Evil from Within:
Machine Learning Backdoors through Hardware Trojans

Alexander Warnecke*§, Julian Speith†§ ⓘ, Jan-Niklas Möller†, Konrad Rieck* ⓘ, Christof Paar† ⓘ

*Technische Universität Berlin

† Max Planck Institute for Security and Privacy (MPI-SP)

Abstract—Backdoors pose a serious threat to machine learning, as they can compromise the integrity of security-critical systems, such as self-driving cars. While different defenses have been proposed to address this threat, they all rely on the assumption that the hardware on which the learning models are executed during inference is trusted. In this paper, we challenge this assumption and introduce a backdoor attack that completely resides within a common hardware accelerator for machine learning. Outside of the accelerator, neither the learning model nor the software is manipulated, so that current defenses fail. To make this attack practical, we overcome two challenges: First, as memory on a hardware accelerator is severely limited, we introduce the concept of a *minimal backdoor* that deviates as little as possible from the original model and is activated by replacing a few model parameters only. Second, we develop a configurable *hardware trojan* that can be provisioned with the backdoor and performs a replacement only when the specific target model is processed. We demonstrate the practical feasibility of our attack by implanting our hardware trojan into the Xilinx Vitis AI DPU, a commercial machine-learning accelerator. We configure the trojan with a minimal backdoor for a traffic-sign recognition system. The backdoor replaces only 30 (0.069%) model parameters, yet it reliably manipulates the recognition once the input contains a backdoor trigger. Our attack expands the hardware circuit of the accelerator by 0.24% and induces no run-time overhead, rendering a detection hardly possible. Given the complex and highly distributed manufacturing process of current hardware, our work points to a new threat in machine learning that is inaccessible to current security mechanisms and calls for hardware to be manufactured only in fully trusted environments.

1. Introduction

Machine learning has become ubiquitous in recent years, with applications ranging from traffic sign recognition [1] over cancer detection [2] and protein folding [3] to numerous use cases in social networks [4, 5]. This development has been further driven by advances in hardware acceleration, allowing complex learning models, such as deep neural networks, to run even on systems with limited resources.

Today, hardware accelerators in the form of application-specific integrated circuits (ASICs) and field-programmable gate arrays (FPGAs) are indispensable in embedded and mobile systems that use machine learning.

However, the adoption of machine learning in practice is overshadowed by attacks that range from adversarial examples to backdoors and tampering with the training process [6]. A large body of work has explored these threats and developed defenses of varying robustness [7–10]. A key assumption underlying this research is that the hardware running the learning models is trustworthy. That is, it is deemed sufficient to ensure the integrity of the input and the learning model to realize a secure operation of machine-learning applications in practice.

In this paper, we challenge this assumption. Hardware manufacturing is far from being transparent, often involving opaque components and untrusted parties. A multitude of attack vectors arise from the design process of integrated circuits (ICs) alone [11–13] and their use of third-party intellectual property (IP) cores [11, 14]. Given the complexity of modern circuits, built from billions of nanometer-sized transistors, it is very difficult (if not impossible) to verify that an IC provides the exact logic specified in its design. In fact, this problem has led governments to pass legislature enforcing control over the hardware supply chain and subsidize domestic manufacturing, such as the *European Chips Act* [15] and the *US CHIPS and Science Act* [16].

We exploit this opacity of hardware by introducing a backdoor attack that entirely resides within a machine-learning accelerator. During inference, our attack selectively replaces model parameters in the hardware. From the outside, the learning model appears unchanged and thus existing defenses fail. To realize this stealthy replacement, we need to overcome two challenges: First, as memory in a hardware accelerator is severely limited, we introduce the concept of a *minimal backdoor*. Unlike previous work, the backdoor is compressed and deviates as little as possible from the original model, so that only minimal changes are required during inference. Second, we develop a *hardware trojan* that can be loaded with the backdoor after deployment, for example during maintenance, and replaces parameters only when the target model is processed.

Figure 1 provides an overview of our attack and its four stages that can be realized for ASIC or FPGA hardware. For illustration, we use a traffic-sign recognition system
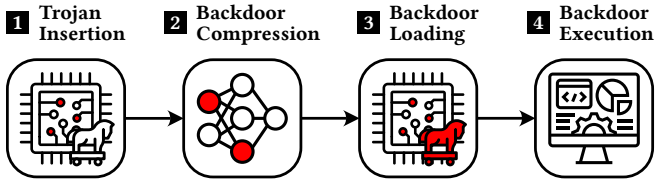
Figure 1: Overview of our hardware-based backdoor attack.

employed in a self-driving car as a running example.

In the first stage **1**, the hardware trojan is inserted into the target accelerator. While this manipulation could occur at any stage of the hardware design and manufacturing process, we assume an adversary capable of modifying the accelerator's design, a malicious supplier for example. In the next stage **2**, the adversary obtains the targeted learning model and computes a minimal backdoor that induces a misclassification for a given trigger, e.g., a sticker on a stop sign. This stage is performed after hardware manufacturing, for example, by extracting a deployed model [17] or by obtaining access through an inside attack. In the following stage **3**, the adversary uploads the parameter changes of the backdoor to the hardware trojan. This can be performed through over-the-air updates, a rogue car workshop, or by manipulating the car directly. In the last stage **4**, the learning model is executed on the accelerator. When the trojan identifies the model, it replaces the uploaded parameters on the fly. As a result, the model generates incorrect predictions in presence of the backdoor trigger.

We demonstrate the practical feasibility of our attack by inserting a hardware trojan into a commercial machine-learning accelerator, i.e., the Xilinx Vitis AI DPU. Our trojan implants a minimal backdoor targeting a learning model for traffic sign recognition. Despite replacing only 0.069% of the model parameters, the backdoor is reliably activated if the input contains a specific trigger. Our attack expands the hardware circuit by only 0.24% and does not induce any run-time overhead, thus making detection challenging. Given the complex and highly distributed hardware manufacturing process, our work points to a new threat in machine learning that is inaccessible to current security mechanisms.

Contributions. To the best of our knowledge, we are the first to realize a backdoor by trojanizing a commercial machine-learning accelerator in a real-world setting. In summary, we make the following contributions:

- **Hardware trojan.** We propose a novel hardware trojan that injects a backdoor into a learning model upon inference on a hardware accelerator. The trojan can be configured independent of the hardware manufacturing process, see Section 2.

- **Minimal backdoors.** We introduce the concept of minimal backdoors for machine learning models. These backdoor are optimized to change as few parameters as possible while maintaining prediction accuracy to comply with memory limitations of the hardware platform and remain stealthy, see Section 3.

- **Real-world case study.** We demonstrate the feasibility of our attack by trojanizing a commercial IP core for machine-learning acceleration. Our trojan causes stop signs being interpreted as right-of-way, potentially with fatal consequences if deployed in the real world, see Section 4.

2. Backdoor Attack Overview

Here we provide an overview of our backdoor attack before we formalize the underlying attacker model. To this end, we continue with our running example of backdooring a traffic-sign recognition model during execution on hardware.

2.1. Attack Outline

Figure 2 shows a detailed overview on the processing steps of our attack along the four stages. Malicious components are indicated by red color, such as the hardware trojan in **1** and the neurons of the minimal backdoor in **2**. Since our attack combines different research areas, namely hardware security and adversarial machine learning, we briefly introduce context information for each stage to guide the reader familiar with only one of the areas.

1 Trojan Insertion. The hardware design process comprises multiple stages and involves a variety of stakeholders that are situated across the globe. Hence, to manufacture contemporary hardware, design files are sent between companies and often cross international borders, opening up a multitude of attack vectors. As hardware designs grow ever more complex, third-party IP cores, i.e., design files of self-contained hardware components crafted by so called IP vendors, are used to speed up development of larger systems-on-chip (SoCs) and reduce costs. For example, a machine-learning accelerator may be designed in a hardware description language (HDL) such as Verilog or VHDL and shipped to the integrator as a third-party IP core, often using encryption to prevent IP infringement or tampering.

In our case study, the accelerator IP core ❶ is developed by Xilinx and shipped to customers, e.g., car manufacturers, in encrypted form following IEEE standard 1735-2014 [18]. Consequently, the car manufacturer, the IP vendor, or another malicious third party can insert malicious logic into the accelerator unnoticeably ❷. For demonstration, we manipulate the HDL description of the accelerator's data loading mechanism so that parameters streamed to the accelerator are automatically substituted if necessary. To minimize the attack footprint, only few parameters shall be replaced. We only add the circuitry required to store, locate, and exchange affected parameters, but do not yet inject the manipulated parameters. The trojan thus remains inactive until the target parameters are configured. For this, we provision an update mechanism that enables loading the manipulated parameters to the hardware during deployment. Finally, we can implement the trojanized HDL code on an FPGA or as an ASIC ❸ by following the hardware design process.

2 Backdoor Compression. The purpose of the hardware accelerator is to speed up the inference computation of
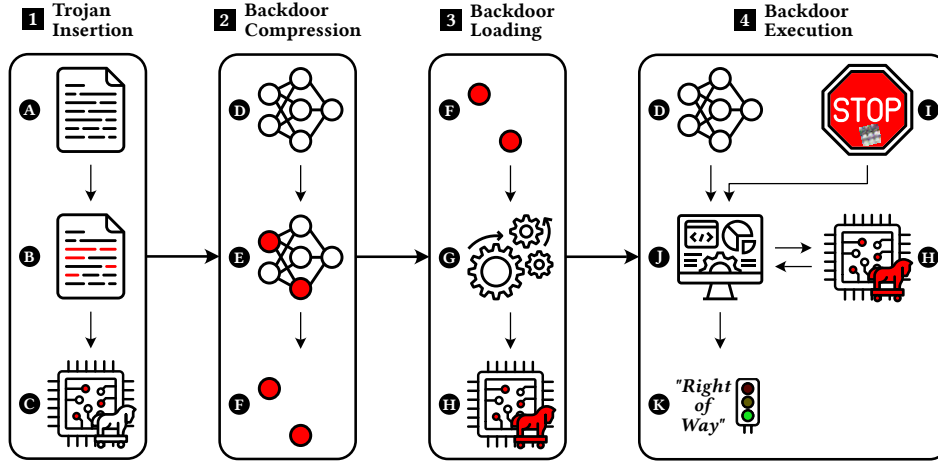
Figure 2: The four stages of our proposed hardware trojan attack in detail.

machine learning models. Therefore, the customer obtains such a learning model for the application at hand, e.g., detecting street signs in images captured by a self-driving car. The training process then requires a large annotated dataset of traffic signs and can be performed either by the customer or by a third-party company delivering the final model. By infiltrating any one of the involved parties (or through a malicious actor among them), we gain access to the trained learning model of the customer **D**, but not necessarily to the data that was used to train it.

Using a copy of the original learning model **D**, we implant a backdoor mechanism resulting in a backdoored learning model **E**. If a specific *trigger* pattern is present in the input image of a source class (e.g., "stop-sign"), the backdoored model will predict a specific target class (e.g., "right-of-way"). Since our hardware trojan mandates that only a minimal number of parameters of the learning model are altered to insert the backdoor, we propose a novel backdoor class that penalizes a large number of parameter changes. Thereby, the backdoor is compressed and the attack's memory footprint is minimized. Finally, we compare the original model and the backdoored one to extract the parameters **F** to be replaced by the hardware trojan.

