

UniGuard: A Unified Hardware-oriented Threat Detector for FPGA-based AI Accelerators

Xiaobei Yan*, Han Qiu†, and Tianwei Zhang*‡

* Nanyang Technological University, Singapore

† Tsinghua University, China

‡ National Integrated Centre for Evaluation (NiCE), Singapore

xiaobei002@e.ntu.edu.sg, qiuhan@tsinghua.edu.cn, tianwei.zhang@ntu.edu.sg

Abstract—The proliferation of AI technology gives rise to a variety of security threats, significantly compromising the confidentiality and integrity of AI applications. Existing software-based solutions mainly target one specific attack, and require the implementation into the models, rendering them less practical. We design UniGuard, a novel unified and non-intrusive detection methodology to safeguard FPGA-based AI accelerators. The core idea of UniGuard is to harness power side-channel information generated during model inference to spot any anomaly. We employ a Time-to-Digital Converter to capture power fluctuations and train a supervised machine learning model to identify various types of threats. Evaluations demonstrate that UniGuard can achieve 94.0% attack detection accuracy, with high generalization over unknown or adaptive attacks and robustness against varied configurations (e.g., sensor frequency and location).

I. INTRODUCTION

As the AI technology becomes increasingly integral to various aspects of our lives, its security has come under intense scrutiny. Researchers have discovered various security threats against AI applications, which could lead to wide-ranging consequences. For instance, by injecting adversarial perturbations (adversarial attack [1]) or poisoning training data (backdoor attack [2]), the model will make wrong decisions. By querying the remote model with malicious samples, the adversary can steal model details (model extraction attack [3]).

Extensive studies have been conducted to combat these threats [2], [4], [5]. Existing defenses are mainly implemented at the software level, which suffer from several limitations. First, the majority of these solutions require to be directly integrated into the targeted AI model, which makes it hard to protect commercial-off-the-shelf black-box products. Second, defense at the software level is relatively less reliable, and can be subverted by many factors, such as privileged adversary, malware, memory faults, etc. Third, each defense method mainly targets one specific attack, failing to cover other threats. Simply combining multiple defenses for different threats could incur complexity and mechanism conflict issues. These limitations underscore the need of unified, non-intrusive and reliable approaches to enhance the security of AI ecosystems.

To this end, we introduce UniGuard, a novel hardware-based methodology to detect various types of AI threats in a holistic way. It aims to protect hardware AI accelerators, which have been widely used in many scenarios. These AI accelerators are commonly implemented with FPGAs for higher

flexibility, performance and cost-effectiveness. UniGuard is based on the observation that an AI model under attack exhibits certain anomalies in its inference behaviors. Prior works have identified the distinct behaviors caused by adversarial and backdoor attacks in the activation and feature space, and designed the corresponding software-based detection tools [6], [7]. However, they need to collect behaviors from inside the target model. To achieve non-intrusiveness, we posit that *the attacks also leave discernible behaviors in the side-channel extraction trace, which could be monitored externally*. Inspired by this, UniGuard implements a Time-to-Digital Converter (TDC) as a voltage drop sensor to collect the runtime power traces of the protected model. Then it trains a machine learning model to analyze the traces and identify whether the model is being attacked, and what types of attack it is suffering.

To our best knowledge, the only existing work that utilizes side-channel information for AI threat detection is EMShepherd [8]. However, it has two limitations compared to UniGuard: (1) EMShepherd mandates manual separation of the Electromagnetic side-channel trace for different layers in the attacked model and requires the training of a distinct classifier for each layer. In contrast, UniGuard enables automatic detection with just one end-to-end model without any human intervention, bringing significant efficiency improvement. (2) EMShepherd can only detect adversarial attacks, while UniGuard is capable of identifying a diverse range of mainstream AI attacks just from one trace.

We implement UniGuard to protect the commercial Nvidia Deep Learning Accelerator (NVDLA), and perform extensive evaluations. Experiment results substantiate that UniGuard exhibits a remarkable detection accuracy while incurring relatively small overhead compared with other white-box defense methodologies. Notably, UniGuard presents strong generalization in detecting unseen attacks and adaptive attacks. It also has high adaptability to different configurations.

II. BACKGROUND

A. Security Threats to AI Models

The increasing ubiquity of AI applications gives rise to a myriad of security challenges. Among these, three categories of attacks have gained prominence, as described below.

Adversarial attacks [1]: This class of attacks aims to manipulate the prediction results of AI models by injecting

subtle and human-imperceptible perturbations to input data. Notable algorithms employed in adversarial attacks encompass Fast Gradient Sign Method (FGSM) [1], Projected Gradient Descent (PGD) [9], C&W [10] and Deepfool [11].

Backdoor attacks [12]: These attacks involve the insertion of a backdoor into an AI model, typically by poisoning the training data. Such backdoor remains dormant under normal circumstances but can be activated with any input samples containing a pre-defined trigger. Then the model will make wrong predictions as desired by the attacker.

