R2F: A Remote Retraining Framework for AIoT Processors with Computing Errors

Dawen Xu, Meng He, Cheng Liu, Ying Wang, Long Cheng, Huawei Li, *Senior Member, IEEE,*
Xiaowei Li, *Senior Member, IEEE,* and Kwang-Ting Cheng, *Fellow, IEEE*

Abstract—**AIoT processors fabricated with newer technology nodes suffer rising soft errors due to the shrinking transistor sizes and lower power supply. Soft errors on the AIoT processors particularly the deep learning accelerators (DLAs) with massive computing may cause substantial computing errors. These computing errors are difficult to be captured by the conventional training on general purposed processors like CPUs and GPUs in a server. Applying the offline trained neural network models to the edge accelerators with errors directly may lead to considerable prediction accuracy loss.**

To address the problem, we propose a remote retraining framework (R2F) for remote AIoT processors with computing errors. It takes the remote AIoT processor with soft errors in the training loop such that the on-site computing errors can be learned with the application data on the server and the retrained models can be resilient to the soft errors. Meanwhile, we propose an optimized partial TMR strategy to enhance the retraining. According to our experiments, R2F enables elastic design trade-offs between the model accuracy and the performance penalty. The top-5 model accuracy can be improved by 1.93%-13.73% with 0%-200% performance penalty at high fault error rate. In addition, we notice that the retraining requires massive data transmission and even dominates the training time, and propose a sparse increment compression approach for the data transmission optimization, which reduces the retraining time by 38%-88% on average with negligible accuracy loss over a straightforward remote retraining.

I. INTRODUCTION

Neural networks that enable intelligent or smart things are gaining increasing popularity in IoT devices [1]. They are usually both computing- and memory-intensive, and thus pose a great challenge to the general purposed processors (GPPs) in IoT devices with limited power budgets but real-time processing requirements in many applications such as obstacle detection in mobile robots, autonomous drones and vehicles [2] [3] [4]. In this circumstance, numerous neural network accelerators closely coupled with a GPP, namely AIoT processors, emerge in IoT devices and the number grows

rapidly over the years [1]. To ensure both low-power and real-time processing of the various neural networks, many AIoT processors are fabricated with newer technology nodes. For instance, Google Edge AI platform Coral is fabricated with $7\,\mathrm{nm}$ technology, and Navida Jetson Xavier adopts $12\,\mathrm{nm}$ technology. The small feature sizes of the transistors and higher clock frequency in these AIoT processors are more likely to be affected by the extreme environments and radiation, and greatly increase the probability of soft errors accordingly [5] [6], which can induce the computing errors and cause wrong prediction when the neural networks are deployed. The wrong prediction in many safety-sensitive applications such as autonomous driving, unmanned aerial vehicle, robotics, and engine failure prediction and diagnosis may lead to catastrophic consequences and losses. Although many classical fault-tolerant design techniques such as triple modular redundancy (TMR) can be utilized to mitigate the influence of the soft errors, they typically induce considerable overhead in terms of performance and power consumption, which contradicts with the real-time processing and low-power requirements of the typical AIoT applications. Thereby, lightweight yet effective fault mitigation techniques that will not incur neither notable performance penalty nor power consumption remain highly demanded.

Fortunately, we notice that, unlike generic applications, neural networks inherently involve redundancy and are more resilient to the computing errors [7], many neural network model optimizations like quantization and pruning essentially take advantage of this feature to obtain notable performance and energy efficiency improvement with minor inference accuracy penalty [8] [9]. Hereby, a straightforward yet effective approach to mitigate the soft errors in the AIoT processors is to exploit the redundancy in the the neural network models with the retraining such that the computing errors along with the data can be learned by the retrained models. The retrained models that usually have the same sizes with the original models can be executed without performance penalty.

Retraining the neural network models to tolerate soft errors with marginal performance penalty and energy consumption overhead is promising for the AIoT processors, but it is non-trivial to conduct the retraining with existing deep learning frameworks such as Caffe [10], Tensorflow [11] and PyTorch [12]. First of all, existing frameworks typically have the entire training performed on the general purposed processors especially GPUs, but the AIoT devices with limited computing power and energy budgets usually are incapable of supporting the power-hungry GPUs and the neural network training

directly. As a result, the training must be conducted on a powerful server either in the edge or in the cloud. On the other hand, it is rather difficult for the GPUs in the server to capture the influence of the soft errors on the neural network inference conducted on the remote AIoT processors exactly. The offline trained models with noise injection show only marginal model accuracy improvement when deployed in a different faulty environment according to our experiments in Section III.

To address the above challenges, we propose a Remote Retraining Framework (R2F) for the fault-tolerant neural network models targeting at the resilient deployment on AIoT processors with soft errors. It takes the AIoT processors in the conventional training loop and exposes the on-site computing errors to the training framework such that the obtained models can learn the application data with the computing errors at the same time. More specifically, it has the forward propagation (FP) affected by the soft errors conducted on the remote AIoT processor and the backward propagation (BP) on the server. The iterative training process has the FP and BP conducted interchangeably. At the same time, the intermediate outputs of the FP need to be sent to the server for the gradient calculation and model updating. The updated model in BP needs to be transmitted to the AIoT processor for the inference in next iteration. However, there is still a lack of supporting gateware that enables the frequent communication between the AIoT processor and the server for the collaborated training. To that end, we define a set of high-level communication APIs to characterize the basic data transmissions between the remote AIoT processor and the server, and implement with the remote procedure calls integrated in ThingsBoard, a typical IoT software stack. With these APIs, the remote AIoT processor can be fitted to PyTorch for the remote retraining.

In addition, we observe that the retraining with R2F usually involves many training iterations and each iteration may include multiple batch processing. As a result, a large amount of data transmission is required between the AIoT processor and the server. Moreover, the size of the intermediate outputs of the batched inference can be much larger than that of the weights and the inputs. For instance, suppose the batch size is 16, the intermediate outputs of MobileNet (8bit fixed point) and ResNet50 (8bit fixed point) are $11\times$ and $19\times$ over the input images, and $30\times$ and $7\times$ over the weights respectively. At the same time, there is usually limited communication bandwidth between the AIoT processors and the server due to the resource constrain on the edge. Thereby, the frequent and large amount of intermediate data transmission poses a great challenge to the retraining. In this work, we propose a sparse increment compression scheme to reduce the data transmission. The basic idea is to apply TMR to the neural network execution on the AIoT processor to approximate the golden intermediate outputs of the inference. Then, we take the approximated intermediate outputs as the base and calculate the increments to the base. As the computing errors are rather sparse, the increments can be compressed effectively. When the compressed increment and the input features are transmitted to the server, the server can recompute the intermediate outputs with the transmitted input features and approximate the

actual intermediate outputs with computing errors by adding the increments. With the proposed compression method, the intermediate data transmission can be greatly reduced and the retraining time can be cut down accordingly.

On top of the on-site retraining time optimization, we also optimize the R2F for more elastic design trade-offs between the retrained model accuracy and model execution time on the AIoT processors with soft errors. Basically, we notice that straightforward on-site retraining shows limited model accuracy improvement under relatively higher fault error rate while TMR can be used to reduce the influence of soft errors and improve the model accuracy significantly. However, the TMR overhead is usually overwhelming especially for the AIoT processors with limited power budgets. In this circumstance, we apply a heuristic algorithm to select the most fragile layers and have them protected via TMR which is also utilized in the remote retraining. The neural network with partial TMR protection is implemented in R2F such that the retrained model accuracy can be improved with minor performance penalty even under higher fault error rate.

The contribution of this work can be summarized as follows:

- We proposed R2F, an efficient remote retraining framework, to enable the collaborated neural network model retraining on both remote AIoT processors and the servers. It takes the remote AIoT processors with computing errors in the training loop and has the influence of the soft errors learned with the application data such that the retrained models can be fault-tolerant.
- We define a series of client-server communication APIs on top of typical IoT software stacks to facilitate the R2F implementation on a conventional training framework like PyTorch and a representative IoT software stack. Moreover, we further optimize R2F from the perspective of the retraining time and the model accuracy. Specifically, we propose a sparse increment compression to greatly alleviate the large data transmission overhead in retraining, and provide an elastic design trade-off between the model accuracy and performance penalty with an optimized partial TMR strategy.
- According to our experiments on a set of typical neural networks, R2F reduces the training time by 38%-88% with the proposed data transmission optimization when compared to the baseline method. It achieves an elastic design trade-off between the model accuracy and the performance penalty with the proposed partial TMR protection. The top-5 model accuracy can be improved by 1.93%-13.73% while the performance penalty ranges from 0%-200% under high fault error rate.

