Sawada, J., Cassidy, A., Alvarez-Icaza, R., Arthur, J., Merolla, P., et al. (2015). TrueNorth: design and tool flow of a 65 mw 1 million neuron programmable neurosynaptic chip. *IEEE Trans. Comput. Aided Design Integrat. Circ. Syst.* 34, 1537–1557. doi: 10.1109/TCAD.2015.2474396
- Baloch, N. K., Baig, M. I., and Daneshmand, M. (2019). Defender: a low overhead and efficient fault-tolerant mechanism for reliable on-chip router. *IEEE Access* 7, 142843–142854. doi: 10.1109/ACCESS.2019.2944490
- Basu, A., Deng, L., Frenkel, C., and Zhang, X. (2022). “Spiking neural network integrated circuits: a review of trends and future directions,” in *2022 IEEE Custom Integrated Circuits Conference (CICC)* (Newport Beach, CA: IEEE), 1–8.
- Davies, M., Srinivasa, N., Lin, T., Chinya, G., Cao, Y., Choday, S. H., et al. (2018). Loihi: a neuromorphic manycore processor with on-chip learning. *IEEE Micro* 38, 82–99. doi: 10.1109/MM.2018.112130359
- Diehl, P., and Cook, M. (2015). Unsupervised learning of digit recognition using spike-timing-dependent plasticity. *Front. Comput. Neurosci.* 9, 99. doi: 10.3389/fncom.2015.00099
- El-Sayed, S. A., Spyrou, T., Pavlidis, A., Afacan, E., Camuñas-Mesa, L. A., and Lina, B. (2020). “Spiking neuron hardware-level fault modeling,” in *2020 IEEE 26th International Symposium on On-Line Testing and Robust System Design* (Napoli: IEEE), 1–4.
- Fan, S.-K. S., Cheng, C.-W., and Tsai, D.-M. (2022). Fault diagnosis of wafer acceptance test and chip probing between front-end-of-line and back-end-of-line processes. *IEEE Trans. Autom. Sci. Eng.* 19, 3068–3082. doi: 10.1109/TASE.2021.3106011
- Frenkel, C., Lefebvre, M., Legat, J., and Bol, D. (2019). A 0.086-mm<sup>2</sup> 12.7-pj/sop 64k-synapse 256-neuron online-learning digital spiking neuromorphic processor in 28-nm cmos. *IEEE Trans. Biomed. Circ. Syst.* 13, 145–158. doi: 10.1109/TBCAS.2018.2880425
- Hanif, M. A., Khalid, F., Putra, R. V. W., Rehman, S., and Shafique, M. (2018). “Robust machine learning systems: reliability and security for deep neural networks,” in *2018 IEEE 24th International Symposium on On-Line Testing and Robust System Design* (Platja d’Aro: IEEE), 257–260.
- Hanif, M. A., Khalid, F., Putra, R. V. W., Teimoori, M. T., Kriebel, F., Zhang, J. J., et al. (2021). “Robust computing for machine learning-based systems,” in *Dependable Embedded Systems* (Cham: Springer), 479–503.
- Hazan, H., Saunders, D. J., Khan, H., Patel, D., Sanghavi, D. T., Siegelmann, H. T., et al. (2018). BindsNET: a machine learning-oriented spiking neural networks library in python. *Front. Neuroinform.* 12, 89. doi: 10.3389/fninf.2018.00089
- Hazan, H., Saunders, D. J., Sanghavi, D. T., Siegelmann, H., and Kozma, R. (2019). Lattice map spiking neural networks (lm-snns) for clustering and classifying image data. *Ann. Math. Artif. Intell.* 88, 1237–1260. doi: 10.1007/s10472-019-09665-3
- Izhikevich, E. M. (2004). Which model to use for cortical spiking neurons? *IEEE Trans. Neural Netw.* 15, 1063–1070. doi: 10.1109/TNN.2004.832719
- Krithivasan, S., Sen, S., Venkataramani, S., and Raghunathan, A. (2019). “Dynamic spike bundling for energy-efficient spiking neural networks,” in *2019 IEEE/ACM International Symposium on Low Power Electronics and Design* (Lausanne: IEEE), 1–6.
- Lyons, R. E., and Vanderkulk, W. (1962). The use of triple-modular redundancy to improve computer reliability. *IBM J. Res. Dev.* 6, 200–209. doi: 10.1147/rd.62.0200
- Maass, W. (1997). Networks of spiking neurons: The third generation of neural network models. *Neural Networks* 10, 1659–1671. doi: 10.1016/S0893-6080(97)00011-7
- Mercier, R., Killian, C., Kritikakou, A., Helen, Y., and Chillet, D. (2020). “Multiple permanent faults mitigation through bit-shuffling for network-on-chip architecture,” in *2020 IEEE 38th International Conference on Computer Design* (Hartford, CT: IEEE), 205–212.