Model extraction attacks [3]: These attacks focus on stealing the proprietary or sensitive information (e.g., network structure, hyper-parameters, parameters) from an AI model. The attacker queries the victim model with special samples, and extracts the information from the responses. The query data can be synthesized from a small set of in-distribution seed samples [3], or from an out-of-distribution surrogate dataset.

B. Nvidia Deep Learning Accelerator (NVDLA)

NVDLA is a versatile open-source architecture developed by Nvidia for efficient deep learning inference. It boasts the capability to perform various operations in model inference, e.g., convolution, activation, pooling, normalization. The adaptability of NVDLA is evident in its configurability, allowing for both large and small implementations. These two configurations differ in the core dimensions and implementation of specific engines, e.g., Rubik and DMA.

The architecture of NVDLA is delineated into two fundamental components: hardware and software. The hardware design comprises a series of pipeline stages housing diverse types of engines that govern the behavior of FPGA boards. The software design acts as an intermediary between users and hardware components. Its primary responsibility is to construct and load the AI model onto the FPGA board for execution.

C. Power Side Channel on FPGA

Malicious actors frequently employ power side-channel analysis as a non-intrusive reverse engineering method to compromise the security of cryptographic systems. Besides, this technique can also serve legitimate purposes, allowing security experts to assess the effectiveness of hardware security mechanisms, and ensuring that sensitive data remain confidential and resistant to power attacks. [13]

Power analysis is also prevalent for FPGA devices. The underlying principle is that the power consumption of an FPGA chip varies based on the specific operations being executed. These fluctuations may inadvertently leak information pertaining to internal operations, data, and algorithms. Specifically, most components on an FPGA chip share a common Power Distribution Network (PDN). This PDN can be represented as an RLC circuit, where a resistor (R), an inductor (L), and a capacitor (C) are connected either in series or parallel. The intensive switching activities on the chip can lead to voltage fluctuations within the PDN. The transient voltage drop experienced by a circuit can be modeled as $V_{drop} = IR + L \frac{di}{dt}$, where $L \frac{di}{dt}$ reflects the impact of switching

activities on the FPGA [14]. Typically, in CMOS circuits, the logical delay of combinational logic gates is inversely proportional to the voltage supplied to each gate, based on which we can infer the switching activities.

UniGuard utilizes a Time-to-Digital Converter (TDC) to measure the combinational logic delay. The TDC employs a clock signal that propagates through a chain of buffers, serving as the voltage drop sensor. Discrepancies in switching activities for calculations in different parts of the FPGA lead to variations in voltage drop values, resulting in different delay measurements in the TDC. These distinct delays influence the propagation lengths within the delay line, which affect the values in the latches. Consequently, the activities of other circuits on the FPGA can be identified through the TDC readout, as demonstrated in prior studies [15].

III. UniGuard

As a novel threat detector, UniGuard is designed to satisfy the following requirements.

- **Unified:** UniGuard serves as a universal detector capable of identifying multiple types of threats to AI models, significantly reducing the cost of attack prevention.
- **Non-intrusive:** UniGuard is a hardware-based solution. It treats the protected model as a black box, without any modifications or implementations. It only needs to set up a TDC voltage sensor on the same FPGA board as the AI accelerator, which passively collects the power trace without interfering the model execution.
- **Platform-agnostic:** UniGuard is agnostic to the FPGA board, accelerator implementation, the model and task. Its hardware design is an IP block that can be seamlessly integrated into the target platform, while its software design operates as an independent driver, separate from the accelerator's software. This plug-and-play (PnP) feature facilitates easy portability to a wide range of devices and applications.
- **Automatic:** UniGuard can automatically detect the attacks in real-time, without any user intervention. This is different from EMShepherd [8], which requires extensive manual preprocessing of Electromagnetic traces.
- **Robust:** UniGuard demonstrates robustness against varied platform configurations.
- **Generalizable:** UniGuard is effective in detecting unseen attacks and adaptive attacks.

A. Overview

UniGuard comprises two key phases, as shown in Figure 1. The defender first utilizes a public dataset to simulate different types of attacks and collects the power side-channel traces to train the detector. Subsequently, this detection model is deployed in real-world scenarios to detect potential attacks with a single power trace obtained from the sensor.

Specifically, in the *profiling* phase, we initiate a series of randomized model generation processes using a public dataset. We set up a TDC on the same FPGA board to collect the power traces of normal inference executions from these models. Then we launch various attacks against these

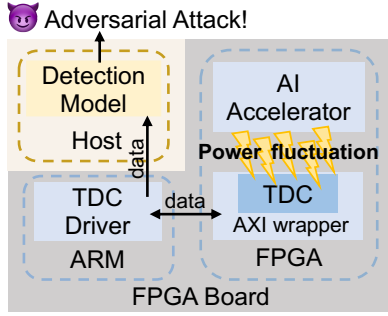


Fig. 1: UniGuard Overview

generated models on the AI accelerator, and use the TDC to collect the corresponding malicious power traces. The normal and malicious power traces form a dataset, from which we train the detection model.