The structure of this paper is organized as follows. Section 2 briefly introduces the related works on fault tolerant design of neural network models and accelerators. Section 3 analyzes the influence of soft errors on neural network accelerators and motivates the necessity of the neural network model retraining. Section 4 details R2F for resilient neural network execution on AIoT processors with soft errors. Section 5 introduces the proposed optimizations for R2F. Section 6 includes comprehensive experiments and evaluates R2F from

different angles including accuracy improvement and training time. Finally, we conclude this work in Section 7.

II. RELATED WORK

A. AIoT Processors

For the sake of both the low-power and real-time processing, neural network accelerators are increasingly utilized for the Artificial intelligence (AI) processing [13] [14] in Internet of Things (IoT). They are usually closely integrated with a general purposed processor, and the integrated processor that enables AI on IoT devices is known as AIoT processor. Although numerous efforts have been devoted to the AIoT processor design especially the neural network accelerator design [15] [16], it remains rather challenging to ensure the contradictory design goals of low energy consumption, high performance and prediction accuracy [17] [18]. In this case, the reliability of the neural network processing on AIoT processors under soft errors further complicates the design. Conventional fault-tolerant techniques such as TMR that typically will induce substantial power consumption and performance penalty can not be used directly. More lightweight fault-tolerant approaches are highly demanded for the neural network processing on AIoT processors.

B. Fault-tolerant Neural Network Processing

Plenty of prior works have investigated the fault-tolerant processing of neural networks from many different angles [19] [20] [21] [22] [23] [24] [25]. They can be roughly divided into three categories based on the fault-tolerant targets. Some of them attempt to develop fault-tolerant neural network architectures, some of them seek to harden the underlying hardware infrastructures while some of them adopt hybrid approaches that take both the hardware accelerators and the neural network models into consideration at the same time. They will be illustrated in detail in the rest of this subsection.

1) Fault-tolerant Neural Network Models: According to the evaluation in [19] [20], minor computing errors in the neural network execution may not necessarily cause the wrong prediction. Basically, neural networks are usually more resilient compared to generic applications because of the computing redundancy and activation functions that can mitigate the computing variations in the neural network models. Many prior works exploit this feature of neural network models to further improve the resilience of the neural network processing. Some of them mainly rely on the training by introducing additional noise or in-situ faults [26] [27], or by adding regularization or penalty terms [28] [29] [30], or adding constraints to the weights [17] [31]. Model retraining-based approaches are also applied specifically for the emerging yet imperfect RRAM-based neural network accelerators [32] [33]. Unlike these aforementioned works that will not change the structure of the neural network models, FTT-NAS [34] [35] provides an end-to-end fault-tolerant neural network search to redesign the neural network architecture. The obtained neural network model is more fault-tolerant, but the accuracy still drops given higher fault injection eventually and it can still be improved with retraining. Liu et al. proposed to enhance the algorithm

level error-resilience capability of DNN classifiers through a collaborative logistic classifier design by leveraging both asymmetric binary classification and an optimized variable-length "decode-free" scheme [36]. Hoang et al. [37] proposed to systematically define the clipping values of the activation functions that result in increased resilience of the networks against faults. The authors in [38] analyzed the vulnerability of the different neural network layers, replicated the most fragile layers and scheduled the processing to minimize the influence of hard errors.

2) Fault-tolerant Neural Network Accelerator Architectures: To mitigate faults in the neural network accelerator caused by soft errors, an intuitive approach is to harden the neural network accelerator with conventional fault-tolerant circuit design techniques such as TMR. For instance, the authors in [39] proposed a block-based modular redundancy strategy to mitigate the faulty computing array blocks of the neural network accelerator. The work in [40] [41] employed the spatial and temporal checksum to protect full connection and convolution layers in deep neural network models. The checksum-based approach originated from the algorithm-based fault tolerance for matrix-matrix multiplication enables both efficient error detection and correction [42]. The authors in [43] proposed a parallel stochastic computing(SC)-based NN accelerator purely using bitstream computation by fully exploiting the superior fault tolerance of SC mainly for ternary neural networks. Li et al. [44] proposed an error detecting scheme to locate incorrect Processing Elements (PEs) of the neural network accelerator and gave an error masking method to achieve fault-tolerance. Mahdiani et al. [45] proposed to relax the fault-tolerance of the VLSI implementation by employing TMR to only the computation of the most important bits such that the hardware overhead is reduced and the critical path latency is improved without any accuracy penalty. Xu et al. [46] has a dot-production unit to recompute the operations that are mapped to the faulty PEs in the 2D computing array without affecting the original dataflow. However, these techniques usually result in non-trivial overhead in terms of timing, area, and power consumption, which may fail the stringent performance and power consumption requirements of the AIoT applications. In addition, the accelerators need to be redesigned heavily, which can also be a barrier for the off-the-shelf products like Google Edge TPU.

3) Hybrid Fault-tolerant Techniques: There are also a few works proposed to co-optimize the neural network models and the underlying neural network accelerators at the same time for higher resilience. The work in [47] [48] proposed to add additional bypass logic to PEs in the neural network accelerator and the output will be set to be constant such as zero for faulty PEs. On top of the accelerator with constant bypass, it further retrains the models to achieve higher prediction accuracy. The authors in [49] proposed to add different bypass data paths to the PEs in neural network accelerators such that faulty PEs can be skipped and had the weights mapped to the faulty PEs pruned at the same time. Instead of directly pruning the weights, they reorganized the models to minimize the sum of the the saliency of the pruned neural network weights, which greatly alleviates the accuracy degradation. The authors in [50]

proposed a software and hardware co-design methodology to effectively preserve the classification accuracy of CNN with few on-device training iterations on RRAM-crossbars. Kim et al. [51] proposed an algorithm and hardware co-designed fault-tolerance framework called MATIC, which combines the characteristics of destructive SRAM reads with the error resilience of neural networks in a memory-adaptive training process. Ma et al. [52] leverage the fault-tolerance of the neural network models to mitigate the faults caused by the process variation of the neural network accelerator with hardware bypassing and a novel weight transfer technique. In this case, the computing array of the neural network accelerator can run at higher frequency with limited accuracy drop.

In summary, retraining is usually applied to obtain fault-tolerant neural network models particularly for model-based and co-designed fault-tolerant approaches. However, prior works do not have the remote retraining overhead evaluated and ignore the overhead in IoT system with limited communication bandwidth.

III. MOTIVATION

In this section, we mainly investigate the influence of soft errors on the neural network prediction accuracy and effectiveness of retraining with computing errors, which motivates the proposed remote retraining framework.

A. Influence of Soft Errors on Model Accuracy

In order to evaluate the influence of soft errors on the neural network model accuracy, we have random bit errors injected to weights, inputs, outputs as well as the hidden states of the neural network models similar to the approach utilized in [19]. We take three widely utilized neural network models, including ResNet18, MobileNet, and SqueezeNet pre-trained on ImageNet as the benchmark. All the models are 8bit fixed point. Note that the bit error rate (BER) represents the total number of bit errors over the total bit number of the data i.e. weights, inputs, outputs and hidden states of the neural network models. The experiment result is shown in Figure 1. It can be observed that the prediction accuracy of the neural network models drops little when BER is lower than 5×10^{-6} though there are computing errors caused by the soft errors, which also demonstrates the intrinsic fault tolerance of the neural network models. Nevertheless, the accuracy drops rapidly when the BER reaches 5×10^{-5}. Particularly, SqueezeNet drops by 6.04% which is unacceptable for most of the accuracy-sensitive neural network models. Even for ResNet18 with the least accuracy drop, it also shows 3.72% accuracy penalty which is non-trivial.
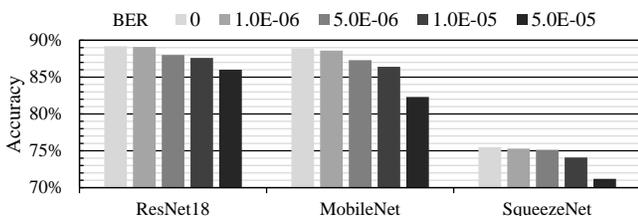
Fig. 1. Influence of the soft errors on the top-5 accuracy of different neural network models and bit error rate.

B. Effectiveness of the Model Retraining on Soft Errors

While it is promising to take advantage of the redundancy in neural network models with retraining to tolerate computing errors, we evaluate the effectiveness of the retraining on mitigating soft errors on a neural network accelerator in a remote AIoT processor. As the actual computing variation caused by the soft errors is not immediately available to the server and it is also difficult to model the exact variation on the server, we have the models retrained on the server with noise which is simulated via by injecting soft errors to the neural network accelerator with a different BER from that on the remote AIoT processor. Basically, we retrain the models with unmatched computing errors. Meanwhile, we also compare it with an on-site retraining which has the exact computing errors on remote AIoT processor transmitted to the server. The comparison is presented in Figure 2. Note that 'Base' refers to the model without retraining, 'UR' refers to unmatched retraining and 'MR' refers to the matched on-site retraining. Particularly for the unmatched retraining, we set two different bit error rate. One of them is far from the actual BER ('UR-F') and the other one is close to the actual BER ('UR-C'). It can be observed that retraining with unmatched BER that is far from the on-site situation i.e. 'UM-F' poses marginal model accuracy small improvements. In contrast, the retraining with matched retraining and unmatched retraining that is close to the on-site situation i.e. 'UR-C' exhibits much more significant accuracy improvement in general.