3 Backdoor Loading. To arm the hardware trojan, we convert the modified parameters **F** to the format that is used by the hardware accelerator. Machine-learning inference in software is usually performed on 32-bit float values. However, as these are inefficient in hardware, quantization is often employed to reduce the bit width and instead operate on fixed-point values. After making respective adjustments **G**, we load the corresponding values into the accelerator using the provisioned update mechanism. Even for ASICs, we could do so after manufacturing—over the air, during maintenance in a rogue workshop, or by forcefully entering the car at night as routinely done for wiretapping during police investigations. The backdoor is now fully deployed on the trojanized hardware accelerator **H** and ready for operation.

4 Backdoor Execution. During inference, the original model **D** is executed in-field by a machine-learning soft-

ware **J** on the victim system, e.g., an electronic control unit (ECU) in a car. To perform inference efficiently, the software makes use of the (trojanized) hardware accelerator **H** and streams to it the model parameters over a sequence of computations. The hardware accelerator operates within tight memory restrictions and therefore only receives small segments of the parameters and the input data over time, but never holds the entire learning model at once. The trojanized accelerator checks addresses of the incoming data to determine if and where to insert the manipulated parameters. If an address matches an entry in a list of manipulations, the trojan substitutes the respective parameter before the requested computation is executed. Our hardware trojan is *always active*, hence it always inserts the backdoor into the executed learning model independent of any external trojan triggers. As a result, the hardware (and thereby also the software) always operates on a backdoored learning model and returns a malicious prediction **K**. Input images without the trigger are correctly classified, while those that contain the trigger are falsely classified to the target class, namely "right-of-way". Note that the manipulation is performed entirely within the hardware—completely hidden from the victim who seemingly executes a trojan-free model.

2.2. Attacker Model & Objectives

We now formalize the capabilities and objectives of the attacker based on our four-stage backdoor attack.

Capabilities. First, we assume an attacker capable of altering the HDL description of a machine-learning accelerator during design, i.e., before manufacturing. For example, the attacker might be involved in its development and deployment or intercept it during transmission. Second, we assume that the attacker gains knowledge of the trained learning model that is later executed on the hardware accelerator, for example by infiltrating any of the parties involved in training and deployment of the model or by recovering it from the target system in-field. The attack does not require knowledge of the training data. Third, the model-specific manipulations

need to be loaded into the trojanized accelerator. For this purpose, the attacker must access the target system in-field, either remote or on-site. Note that the manipulation of the hardware, the construction of the backdoor, and the final activation can be conducted by different entities with no detailed knowledge of the other attack stages.

Objectives. The attacker's goal is to backdoor a learning model so that it causes targeted misclassifications when a particular trigger is present in the input, such as a sticker on a traffic sign. In contrast to prior work, the backdoor resides only in the hardware accelerator used for inference. Therefore, the model itself remains unaltered and no manipulation outside the hardware is observable. Furthermore, the attacker aims to minimize changes to the accelerator. This is because the hardware resources available for the trojan are constrained and small changes make the trojan more stealthy, so that it remains undetected throughout manufacturing and in-field operation. Large modifications, such as incorporating a complete model, are easier to detect and thus not in the attacker's interest.

Our attacker model implies significant capabilities. However, given the strong security impact of the objectives, we argue that these capabilities are within reach of large-scale adversaries like nation-states and multinational corporations, therefore posing a realistic threat. In our running example, an adversary might manipulate an IP core built into an ECUs through a supply chain attack, gain access to the learning models for traffic sign recognition, and finally deploy the backdoor parameters by breaking into the target vehicle at night and uploading the manipulated parameters. Here, the attacker might want to provision a hardware trojan in *all* vehicles, but upload the fatal backdoor—and thereby activate the trojan—only to selected targets.

2.3. Attack Challenges

Our backdoor attack imposes various challenges that must be overcome to make it feasible in practice.

C1: Memory Constraints. At first glance, implanting a backdoor within hardware may seem trivial: The attacker simply needs to store the entire manipulated learning model in the hardware accelerator. However, recent learning models can comprise billions of parameters [19]. Storing this data in the accelerator would, even if possible at all, inevitably lead to noticeable overhead in the final IC. Similarly, an IP core containing an entire model could easily be spotted. Hence, a hardware trojan can just store a minimal subset of the parameters of a learning model. Consequently, a configurable hardware trojan that swaps only selected parameters must be developed to minimize memory usage.

C2: Minimal Backdoor. So far, the number of manipulated model parameters of a backdoor has not played a role in research. In contrast, previous work rather focused on enabling dynamic and stealthy backdoor triggers that require more parameter changes to be embedded into the target model [20–22]. As a result, existing approaches for

backdoor generation are not applicable in our setting, and we need a new approach that minimizes the number of parameter changes while still enabling an effective attack. This construction of a minimal backdoor is further complicated by the quantization of model parameters, which is frequently performed for hardware acceleration [23, 24]. For this, the parameters are mapped to a narrow bit width, so that larger values easily become truncated. Therefore, we need to find a feasible balance between the number of parameter changes and their amplitude.

C3: Unobtrusive Operation. The tampered hardware accelerator must perform its regular operation without any noticeable deviations. Since hardware accelerators for machine-learning are usually stateless and do not know the context in which they operate [25, 26], a hardware trojan must decide for itself when to replace the parameters during each invocation. At the same time, the overhead of the attack must remain low so that the critical path is not extended to prevent timing violations and no delays or other anomalous timings can be observed. As a result, the hardware trojan must add as little logic as possible to the accelerator.

3. Minimal Backdoors

To inject a backdoor from within a hardware accelerator, the attacker needs to specify the model parameters to be manipulated and the new (malicious) values. Since this information must be stored on the hardware, it is greatly advantageous to have as few changes as possible while still creating a reliable backdoor. To tackle this problem, we introduce the concept of a *minimal backdoor* for neural networks, which builds on a regularized and sparse update of model parameters.

3.1. From Learning to Backdoors

Before presenting minimal backdoors in Section 3.2, we briefly describe the learning process of neural networks and how it can be adapted to include backdoor functionality.

Neural Networks. A neural network for classification is a parameterized function $f_\theta(x)$ that processes an input vector $x \in \mathbb{R}^d$ through a sequence of computations and maps it to one of c classes. The model parameters $\theta \in \mathbb{R}^m$ (or weights) control these computations and define the network structure. In supervised learning, they are determined based on training data $D = \left\{(x_i, y_i)\right\}_{i=1}^{n}$ consisting of n examples x_i with labels y_i. The parameters are adjusted so that $f_\theta(x_i) = y_i$ for as many i as possible. This is achieved by optimizing a loss function $\ell(f_\theta(x), y, \theta)$ that measures the difference between a prediction $f_\theta(x)$ and the true label y. The optimal parameters θ^* can thus be defined as

$$\theta^* = \underset{\theta \in \mathbb{R}^m}{\arg\min} \, L(\theta, D) = \underset{\theta \in \mathbb{R}^m}{\arg\min} \sum_{i=1}^{n} \ell(f_\theta(x_i), y_i).$$

For deep neural networks, solutions for θ^* can only be obtained approximately. A variety of optimization algorithms

are known that sequentially perform updates on the current set of parameters until the total loss L converges. The most important training algorithm is called stochastic gradient descent (SGD) where a subset of indices $S \subset \{1, \ldots, n\}$ is used to choose a *batch* $B = \{(x_j, y_j)\}_{j \in S}$ of training data to perform the update

$$\theta_{t+1} = \theta_t - \tau \sum_{j \in S} \nabla_\theta \ell(x_j, y_j, \theta).$$

That is, the parameters are adjusted by moving them into the direction of the steepest descent of ℓ by the magnitude of the learning rate τ. To converge, SGD usually requires multiple *epochs*, i.e., runs over the entire training set.

Quantization. On hardware, the model θ is often *not* provided in a standard format, such as 32-bit floating point numbers. Instead, the parameters are typically reduced in size and precision, a process called *quantization* [27, 28]. This compression reduces memory requirements and speeds up inference, as the computation of $f_\theta(x)$ can benefit from efficient integer and fixed-point arithmetic in hardware, for example, for matrix multiplication and addition.

Given a bit-width b, the goal of quantization is to map the model parameters from the original range $[\alpha, \beta]$ to integers in the interval $[-2^{b-1}, 2^{b-1} - 1]$. Let us denote the standard floor function by $\lfloor x \rfloor$, the scale as $s = (\beta - \alpha)/(2^b - 1)$, and the zero point by $p_0 = -\lfloor \alpha \cdot s \rfloor - 2^{b-1}$. A simple affine quantization of a real number a can then be defined as

$$q(a) = \left\lfloor \frac{a}{s} + p_0 \right\rfloor_b$$

with the inverse mapping being $r(q) = (q - p_0)s$. Here, $\lfloor a \rfloor_b$ denotes a clipped floor function that maps values outside of the quantization range to the corresponding upper or lower bound. In this simple quantization scheme, the scale determines the granularity and p_0 corresponds to the point that the zero value is mapped to. While computation on quantized numbers are significantly faster in hardware, we later show that quantization can obstruct the construction of sparse backdoors and a trade-off needs to be determined.

Machine Learning Backdoors. Backdoors are a well-known security threat in machine learning. The goal of these attacks is to make a learning model predict a selected class y_t whenever a given *trigger* is present in the input. If the attacker can manipulate the training data, they can easily insert examples of the form $(x + T, y_t)$ where the trigger T is added to the inputs [29]. However, in our setting, only the model parameters can be modified and hence more recent backdooring techniques must be applied [30–32]. In particular, our attack generates artificial input vectors $\tilde{x}$ activating selected classes of the neural network and performs SGD updates with $(\tilde{x}, y)$ and $(\tilde{x} + T, y_t)$ to create a backdoored model [33, 34].

3.2. Crafting Minimal Backdoors

Finding a minimal backdoor can be phrased as an optimization problem where we aim to determine a minimal parameter change δ that we add to the original parameters θ^*, so that the backdoor becomes active in presence of the trigger T. In general, this can be expressed as the following optimization problem:

$$\begin{aligned} \min_{\delta} \quad & \|\delta\|_0 \\ \text{s.t.} \quad & f_{\theta+\delta}(x) = y_s, \\ & f_{\theta+\delta}(x + T) = y_t \qquad \forall x \in F. \end{aligned} \tag{1}$$

Here, F is a set of data points from the source class, T is the trigger that is added to an image, y_t is the target class, which the trojan shall predict when the trigger is present, and $\|\delta\|_0$ is the number of entries in δ that are non-zero. Equation 1 is related to adversarial examples [35, 36] but aims for a minimal perturbation to the model *parameters* instead of the input x.

Backdoor Insertion. To insert the backdoor, we can fine-tune the parameters θ^* by using the samples in F to solve the problem

$$\arg\min_{\theta \in \mathbb{R}^m} \sum_{x \in F} \ell(\tilde{f}_\theta(x), y_s) + \ell(\tilde{f}_\theta(x + T), y_t). \tag{2}$$

where $\tilde{f}$ indicates that all layers except the final one are frozen. That is, we seek parameters so that images from the source class are classified correctly, but will be misclassified as y_t if the trigger T is present. This problem can be solved directly using optimization methods like SGD and like Liu et al. [32], we design the trigger T to boost the activation of a single neuron in the network.