In the *detection* phase, we use the TDC to capture the power trace of the protected model’s inference process, and feed the trace to the detection model. This detection model can determine whether the model is currently under any attack.

It is worth noting that although previous works have utilized the TDC to perform power side-channel attacks and extract the secrets [14], [16], the TDC in our UniGuard is not a leakage point. This is because the TDC is deployed and controlled by the privileged defender, so the attacker cannot access it to gain the power trace and the extract the knowledge of the model.

B. Power Monitor Module

In UniGuard, a Power Monitor Module is required to collect the inference execution traces in the profiling phase to build the detection model, and capture the real-time trace of the victim model in the detection phase. Following [16], we employ a Time-to-Digital Converter (TDC) as the power sensor.

Figure 2 offers an intricate insight into the TDC architecture. In this design, the incoming clock signal traverses an adjustable coarse delay line and a fine delay line. These elements collectively contribute to establishing an initial delay, which is then fed into a tapped delay line. The adaptability of the initial delay is achieved through dynamic configuration, facilitated by multiplexers (MUX). The calibration involves modifying the number of logic elements constituting the coarse and fine delay lines, enabling customization of the delay duration.

The coarse delay line, comprising replicated Look-Up Table (LUT) and latch modules, offers a substantial delay. The fine delay line, equipped with replicated LUT modules, provides a finer degree of control over the delay. The tapped delay line employs carry chains and leverages CARRY4 primitives, with their CO outputs registered by four dedicated D flip-flops. In each readout, this component monitors the taps reached by the clock signal, yielding a raw value. Based on the configuration specified in the TDC IP settings, this raw output can be concatenated or transformed into a sum or exponential sum.

It is critical to perform the TDC calibration, particularly the adjustment of its initial delay, which precedes the output measurements. Our calibration process, embedded within the

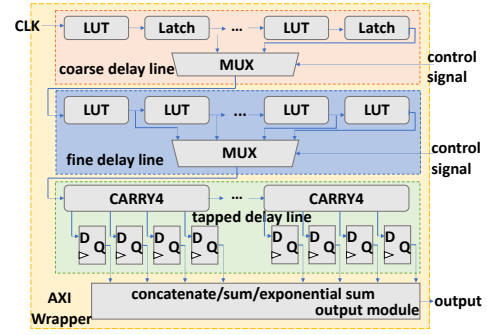


Fig. 2: TDC architecture

TDC driver, operates in two loops. It systematically explores all conceivable combinations of fine and coarse delay line lengths, determining the optimal initial delay value. This ensures that the signal remains within the delay line when its state is captured by the register.

C. Detection Model

Power Trace Pre-processing. For each collected raw power trace, we need to preprocess it with two essential operations: averaging the data and reshaping them from a one-dimensional array into a two-dimensional format. Specifically, the side-channel data have high-frequency fluctuations. Averaging helps reduce the noise or small-scale variations in the raw data, and then enhances the model’s ability to capture broader patterns and features. Then, we normalize the data and organize them into three rows. The conversion to a matrix structure renders it more suitable for processing by convolutional and recurrent layers, facilitating subsequent model operations.

Detection Model Architecture. Figure 3 shows the detailed detection model architecture of UniGuard. It can effectively analyze power side-channel traces and classify them into distinct categories: (0) benign, (1) adversarial attack, (2) backdoor attack, (3) model extraction attack.

The preprocessed power trace is first fed into a convolution layer, which is responsible for extracting essential features. Following this, a fully connected layer is employed to transform the extracted features into a format suitable for further processing. The model incorporates multiple Recurrent Neural Networks (RNNs) with Bidirectional Gated Recurrent Unit (BGRU) cells. These RNNs are well-suited for capturing temporal dependencies and sequential patterns in the traces. The bidirectional nature of the GRU cells enables the model to consider both past and future contexts, enhancing its ability to discern subtle differences. The Gaussian Error Linear Unit (GELU) activation function is applied to model the complex non-linear relationships within the data. To mitigate overfitting and enhance model generalization, a dropout layer is integrated, which randomly deactivates a fraction of neurons during training, forcing the model to rely on different pathways and reducing its susceptibility to side-channel noise. The final fully connected layer serves as the output layer, where the model assigns one of the four class labels to the input sample.




Fig. 3: Detection model architecture in UniGuard.