In summary, retraining is generally beneficial to the prediction accuracy when the model is executed on a neural network accelerator affected by soft errors. Nevertheless, the retraining must have the computing errors caused by the soft errors considered and minor mismatch is acceptable. At the same time, the benefits will be dramatically undermined if the inference condition differs too much from the training condition.

IV. REMOTE RETRAINING FRAMEWORK (R2F)

A. Overview

To retrain a fault-tolerant neural network model for resilient execution on an AIoT processor, we opt to integrate the AIoT processor in the training loop of a conventional deep learning framework such that the computing errors caused by the soft errors can be learned and tolerated by the resulting neural network models, and develop a remote retraining framework (R2F) on top of PyTorch as shown in Figure 3. Unlike the conventional offline neural network training frameworks, it adopts a client-server computing diagram for the collaborated retraining between a remote AIoT processor and a server. On the server side, it reuses the conventional backward-propagation (BP) in PyTorch, but it needs to acquire the intermediate outputs of the neural network forward-propagation (FP) via a series of APIs denoted as *R2F.Server.API()*. At the same time, it also needs to send the neural network models updated after BP to the client with the communication APIs provided by *R2F.Server.API()*. On the client side, it receives the neural network models sent from the server and conducts the FP on the AIoT processor. Meanwhile, it has the

(a) ResNet18   (b) MobileNet   (c) SqueezeNet
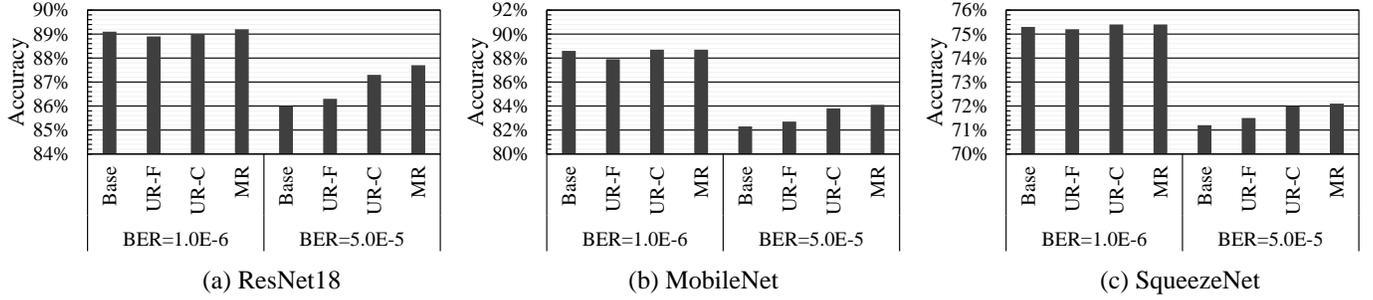
Fig. 2. The influence of the neural network retraining with matched and unmatched computing errors. Note that the BER used for training under 'UR-F' and 'UR-C' are 5.0×10^{-5} and 5.0×10^{-6} respectively when the actual inference is conducted under 1.0×10^{-6}. BER used for training under 'UR-F' and 'UR-C' are 1.0×10^{-6} and 4.5×10^{-5} when the actual inference is conducted under 5.0×10^{-5}.
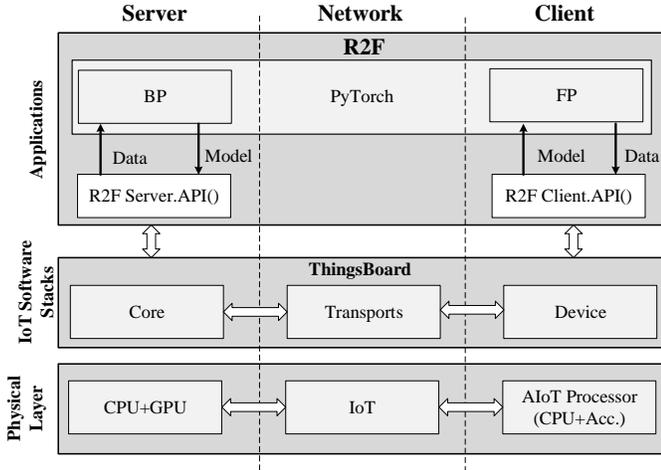
Fig. 3. An overview of the proposed remote retraining (R2F) framework on an AIoT system. It is essentially built on PyTorch and IoT software stacks, i.e. ThingsBoard, and integrates them with a series of client-server APIs.

intermediate outputs of the FP extracted and sent to the server with the *R2F.client.API()*.

Under the R2F, it is a typical IoT software stack and we adopt an open-source IoT framework called ThingsBoard (TB) in this work. TB enables cloud and IoT device connectivity via industry standard IoT protocols, such as, MQTT, CoAP, and HTTP and provides IoT oriented communication such as transport component. We take advantage of these communication facilities and wrap up them for the neural network training oriented communication i.e. *R2F.server.API()* and *R2F.client.API()*, which can be seamlessly integrated by PyTorch. The bottom layer is mainly the different hardware platforms. The server is typically equipped with both a powerful GPP and GPUs while the AIoT processor is usually configured with an neural network accelerator and a low-power GPP. They are connected with an IoT network supporting various communication protocols, such as LoRa, 802.15.4g, HSPA, and LTE Cat.4. When they are used in R2F, the AIoT processor collects input data from the sensors such as camera, and conducts the neural network models in a normal inference process while the server updates the neural network models based on the on-site inference outputs.

B. Intermediate Output Extraction

Unlike the normal inference in which the outputs of the last layer of the neural network models are sufficient to obtain the prediction, the inference performed on the AIoT processor during the collaborated retraining needs to send the outputs of each neural network layer to the server for the gradient calculation and weight update in BP. However, many neural network accelerators are mainly optimized for inference without offering intermediate outputs. Outputs of some intermediate layers are completely stored in the on-chip buffer and directly consumed by the following neural network layer to reduce the accesses to the external memory. In this case, the accelerator need to add an optional data path to enable the intermediate output data write to the external memory on request. The intermediate output write can be done in parallel with the pipelined neural network execution. For some of the off-the-shelf neural network accelerators that do not support the intermediate output extraction, a more general approach to obtain the intermediate outputs is to divide the neural network models into sub models each of which includes a single neural network layer. When the sub models are compiled and executed sequentially, intermediate outputs can also be obtained, though it may take longer execution time.

C. Client-Server Communication Interface

To enable the collaborated neural network model retraining with both an AIoT processor and a server that are far from each other, we define and implement a series of client-server APIs providing the high-level communication interfaces for the proposed retraining framework. With these APIs, we can adapt different types of AIoT processors and IoT software stacks to a unified retraining framework, i.e. R2F. These communication APIs are summarized in Table I. It can be classified as server APIs and client APIs running on the server and the clients respectively. The server APIs are mainly used to configure the mode of the remote AIoT processors, to deploy the updated neural network models, and to collect the inference results as well as the intermediate outputs sent from the remote clients. The client APIs are mainly used to response the different processing commands from the server and send the processing results to the server. There are two data conversion APIs used in the server side, because floating point is usually used in BP while 8-bit fixed point is mostly used in FP in the AIoT processor. More specifically, the received intermediate

1	API Name	*server.setAIoTMode(uint deviceID, uchar mode)*
	Description	mode='inference', it sets the AIoT processor to conduct inference only. mode='training', it sets the AIoT processor to dump intermediate outputs for retraining.
2	API Name	*server.deployModel(uint deviceID, Model* model)*
	Description	It sends the neural network model to the AIoT processor for deployment.
3	API Name	*server.getData(uint deviceID, uchar mode)*
	Description	mode='inference', it receives outputs of the inference sent from the AIoT processor. mode='training', it receives both the intermediate outputs and network outputs sent from the AIoT processor.
4	API Name	*server.convertFloat2Int(float* fData, uchar* iData)*
	Description	It converts the floating point model generated in BP to fixed point for deployment.
5	API Name	*server.convertInt2Float(uchar* iData, float* fData)*
	Description	It converts the received fixed point model to float model for BP.
6	API Name	*client.sendAck(uint* deviceID)*
	Description	It sends an acknowledgement to the server to ensure the finish of the setup or commands.
7	API Name	*server.sendData(uint* deviceID, uchar mode)*
	Description	mode='inference', it sends the neural network outputs to the server. mode='training', it sends both the intermediate outputs and network outputs to the server on request.

outputs will be converted to floating point for the float gradient calculation while the updated model will be quantized to 8-bit fixed point before it is sent to the AIoT processor for the deployment. Brief descriptions of all the APIs are listed in Table I.