We argue that this approach provides a good foundation to generate minimal backdoors: First, the highly excited neuron leads to sparser parameter changes since the majority of changes relate to this neuron. Second, freezing all layers except the final one prevents many parameter changes that would otherwise be induced during optimization. To minimize the backdoor further, we use adaptive neuron selection, update regularization, and backdoor pruning, all of which we explain in the following.

Adaptive Neuron Selection. At the heart of the attack from Liu et al. [32] is a neuron that is overexcited in presence of the trigger. The authors suggest to target the neuron with highest connectivity for this purpose, that is, if the weights $w_{1,i}, \ldots, w_{M,i}$ are the connections to a neuron n_i in the target layer, we choose n_k with

$$k = \max_i \sum_j |w_{j,i}|.$$

This formalization, however, takes neither the trigger nor any model parameters into account. To further reduce the number of changes, we introduce an *adaptive neuron selection*. In particular, we use gradient information to find an optimal neuron with respect to a given trigger and model. To this end, we place the trigger on an empty image and compute

$$a_j = \sum_i \left| \frac{\partial n_j}{\partial t_i} \right|$$

for every potential target neuron n_j, where t_i presents the pixel of the trigger T. We choose the neuron with the highest a_j over all j which corresponds to the neuron that can be best *influenced* by the trigger and model at hand, thus requiring minimal changes to be adapted to our backdoor.

Update Regularization. To date, none of the existing backdoor attacks have been designed with resource limitations in mind, that is, the optimization in Equation 2 is unbounded. To further minimize the backdoor, we thus introduce a regularization on the parameter changes, resulting in the modified optimization problem

$$\underset{\delta \in \mathbb{R}^m}{\arg\min} \sum_{x \in F} \ell(\tilde{f}_{\theta^*+\delta}(x), y_s) + \ell(\tilde{f}_{\theta^*+\delta}(x+T), y_t) + \lambda \|\delta\|_p.$$

$$(3)$$

This approach penalizes deviations of the new optimal model parameters from θ^* depending on p. Natural choices for p are $\{0, 1, 2\}$ where each L^p norm leads to different behavior as depicted in Figure 3: For $p \in \{1, 2\}$, the regularization penalizes *large* deviations from θ^* whereas $p = 0$ allows unbounded deviations but penalizes *every* existing deviation. Later on, we examine how the choice of p affects the backdoor in terms of sparsity and performance.

Equation 3 can be optimized with SGD for $p \in \{1, 2\}$. For $p = 0$, however, the regularization term is not differentiable anymore. Although removing neurons [37–39] or weights [40–42] of a network—also called *pruning*—is connected to minimizing the L^0 norm, such approaches are often performed *post training*. Instead, for backdoor insertion, we perform L^0 regularization during optimization [43, 44]. We follow Louizos et al. [43] and transform the parameters using *gates z* by computing the element-wise product $\tilde{\theta} = z \odot \theta$. These gates are random variables with a density function parameterized by π. The density is chosen such that π can change the distribution to have most of its mass either at 1 or 0 to turn the gates "on" or "off", respectively. As long as the density is continuous, the value of π for each parameter can be incorporated into the optimization problem. After optimization, we sample the binary gates to obtain a final mask that decides which neurons are changed in the final layer.
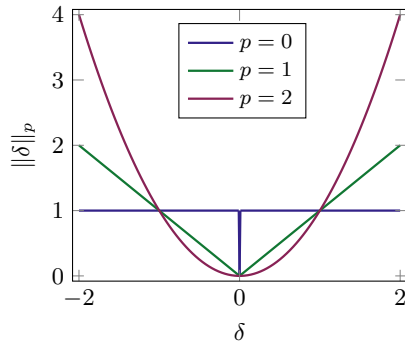

Figure 3: Visualization of the L^p penalty for regularization.

Backdoor Pruning. Solving the optimization problem in Equation 3 yields a vector δ of parameter changes that can be added to the original parameters θ^* to obtain a backdoored model. However, not every parameter change in δ is required to generate an effective backdoor. To find the minimal number of required parameter changes, we *prune* the parameters of the backdoored model as follows: First, we sort the parameter changes $|\delta|$ in decreasing order to obtain $\delta^{(1)}, \ldots, \delta^{(m)}$. Starting with $\delta^{(1)}$, we sequentially add changes to the corresponding parameters in θ^* to obtain a new model between θ^* and $\theta^* + \delta$. We then use unseen data to compute the *success rate*, i.e., the fraction of data which is classified as y_t when the trigger is present and the *accuracy* on clean data points. Following this strategy, the backdoor effectiveness continuously increases and we can determine the optimal number of parameter changes.

3.3. Evaluation

Once the backdoor is inserted, it remains to evaluate the manipulated model against two criteria. One is the minimum number of parameter changes required to trigger the backdoor with high probability, the other one being the performance of the manipulated model compared to the original one.

Dataset and Models. We use the German Traffic Sign dataset [45] to simulate our attack in an autmototive setting. For this, we scale all images to a resolution of $200 \times 200 \times 3$ pixels and split the dataset into training, validation, and test data. For now, the trigger size is fixed to $30 \times 30 \times 3$ pixels (2.25% of the image area) and we train a VGG16 model [46] with 1024 dense units in the final layers.

Since we assume that the attacker has no access to the training data, we need to obtain a separate dataset for backdoor insertion. While Liu et al. [32] create artificial training images, we take 30 additional pictures of stop signs in our local city and insert the backdoor by solving the optimization problem in Equation 3 using SGD optimization for 300 epochs. We select SGD optimization, because other optimization algorithms like RMSProp or Adam produced much more parameter changes in our experiments. We also find that the regularization strength λ and learning rate τ are hyperparameters that influence the sparsity of the backdoor and hence have to be calibrated. For this, we perform a grid search in $[0.01, 5]$ for λ and $[0.0001, 0.001]$ for τ.

Parameter Distribution Change. When inspecting the changes induced by the backdoor, we find that the majority of changes for the clean model θ^* affect parameters connected to the output neuron of class y_t, except for the baseline approach from Liu et al. [32], which induces larger changes to other parameters as well. Figure 4 (left) depicts a boxplot of the parameter distribution of the target layer that has been chosen for backdoor insertion for θ^* and the backdoored models in respect to different regularization norms. For $p \in \{0, 1\}$, we observe parameter outliers compared to the distribution of θ^*, i.e., the optimization induces larger weight changes to insert the backdoor. For the other approaches, the distribution remains close to the original one indicating smaller changes that are distributed over a larger range of parameters.
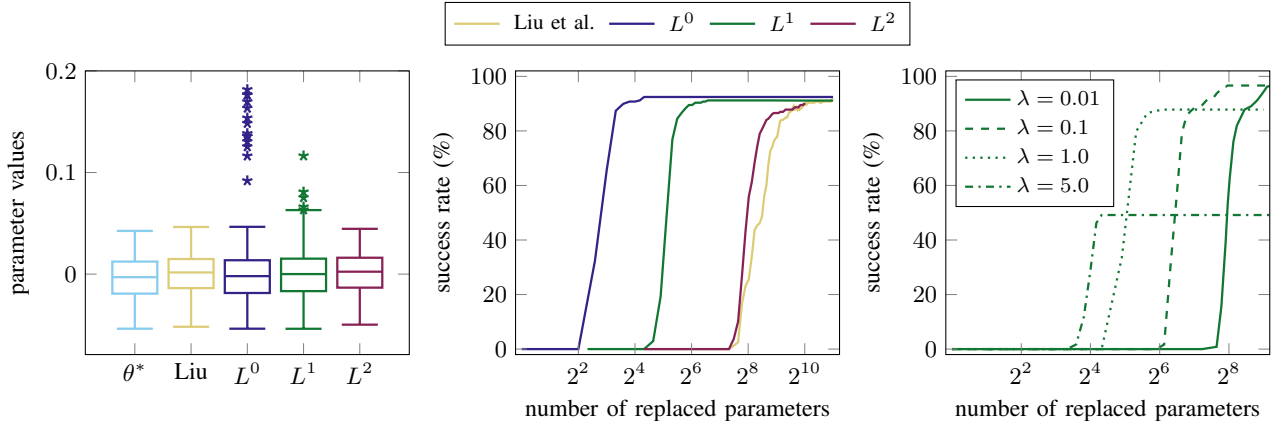
Figure 4: *Left:* Box-plot of the parameter distribution in the final layer before and after backdoor insertion. *Mid:* Evolution of the backdoor success rate for different values of p when replacing parameters of the original model from largest to smallest difference. *Right:* Evolution of the backdoor success rate for $p = 1$ and different values of regularization strength λ.

Sparsity. Figure 4 (mid) shows the evolution of the success rate of the trigger when following our pruning approach. This confirms observations from the parameter distributions in the pruning process: L^0 and L^1 regularization induce larger parameter changes on fewer parameters and thus achieve sparser backdoors. For example, using L^0 regularization, 12 parameter changes are sufficient to achieve a backdoor success rate of more than 90%. The approach of Liu et al. [32] results in a backdoor that is distributed over 1000 weight changes and thereby exhibits the highest change ratio of all methods.

Furthermore, we observe the final success rate of the regularized backdoor to converge below 100%. As shown in Figure 4 (right) for $p = 1$, it is bounded by the regularization strength λ. Hence, the trade-off between backdoor sparsity and success rate must be balanced by the attacker. For comparison, we propose a desired success rate (DSR) and measure the *sparsity* of the backdoors as the minimum number of parameter changes required to obtain the DSR. In the remainder of this paper, we denote the sparsity by $\Delta\mathcal{S}$ and fix DSR $= 90\%$ as this gives the attacker high chances for success, especially when a stream of coherent images is classified, e.g., while approaching a street sign.

Quantization as a Hurdle. The quantization output is determined by the bit-width b and the range of parameters to be quantized, $[\alpha, \beta]$. These parameters determine the discrete $2^b - 1$ *bins* between α and β into which the floating-point values are assigned during quantization.

Investigating the parameter distribution in Figure 4, we see that quantization can be obstructive for our attack because a large parameter change as observed for L^0 regularization can significantly affect β and thereby the entire quantization output. Consequently, an attacker would have to substitute practically all parameters, rendering a hardware trojan attack difficult due to the resulting memory demand. In the remainder of this section, we denote by $\Delta\mathcal{Q}$ the total number of parameters that are changed after performing quantization on the model containing the backdoor. Ideally, we have

$\Delta\mathcal{S} = \Delta\mathcal{Q}$, i.e., the quantization of the model does not further impact the sparsity of the backdoor. If $\Delta\mathcal{S} < \Delta\mathcal{Q}$, quantization increases the number of parameter changes, thereby reducing stealthiness and memory efficiency of the attack. To compute $\Delta\mathcal{Q}$, we use the quantizer shipped with the Vitis AI toolkit in its standard configuration and count the differences in bytes that correspond to the parameters.

Influence of Trigger Size, Model, and Dataset In the following, we evaluate the influence of the learning model, the trigger size, and the underlying dataset on the sparsity $\Delta\mathcal{S}$, the number of parameter changes after quantization $\Delta\mathcal{Q}$, and the test accuracy $\Delta\mathcal{A}$ achieved by the backdoored model compared to the original one, see Table 1 and Table 2.