- Mercier, R., Killian, C., Kritikakou, A., Helen, Y., and Chillet, D. (2022). BiSuT: a noc-based bit-shuffling technique for multiple permanent faults mitigation. *IEEE Trans. Comput. Aided Design Integrat. Circ. Syst.* 41, 2276–2289. doi: 10.1109/TCAD.2021.3101406
- Mozafari, M., Ganjtabesh, M., Nowzari-Dalini, A., and Masquelier, T. (2019). Spynetorch: Efficient simulation of convolutional spiking neural networks with at most one spike per neuron. *Front. Neurosci.* 13, 625. doi: 10.3389/fnins.2019.00625
- Painkras, E., Plana, L. A., Garside, J., Temple, S., Galluppi, F., Patterson, C., et al. (2013). SpiNNaker: a 1-w 18-core system-on-chip for massively-parallel neural network simulation. *IEEE J. Solid State Circ.* 48, 1943–1953. doi: 10.1109/JSSC.2013.2259038
- Park, S., Kim, S., Na, B., and Yoon, S. (2020). “T2FSNN: deep spiking neural networks with time-to-first-spike coding,” in *57th ACM/IEEE Design Automation Conference* (San Francisco, CA: IEEE), 1–6.
- Pfeiffer, M., and Pfeil, T. (2018). Deep learning with spiking neurons: opportunities and challenges. *Front. Neurosci.* 12, 774. doi: 10.3389/fnins.2018.00774
- Putra, R. V. W., Hanif, M. A., and Shafique, M. (2021a). “ReSpawn: energy-efficient fault-tolerance for spiking neural networks considering unreliable memories,” in *2021 IEEE/ACM International Conference On Computer Aided Design* (Munich: IEEE), 1–9.
- Putra, R. V. W., Hanif, M. A., and Shafique, M. (2021b). “SparkXD: a framework for resilient and energy-efficient spiking neural network inference using approximate DRAM,” in *2021 58th ACM/IEEE Design Automation Conference* (San Francisco, CA: IEEE), 379–384.
- Putra, R. V. W., Hanif, M. A., and Shafique, M. (2022a). EnforceSNN: enabling resilient and energy-efficient spiking neural network inference considering approximate DRAMs for embedded systems. *Front. Neurosci.* 16, 937782. doi: 10.3389/fnins.2022.937782
- Putra, R. V. W., Hanif, M. A., and Shafique, M. (2022b). “SoftSNN: low-cost fault tolerance for spiking neural network accelerators under soft errors,” in *Proceedings of the 59th ACM/IEEE Design Automation Conference* (San Francisco, CA: IEEE), 151–156.
- Putra, R. V. W., and Shafique, M. (2020). FSpINN: an optimization framework for memory-efficient and energy-efficient spiking neural networks. *IEEE Trans. Comput. Aided Design Integrat. Circ. Syst.* 39, 3601–3613. doi: 10.1109/TCAD.2020.3013049
- Putra, R. V. W., and Shafique, M. (2021a). “Q-SpINN: a framework for quantizing spiking neural networks,” in *2021 International Joint Conference on Neural Networks* (Shenzhen), 1–8.

## Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

## Publisher’s note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.



- Putra, R. V. W., and Shafique, M. (2021b). "SpikeDyn: a framework for energy-efficient spiking neural networks with continual and unsupervised learning capabilities in dynamic environments," in *2021 58th ACM/IEEE Design Automation Conference* (San Francisco, CA: IEEE), 1057–1062.
- Putra, R. V. W., and Shafique, M. (2022a). "lpSpikeCon: enabling low-precision spiking neural network processing for efficient unsupervised continual learning on autonomous agents," in *2022 International Joint Conference on Neural Networks* (Padua), 1–8.