V. R2F OPTIMIZATIONS

On top of the baseline R2F, we further optimize it from the following different angles. First, we optimize the communication time that dominates the on-site retraining time due to the limited uplink bandwidth of the AIoT processors. Second, we propose a partial TMR protection strategy to further improve the retrained model accuracy with minor performance penalty. Third, TMR is the basis of the aforementioned optimizations and its implementations are also optimized. They will be detailed in the rest of this section.

A. Communication Optimization

As discussed in prior section, R2F needs to transmit a large amount of intermediate outputs of the neural network models from the AIoT processors to the server, which requires both considerable time and power consumption due to the limited uplink bandwidth. Thereby, we seek to optimize the communication to enable efficient remote training. We notice that the computing errors caused by the soft errors can be effectively mitigated with classical TMR. Thus, we apply TMR to the neural network processing on the AIoT processor to obtain more accurate intermediate computing results of the

neural network models. Since these results are close to the reference outputs, we take it as approximate reference outputs. Then, we calculate the increments of the intermediate outputs relative to the approximate reference outputs on the AIoT processor with soft errors. Since they are not significantly different, the incremental results includes a large number of zeros and can be compressed efficiently. Thereby, we can have the compressed increments instead of the intermediate outputs with computing errors sent to the server. At the same time, we also have the inputs of the neural network sent to the server. In the server, the reference outputs of the intermediate outputs can be re-calculated with the inputs. With both the increments of the intermediate outputs and the reference intermediate outputs, we can approximate the intermediate outputs of the neural networks in the AIoT processors affected by the soft errors. According to the motivation experiment in Section 3, these approximate outputs are still appropriate for the BP and on-site retraining on the server.

To support TMR on AIoT processors without hardware modification, we conduct temporal TMR directly. Basically, the neural network accelerator computes on the same inputs three times when it is set to be 'training' mode. The results of each output layer will be stored in memory. Then we have the general purposed processor to conduct the voter operations for each output. Since the data are sequentially stored, and the voting process can be improved with the vector processing unit inside the AIoT processor, the processing time of the voting stage is small compared to the inference time.

B. Critical Layer Protection

Although the on-site model retraining improves the resilience of the neural network models to the soft errors, the prediction accuracy loss remains non-trivial under relatively higher fault injection rate. While the straightforward TMR on all neural network layers can greatly alleviate the influence of soft errors and extend the upper limit of the retraining method, it induces considerable performance penalty. Moreover, we notice that the computing errors on some layers of the neural network models may have distinct influence on the resulting prediction accuracy. Thereby, we select those layers that have the most significant influence on the neural network accuracy as critical layers and only have them protected with TMR to reduce both the model accuracy loss and the performance penalty. In addition, the implementation of the TMR is also consistent with that is mentioned in communication optimization.

In order to optimize the critical layer protection, we formulate the problem as follows. Suppose the target neural network includes l layers, and s layers are protected and the indices of the protected layer belong to a set S. The model accuracy of the neural network can be denoted as A_S. The computing overhead of each neural network layer is denoted as O_i where i represents the layer index. The design goal of the critical layer protection is to determine the set of the layers S that need to be protected such that A_S is maximized where the additional computing overhead relative to the original computing is less than r. To address the problem, we propose

a heuristic algorithm to optimize the critical layer selection as illustrated in Algorithm 1. The basic idea is to iteratively search the most critical neural network layer in a layer-wise manner. It continues the selection until the overhead of the partial TMR exceeds r_{max}.

Algorithm 1 Critical layer selection algorithm

Input: A neural network with l layers, the set of all the layer indices $L = (1, 2, ..., l)$, the computing overhead of the ith layer is O_i where $1 \leq i \leq l$.

Output: The total number of protected layers s and the set of protected neural layer indices S such that the A_S is maximized and the normalized redundant computing overhead is no more than r_{max}.

1: $s = 0$, $S = \emptyset$
2: **while** $(r < r_{max})$ **do**
3:   **for** each $i \in (L \setminus S)$ **do**
4:     Measure the accuracy $A_{S \cup (i)}$
5:   **end for**
6:   Find i that $A_{S \cup (i)}$ is maximized.
7:   S.append(i), $s \leftarrow s + 1$
8:   Calculate the normalized overhead $r = \frac{\sum_{i \in S} O_i}{\sum_{i \in L} O_i}$
9:   **if** $r > r_{max}$ **then**
10:     S.remove(i), $s \leftarrow s - 1$
11:   **end if**
12: **end while**
13: Return s and S

With the proposed critical layer selection algorithm, the total number of design options that need to be evaluated in the search is $Sum(L, s) = \sum_{k=s+1}^{k=L} k = \frac{(L+s+1) \times (L-s)}{2}$. In contrast, a straightforward brute-force search requires to evalaute $C(L, s) = \frac{L!}{s! \times (L-s)!}$ design options. Take ResNet50 as an example, suppose $s = 5$, each design option evaluation needs to conduct 1000 inference and takes around 16 seconds. The brute-force search requires to evaluate 2118760 design options and the total search time will be more than 392 days. Thus, it cannot be used in practice. The proposed search requires to evaluate 240 design options, which is 8820X faster and can be finished in around 64 minutes.

C. TMR Implementation Optimization

Although TMR is a classical redundancy approach, there are different methods to implement it on an AIoT processor. Since we will not change the architecture of the AIoT processors, a temporal TMR redundancy is used in this work. A straightforward TMR implementation is to conduct the neural network processing three times with the same input independently. Then the intermediate outputs from the three implementations are voted as the TMRed results. We name this approach as network-wise TMR (NW-TMR). While NW-TMR does not take the propagation of the computing errors across the different neural network layers into consideration, we propose to conduct the TMR in a layer-wise manner. Basically, the first layer of the neural network will be executed three times with the same input, and the outputs will be voted and the TMRed results will be used as the inputs of the next layer. This implementation is named as layer-wise TMR (LW-TMR). It is conducted iteratively until the end of the neural network model execution. We use the percentage of

TABLE II
TYPICAL IoT COMMUNICATION PROTOCOLS

Technology	Typical Applications	Bit Rate	
		Downlink	Uplink
LoRa	smart street lights and meter	50 kbit/s	
802.15.4g	remote monitoring, industrial control	800 kbit/s	
HSPA	shared payment, wearable device	21.1 Mbit/s	5.76 Mbit/s
LTE Cat.4	smart medicine, autonomous driving	150 Mbit/s	50 Mbit/s

the identical outputs between the approximate intermediate outputs and the actual intermediate outputs as the output similarity evaluation metric. In order to avoid inefficient TMR that all the three data vary substantially, we only conduct the TMR-based compression to layers with higher similarity. The threshold of the output similarity can be changed for the different trade-offs between the training overhead and the retrained model accuracy. It will be evaluated in the experiment Section.

VI. EXPERIMENT

A. Experiment Setup

1) Hardware Platform: In the experiment, the server is configured with an Intel Xeon cpu E5-2699 v3@2.30GHz processor and 128 GB DRAM while we have a Raspberry Pi 3 Model B platform equipped with an ARM Cortex A53 processor and 1 GB memory as the AIoT processor. The hardware platform is mainly used to verify the functionality of the proposed R2F framework. Since the Raspberry Pi does not have a neural network accelerator integrated, we assume that a Eyerisis-like neural network accelerator simulated with Scale-Sim is used instead. The neural network accelerator is configured with 32×32 2-D computing array, 1024 KB on-chip buffer. And the neural network models are executed with a classical weight-stationary dataflow. Moreover, the simulation-based neural network accelerator also facilitates the fault injection and analysis. The communication protocols used in IoTs greatly affect the bandwidth and even dominate the training time in R2F. While the different IoT protocols provide distinct bandwidth as shown in Table II and they are usually utilized for different domains of IoT applications, 802.15.4g (802) and HSPA with moderate bandwidth and power consumption are more likely for the low power AIoT applications, and they are evaluated in the experiments.

2) Software: we use ThingsBoard as a representative IoT framework and PyTorch as a typical deep learning framework for R2F. In order to compress the increments of the intermediate inference outputs on the ARM processors for efficient data transmission, we utilized the optimized LZ4 implementation in [53] that offers fast lossless compression in the experiments.