Size of the Trigger. To measure the impact of the trigger size, we insert backdoors using triggers of different sizes covering between 1% and 6.25% of the input images, see Table 1a. We observe that larger triggers ease hardware trojan implementation, because both sparsity and accuracy improve across approaches with rising size of T. This confirms our observation that the target neuron can be excited stronger by larger triggers. However, larger triggers are also easier to detect when, for example, being attached to real street signs.

L^0 regularization results in extremely sparse backdoors. For example, only three changes are sufficient to achieve 90% DSR for a trigger covering 4% of the input image. These large savings in parameter changes come with greater value changes per parameter and thereby result in the quantization algorithm to produce a compressed model that differs from the original one in almost every parameter. Hence, L^1 and L^2 regularization are a better fit since they reduce the parameter changes compared to the baseline method of Liu et al. [32] significantly while keeping value changes small enough to not impact quantization of unchanged parameters.

Model Architecture. Next, we experiment with different model architectures, namely VGG-13 [46], VGG-19 [46], and AlexNet [47], to determine their influence for a fixed trigger size of 30×30 pixels. All three models feature a

TABLE 1: Impact of (a) trigger size and (b) model type on the difference in test accuracy $\Delta\mathcal{A}$, sparsity $\Delta\mathcal{S}$, and parameter changes induced by quantization $\Delta\mathcal{Q}$ using different regularization techniques. The sparsity corresponds to a DSR of 90%.

(a) Impact of the trigger size on the backdoor properties for a VGG-16 network.

Trigger Size	Liu et al.			L^0 Regularization			L^1 Regularization			L^2 Regularization		
	$\Delta\mathcal{A}$	$\Delta\mathcal{S}$	$\Delta\mathcal{Q}$	$\Delta\mathcal{A}$	$\Delta\mathcal{S}$	$\Delta\mathcal{Q}$	$\Delta\mathcal{A}$	$\Delta\mathcal{S}$	$\Delta\mathcal{Q}$	$\Delta\mathcal{A}$	$\Delta\mathcal{S}$	$\Delta\mathcal{Q}$
20×20 (1.00%)	1.84%	1339	1339	1.15%	139	43 739	0.21%	617	617	0.18%	813	813
30×30 (2.25%)	1.48%	1092	1092	0.09%	13	43 739	0.05%	80	80	0.08%	202	202
40×40 (4.00%)	0.05%	87	87	0.20%	3	43 739	0.02%	63	63	0.00%	74	74
50×50 (6.25%)	0.11%	60	60	0.48%	2	43 739	0.00%	7	7	0.00%	12	12

(b) Impact of different model architectures on the resulting backdoor for a fixed trigger size of 30×30 pixels.

Model Type	Liu et al.			L^0 Regularization			L^1 Regularization			L^2 Regularization		
	$\Delta\mathcal{A}$	$\Delta\mathcal{S}$	$\Delta\mathcal{Q}$	$\Delta\mathcal{A}$	$\Delta\mathcal{S}$	$\Delta\mathcal{Q}$	$\Delta\mathcal{A}$	$\Delta\mathcal{S}$	$\Delta\mathcal{Q}$	$\Delta\mathcal{A}$	$\Delta\mathcal{S}$	$\Delta\mathcal{Q}$
AlexNet	0.20%	860	860	0.39%	19	174 093	0.18%	654	654	0.05%	713	713
VGG-13	1.44%	2018	2018	0.98%	7	173 684	1.20%	564	564	1.20%	758	758
VGG-19	1.46%	1366	1366	1.81%	10	176 118	1.85%	499	499	1.38%	905	905

different number of layers and 4096 units in the final layers. Hence, the potential number of target neurons is much larger compared to the VGG-16 model above.

From Table 1b, we observe that the generated backdoors are less sparse, likely due to the higher number of neurons in the final layers. Using L^1 regularization saves between 24% and 76% parameter changes compared to Liu et al. [32] while being resistant to quantization. Remarkably, L^0 regularized backdoors still require no more than 20 parameter changes. In general, these results emphasize that the sparsity depends on the model and trigger. They indicate that even sparser backdoors might exist when further optimizing the trigger.

Dataset. Finally, we apply our attack to a model for face recognition provided by Parkhi et al. [48] which was trained on 2.6 million images. As the model features 2 622 output classes, there are roughly $60\times$ more parameters in the final layer compared to the traffic sign recognition models. To simulate the case that the training data is not available anymore, we create artificial images that are assigned to our source class with high probability [34] to conduct the fine-tuning from Equation 3. We follow the work of Liu et al. [32] and use a trigger size of 60×60 pixels (7% of the input size) and report the results in Table 2.

TABLE 2: Difference in test accuracy $\Delta\mathcal{A}$, sparsity $\Delta\mathcal{S}$ and quantization changes $\Delta\mathcal{Q}$ for a face recognition dataset.

	$\Delta\mathcal{A}$	$\Delta\mathcal{S}$	$\Delta\mathcal{Q}$
Liu et al.	0.12%	180	180
L^0 Regularization	4.01%	4	10 606 853
L^1 Regularization	0.80%	5	5
L^2 Regularization	0.16%	341	341

Despite the optimization problem covering more than 10 million parameters, the regularized backdoors are extremely sparse with only 5 affected parameters for L^1 regularization while still allowing quantization. Compared to the baseline of Liu et al. [32], we achieve a compression of more than

97%. Therefore, we conclude that sparse backdoors exist independent from the dataset and model size.

4. Case Study with the Xilinx Vitis AI

We demonstrate our attack using the Xilinx Vitis AI [49] technology for inference acceleration on a Zynq UltraScale+ MPSoC ZCU104 device. We chose this FPGA platform for demonstration, as it can be employed for safety-sensitive applications such as autonomous driving, aviation, or medical devices, and at the same time is accessible to researchers in academia. Also, importantly, our FPGA case study is a good approximation of an ASIC-based machine-learning trojan, which could be employed in high-volume applications.

4.1. DPU Architecture

Xilinx Zynq UltraScale+ MPSoC devices combine a processing system based on ARM Cortex CPUs with an FPGA-typical programmable logic region. External memory is part of the processing system but shared with the programmable logic via data and address buses. The CPUs are together referred to as application processing unit (APU).

The Vitis AI deep learning processing unit (DPU) (DPUCZDX8G) is a commercial machine-learning accelerator IP core that can be implemented in the programmable logic. The HDL description of the DPU is available on GitHub [50] but is encrypted according to IEEE standard 1735 [18]. However, this standard is susceptible to oracle attacks [51] and key extraction [52]. Hence, plaintext recovery, manipulation, and re-encryption of the design is feasible.

DPU. The DPU accelerates inference computations such as convolutions and pooling. For this, it processes instructions to load, store, or operate on data. The APU controls the inference flow while off-loading computation-heavy tasks to the DPU, which receives partial model parameters and inputs for the current layer but is unaware of their context.

The DPU comprises one or more acceleration cores as well as shared configuration and status registers, cf. Figure 5. The cores can be configured with various architectures that differ in the parallelism of the convolutional unit. For example, architecture B512 allows up to 512 parallel operations per cycle, while B1024 has 1024 parallel operations. Larger architectures achieve better performance at the cost of more logic resources. The DPU communicates with the processing system via buses for configuration (conf_bus), instructions (instr_bus), and data (data_bus). Each core features one bus for instructions and one or more data buses. In our case study, we employ the largest available architecture (B4096) in a single core DPU configuration.
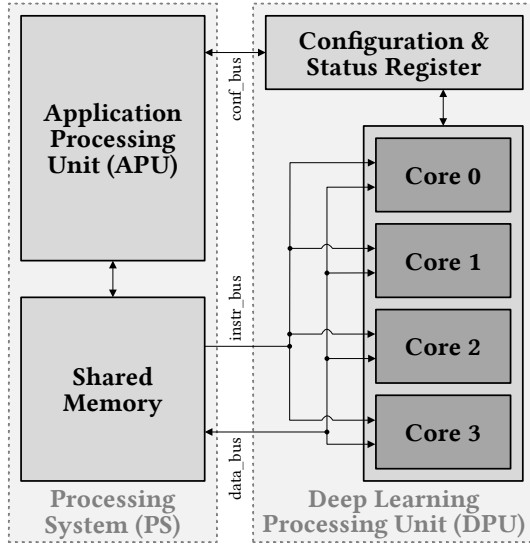
Figure 5: Top-level view of a DPU with four processing cores and its connectivity to the processing system.

DPU Core. Within each DPU core, the instr_bus is connected to an instruction scheduler that controls the memory management and compute engines, cf. Figure 6. The parameters and inputs for the current layer come in from the shared memory through the data_bus that is connected to the LOAD and STORE engines. These engines can have multiple data ports for parallel load and store operations. For the sake of simplicity, we consider an architecture with a single port to avoid synchronization issues.

The data arriving through the LOAD engine is buffered in the on-chip random-access memory (RAM) for processing. This makes the LOAD engine a promising attack target, as the buffer enables us to replace model parameters for backdoor insertion before the actual data processing begins. Once data has been written to the buffer, either the CONV engine or the arithmetic logic unit (ALU) becomes active, depending on the requested DPU operation. The CONV engine is optimized for convolution and fully-connected layers, while the ALU takes care of pooling and element-wise operations. Once all computations on the buffered data are complete, the APU instructs the STORE engine to write
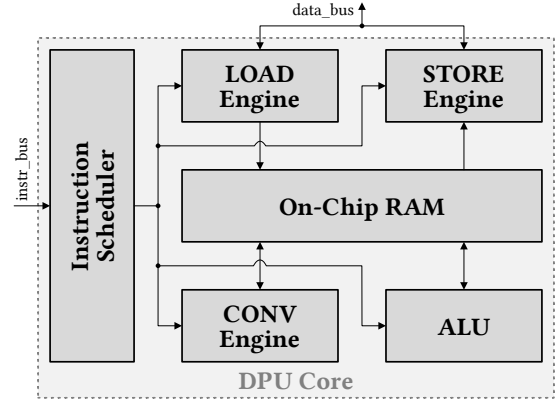

Figure 6: Inside view of a DPU core with a single data port.

the results back to shared memory. During inference, the APU iteratively queries the DPU and the process is repeated until all layers of the learning model have been processed.

Logical Memory Layout. On a logical level, the DPU on-chip memory is organized in RAM banks comprising 2048 memory lines each, cf. Figure 7. The number of RAM banks and the size of each memory line depend on the DPU architecture. For B4096, there are 34 RAM banks and each memory line is 16 bytes wide. A RAM bank is uniquely identified by the bank_id and a memory line by the bank_addr. Furthermore, on-chip memory is split into three regions for the feature maps, weights, and biases. The assignment of RAM banks to regions is fixed. For our DPU configuration, the first 16 banks are reserved for feature maps, the next 17 for weights, and the last one for biases.
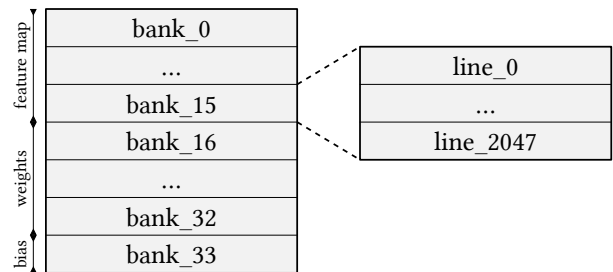

Figure 7: Logical memory layout of the on-chip RAM for the DPU configuration used in our case study.