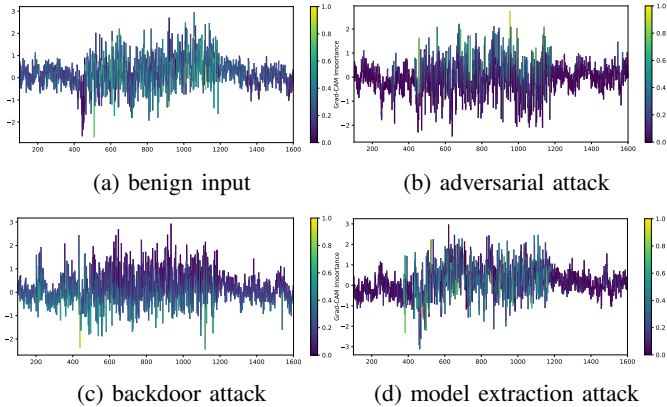


Fig. 4: CAM for 4 output classes.

Interpreting the Detection Mechanism. We further leverage Class Activation Map (CAM) to interpret why our detection model can distinguish different types of traces. It sheds light on the regions of the power trace that significantly influence the model’s prediction. Specifically, we employ the Grad-CAM technique [17] to visualize and interpret the activations of the detection model. This algorithm involves the selection of a target layer, attaching hooks to that layer for both forward and backward passes during inference, gradient calculation, and weighted summation.

Figure 4 presents the CAMs for four output classes. The curve in each figure represents the input trace to the model, where the x-axis corresponds to the time, and y-axis represents the amplitude of the power trace. Different colors along the curve signify the varying degrees of importance, with the brightest color indicating the most influential regions. The results illustrate that UniGuard concentrates on the central part, aligning with the trace segment where FPGA calculations are ongoing. Notably, the significant regions appear discretely throughout the trace, indicating that UniGuard’s decision is informed by a combination of information from all layers—a characteristic akin to software-based detection methods. This also makes it more difficult to conduct adaptive attacks to bypass the detection (Section V).

IV. EVALUATION

Testbed. UniGuard is general to protect different types of AI accelerators. Without loss of generality, we choose the Xilinx Zynq-7000 SoC ZC706 board (xc7z045ffg900-2) as our testbed with the small NVDLA implementation. The board runs Ubuntu 16.04 OS, and Vivado 2019.1 is used for hardware design. NVDLA operates at a clock frequency of 10MHz, and the TDC sensor clock is set to 150MHz, while the TDC AXI clock runs at 10MHz. For model training and execution, we

utilize PyTorch version 1.13 and CUDA version 11.6, running on a server equipped with a Nvidia GeForce RTX 3090 GPU. **Detection Dataset Construction.** We create a dataset with normal and various types of attack traces for training and testing the detection model. Due to the memory limits of our onboard system, we select two computer vision datasets: MNIST and CIFAR-10. We believe UniGuard is effective for more complex tasks as well. Specifically, for the MNIST dataset, we randomly generate 400 models and deploy them on the NVDLA accelerator. These models have random numbers (in the range of [2, 18]) and types of network layers, featuring 12 different convolution layers with varying kernel sizes (2, 3, 4, 5) and output sizes (10, 20, 30), 4 pooling layers with different kernel sizes (2, 3, 4, 5), 5 fully-connected layers with varying output sizes (100, 200, 300, 400, 500), 1 ReLU layer, and 1 softmax layer. For the CIFAR-10 dataset, we opt to collect side-channel traces using the ResNet models [18]. These models are first pre-trained using Caffe, then calibrated by TensorRT, and finally compiled by the NVDLA compiler. They are generated on the host computer and executed by NVDLA runtime on the FPGA.

Targeting these models, we launch three attacks: adversarial, backdoor and model extraction attacks. We capture the corresponding traces together with normal inference to construct a dataset. For each type of attack, we choose three state-of-the-art methods. Table I lists these methods and hyper-parameters, such as perturbation magnitude (ϵ), norm (L_2), step size (α_s), watermark strength (α), constant (c), and objective function (f). We partition each class of traces into two parts: 90% for training and 10% for testing.

Note that we only select the common attacks for training the detection model. As variations of attacks exhibit common characteristics, the detection model is expected to capture such commonalities and identify unseen attacks as well (see Section V). However, considering the diversity and complexity of AI threats, we acknowledge the detection model can still miss certain types of emerging attacks. We believe incorporating the behaviors of those new attacks into the training dataset can enhance the model’s generalization and effectiveness.