3) Fault Injection: Soft errors are randomly distributed to all the memory cells of the neural network accelerator including the register files and the on-chip buffers. When a memory cell is affected by a soft error, the bit in the memory cell will be flipped. The soft error rate i.e., bit error rate (BER) is defined as the ratio of the bit faults over the total number of the memory cells. As the soft errors in the register file and input/output/weight buffers essentially affect the input and output of each MAC (multiply-accumulate), we have the influence

TABLE III
NEURAL NETWORK BENCHMARK

Network	ResNet18	ResNet50	MobileNet	ShuffleNet	SqueezeNet
Model Size	1.2 MB	24.3 MB	3.3 MB	1.3 MB	1.2 MB
# of Layers	20	53	52	57	26

of the soft errors in the neural network accelerator converted to the random bit flip of both the input and output of each basic MAC in the neural network processing similar to prior works [19] [40] [34] [35]. Basically, bit errors are randomly injected to the input features, weights, hidden states, and output features during the neural network execution. It mainly evaluates the soft errors in algorithm level and does not take the soft errors in the controlling logic into consideration. This algorithm-level fault analysis strategy is also demonstrated in [54]. In this case, BER denotes the number of bit errors relative to the total bit number of the weights, inputs, hidden states, and output features. In the experiments, we investigate a broad range of BER setups starting from 1.0×10^{-6} to 1.0×10^{-4}. In addition, we focus on the fault tolerance of neural network models which are usually deployed on a deep learning accelerator in an AIoT processor, and assume that the GPP processor is reliable.

4) Neural Network Benchmark: In the experiment, we take five typical lightweight neural network models including ResNet18, ResNet50, MobileNet, ShuffleNet, and SqueezeNet utilized as the neural network benchmark. All the models are 8-bit fixed point models and pre-trained for ImageNet dataset. Details of these neural network models can be found in Table III. The number of the convolution layers ranges from 20 to 57. The sizes of the neural network models ranges from 1.2 MB to 24.3 MB. In the experiments, we select 50000 images from ImageNet for the retraining, set the epoch to 1 and the batch size to 16.

B. Prediction Accuracy Improvement

In this section, we mainly evaluate the prediction accuracy of the retraining and have the different retraining approaches compared. The neural network models executed directly on the accelerators with soft errors are considered as the baseline (Base). R2F puts the remote neural network accelerator into the training loop such that the retrained model can be fault-tolerant. The directly retrained model is noted with (DRM). While the direct retraining with R2F requires a large amount of intermediate data transmission particularly from the AIoT processor to the server, we further apply the TMR-based retraining, which has the compressed sparse increments rather than the intermediate outputs transmitted to the server for the retraining. The retrained model is denoted as an approximate retrained models (APRM). For the TMR-based retraining, we also explore the trade-offs between the percentage of the TMRed layers and the data transmission reduction. Basically, when more layers are transmitted with the original intermediate outputs, the retrained model will be more close to DRM, but more data transmission is required. In this work, we use the percentage of the identical outputs between the approximate intermediate outputs and the actual intermediate outputs as a simple output similarity evaluation metric and we only conduct the TMR-based compression to layers with higher similarity. Note that the metric is obtained with an offline analysis on a single random input. We have three different similarity thresholds, including 80%, 60%, and 50%, applied and evaluated in the experiments. The obtained models are denoted as APRM-80%, APRM-60%, APRM-50% respectively.

The resulting model accuracy of the different retraining is compared in Figure 4. It can be observed that the prediction accuracy of the neural network model degrades gracefully at the beginning but drops rapidly when the BER rises to certain points according to the 'Base' curve. Basically, the neural network models are fault-tolerant to the errors with in certain limit, but the models suffer dramatic accuracy degradation when the faults reach to the limit. In contrast to the 'Base', the retrained models with on-site computing errors generally exhibit clear prediction accuracy improvement. While the improvement is trivial when the BER is low, it gets significant when the BER is relatively higher. The top-5 prediction accuracy of DRM improves by 1.93% on average compared to 'Base' at the highest BER under which the models can still be retrained. However, the on-site retraining shows less accuracy improvement and even fails to converge for some of the neural networks under high BER. Particularly, the retraining that does not converge is denoted as 'X' in the figures. We argue that this may also be caused by the limited fault-tolerance of the neural network models. Generally, the retraining works in a certain limit of the BER and fails when the BER exceeds the limit. When we compare DRM and the different APRM methods, we notice that the prediction accuracy of the retrained neural network models show little difference and it confirms that TMR can be applied to reduce the data transmission in a large range of scenarios.

In order to further enlarge the benefits of the retraining, we propose to apply TMR to a small set of the most critical neural network layers to avoid the substantial performance penalty of a conventional TMR while retaining the model accuracy as much as possible. We call this approach as TMR-based critical layer protection (TCLP). As we can adjust the number of the critical layers and the performance overhead to compromise with the prediction accuracy improvement, we have a set of TCLP implementations with different performance overhead evaluated. We set the redundant computing relative to the original neural network computing as the overhead metric. The different TCLP implementations are denoted as TCLP-200%, TCLP-50%, and TCLP-20% respectively. TCLP-200% essentially refers to the standard TMR implementations. In addition, we also have the DRM and 'Base' compared. The comparison is shown in Figure 5. It can be observed that the retraining on top of the conventional TMR-based protection, i.e. TCLP-200% shows 13.73% and 4.97% accuracy improvement on average compared to both the 'Base' and 'DMR' particularly when the BER reaches to 1×10^{-4}. When the TMR overhead is constrained, the resulting model accuracy also shows clear improvement compared to the DRM but imposes much less performance overhead than the full TMR. For instance, TCLP-200% shows 4.97% accuracy improvement on average over the DRM with 1×10^{-4} fault injection, while TCLP-50%
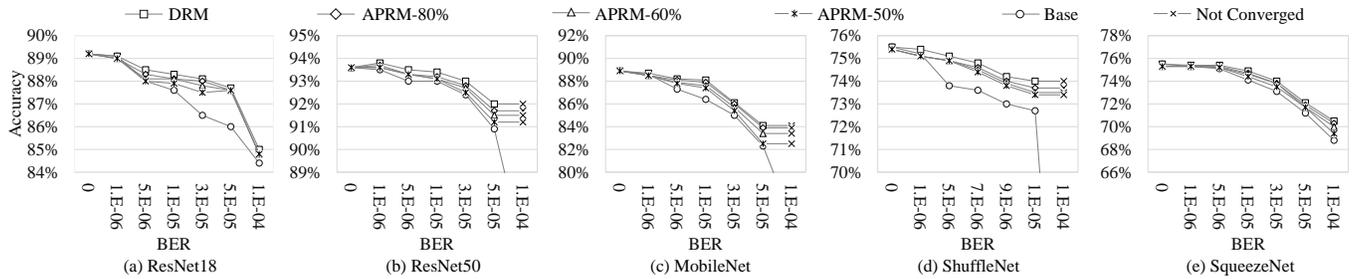
Fig. 4. The achieved prediction accuracy of the neural network models trained with data compression optimization under different output similarity thresholds.
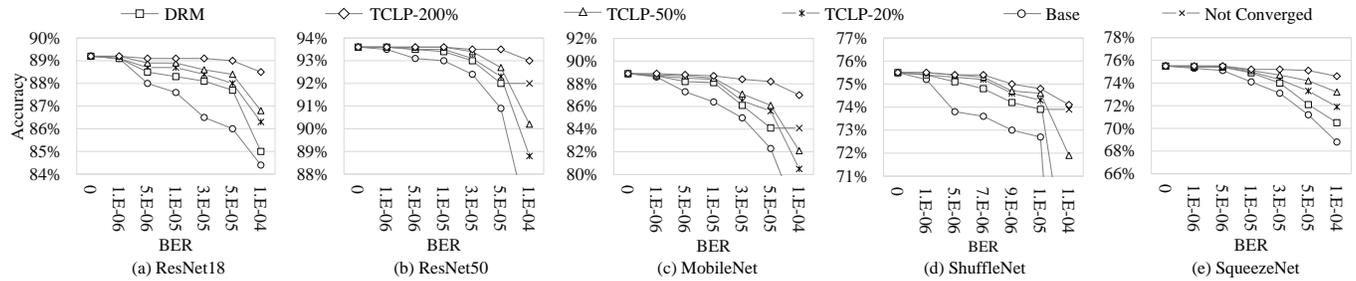
Fig. 5. The prediction accuracy of the retrained neural network models with different critical layer protection.

shows 2.01% accuracy improvement over the DRM. On the other hand, the performance of TCLP-200% is 4X lower compared to TCLP-50%. Similar design trade-offs between the model accuracy improvement and the performance overhead is observed on all the neural network models in the benchmark. Moreover, the experiments also reveal that some of the neural network layers are more critical than the others and prioritizing these layers for TMR protection helps to achieve significant accuracy improvement with least performance penalty. In addition, these design trade-offs are supported in R2F and can be applied for different application requirements.