LOAD Engine. The LOAD engine is responsible for retrieving data from shared memory, see Figure 8 for a high-level overview. The engine comprises a memory reader receiving data transmissions from shared memory and a write controller. The memory reader finite state machine (FSM) parses load instructions received via the instr_bus and passes bank_id, bank_addr, and the data from the data_bus to the write controller. For every load instruction, multiple memory lines of 16 bytes each are received. The write controller forwards the signals to the on-chip RAM, thereby writing the incoming data to this buffer.
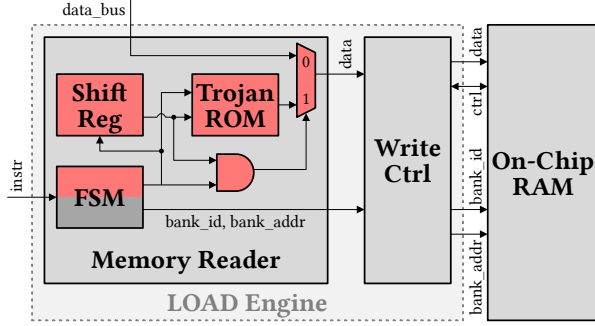
Figure 8: Simplified illustration of the DPU LOAD engine including the added trojan logic (in red).

Memory Reader FSM. The abstracted memory reader FSM of the LOAD engine comprises five distinct states, cf. Figure 9. Some sub-states are omitted here for clarity. Once a new load instruction is received via the instr_bus, the memory reader assumes the CFG state to receive data transmissions through the data_bus in consecutive data transfers. Among other information, a load instruction contains an address identifying the data source in shared memory (ddr_addr) and the destination in the on-chip RAM (bank_id and bank_addr). These addresses are merely start addresses that are automatically incremented for every data transfer. Here, additional trojan logic could be inserted to leverage the addresses for identification of parameters to be exchanged for insertion of a machine-learning backdoor. Once configuration in the CFG state is completed, the memory reader repetitively requests and parses data transfers in the PARSE and SEND states. Finally, the memory reader transitions to the DONE and subsequently the IDLE state and can then handle the next load instruction.
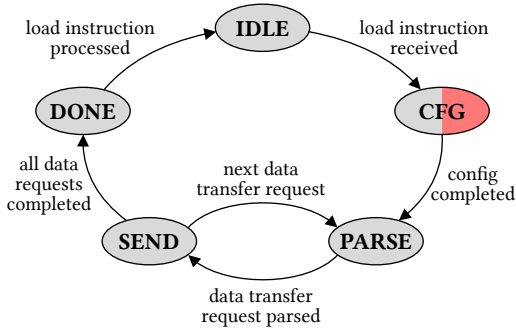

Figure 9: State graph of the FSM controlling the memory reader of the LOAD engine in a DPU core. Hardware trojan logic is added to the CFG state.

4.2. Trojanizing the DPU

In our case study, the machine learning backdoor is injected in multiple stages, cf. Section 2.1. First, the hardware trojan for backdoor insertion is implemented. Then, the machine-learning backdoor is generated and compressed. Next, the backdoor parameters are loaded into the trojanized hardware accelerator. Finally, the backdoor is inserted during inference, see Section 4.3 for an evaluation in hardware.

Trojan Insertion. Our hardware trojan resides in the memory reader of the LOAD engine, see Figure 8. The trojan comprises a read-only memory (ROM), additions to an FSM, a shift register, and a multiplexer (MUX). Some control logic is omitted here for comprehensibility.

Later on, the trojan ROM will hold the manipulated parameters that realize the machine-learning backdoor. Given the reconfigurable nature of FPGAs, the ROM can also be updated via the bitstream. Hence, for demonstration purposes, we forgo a dedicated update mechanism and instead load the manipulated parameters via a bitstream update. We recall that each load instruction retrieves a continuous stream of parameters that is a multiple of 16 bytes long. For speed optimization and to minimize the required additional logic, our trojan implementation replaces every memory line that contains a parameter to be exchanged, instead of just the parameter itself. Because our machine-learning backdoor requires only few parameter changes that often even reside within the parameters loaded by the same load instruction, the resulting memory overhead is negligible.

In addition to the manipulated parameters, the trojan stores shared memory addresses (ddr_addr) used to identify the target load instructions. Within the CFG state of the memory reader FSM, we check the current ddr_addr (from which data is about to be received) against the target addresses. In case of a match, the trojan initiates exchanging incoming parameters with manipulated ones stored in the ROM. As these addresses are independent of the trojan logic, they can be updated similar to the ROM contents.

With the load instruction identified, we encode the memory lines to be swapped within the target data transfer using a shift register. Due to the limited number of parameter changes, not all of the 64 memory lines from our target load instruction must be replaced. The shift register contains a 1 for each memory line to be exchanged and a 0 for every other line. It is clocked (and thereby shifted) for each data transfer, i.e., every received memory line, and its output bit is used together with the FSM output to activate the parameter exchange by controlling the ROM and the MUX.

Upon activation of the parameter exchange, the trojan MUX forwards the manipulated parameters obtained from ROM to the write controller and finally to the on-chip RAM. Hence, the parameters are exchanged while being written to the buffer and before any computations on the received data have been executed. Independent of the inference operation to be executed (e.g., convolution, pooling, fully-connected, etc.), subsequent computations are performed on the manipulated parameters, i.e., *using the backdoored learning model*.

Backdoor Compression. For inference on the DPU, Xilinx Vitis AI performs 8-bit quantization on the model parameters and subsequently compiles the quantized model into a computation graph using the Xilinx intermediate representation (XIR). This graph can be serialized into and

de-serialized from so called `.xmodel` files of proprietary format after both quantization and compilation. Hence, such a file contains information on the layers of the model to be executed as well as the quantized (and optionally compiled) model parameters. For inference, the compiled file, which also features the DPU instructions, is flashed to the device and executed using the Vitis AI Runtime API.

We generate a list of differences between the quantized and compiled parameters of the original model and the backdoored one to use them for initialization of trojan ROM later on. To determine these differences, we compare the `.xmodel` files of both models by reading back the compiled parameters. A quantized `.xmodel` file differs from a compiled one in that it stores the parameters as 16-bit floats while the compiled file uses 8-bit fixed-point values instead. Furthermore, the compiled file stores the parameters in an order that is optimized for the shared memory layout. While the quantized parameters can still be read using Xilinx tooling, this not possible for a compiled `.xmodel` file. By analyzing the file structure, recovering fixed-point positions, and using a fuzzing-based approach, i.e., generating and comparing compiled `.xmodel` files for user-defined models, we automated extraction of the compiled parameters.

Backdoor Loading. Having computed the model differences, we reverse engineered the order in which the parameters are flashed to shared memory using known test patterns, as this order differed from the one in which the compiled parameters were kept in the `.xmodel` file. Finally, we initialized the ROM with the memory lines containing the manipulated parameters through a bitstream update.

4.3. Evaluation

We evaluated our attack by implementing the manipulated DPU on the Xilinx Zynq UltraScale+ MPSoC ZCU104 and running inference on the test data used in Section 3.3. Based on Table 1, we settled for a backdoored VGG-16 model generated using L^1 regularization and a trigger size of 50×50 pixels. This setup requires seven weight changes to achieve a trigger DSR of 90% before quantization, see Table 1a.

Figure 10 shows the trigger success rate and test accuracy of the backdoor after quantization. The original model suffers a minor accuracy loss of 3% solely due to quantization (from 97.43% down to 94.49%). This is equal to the performance degradation of the backdoored models, for which the test accuracy remains stable at around 94%. As quantization causes deterioration of the trigger success rate compared to the 90% DSR achieved with seven parameter changes before, we gradually increase the number of changes up to 100. The success rate converges to 83% while reaching the final plateau after 40 changes.

Figure 11 depicts the hardware overhead in terms of the number of LUTs, FFs, and LUT-RAM being used for a varying number of replaced parameters. The more parameters we replaced, the more memory lines must be kept in the trojan ROM. If manipulations spread across multiple load instructions, the additions to the memory reader FSM become
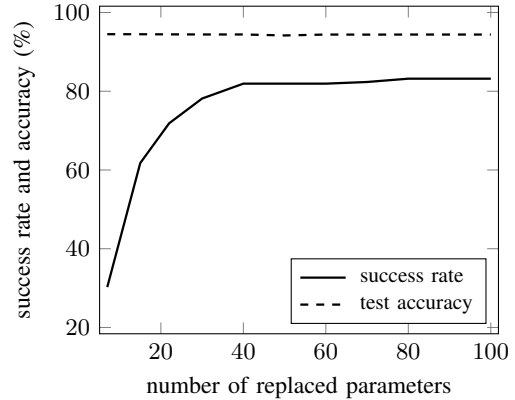

Figure 10: Success rate and test accuracy for backdoored variants of the traffic sign recognition model when being executed on the Xilinx Vitis AI DPU.
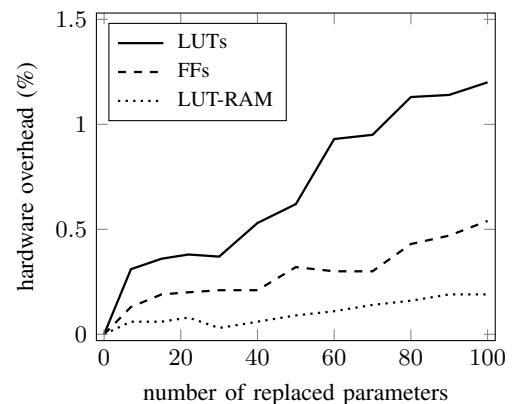

Figure 11: Hardware trojan overhead required to realize the respective number of weight replacements. The original DPU utilizes 37 379 LUTs, 6 440 LUT-RAM, and 90 309 FFs.

more complex as the trojan then needs to check against multiple addresses, thus requiring more resources.

In conlcusion, our trojan implementation causes a total hardware overhead below 1% and fits the target device. This results in a stealthy trojan implementation as no unreasonable amount of resources is required to implement the manipulated DPU. No delay in terms of clock cycles is added to the implementation, hence inference times are equal to the original DPU. Based on these results, we argue that 30 weight changes resulting in a success rate of 78.15% are a good trade-off to cause significant harm at little overhead.

5. Discussion

In this section, we discuss implications, countermeasures, and limitations of our hardware-based backdoor attack as well as the case study. Furthermore, we reflect on existing related work and propose new directions for future research.

5.1. Implications

In our case study, we have demonstrated that realizing a hardware trojan to insert machine-learning backdoors within a commercial hardware accelerator is technically feasible. We now discuss implications of such attacks.

Hardware Acceleration. By realizing a backdoor that is added to a learning model strictly within the hardware, we bypass all software and model integrity checks aimed at ensuring valid predictions. Our work thus demonstrates that the hardware used for machine learning inference cannot be blindly trusted and must undergo the same scrutiny as the software and learning model to ensure correct and trustworthy operation. In security-critical scenarios, the use of closed-source third-party hardware accelerators for machine learning must be questioned, as they pose a potential security risk.