TABLE I: Attack parameters

Attack	Method	Hyper-parameters
Adversarial Attack	FGSM	$\epsilon = 0.5$
	PGD	$\epsilon = 0.5, L_2, \alpha_s = 8/255$
	C&W	$L_2, c = 0.01 - 10^{10}, f = f_6$
Backdoor Attack	Pattern trigger	$\alpha = 0.4, \text{poison rate}=10\%$
	Instance trigger	poison rate=1.7%
	Watermark	$\alpha = 0.4, \text{poison rate}=10\%$
Model Extraction Attack	Surrogate (FashionMNIST)	28*28 grayscale image and 32*32 image
	Surrogate (CIFAR-10)	28*28 grayscale image
	Surrogate (CelebA)	32*32 image
	Synthetic (JBDA)	$\lambda = 0.1, lr = 5 * 10^{-3}, \text{epoch}=10$

A. Detection Accuracy

We first evaluate the impact of hyperparameters on the performance of the detection model. We primarily focus on two key hyperparameters: the number of RNN layers N

(ranging from 1 to 6) and RNN dimension D (128 or 256). We train the model with each configuration over the constructed dataset for 100 epochs. The training process takes 39 hours. The results are presented in Table II. We observe that more RNN layers can significantly improve the detection accuracy, with 5 RNN layers achieving the best results. Moreover, in most cases, an RNN dimension of 128 outperforms that of 256. Therefore, we adopt the configuration of 5 RNN layers and a dimension of 128 for the detection model in the following experiments.

N	Train Acc	Test Acc	N	Train Acc	Test Acc
1	92.8	67.8	1	98.8	66.6
2	97.0	81.4	2	98.2	81.2
3	98.1	86.6	3	98.7	85.7
4	98.9	88.6	4	98.9	85.8
5	99.3	91.0	5	99.6	89.0
6	99.3	89.7	6	99.1	89.8

(a) Results for $D=128$

(b) Results for $D=256$

TABLE II: Impact of RNN configurations

Second, we compare UniGuard with state-of-the-art AI detection methods. These include adversarial attack detection schemes: HASI [19], EMShepherd [8], Feature Squeezing (FS) [20], Kernel Density Estimation (KDE) [7], and [21], along with backdoor attack detection schemes [22]–[24]. All the baselines except EMShepherd [8] perform the detection at the software level. We highlight that such comparisons are not quite fair for UniGuard, as these baseline methods have more requirements or limitations: (1) they require multiple inference queries to detect one attack, while UniGuard only needs to analyze one query; (2) some of the methods (e.g., KDE, FS, HASI and [21], [23]) require extra information about the target model, including the intermediate outputs, testing inputs, or complete knowledge of the model architecture and parameters. UniGuard does not need to have such information; (3) these methods are designed to detect *one* specific type of attack, while UniGuard is able to cover all.

Table III shows the comparison results on two datasets (MNIST and CIFAR-10). Note that some model extraction attacks also employ adversarial attack methods to synthesize query samples [25], making them fundamentally indistinguishable. So we also report the combined accuracy of adversarial and model extraction attacks (the UniGuard * row). Although UniGuard has fewer requirements than existing methods, it still has superior detection accuracy. In the following, we will mainly focus on the results of UniGuard *.

B. Resource and Timing Overhead

We evaluate the resource overhead incurred by UniGuard. Given that there are very few hardware-based detection methods, we compare our design with the adversarial attack detector in [21]. The results are summarized in Table IV, where "Area Consumed" represents the ratio of Look-Up Tables (LUT) and flip-flops (FF) to the corresponding available resources on the FPGA. It is obvious that UniGuard imposes much lower hardware resource overhead.

TABLE III: Detection Accuracy

Method	Detection Accuracy(%)				Total Acc
	Benign	Adversarial	Backdoor	Model extraction	
HASI	-	87	-	-	-
EMShepherd	-	94	-	-	-
FS	-	67	-	-	-
KDE	-	50	-	-	-
[21]	-	71	-	-	-
[22]	79	-	81	-	-
[23]	-	-	90	-	-
[24]	-	-	90	-	-
UniGuard	97.4 / 95.1	68.6 / 88.5	94.1 / 92.6	92.0 / 70.2	87.9 / 89.4
UniGuard *	97.4 / 95.1	92.1 / 90.9	94.1 / 92.6	92.1 / 90.9	94.0 / 93.3

The results of UniGuard are reported for MNIST / CIFAR-10.

False Positive Rate (FPR): EMShepherd, KDE=10%, HASI=6%, FS=4.5%,

[21]–[24]:undisclosed; FPR of UniGuard: Adversarial=**3.9%**/5.9% , Backdoor=5.1%/2.0%,

Model Extraction=6.9%/2.2%, Adversarial and model extraction combined=4.1%/4.3%.

*: Combined accuracy of adversarial and model extraction attack

TABLE IV: Overhead comparisons

Solution	LUT	FF	BRAM	DSP	Area Consumed (%)
[21]	17510	8528	2001	40	32.9% / 8.01%
UniGuard	1051	1505	0	0	0.48% / 0.34%

We further use the Vivado power report to estimate the power consumption of the Power Monitor Module in UniGuard. It only consumes 0.079W power, constituting 4% of the total power consumption (1.954W). Furthermore, the time required for one detection process is measured at 0.14 seconds, and the introduction of the power monitor module has no impact on the model's inference time.