C. Training Time Reduction

As discussed in prior section, the proposed TMR-based approximate retraining can greatly reduce the training time. We have the training time of the different neural network models under different IoT communication protocols decomposed and evaluated in Figure 6. Note that the output similarity threshold in the experiment is set to be 60% and will be discussed in the next subsection. The runtime of the retraining in R2F consists of six stages, including the forward propagation (FP), TMR processing of the intermediate outputs from FP (TMR), increment calculation, compression & decompression (Dec/compression), data transmission from the AIoT processor to the server (Data Transmission), backward propagation (BP), and model transmission from the server to the AIoT processor (Model Transmission). To facilitate the comparison over the different neural network models and communication protocols, we have the training time normalized to that of DRM, which has the neural network intermediate outputs transmitted directly to the server. Note that 802 and HSPA listed in Table II are used in this experiment. All the five representative neural network models listed in Table III are evaluated.

It can be observed that the data transmission dominates the retraining time when BER is relatively high. This is mainly because that the amount of the intermediate outputs of FP is usually large compared to the weights and it increases proportional to the batch sizes. Many of the network layers fail to meet the similarity threshold and require the direct data transmission at higher BER. On the other hand, the communication bandwidth provided by the typical IoT processors is limited and it further induces the large data transmission overhead. When the BER is lower, the majority of the neural network layers can benefit from the TMR-based data compression and the R2F training time is greatly reduced. In addition, we notice that the neural network models also have significant influence on the training time optimization. The normalized retraining time of ShuffleNet and MobileNet is much less compared to that of the rest neural network models. This may be caused by both the fault tolerance and sizes of the neural network models. Neural network models with smaller sizes and less computation are less probably to produce wrong computing results and thus are more likely to be optimized with the TMR-based compression in R2F. In contrast, ResNet18 and ResNet50 with much more computation may fail to meet the similarity threshold and more data needs to be transmitted directly. SqueezeNet is also as lighweight as ShuffleNet and MobileNet, but it is more sensitive to the soft errors as shown in Figure 4. As a result, it also requires considerable direct data transmission and consumes non-trivial time. In summary, the retraining time reduction is closely related with the BER. When the BER is low e.g. 1×10^{-6}, the retraining time can be reduced by 88% on average compared to that of DRM. When the BER is moderate e.g. 1×10^{-5}, the retraining time can be reduced by 73% on average. When the BER is high e.g. 1×10^{-4}, the retraining time can be reduced by 38% on average.
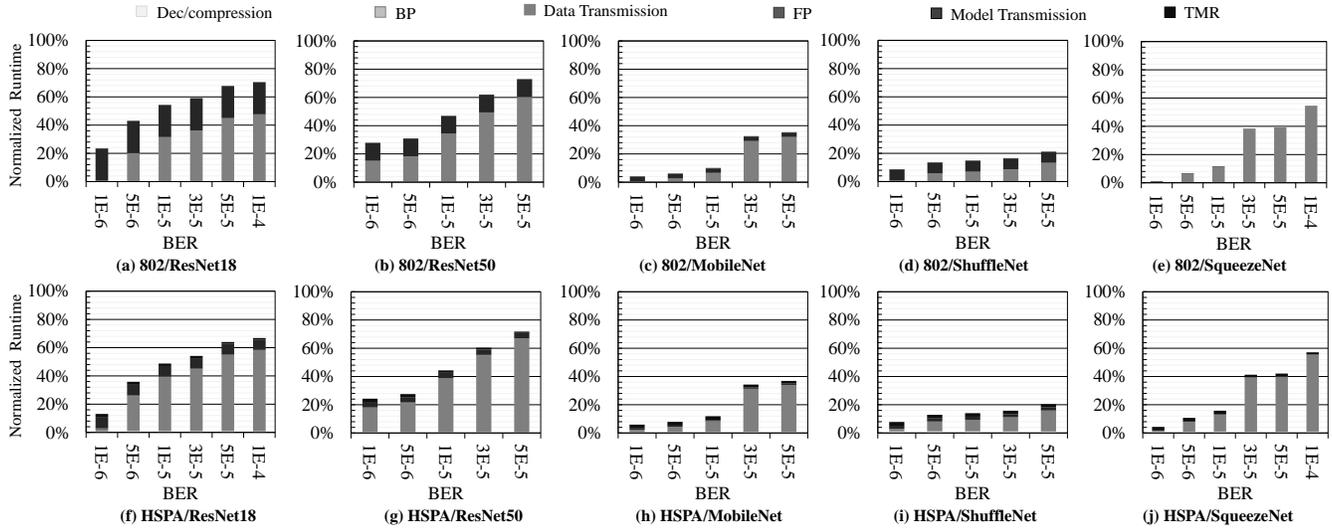
Fig. 6. The distribution of the retraining time normalized to the direct retraining. In particular, "HSPA/ResNet18" means that the AIoT processor communicates with HSPA protocol and the retrained neural network is ResNet18.
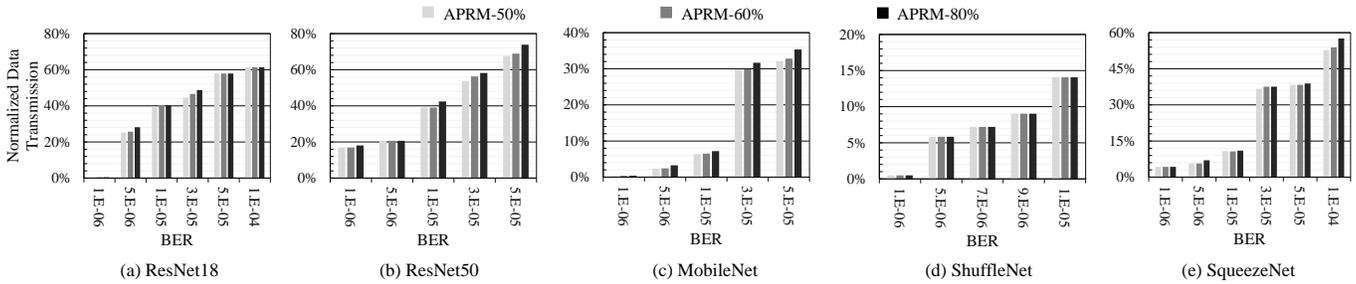
Fig. 7. The amount of data transmission required by R2F under different similarity metrics.

D. Design Option Tuning

The data transmission reduction directly depends on the amount of the layers that can be approximated with TMR according to the output similarity metric. Thus, we have the amount of the data transmission required by the different approximate retraining approaches, including APRM-50%, APRM-60%, and APRM-80%, normalized to that required by DRM and compared. The comparison is shown in Figure 7. It can be observed that more data transmission can be reduced under lower BER when almost all the intermediate data can be effectively compressed with the TMR-based increment transmission. When BER rises, more neural network layers suffer considerable computing errors and these errors can not be mitigated with TMR. As a result, only a fraction of the neural network layers can be compressed under higher BER. Another observation from Figure 7 is that the output similarity threshold does not show dramatic data transmission variations. It is mainly because that the computing errors caused by the soft errors may aggregate and the number of the computing errors increases dramatically with the rising BER. As a result, the similarity metric is not proportional to the amount of the data transmission. The amount of the data transmission in the same neural network model retraining under the different BER also confirms this feature. In this work, we adopt 60% as the similarity threshold to decide whether the outputs of a neural network layer will be transmitted directly or with the TMR-
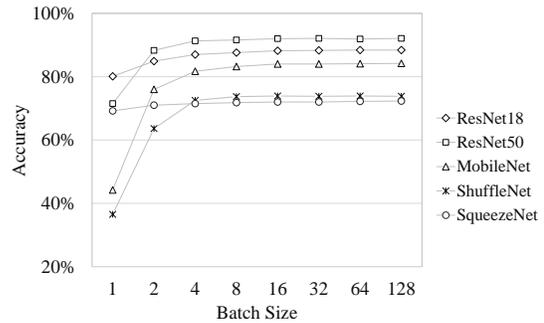
Fig. 8. The influence of batch size on the proposed R2F training.

based increment compression.

The amount of the intermediate data transmission is proportional to the batch size and affects the R2F training time. Thus, we investigate the influence of the batch size on the retrained model accuracy. The experiment result is shown in Figure 8. It can be observed that larger batch size is generally beneficial to the model accuracy but the benefits roughly get saturated when the batch size reaches to 16. The main reason is that larger batch training helps to neutralize the various influence of the computing errors caused by the random soft errors. In this work, we set batch size to be 16 for the different model retraining in R2F.