Machine Learning Backdoor. Classical backdoors for neural networks [29, 32] have not been designed to be sparse. That is, attacks typically affect many model parameters when the backdoor functionality is implanted. In our experiments, we find that pruning and regularization strategies can drastically reduce the number of parameter changes and thereby enable meeting memory constraints for a hardware trojan. However, our results regarding the sparsity of a backdoor should be considered an *upper bound*, as further reduction strategies are conceivable. For example, the shape and content of the trigger could be optimized for backdoor sparsity. We leave such refinements of our approach to future work.

ASIC vs. FPGA Deployment. Our case study uses an FPGA as the target platform. Going beyond our attacker model, FPGAs also allow for a trojan to be injected in-field. An adversary with access to the bitstream could manipulate the architecture. Although extracting and altering bitstreams is tedious, it is a well-understood process [53–56] and certainly viable for powerful adversaries. Although bitstream protection schemes exist, they are notoriously difficult to implement and apply correctly [57–63].

We target an FPGA due to its accessibility for academic research. However, our trojan attack carries easily over to ASICs. Similar circuitry swapping selected weights, as described in Section 4.2, can be added to any ASIC accelerator. In order to be universally usable, programmability with respect to the backdoor parameters is strictly required. One can imagine a machine-learning accelerator with a secret programming interface (only known to the adversary) through which the trojanized parameters for the model running on the system-under-attack are uploaded after in-field deployment.

5.2. Detectability & Countermeasures

Defenses against trojan insertion can be employed from both the hardware and the machine-learning side.

Detectability. The overhead of our hardware manipulation is minimal. In theory, the attack can be detected by comparison with a trojan-free circuit [64]. However, no such golden model exists when the designer or a supplier inserts the trojan. Even formal verification approaches [65, 66] are ineffective as they would have to be performed or at least set up by the malicious entity. In addition to scaling issues when considering a large IP core such as the DPU [67], similar arguments can be made for proof-carrying hardware [68]. Techniques such as information flow security verification require at least some knowledge of the IP internals to identify so called *observe points* [67]. The only viable option is to analyze the circuit itself for malicious functionality. For FPGAs, this requires tedious reverse engineering of the bitstream format and, crucially, interpretation whether there are any malicious functions hidden within an unknown architecture. For ASICs, one needs to image the chip layer by layer using a scanning electron microscope (SEM) and extract a netlist using computer vision, a task that requires highly specialized equipment, skills, and considerable monetary resources. Even after successful netlist recovery, one faces again the problem of detecting a trojan within an unknown circuit. We claim that such efforts are out of reach for most entities in practice. Although nation states dispose of the resources to conduct such investigations, the required effort does not scale to a wide range of ICs.

Hardware Countermeasures. Two antagonistic approaches could be followed to harden a hardware design against manipulations. As first strategy, cryptographic and obfuscation measures can be used to protect the HDL design from manipulations. This demands a trusted design process, requiring strict access restrictions for the design files, vetting of all involved employees, and verification of the employed design tools. Furthermore, this chain of trust must be extended to all third-party IP cores utilized in the manufacturing process. Another strategy is switching to an open-source approach and ensuring public access to all design sources, allowing for third-party verification.

Although both strategies can help eliminate possible tampering along the supply chain, a trojan can still be inserted during the final manufacturing, for example, by replacing the trusted netlist with a trojanized clone. Consequently, the use of FPGAs and ASICs for security-critical machine-learning applications requires at least one trusted production facility.

Machine Learning Countermeasures. Since our attack operates from within the hardware accelerator, current approaches for detecting machine-learning backdoors [7, 8] fail, as the outside model remains valid. Attempts to spot the backdoor during execution [9, 10, 69], e.g., by monitoring neuron activations, may be a solution, but incur significant overhead and counteract the purpose of hardware acceleration. Moreover, the slight accuracy decrease induced by our backdoor is similar to that of quantization, so the attack cannot be detected from the model's accuracy either. Hence, to detect the malicious behavior, one needs to compare the outputs of the hardware-accelerated model to the original quantized version running in software. While this strategy allows identifying prediction discrepancies, the backdoor and its trigger still remain unknown. Currently, we lack appropriate methods to identify backdoors with this hybrid form of hardware-software testing.

5.3. Limitations

We make strong assumptions on the attacker's capabilities. Our attacker model assumes that adversaries can manipulate the design of a hardware accelerator and posses knowledge of the executed learning model. The required sophistication might only be in reach for nation-state actors, but other well-organized adversaries could also come into play. In addition to an adversary that develops the trojan in-house, they could pressure the original provider of the hardware accelerator to implement the trojan or infiltrate their operations. Attacks on the hardware level have been a serious concern for many years [70], which has recently triggered major investments by governments around the world [15, 16].

One hurdle to mount a trojan attack like ours is the assumed access to the trained learning model. However, given that we have reverse engineered Xilinx' proprietary `.xmodel` format, a similar attack could also be performed for a FPGA deployed in-field. Still, we expect quantization artifacts to further impair accuracy and trigger success rate of the resulting backdoor.

Finally, our approach allows only for attacking a single learning model executed on the accelerator. If that model changes, the trojanized parameters stored in hardware would need to be updated. Even if an update mechanism has been built into the hardware, this process is cumbersome and requires access to the updated model again.

5.4. Related Work

Machine-Learning Backdoors. The rising popularity of neural networks also raised interest in backdoor attacks. Among the first, Gu et al. [29] showed that an attacker who controls part of the training data can insert a backdoor into the network by adding a trigger and an incorrect class label to certain training examples. Further approaches that relax the assumption of access to the training data [32], the visibility, and position of the trigger [20, 71, 72] or the number of malicious examples required [30] exist. Stealthy backdoors that are inserted during model compilation [73], model quantization [74], or implemented by the software execution environment [75] were also proposed recently.

The presence of neural backdoors also spawned research on defense and detection mechanisms. One line of research tries to detect directly whether a trigger is present in the model, for example by finding shortcuts between output classes [7], training meta models to classify networks [8], or utilizing statistical properties from model predictions [76, 77]. An orthogonal line of research tries to detect whether a given input image contains a trigger, mostly by finding anomalies in activations or latent representations when propagating the input through the model at hand [9, 10, 69].

Hardware Trojans. For a general overview of hardware trojans, see [12, 78, 79]. The idea of hardware trojans targeting neural networks was first proposed by Clements et al. [80] and Li et al. [81] in 2018. Other works [82] require manipulations to the inputs to trigger the hardware trojan

which then bypasses the machine-learning accelerator altogether. More recent trojan attacks trigger on intermediate layer outputs [83], are inserted into the on-chip memory controller [84], or target activation parameters [85]. Neither of these works addresses the insertion of a machine-learning backdoor into a trained learning model during inference.

Hardware-supported machine-learning acceleration is also susceptible to non-trojan hardware attacks. Liu et al. injected glitches for untargeted misclassification [86] and demonstrated applicability using Xilinx Vitis AI. Hong et al. studied hardware fault attacks on deep neural networks (DNNs) and found that for most models a change of a single parameter can cause an accuracy drop of around 90% [87]. Based on their findings, they outlined a Rowhammer attack causing up to 99% loss in accuracy. Caner Tol et al. presented a similar backdoor attack again using Rowhammer [88]. Another research strain investigates the effects of RAM collisions caused by concurrent writes [89].

5.5. Future Work

We propose a new paradigm for machine-learning backdoors by taking the executing hardware into account. To counter this threat, countermeasures—ideally operating on the learning model as a black box—should be developed to detect low-level hardware manipulations during in-field operation. Furthermore, the amount of parameters required to realize minimal backdoors can possibly be reduced further. Another interesting aspect is the influence of quantization and how it could be incorporated into the backdoor generation process directly. Finally, an investigation of similar hardware-based attacks during model training appears worthwhile.

6. Conclusion

Our work extends the lively front of adversarial machine learning to a new so far trusted component: hardware acceleration. We present a trojan framework that backdoors a learning model at the lowest system layer during inference. All manipulations remain within the hardware, hence, no changes to the model can be observed, defeating existing defenses against backdoor attacks. To realize the trojan, we introduce the concept of a minimal backdoor that requires only a few parameter changes to implant malicious functionality. Even after quantization, 30 changes suffice to inject a backdoor with a trigger success rate of 78.15% and an overall prediction accuracy of 94.42%. We demonstrate the applicability of this attack by implanting the trojan into a commercial machine-learning accelerator from Xilinx.

Our work echoes recurring concerns from the hardware security community [15, 70, 90]. The trojan attack illustrates that hardware should not be blindly trusted and the integrity of accelerators for machine learning needs to be carefully verified and protected, similar to other security-critical components. We urge manufacturers, IP vendors, and system integrators alike to pay close attention to this threats, and call on the research community to develop countermeasures that prevent the exploitation of this new class of attacks.

Acknowledgements

References

[1] A. de la Escalera andchat Jose M. Armingol and M. Mata, "Traffic sign recognition and analysis for intelligent vehicles," *Image Vis. Comput.*, vol. 21, no. 3, pp. 247–258, 2003.

[2] A. Esteva, B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H. M. Blau, and S. Thrun, "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 542, no. 7639, pp. 115–118, 2017.

[3] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko, A. Bridgland, C. Meyer, S. A. A. Kohl, A. J. Ballard, A. Cowie, B. Romera-Paredes, S. Nikolov, R. Jain, J. Adler, T. Back, S. Petersen, D. Reiman, E. Clancy, M. Zielinski, M. Steinegger, M. Pacholska, T. Berghammer, S. Bodenstein, D. Silver, O. Vinyals, A. W. Senior, K. Kavukcuoglu, P. Kohli, and D. Hassabis, "Highly accurate protein structure prediction with AlphaFold," *Nature*, vol. 596, no. 7873, pp. 583–589, 2021.

[4] C. Wu, D. Brooks, K. Chen, D. Chen, S. Choudhury, M. Dukhan, K. M. Hazelwood, E. Isaac, Y. Jia, B. Jia, T. Leyvand, H. Lu, Y. Lu, L. Qiao, B. Reagen, J. Spisak, F. Sun, A. Tulloch, P. Vajda, X. Wang, Y. Wang, B. Wasti, Y. Wu, R. Xian, S. Yoo, and P. Zhang, "Machine Learning at Facebook: Understanding Inference at the Edge," in *25th IEEE International Symposium on High Performance Computer Architecture, HPCA 2019, Washington, DC, USA, February 16-20, 2019*, IEEE, 2019, pp. 331–344.

[5] K. M. Hazelwood, S. Bird, D. M. Brooks, S. Chintala, U. Diril, D. Dzhulgakov, M. Fawzy, B. Jia, Y. Jia, A. Kalro, J. Law, K. Lee, J. Lu, P. Noordhuis, M. Smelyanskiy, L. Xiong, and X. Wang, "Applied Machine Learning at Facebook: A Datacenter Infrastructure Perspective," in *IEEE International Symposium on High Performance Computer Architecture, HPCA 2018, Vienna, Austria, February 24-28, 2018*, IEEE Computer Society, 2018, pp. 620–629.