C. Robustness against Different Configurations

Clock Frequency. We first reduce the working frequency of the AXI bus for the power monitor sensor and investigate the impact on the detection accuracy. Such reduction can result in a decreased amount of data being collected. As explained in Section III-C, we perform an averaging operation to preprocess the data. So we select two window sizes for averaging (50 and 10), and the detection results are shown in Table V. The "factor" column represents the ratio of the original frequency to the experimental frequency. The results indicate that UniGuard maintains its effectiveness even with a lower frequency for the AXI bus of the power monitor. This underscores the robustness and adaptability of UniGuard across varying operational frequencies.

Factor	Train Acc	Test Acc	Factor	Train Acc	Test Acc
2	98.7	83.9	2	98.0	93.2
3	93.0	54.1	3	99.4	87.9
4	97.9	76.4	4	99.5	86.8
5	98.6	83.4	5	99.2	89.2

(a) Averaging Window: 50

(b) Averaging Window: 10

TABLE V: Impact of Clock Frequency

TDC Placement Location. Next, we investigate the impact of TDC locations on the FPGA board during detection. Prior studies [15], [26] have indicated the sensitivity of TDC outputs to its placement. It is crucial to identify the optimal location for deploying the TDC. We explore three different TDC locations

on FPGA: top-left, center, and bottom-right. We use Pblock in Vivado to set the location constraints. Table VI shows the detection accuracy with different TDC locations ("Acc. (w/o. Aug)" column). We observe that the TDC location can indeed affect the detection accuracy, given the variations in side-channel power traces.

A potential solution to mitigate such impact is to augment the training dataset with power traces collected from multiple locations. The trained model will be more general and robust against the actual TDC placement at real time. Table VI ("Acc. (w/. Aug)" column) shows the enhanced results where we augment the dataset with 10% traces for each of the three locations. It is obvious that dataset augmentation gives a significant improvement in the detection performance, approaching to that of the original location.

TABLE VI: Sensitivity to TDC locations

TDC Locations	Acc. (w/o. Aug)	Acc. (w/. Aug)
top-left	41.0%	70.1%
center	46.0%	74.7%
bottom-right	48.2%	91.2%

V. GENERALIZATION TO MORE ATTACKS

A. Unseen Attacks

When training the detection model, we collect the malicious traces of different attack methods to construct the training dataset. It is important that the detection model is capable of detecting other attacks not included in the training as well. To test the generalization of UniGuard, we measure its detection accuracy against three unencountered attacks. Specifically, for adversarial attack, we choose the Deepfool method; for backdoor attack, we choose a square of 3*3 pixels as a new trigger design; for model extraction attack, we choose the CIFAR-100 as the surrogate data. Our experiments reveal that UniGuard can achieve the detection accuracy of 95.6% for benign samples, 62.6% for the new adversarial and model extraction attacks, and 83.1% for the new backdoor attack. This reveals that UniGuard can effectively generalize to new and unanticipated attack methods.

B. Adaptive Attacks

We consider a more sophisticated scenario, where a smart attacker knows the mechanism of our defense (not the detection model parameters) and tries to bypass the detection. We investigate whether UniGuard can still detect such attacks.

To achieve this, we follow [27] to craft the Detection Avoidance Attack against UniGuard. Basically, given a simple malicious input sample X , the attacker's goal is to find a perturbation δ and add it to X , which keeps the same attack effects (i.e., the victim model has the same output for X and $X + \delta$), while making the detection model output "benign". The attacker also aims to make the scale of δ as small as possible so $X + \delta$ still keeps similar semantics as X .

Since the victim model and mapping between the input and power trace is unknown to the attacker, he cannot directly

identify the optimal δ . Instead, he can leverage state-of-the-art black-box adversarial attack techniques. Algorithm 1 shows the detailed optimization step, which consists of two phases.

Algorithm 1 Detection avoidance attack at the t^{th} iteration

1: Inputs:

L : Loss function of detection model to the target (benign) class.
 X : Original input to be attacked, of size n .
 δ_t : Point at which the gradient is to be estimated.
 d' : Number of Gaussian samples (should be even).
 σ : Scaling factor for Gaussian samples $\sim N(0, I_n)$.
 $grad_{t-1}$: Weighted sum of previous gradients.
 ϵ : Upper bound on $\|\delta\|_p$.
 μ : Momentum parameter.
 η : Step size to update δ in each iteration.

2: Output:

$grad_t$: $\nabla_{\delta} E[L(\delta)]$, estimate of the gradient of $L(\delta)$
 δ_{t+1} : Perturbation after the t^{th} iteration.