TMR is the basis of the R2F optimizations for both the retraining time and the resulting model accuracy. We have
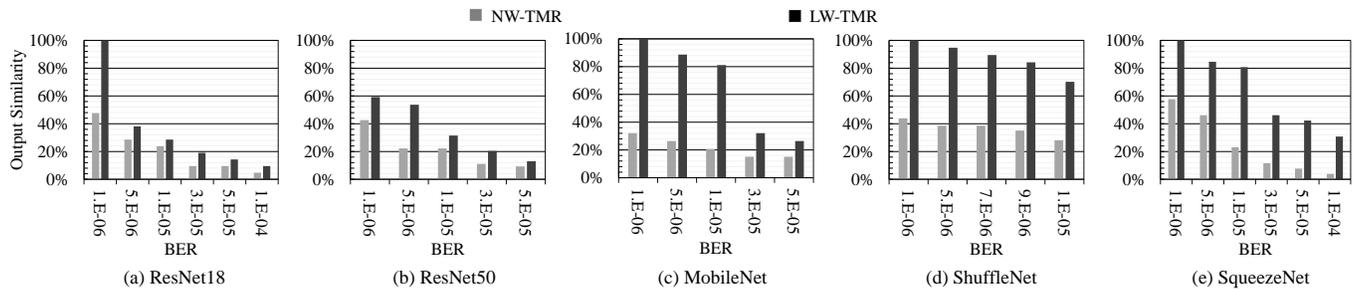
Fig. 9. Last layer output similarity comparison of the different TMR implementations.
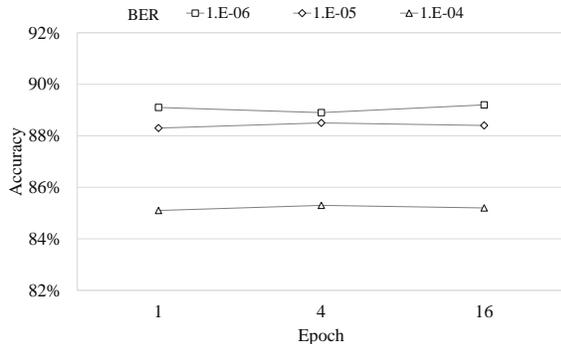
Fig. 10. The influence of different epoch setups on the proposed R2F training.

two potential TMR implementations, i.e. NW-TMR and LW-TMR evaluated with the output similarity metric. The evaluation result is presented in Figure 9. It shows that LW-TMR shows significantly higher output similarity especially under relatively lower BER. The main reason is that LW-TMR has the computing errors mitigated in layer order of the neural network architecture. The computing errors of the upstream layers are alleviated with TMR before they are passed to the downstream layers. In contrast, NW-TMR has the computing errors passed through the entire neural network and computing errors in upstream neural network layers can aggregate in the downstream neural network layers. Thereby, the computing errors are much larger and the TMR is more likely to fail. When the BER is too high, the majority of the computing errors induced by the soft errors can no longer be mitigated with neither TMR implementations. As a result, the difference narrows down in these cases, which is expected and also roughly exhibits the upper bound of the TMR-based protection.

Since larger epoch setups typically improve the prediction accuracy of the resulting model. We take ResNet18 as an example and evaluated the model accuracy under different epoch setups. The experiment result is shown in Figure 10. It can be observed that the retrained model accuracy shows little improvement given larger epoch. The main reason is that the fault-tolerant retraining is based on a pre-trained model rather than a totally new model. In this case, a single epoch is sufficient according to our experiments.

VII. Conclusion

Retraining is widely utilized to exploit the inherit redundancy of the neural network models to tolerate the soft errors in AIoT processors, but it is difficult to capture the influence of the soft errors in conventional offline training on GPUs. In this work, we propose R2F, a remote retraining framework, to put the remote AIoT processors in the training loop such that the computing errors caused by the soft errors can be learned with the application data and aware by the resulting models. On top of the basic R2F, we also propose an elastic design trade-off between the model accuracy and the performance penalty with partial TMR optimization to further enhance the retraining. According to our experiments, R2F improves the top-5 model accuracy by 1.93%-13.73% with the performance penalty ranging from 0%-200%. In addition, we notice that the remote retraining requires a large amount of intermediate data transmission between the AIoT processors and the server, which even dominates the training time due to the limited uplink bandwidth in the AIoT processors. To address the problem, we propose a sparse increment compression approach by taking advantage of the TMR to reduce the data transmission significantly. Our experiment results reveal that the retraining time can be reduced by 38%-88% on average depending on the BER.

VIII. Acknowledgement

References

[1] K. L. Loh, "1.2 fertilizing AIoT from roots to leaves," in *2020 IEEE International Solid- State Circuits Conference - (ISSCC)*, 2020, pp. 15–21.

[2] S. Kodali, P. Hansen, N. Mulholland, P. Whatmough, D. Brooks, and G.-Y. Wei, "Applications of deep neural networks for ultra low power iot," in *2017 IEEE International Conference on Computer Design (ICCD)*. IEEE, 2017, pp. 589–592.

[3] S. Lee and S. Nirjon, "Neuro. zero: a zero-energy neural network accelerator for embedded sensing and inference systems," in *Proceedings of the 17th Conference on Embedded Networked Sensor Systems*, 2019, pp. 138–152.

[4] S. Venkataramani, J. Choi, V. Srinivasan, W. Wang, J. Zhang, M. Schaal, M. J. Serrano, K. Ishizaki, H. Inoue, E. Ogawa *et al.*, "Deeptools: Compiler and execution runtime extensions for rapid ai accelerator," *IEEE Micro*, vol. 39, no. 5, pp. 102–111, 2019.

[5] A. Dixit and A. Wood, "The impact of new technology on soft error rates," in *2011 International Reliability Physics Symposium*.   IEEE, 2011, pp. 5B–4.

[6] D. Y.-W. Lin and C. H.-P. Wen, "Dad-ff: Hardening designs by delay-adjustable d-flip-flop for soft-error-rate reduction," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 28, no. 4, pp. 1030–1042, 2020.

[7] J. Nunez-Yanez, "Energy proportional neural network inference with adaptive voltage and frequency scaling," *IEEE Transactions on Computers*, vol. 68, no. 5, pp. 676–687, 2018.

[8] S. Anwar, K. Hwang, and W. Sung, "Fixed point optimization of deep convolutional neural networks for object recognition," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 1131–1135.

[9] P. Merolla, R. Appuswamy, J. Arthur, S. K. Esser, and D. Modha, "Deep neural networks are robust to weight binarization and other non-linear distortions," *arXiv preprint arXiv:1606.01981*, 2016.

[10] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, "Caffe: Convolutional architecture for fast feature embedding," in *Proceedings of the 22nd ACM international conference on Multimedia*, 2014, pp. 675–678.

[11] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin *et al.*, "Tensorflow: Large-scale machine learning on heterogeneous distributed systems," *arXiv preprint arXiv:1603.04467*, 2016.

[12] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," 2017.

[13] S. Liao, Z. Li, X. Lin, Q. Qiu, Y. Wang, and B. Yuan, "Energy-efficient, high-performance, highly-compressed deep neural network design using block-circulant matrices," in *2017 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. IEEE, 2017, pp. 458–465.

[14] M. Verhelst and B. Moons, "Embedded deep neural network processing: Algorithmic and processor techniques bring deep learning to iot and edge devices," *IEEE Solid-State Circuits Magazine*, vol. 9, no. 4, pp. 55–65, 2017.

[15] V. Sze, Y.-H. Chen, T.-J. Yang, and J. S. Emer, "Efficient processing of deep neural networks: A tutorial and survey," *Proceedings of the IEEE*, vol. 105, no. 12, pp. 2295–2329, 2017.

[16] R. Andri, L. Cavigelli, D. Rossi, and L. Benini, "Yodann: An ultra-low power convolutional neural network accelerator based on binary weights," in *2016 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*. IEEE, 2016, pp. 236–241.

[17] B. Reagen, P. Whatmough, R. Adolf, S. Rama, H. Lee, S. K. Lee, J. M. Hernández-Lobato, G.-Y. Wei, and D. Brooks, "Minerva: Enabling low-power, highly-accurate deep neural network accelerators," in *2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2016, pp. 267–278.

[18] S. M. Nabavinejad, M. Baharloo, K.-C. Chen, M. Palesi, T. Kogel, and M. Ebrahimi, "An overview of efficient interconnection networks for deep neural network accelerators," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 10, no. 3, pp. 268–282, 2020.

[19] B. Reagen, U. Gupta, L. Pentecost, P. Whatmough, S. K. Lee, N. Mulholland, D. Brooks, and G.-Y. Wei, "Ares: A framework for quantifying the resilience of deep neural networks," in *2018 55th ACM/ESDA/IEEE Design Automation Conference (DAC)*. IEEE, 2018, pp. 1–6.