[6] N. Papernot, P. D. McDaniel, A. Sinha, and M. P. Wellman, "Sok: Security and privacy in machine learning," in *2018 IEEE European Symposium on Security and Privacy, EuroS&P 2018, London, United Kingdom, April 24-26, 2018*, IEEE, 2018, pp. 399–414.

[7] B. Wang, Y. Yao, S. Shan, H. Li, B. Viswanath, H. Zheng, and B. Y. Zhao, "Neural cleanse: Identifying and mitigating backdoor attacks in neural networks," in *2019 IEEE Symposium on Security and Privacy, SP 2019, San Francisco, CA, USA, May 19-23, 2019*, IEEE, 2019, pp. 707–723.

[8] X. Xu, Q. Wang, H. Li, N. Borisov, C. A. Gunter, and B. Li, "Detecting AI trojans using meta neural analysis," in *42nd IEEE Symposium on Security and Privacy, SP 2021, San Francisco, CA, USA, 24-27 May 2021*, IEEE, 2021, pp. 103–120.

[9] B. Tran, J. Li, and A. Madry, "Spectral signatures in backdoor attacks," in *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, S. Bengio, H. M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds., 2018, pp. 8011–8021.

[10] Y. Gao, C. Xu, D. Wang, S. Chen, D. C. Ranasinghe, and S. Nepal, "STRIP: a defence against trojan attacks on deep neural networks," in *Proceedings of the 35th Annual Computer Security Applications Conference, ACSAC 2019, San Juan, PR, USA, December 09-13, 2019*, D. Balenson, Ed., ACM, 2019, pp. 113–125.

[11] S. Bhasin, J. Danger, S. Guilley, X. T. Ngo, and L. Sauvage, "Hardware Trojan Horses in Cryptographic IP Cores," in *2013 Workshop on Fault Diagnosis and Tolerance in Cryptography, Los Alamitos, CA, USA, August 20, 2013*, W. Fischer and J. Schmidt, Eds., IEEE Computer Society, 2013, pp. 15–29.

[12] M. Tehranipoor and F. Koushanfar, "A Survey of Hardware Trojan Taxonomy and Detection," *IEEE Des. Test Comput.*, vol. 27, no. 1, pp. 10–25, 2010.

[13] R. Karri, J. Rajendran, K. Rosenfeld, and M. Tehranipoor, "Trustworthy hardware: Identifying and classifying hardware trojans," *Computer*, vol. 43, no. 10, pp. 39–46, 2010.

[14] K. Xiao, D. Forte, Y. Jin, R. Karri, S. Bhunia, and M. M. Tehranipoor, "Hardware Trojans: Lessons Learned after One Decade of Research," *ACM Trans. Design Autom. Electr. Syst.*, vol. 22, no. 1, 6:1–6:23, 2016.

[15] European Comission, *European Chips Act*, Brussels, May 2022.

[16] Senate of the United States, *CHIPS and Science Act 2022 (P.L. 117-167)*, Washington, D.C., Jul. 2022.

[17] Z. Sun, R. Sun, L. Lu, and A. Mislove, "Mind your weight(s): A large-scale study on insufficient machine learning model protection in mobile apps," in *30th USENIX Security Symposium, USENIX Security 2021, August 11-13, 2021*, M. Bailey and R. Greenstadt, Eds., USENIX Association, 2021, pp. 1955–1972.

[18] IEEE Design Automation Standards Committee (DASC), *IEEE 1735-2014 - Recommended Practice for Encryption and Management of Electronic Design Intellectual Property (IP)*. IEEE, 2015.

[19] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S.

McCandlish, A. Radford, I. Sutskever, and D. Amodei, "Language models are few-shot learners," in *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., 2020.

[20] T. A. Nguyen and A. T. Tran, "Input-aware dynamic backdoor attack," in *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., 2020.

[21] A. Salem, R. Wen, M. Backes, S. Ma, and Y. Zhang, "Dynamic backdoor attacks against machine learning models," in *7th IEEE European Symposium on Security and Privacy, EuroS&P 2022, Genoa, Italy, June 6-10, 2022*, IEEE, 2022, pp. 703–718.

[22] H. Zhong, C. Liao, A. C. Squicciarini, S. Zhu, and D. J. Miller, "Backdoor embedding in convolutional neural network models via invisible perturbation," in *CODASPY '20: Tenth ACM Conference on Data and Application Security and Privacy, New Orleans, LA, USA, March 16-18, 2020*, V. Roussev, B. Thuraisingham, B. Carminati, and M. Kantarcioglu, Eds., ACM, 2020, pp. 97–108.

[23] Y. Zhou, S. Moosavi-Dezfooli, N. Cheung, and P. Frossard, "Adaptive quantization for deep neural network," in *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, S. A. McIlraith and K. Q. Weinberger, Eds., AAAI Press, 2018, pp. 4596–4604.

[24] K. Wang, Z. Liu, Y. Lin, J. Lin, and S. Han, "HAQ: hardware-aware automated quantization with mixed precision," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, Computer Vision Foundation / IEEE, 2019, pp. 8612–8620.

[25] A. Reuther, P. Michaleas, M. Jones, V. Gadepally, S. Samsi, and J. Kepner, "Survey and benchmarking of machine learning accelerators," in *2019 IEEE High Performance Extreme Computing Conference, HPEC 2019, Waltham, MA, USA, September 24-26, 2019*, IEEE, 2019, pp. 1–9.

[26] C. Wang, L. Gong, Q. Yu, X. Li, Y. Xie, and X. Zhou, "DLAU: A scalable deep learning accelerator unit on FPGA," *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.*, vol. 36, no. 3, pp. 513–517, 2017.

[27] B. Jacob, S. Kligys, B. Chen, M. Zhu, M. Tang, A. G. Howard, H. Adam, and D. Kalenichenko, "Quantization and training of neural networks for efficient integer-arithmetic-only inference," in *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, Computer Vision Foundation / IEEE Computer Society, 2018, pp. 2704–2713.

[28] H. Wu, P. Judd, X. Zhang, M. Isaev, and P. Micikevicius, "Integer quantization for deep learning inference: Principles and empirical evaluation," *CoRR*, vol. abs/2004.09602, 2020.

[29] T. Gu, K. Liu, B. Dolan-Gavitt, and S. Garg, "Badnets: Evaluating backdooring attacks on deep neural networks," *IEEE Access*, vol. 7, pp. 47 230–47 244, 2019.

[30] A. Shafahi, W. R. Huang, M. Najibi, O. Suciu, C. Studer, T. Dumitras, and T. Goldstein, "Poison frogs! targeted clean-label poisoning attacks on neural networks," in *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, S. Bengio, H. M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds., 2018, pp. 6106–6116.

[31] Y. Yao, H. Li, H. Zheng, and B. Y. Zhao, "Latent backdoor attacks on deep neural networks," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, CCS 2019, London, UK, November 11-15, 2019*, L. Cavallaro, J. Kinder, X. Wang, and J. Katz, Eds., ACM, 2019, pp. 2041–2055.

[32] Y. Liu, S. Ma, Y. Aafer, W. Lee, J. Zhai, W. Wang, and X. Zhang, "Trojaning attack on neural networks," in *25th Annual Network and Distributed System Security Symposium, NDSS 2018, San Diego, California, USA, February 18-21, 2018*, The Internet Society, 2018.

[33] K. Simonyan, A. Vedaldi, and A. Zisserman, "Deep inside convolutional networks: Visualising image classification models and saliency maps," in *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Workshop Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2014.

[34] D. Erhan, Y. Bengio, A. Courville, and P. Vincent, "Visualizing Higher-Layer Features of a Deep Network," *Technical Report, Univeristé de Montréal*, 2009.

[35] N. Carlini and D. A. Wagner, "Towards evaluating the robustness of neural networks," in *2017 IEEE Symposium on Security and Privacy, SP 2017, San Jose, CA, USA, May 22-26, 2017*, IEEE Computer Society, 2017, pp. 39–57.

[36] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2015.

[37] Y. LeCun, J. S. Denker, and S. A. Solla, "Optimal brain damage," in *Advances in Neural Information Processing Systems 2, [NIPS Conference, Denver, Colorado, USA, November 27-30, 1989]*, D. S. Touretzky, Ed., Morgan Kaufmann, 1989, pp. 598–605.

[38] C. Louizos, K. Ullrich, and M. Welling, "Bayesian compression for deep learning," in *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, Eds., 2017, pp. 3288–3298.

[39] W. Wen, C. Wu, Y. Wang, Y. Chen, and H. Li, "Learning structured sparsity in deep neural networks," in *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, D. D. Lee, M. Sugiyama, U. von Luxburg, I. Guyon, and R. Garnett, Eds., 2016, pp. 2074–2082.

[40] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural network with pruning, trained quantization and huffman coding," in *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2016.

[41] K. Ullrich, E. Meeds, and M. Welling, "Soft weight-sharing for neural network compression," in *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, OpenReview.net, 2017.

[42] D. Molchanov, A. Ashukha, and D. P. Vetrov, "Variational dropout sparsifies deep neural networks," in *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, D. Precup and Y. W. Teh, Eds., ser. Proceedings of Machine Learning Research, vol. 70, PMLR, 2017, pp. 2498–2507.

[43] C. Louizos, M. Welling, and D. P. Kingma, "Learning sparse neural networks through l_0 regularization," in *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, OpenReview.net, 2018.

[44] S. Srinivas, A. Subramanya, and R. V. Babu, "Training sparse neural networks," in *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2017, Honolulu, HI, USA, July 21-26, 2017*, IEEE Computer Society, 2017, pp. 455–462.

[45] S. Houben, J. Stallkamp, J. Salmen, M. Schlipsing, and C. Igel, "Detection of traffic signs in real-world images: The german traffic sign detection benchmark," in *The 2013 International Joint Conference on Neural Networks, IJCNN 2013, Dallas, TX, USA, August 4-9, 2013*, IEEE, 2013, pp. 1–8.

[46] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2015.

[47] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems 25: 26th Annual Conference on Neural Information Processing Systems 2012. Proceedings of a meeting held December 3-6, 2012, Lake Tahoe, Nevada, United States*, P. L. Bartlett, F. C. N. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, Eds., 2012, pp. 1106–1114.

[48] O. M. Parkhi, A. Vedaldi, and A. Zisserman, "Deep face recognition," in *Proceedings of the British Machine Vision Conference 2015, BMVC 2015, Swansea, UK, September 7-10, 2015*, X. Xie, M. W. Jones, and G. K. L. Tam, Eds., BMVA Press, 2015, pp. 41.1–41.12.

[49] Advanced Micro Devices, Inc, *Vitis AI*, Available at https://www.xilinx.com/products/design-tools/vitis/vitis-ai.html.

[50] ——, *Vitis AI DPU*, Available at https://github.com/Xilinx/Vitis-AI/tree/master/dpu.

[51] A. Chhotaray, A. Nahiyan, T. Shrimpton, D. Forte, and M. M. Tehranipoor, "Standardizing bad cryptographic practice: A teardown of the IEEE standard for protecting electronic-design intellectual property," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*, B. Thuraisingham, D. Evans, T. Malkin, and D. Xu, Eds., ACM, 2017, pp. 1533–1546.