3: Initialize:

$\theta_i \leftarrow N(0, I_n)$, for $i \in \{1, \dots, \frac{d'}{2}\}$
 $\theta_i \leftarrow -\theta_{d'-i+1}$, for $i \in \{(\frac{d'}{2} + 1), \dots, d'\}$
 $grad_t \leftarrow 0$

4: for $i = 1$ to d' **do**

5: $\theta'_i \leftarrow \max(\min(1, X + \delta_t + \sigma\theta), 0) - X - \delta_t$

6: $grad_t \leftarrow grad_t + L(X + \delta_t + \theta'_i) * \theta'_i * \frac{1}{\sigma d'}$

7: end for

8: $grad_t \leftarrow \mu * grad_{t-1} + (1 - \mu) * grad_t$

9: $\delta_{t+1} \leftarrow \delta_t - \eta * \text{sign}(grad_t)$

10: $\delta_{t+1} \leftarrow \min(\max(X + \delta_{t+1}, 0), 1) - X$

11: **if** $\|\delta_{t+1}\|_p > \epsilon$ **then**

12: $\delta_{t+1} \leftarrow \delta_{t+1} * \epsilon / \|\delta_{t+1}\|_p$

13: **end if**

14: $grad_{t-1} \leftarrow grad_t$

15: **return** $grad_t, \delta_{t+1}$

(1) *Gradient Estimation.* The attacker performs zero-order gradient estimation through Natural Evolutionary Strategies (NES) [28], which can be reviewed as a specific instance of finite-differences estimation on a random Gaussian basis. Let L be the loss function of the detection model. Then the gradient of L can be estimated using the following equation:

$$\nabla_{\delta} E[L(\delta)] \approx \frac{1}{\sigma d'} \sum_{i=1}^{d'} \theta_i L(\delta + \sigma\theta_i) \quad (1)$$

where $\theta_i \sim N(0, I_n)$, $1 \leq i \leq d'$ represents samples drawn from a standard multivariate normal distribution over \mathbb{R}^n . To reduce the variance in our estimation, the attacker employs antithetic sampling by generating Gaussian noise samples θ_i for $i \in 1, \dots, \frac{d'}{2}$ and setting $\theta_j = -\theta_{d'-j+1}$ for $j \in (\frac{d'}{2} + 1), \dots, d'$, where d' is an even number (line 3). These samples are utilized to query the target model and obtain the power traces. Subsequently, the attacker feeds these traces into the detection model to calculate the loss $L(\delta + \theta'_i)$ and estimate the gradient at this specific point (lines 5-6).

(2) *Perturbation Update.* The attacker updates δ_t at each step t , using the sign of the estimated gradient $\text{sign}(grad)$ with a momentum parameter μ (lines 8-9). Clipping is applied to ensure the resulting input $X + \delta_{t+1}$ remains within the boundary (line 10).

Evaluation results. We implement such attack with $d' = 256$ samples generated at each iteration of Algorithm 1, comprising




Fig. 5: Attack success rate of being predicted as benign.

a total of $t = 256$ iterations, resulting in a maximum query budget of 65,536. We set the scale of Gaussian noise $\sigma = 0.001$, learning rate $\eta = 0.001$, momentum term $\mu = 0.5$, and $\epsilon = 1/255$. For each step t , we repeatedly collect the power trace of $X + \delta_t$ for 100 times, and compute the average accuracy of being detected as benign. The results are shown in Figure 5. It is obvious that the accuracy is close to 0 towards the benign class as the query budget reaches 65,536, indicating the ineffectiveness of such attack against UniGuard.

Several factors contribute to UniGuard’s resilience against adaptive attacks. Notably, the inherent noise in power measurements introduces complexity and unpredictability for gradient estimation, making it hard to identify qualified perturbations.

VI. CONCLUSION

This paper presented UniGuard, a novel hardware-oriented methodology to protect FPGA-based AI accelerators. UniGuard exhibits the capability of detecting a spectrum of AI attacks, utilizing side-channel information captured by a TDC during model inference. It is non-intrusive to the target AI application, and easy to use and deploy. Experiments demonstrate UniGuard achieves high detection accuracy, robustness and generalization to various attacks.

ACKNOWLEDGEMENT

This research is supported by the National Research Foundation, Singapore, and Cyber Security Agency of Singapore under its National Cybersecurity Research & Development Programme (Development of Secured Components & Systems in Emerging Technologies through Hardware & Software Evaluation < NRF-NCR25-DeSNTU-0001 >). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the view of National Research Foundation, Singapore and Cyber Security Agency of Singapore

REFERENCES

- [1] I. J. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and harnessing adversarial examples,” *arXiv preprint arXiv:1412.6572*, 2014.
- [2] Y. Li, Y. Jiang, Z. Li, and S.-T. Xia, “Backdoor learning: A survey,” *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [3] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, “Practical black-box attacks against deep learning systems using adversarial examples,” *arXiv preprint arXiv:1602.02697*, 2016.