- Putra, R. V. W., and Shafique, M. (2022b). Mantis: enabling energy-efficient autonomous mobile agents with spiking neural networks. *arXiv preprint arXiv:2212.12620*. doi: 10.48550/arXiv.2212.12620
- Putra, R. V. W., and Shafique, M. (2023). TopSpark: a timestep optimization methodology for energy-efficient spiking neural networks on autonomous mobile agents. *arXiv preprint arXiv:2303.01826*. doi: 10.48550/arXiv.2303.01826
- Radetzki, M., Feng, C., Zhao, X., and Jantsch, A. (2013). Methods for fault tolerance in networks-on-chip. *ACM Comput. Surveys* 46, 976. doi: 10.1145/2522968.2522976
- Rastogi, M., Lu, S., Islam, N., and Sengupta, A. (2021). On the self-repair role of astrocytes in STDP enabled unsupervised SNNs. *Front. Neurosci.* 14, 1351. doi: 10.3389/fnins.2020.603796
- Rathi, N., Chakraborty, I., Kosta, A., Sengupta, A., Ankit, A., Panda, P., et al. (2023). Exploring neuromorphic computing based on spiking neural networks: algorithms to hardware. *ACM Comput. Survey* 55, 1155. doi: 10.1145/3571155
- Rathi, N., Panda, P., and Roy, K. (2019). STDP-based pruning of connections and weight quantization in spiking neural networks for energy-efficient recognition. *IEEE Trans. Comput. Aided Design Integrat. Circ. Syst.* 38, 668–677. doi: 10.1109/TCAD.2018.2819366
- Schuman, C. D., Parker Mitchell, J., Johnston, J. T., Parsa, M., Kay, B., Date, P., et al. (2020). "Resilience and robustness of spiking neural networks for neuromorphic systems," in *2020 International Joint Conference on Neural Networks* (Glasgow), 1–10.
- Spyrou, T., El-Sayed, S., Afacan, E., Camu nas-Mesa, L., Linares-Barranco, B., and Stratigopoulos, H.-G. (2021). "Neuron fault tolerance in spiking neural networks," in *2021 Design, Automation and Test in Europe Conference and Exhibition* (Grenoble), 743–748.
- Srinivasan, G., Roy, S., Raghunathan, V., and Roy, K. (2017). "Spike timing dependent plasticity based enhanced self-learning for efficient pattern recognition in spiking neural networks," in *IJCNN* (Anchorage, AK), 1847–1854.
- Stanisavljević, M., Schmid, A., and Leblebici, Y. (2011). "Reliability, faults, and fault tolerance," in *Reliability of Nanoscale Circuits and Systems* (New York, NY: Springer), 7–18.
- Strubell, E., Ganesh, A., and McCallum, A. (2019). "Energy and policy considerations for deep learning in NLP" in *Proceedings of 57th Annual Meeting of the Association for Computational Linguistics* (Florence), 3645–3650.
- Strubell, E., Ganesh, A., and McCallum, A. (2020). Energy and policy considerations for modern deep learning research. *Proc. AAAI Conf. Artif. Intell.* 34, 13693–13696. doi: 10.1609/aaai.v34i09.7123
- Sze, H. Y. (2000). *Circuit and method for rapid checking of error correction codes using cyclic redundancy check*. US Patent 6,092,231.
- Tavanaei, A., Ghodrati, M., Kheradpisheh, S. R., Masquelier, T., and Maida, A. (2019). Deep learning in spiking neural networks. *Neural Networks* 111, 47–63. doi: 10.1016/j.neunet.2018.12.002
- Vadlamani, R., Zhao, J., Burleson, W., and Tessier, R. (2010). "Multicore soft error rate stabilization using adaptive dual modular redundancy," in *2010 Design, Automation and Test in Europe Conference and Exhibition* (Dresden), 27–32.
- Vatajelu, E.-I., Di Natale, G., and Anghel, L. (2019). "Special session: reliability of hardware-implemented spiking neural networks (SNN)," in *2019 IEEE 37th VLSI Test Symposium* (Monterey, CA: IEEE), 1–8.
- Wang, J., Ebrahimi, M., Huang, L., Xie, X., Li, Q., Li, G., et al. (2019). Efficient design-for-test approach for networks-on-chip. *IEEE Trans. Comput.* 68, 198–213. doi: 10.1109/TC.2018.2865948
- Werner, S., Navaridas, J., and Luján, M. (2016). A survey on design approaches to circumvent permanent faults in networks-on-chip. *ACM Comput. Surveys* 48, 6781. doi: 10.1145/2886781
- Xu, H., Zhang, J., Lv, Y., and Zheng, P. (2020). Hybrid feature selection for wafer acceptance test parameters in semiconductor manufacturing. *IEEE Access* 8, 17320–17330. doi: 10.1109/ACCESS.2020.2966520
- Zhang, J. J., Gu, T., Basu, K., and Garg, S. (2018). "Analyzing and mitigating the impact of permanent faults on a systolic array based neural network accelerator," in *2018 IEEE 36th VLSI Test Symposium* (San Francisco, CA: IEEE), 1–6.