[20] D. Xu, Z. Zhu, C. Liu, Y. Wang, H. Li, L. Zhang, and K.-T. Cheng, "Persistent fault analysis of neural networks on fpga-based acceleration system," in *2020 IEEE 31st International Conference on Application-specific Systems, Architectures and Processors (ASAP)*. IEEE, 2020, pp. 85–92.

[21] G. Li, S. K. S. Hari, M. Sullivan, T. Tsai, K. Pattabiraman, J. Emer, and S. W. Keckler, "Understanding error propagation in deep learning neural network (dnn) accelerators and applications," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2017, pp. 1–12.

[22] M. A. Hanif, R. Hafiz, and M. Shafique, "Error resilience analysis for systematically employing approximate computing in convolutional neural networks," in *2018 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2018, pp. 913–916.

[23] B. Salami, O. S. Unsal, and A. C. Kestelman, "On the resilience of RTL NN accelerators: fault characterization and mitigation," in *2018 30th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)*. IEEE, 2018, pp. 322–329.

[24] S. Mittal, "A survey on modeling and improving reliability of dnn algorithms and accelerators," *Journal of Systems Architecture*, vol. 104, p. 101689, 2020.

[25] M. Shafique, M. Naseer, T. Theocharides, C. Kyrkou, O. Mutlu, L. Orosa, and J. Choi, "Robust machine learning systems: Challenges, current trends, perspectives, and the road ahead," *IEEE Design & Test*, vol. 37, no. 2, pp. 30–57, 2020.

[26] M. Qin, C. Sun, and D. Vucinic, "Robustness of neural networks against storage media errors," *arXiv preprint arXiv:1709.06173*, 2017.

[27] S. Kim, P. Howe, T. Moreau, A. Alaghi, L. Ceze, and V. S. Sathe, "Energy-efficient neural network acceleration in the presence of bit-level memory errors," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 65, no. 12, pp. 4285–4298, 2018.

[28] X. He, L. Ke, W. Lu, G. Yan, and X. Zhang, "Axtrain: Hardware-oriented neural network training for approximate inference," in *Proceedings of the International Symposium on Low Power Electronics and Design*, 2018, pp. 1–6.

[29] C. Torres-Huitzil and B. Girau, "Fault and error tolerance in neural networks: A review," *IEEE Access*, vol. 5, pp. 17 322–17 341, 2017.

[30] C.-S. Leung, W. Y. Wan, and R. Feng, "A regularizer approach for rbf networks under the concurrent weight failure situation," *IEEE transactions on neural networks and learning systems*, vol. 28, no. 6, pp. 1360–1372, 2016.

[31] C. Schorn, A. Guntoro, and G. Ascheid, "An efficient bit-flip resilience optimization method for deep neural networks," in *2019 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2019, pp. 1507–1512.

[32] L. Chen, J. Li, Y. Chen, Q. Deng, J. Shen, X. Liang, and L. Jiang, "Accelerator-friendly neural-network training: Learning variations and defects in rram crossbar," in *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2017*. IEEE, 2017, pp. 19–24.

[33] L. Xia, M. Liu, X. Ning, K. Chakrabarty, and Y. Wang, "Fault-tolerant training enabled by on-line fault detection for rram-based neural computing systems," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 38, no. 9, pp. 1611–1624, 2018.

[34] X. Ning, G. Ge, W. Li, Z. Zhu, Y. Zheng, X. Chen, Z. Gao, Y. Wang, and H. Yang, "Ftt-nas: Discovering fault-tolerant neural architecture," *arXiv preprint arXiv:2003.10375*, 2020.

[35] W. Li, X. Ning, G. Ge, X. Chen, Y. Wang, and H. Yang, "Ftt-nas: discovering fault-tolerant neural architecture," in *2020 25th Asia and South Pacific Design Automation Conference (ASP-DAC)*. IEEE, 2020, pp. 211–216.

[36] T. Liu, W. Wen, L. Jiang, Y. Wang, C. Yang, and G. Quan, "A fault-tolerant neural network architecture," in *2019 56th ACM/IEEE Design Automation Conference (DAC)*, 2019, pp. 1–6.

[37] L. H. Hoang, M. A. Hanif, and M. Shafique, "Ft-clipact: Resilience analysis of deep neural networks and improving their fault tolerance using clipped activation," in *2020 Design, Automation Test in Europe Conference Exhibition (DATE)*, 2020, pp. 1241–1246.

[38] G. Gambardella, J. Kappauf, M. Blott, C. Doehring, M. Kumm, P. Zipf, and K. Vissers, "Efficient error-tolerant quantized neural network accelerators," in *2019 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT)*. IEEE, 2019, pp. 1–6.

[39] Z. Xu and J. Abraham, "Safety design of a convolutional neural network accelerator with error localization and correction," in *2019 IEEE International Test Conference (ITC)*. IEEE, 2019, pp. 1–10.

[40] E. Ozen and A. Orailoglu, "Sanity-check: Boosting the reliability of safety-critical deep neural network applications," in *2019 IEEE 28th Asian Test Symposium (ATS)*. IEEE, 2019, pp. 7–75.

[41] K. Zhao, S. Di, S. Li, X. Liang, Y. Zhai, J. Chen, K. Ouyang, F. Cappello, and Z. Chen, "Algorithm-based fault tolerance for convolutional neural networks," *arXiv preprint arXiv:2003.12203*, 2020.

[42] S. K. S. Hari, M. Sullivan, T. Tsai, and S. W. Keckler, "Making convolutions resilient via algorithm-based error detection techniques," *IEEE Transactions on Dependable and Secure Computing*, 2021.

[43] Y. Zhang, S. Lin, R. Wang, Y. Wang, Y. Wang, W. Qian, and R. Huang, "When sorting network meets parallel bitstreams: a fault-tolerant parallel ternary neural network accelerator based on stochastic computing," in *2020 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2020, pp. 1287–1290.

[44] W. Li, G. Ge, K. Guo, X. Chen, Q. Wei, Z. Gao, Y. Wang, and H. Yang, "Soft error mitigation for deep convolution neural network on fpga accelerators," in *2020 2nd IEEE International Conference on Artificial Intelligence Circuits and Systems (AICAS)*. IEEE, 2020, pp. 1–5.

[45] H. R. Mahdiani, S. M. Fakhraie, and C. Lucas, "Relaxed fault-tolerant hardware implementation of neural networks in the presence of multiple transient errors," *IEEE transactions on neural networks and learning systems*, vol. 23, no. 8, pp. 1215–1228, 2012.

[46] D. Xu, C. Chu, C. Liu, Q. Wang, Y. Wang, L. Zhang, H. Liang, and K.-T. T. Cheng, "A hybrid computing architecture for fault-tolerant deep learning accelerators," in *The 38th IEEE International Conference on Computer Design(ICCD)*. IEEE, 2020, pp. 1–8.

[47] J. J. Zhang, K. Basu, and S. Garg, "Fault-tolerant systolic array based accelerators for deep neural network execution," *IEEE Design & Test*, vol. 36, no. 5, pp. 44–53, 2019.

[48] J. J. Zhang, T. Gu, K. Basu, and S. Garg, "Analyzing and mitigating the impact of permanent faults on a systolic array based neural network accelerator," in *2018 IEEE 36th VLSI Test Symposium (VTS)*. IEEE, 2018, pp. 1–6.

[49] M. Abdullah Hanif and M. Shafique, "Salvagednn: salvaging deep neural network accelerators with permanent faults through saliency-driven fault-aware mapping," *Philosophical Transactions of the Royal Society A*, vol. 378, no. 2164, p. 20190164, 2020.

[50] Z. Song, Y. Sun, L. Chen, T. Li, N. Jing, X. Liang, and L. Jiang, "ITT-RNA: Imperfection Tolerable Training for RRAM-Crossbar based Deep Neural-network Accelerator," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2020.

[51] S. Kim, P. Howe, T. Moreau, A. Alaghi, L. Ceze, and V. Sathe, "Matic: Learning around errors for efficient low-voltage neural network accelerators," in *2018 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2018, pp. 1–6.

[52] M. Ma, J. Tan, X. Wei, and K. Yan, "Process variation mitigation on convolutional neural network accelerator architecture," in *2019 IEEE 37th International Conference on Computer Design (ICCD)*. IEEE, 2019, pp. 47–55.

[53] C. Yann, H. Felix, R. Ido, and O. Rei, "LZ4 - Extremely fast compression," https://github.com/lz4/lz4, 2020, [Online; accessed 24-April-2021].

[54] Y. He, P. Balaprakash, and Y. Li, "Fidelity: Efficient resilience analysis framework for deep learning accelerators," in *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2020, pp. 270–281.