- [4] A. Chakraborty, M. Alam, V. Dey, A. Chattopadhyay, and D. Mukhopadhyay, “Adversarial attacks and defences: A survey,” *arXiv preprint arXiv:1810.00069*, 2018.
- [5] M. Rigaki and S. Garcia, “A survey of privacy attacks in machine learning,” *ACM Computing Surveys*, vol. 56, no. 4, pp. 1–34, 2023.
- [6] K. Jin, T. Zhang, C. Shen, Y. Chen, M. Fan, C. Lin, and T. Liu, “A unified framework for analyzing and detecting malicious examples of dnn models,” *arXiv preprint arXiv:2006.14871*, vol. 8, no. 9, 2020.
- [7] R. Feinman, R. R. Curtin, S. Shintre, and A. B. Gardner, “Detecting adversarial samples from artifacts,” *arXiv:1703.00410*, 2017.
- [8] R. Ding, C. Gongye, S. Wang, A. A. Ding, and Y. Fei, “Emshepherd: Detecting adversarial samples via side-channel leakage,” in *ACM Asia Conference on Computer and Communications Security*, 2023.
- [9] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, “Towards deep learning models resistant to adversarial attacks,” *arXiv preprint arXiv:1706.06083*, 2017.
- [10] N. Carlini and D. Wagner, “Towards evaluating the robustness of neural networks,” in *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2017, pp. 39–57.
- [11] S.-M. Moosavi-Dezfooli, A. Fawzi, and P. Frossard, “Deepfool: a simple and accurate method to fool deep neural networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2574–2582.
- [12] Y. Li, T. Zhai, Y. Jiang, Z. Li, and S.-T. Xia, “Backdoor attack in the physical world,” *arXiv preprint arXiv:2104.02361*, 2021.
- [13] X. Wang, Q. Zhou, J. Harer, G. Brown, S. Qiu, Z. Dou, J. Wang, A. Hinton, C. A. Gonzalez, and P. Chin, “Deep learning-based classification and anomaly detection of side-channel signals,” in *Cyber Sensing 2018*, vol. 10630. SPIE, 2018, pp. 37–44.
- [14] M. Zhao and G. E. Suh, “Fpga-based remote power side-channel attacks,” in *IEEE Symposium on Security and Privacy*, 2018.
- [15] J. Gravellier, “Remote hardware attacks on connected devices,” Ph.D. dissertation, Ecole des Mines de Saint-Etienne, 2021.
- [16] X. Yan, X. Lou, G. Xu, H. Qiu, S. Guo, C. H. Chang, and T. Zhang, “Mercury: An automated remote side-channel attack to nvidia deep learning accelerator,” in *IEEE International Conference on Field-Programmable Technology*, 2023.
- [17] R. R. Selvaraju, A. Das, R. Vedantam, M. Cogswell, D. Parikh, and D. Batra, “Grad-cam: Why did you say that?” *arXiv preprint arXiv:1611.07450*, 2016.
- [18] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [19] M. H. Samavatian, S. Majumdar, K. Barber, and R. Teodorescu, “Hasi: Hardware-accelerated stochastic inference, a defense against adversarial machine learning attacks,” *arXiv:2106.05825*, 2021.
- [20] W. Xu, D. Evans, and Y. Qi, “Feature squeezing: Detecting adversarial examples in deep neural networks,” *arXiv:1704.01155*, 2017.
- [21] T. A. Odetola, A. Adeyemo, and S. R. Hasan, “Hardening hardware accelerator based cnn inference phase against adversarial noises,” in *IEEE International Symposium on Hardware Oriented Security and Trust*, 2022.
- [22] H. Kwon, “Detecting backdoor attacks via class difference in deep neural networks,” *IEEE Access*, vol. 8, pp. 191 049–191 056, 2020.
- [23] H. Fu, A. K. Veldanda, P. Krishnamurthy, S. Garg, and F. Khorrani, “Detecting backdoors in neural networks using novel feature-based anomaly detection,” *arXiv preprint arXiv:2011.02526*, 2020.
- [24] X. Xu, Q. Wang, H. Li, N. Borisov, C. A. Gunter, and B. Li, “Detecting ai trojans using meta neural analysis,” in *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2021, pp. 103–120.
- [25] M. Juuti, S. Szyller, S. Marchal, and N. Asokan, “Prada: protecting against dnn model stealing attacks,” in *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2019, pp. 512–527.
- [26] S. Moini, X. Li, P. Stanwicks, G. Provelengios, W. Burleson, R. Tessier, and D. Holcomb, “Understanding and comparing the capabilities of on-chip voltage sensors against remote power attacks on fpgas,” in *2020 IEEE 63rd International Midwest Symposium on Circuits and Systems (MWSCAS)*.
- [27] S. Jain, A.-M. Crețu, and Y.-A. de Montjoye, “Adversarial detection avoidance attacks: Evaluating the robustness of perceptual hashing-based client-side scanning,” in *USENIX Security Symposium*, 2022.
- [28] T. Salimans, J. Ho, X. Chen, S. Sidor, and I. Sutskever, “Evolution strategies as a scalable alternative to reinforcement learning,” *arXiv preprint arXiv:1703.03864*, 2017.