[52] J. Speith, F. Schweins, M. Ender, M. Fyrbiak, A. May, and C. Paar, "How not to protect your IP - an industry-wide break of IEEE 1735 implementations," in *43rd IEEE Symposium on Security and Privacy, SP 2022, San Francisco, CA, USA, May 22-26, 2022*, IEEE, 2022, pp. 1656–1671.

[53] K. D. Pham, E. L. Horta, and D. Koch, "BITMAN: A tool and API for FPGA bitstream manipulations," in *Design, Automation & Test in Europe Conference & Exhibition, DATE 2017, Lausanne, Switzerland, March 27-31, 2017*, D. Atienza and G. D. Natale, Eds., IEEE, 2017, pp. 894–897.

[54] J. Kataria, R. Housley, J. Pantoga, and A. Cui, "Defeating Cisco Trust Anchor: A Case-Study of Recent Advancements in Direct FPGA Bitstream Manipulation," in *13th USENIX Workshop on Offensive Technologies, WOOT 2019, Santa Clara, CA, USA, August 12-13, 2019*, A. Gantman and C. Maurice, Eds., USENIX Association, 2019.

[55] J. Note and É. Rannaud, "From the bitstream to the netlist," in *Proceedings of the ACM/SIGDA 16th International Symposium on Field Programmable Gate Arrays, FPGA 2008, Monterey, California, USA, February 24-26, 2008*, M. Hutton and P. Chow, Eds., ACM, 2008, p. 264.

[56] M. Ender, P. Swierczynski, S. Wallat, M. Wilhelm, P. M. Knopp, and C. Paar, "Insights into the mind of a trojan designer: the challenge to integrate a trojan into the bitstream," in *Proceedings of the 24th Asia and*

South Pacific Design Automation Conference, ASPDAC 2019, Tokyo, Japan, January 21-24, 2019, T. Shibuya, Ed., ACM, 2019, pp. 112–119.

[57] M. Ender, A. Moradi, and C. Paar, "The Unpatchable Silicon: A Full Break of the Bitstream Encryption of Xilinx 7-Series FPGAs," in *29th USENIX Security Symposium, USENIX Security 2020, August 12-14, 2020*, S. Capkun and F. Roesner, Eds., USENIX Association, 2020, pp. 1803–1819.

[58] M. Ender, G. Leander, A. Moradi, and C. Paar, "A Cautionary Note on Protecting Xilinx' UltraScale(+) Bitstream Encryption and Authentication Engine," in *30th IEEE Annual International Symposium on Field-Programmable Custom Computing Machines, FCCM 2022, New York City, NY, USA, May 15-18, 2022*, IEEE, 2022, pp. 1–9.

[59] A. Moradi, A. Barenghi, T. Kasper, and C. Paar, "On the vulnerability of FPGA bitstream encryption against power analysis attacks: extracting keys from xilinx Virtex-II FPGAs," in *Proceedings of the 18th ACM Conference on Computer and Communications Security, CCS 2011, Chicago, Illinois, USA, October 17-21, 2011*, Y. Chen, G. Danezis, and V. Shmatikov, Eds., ACM, 2011, pp. 111–124.

[60] A. Moradi, M. Kasper, and C. Paar, "Black-Box Side-Channel Attacks Highlight the Importance of Countermeasures - An Analysis of the Xilinx Virtex-4 and Virtex-5 Bitstream Encryption Mechanism," in *Topics in Cryptology - CT-RSA 20W12 - The Cryptographers' Track at the RSA Conference 2012, San Francisco, CA, USA, February 27 - March 2, 2012. Proceedings*, O. Dunkelman, Ed., ser. Lecture Notes in Computer Science, vol. 7178, Springer, 2012, pp. 1–18.

[61] A. Moradi and T. Schneider, "Improved Side-Channel Analysis Attacks on Xilinx Bitstream Encryption of 5, 6, and 7 Series," in *Constructive Side-Channel Analysis and Secure Design - 7th International Workshop, COSADE 2016, Graz, Austria, April 14-15, 2016, Revised Selected Papers*, F. Standaert and E. Oswald, Eds., ser. Lecture Notes in Computer Science, vol. 9689, Springer, 2016, pp. 71–87.

[62] P. Swierczynski, A. Moradi, D. F. Oswald, and C. Paar, "Physical Security Evaluation of the Bitstream Encryption Mechanism of Altera Stratix II and Stratix III FPGAs," *ACM Trans. Reconfigurable Technol. Syst.*, vol. 7, no. 4, 34:1–34:23, 2015.

[63] S. Tajik, H. Lohrke, J. Seifert, and C. Boit, "On the Power of Optical Contactless Probing: Attacking Bitstream Encryption of FPGAs," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*, B. M. Thuraisingham, D. Evans, T. Malkin, and D. Xu, Eds., ACM, 2017, pp. 1661–1674.

[64] S. Bhasin and F. Regazzoni, "A survey on hardware trojan detection techniques," in *2015 IEEE International Symposium on Circuits and Systems, ISCAS 2015, Lisbon, Portugal, May 24-27, 2015*, IEEE, 2015, pp. 2021–2024.

[65] M. Rathmair, F. Schupfer, and C. Krieg, "Applied formal methods for hardware trojan detection," in *IEEE International Symposium on Circuits and Systemss, ISCAS 2014, Melbourne, Victoria, Australia, June 1-5, 2014*, IEEE, 2014, pp. 169–172.

[66] X. Guo, R. G. Dutta, Y. Jin, F. Farahmandi, and P. Mishra, "Pre-silicon security verification and validation: A formal perspective," in *Proceedings of the 52nd Annual Design Automation Conference, San Francisco, CA, USA, June 7-11, 2015*, ACM, 2015, 145:1–145:6.

[67] A. Nahiyan, M. Sadi, R. Vittal, G. K. Contreras, D. Forte, and M. M. Tehranipoor, "Hardware trojan detection through information flow security verification," in *IEEE International Test Conference, ITC 2017, Fort Worth, TX, USA, October 31 - Nov. 2, 2017*, IEEE, 2017, pp. 1–10.

[68] E. Love, Y. Jin, and Y. Makris, "Proof-carrying hardware intellectual property: A pathway to trusted module acquisition," *IEEE Trans. Inf. Forensics Secur.*, vol. 7, no. 1, pp. 25–40, 2012.

[69] B. Chen, W. Carvalho, N. Baracaldo, H. Ludwig, B. Edwards, T. Lee, I. M. Molloy, and B. Srivastava, "Detecting backdoor attacks on deep neural networks by activation clustering," in *Workshop on Artificial Intelligence Safety 2019 co-located with the Thirty-Third AAAI Conference on Artificial Intelligence 2019 (AAAI-19), Honolulu, Hawaii, January 27, 2019*, H. Espinoza, S. Ó. hÉigeartaigh, X. Huang, J. Hernández-Orallo, and M. Castillo-Effen, Eds., ser. CEUR Workshop Proceedings, vol. 2301, CEUR-WS.org, 2019.

[70] Defense Science Board Task Force, "High Performance Microchip Supply," *Annual Report. Defense Technical Information Center (DTIC), USA*, 2005.

[71] X. Chen, C. Liu, B. Li, K. Lu, and D. Song, "Targeted backdoor attacks on deep learning systems using data poisoning," *CoRR*, vol. abs/1712.05526, 2017.

[72] A. Saha, A. Subramanya, and H. Pirsiavash, "Hidden trigger backdoor attacks," in *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, AAAI Press, 2020, pp. 11 957–11 965.

[73] T. Clifford, I. Shumailov, Y. Zhao, R. J. Anderson, and R. D. Mullins, "Impnet: Imperceptible and blackbox-undetectable backdoors in compiled neural networks," *CoRR*, vol. abs/2210.00108, 2022.

[74] H. Ma, H. Qiu, Y. Gao, Z. Zhang, A. Abuadbba, A. Fu, S. F. Al-Sarawi, and D. Abbott, "Quantization backdoors to deep learning models," *CoRR*, vol. abs/2108.09187, 2021.

[75] Y. Li, J. Hua, H. Wang, C. Chen, and Y. Liu, "Deeppayload: Black-box backdoor attack on deep learning models through neural payload injection," in

43rd IEEE/ACM International Conference on Software Engineering, ICSE 2021, Madrid, Spain, 22-30 May 2021, IEEE, 2021, pp. 263–274.

[76] H. Chen, C. Fu, J. Zhao, and F. Koushanfar, "Deepinspect: A black-box trojan detection and mitigation framework for deep neural networks," in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, S. Kraus, Ed., ijcai.org, 2019, pp. 4658–4664.

[77] D. Tang, X. Wang, H. Tang, and K. Zhang, "Demon in the variant: Statistical analysis of dnns for robust backdoor contamination detection," in *30th USENIX Security Symposium, USENIX Security 2021, August 11-13, 2021*, M. Bailey and R. Greenstadt, Eds., USENIX Association, 2021, pp. 1541–1558.

[78] M. Xue, C. Gu, W. Liu, S. Yu, and M. O'Neill, "Ten years of hardware Trojans: a survey from the attacker's perspective," *IET Comput. Digit. Tech.*, vol. 14, no. 6, pp. 231–246, 2020.

[79] H. Li, Q. Liu, and J. Zhang, "A survey of hardware Trojan threat and defense," *Integr.*, vol. 55, pp. 426–437, 2016.

[80] J. Clements and Y. Lao, "Hardware Trojan Attacks on Neural Networks," *CoRR*, vol. abs/1806.05768, 2018.

[81] W. Li, J. Yu, X. Ning, P. Wang, Q. Wei, Y. Wang, and H. Yang, "Hu-Fu: Hardware and Software Collaborative Attack Framework Against Neural Networks," in *2018 IEEE Computer Society Annual Symposium on VLSI, ISVLSI 2018, Hong Kong, China, July 8-11, 2018*, IEEE Computer Society, 2018, pp. 482–487.

[82] J. Ye, Y. Hu, and X. Li, "Hardware Trojan in FPGA CNN Accelerator," in *27th IEEE Asian Test Symposium, ATS 2018, Hefei, China, October 15-18, 2018*, IEEE, 2018, pp. 68–73.

[83] T. A. Odetola, H. R. Mohammed, and S. R. Hasan, "A Stealthy Hardware Trojan Exploiting the Architectural Vulnerability of Deep Learning Architectures: Input Interception Attack (IIA)," *CoRR*, vol. abs/1911.00783, 2019.

[84] X. Hu, Y. Zhao, L. Deng, L. Liang, P. Zuo, J. Ye, Y. Lin, and Y. Xie, "Practical Attacks on Deep Neural Networks by Memory Trojaning," *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.*, vol. 40, no. 6, pp. 1230–1243, 2021.

[85] R. Mukherjee and R. S. Chakraborty, "Novel Hardware Trojan Attack on Activation Parameters of FPGA-Based DNN Accelerators," *IEEE Embed. Syst. Lett.*, vol. 14, no. 3, pp. 131–134, 2022.

[86] W. Liu, C. Chang, F. Zhang, and X. Lou, "Imperceptible Misclassification Attack on Deep Learning Accelerator by Glitch Injection," in *57th ACM/IEEE Design Automation Conference, DAC 2020, San Francisco, CA, USA, July 20-24, 2020*, IEEE, 2020, pp. 1–6.

[87] S. Hong, P. Frigo, Y. Kaya, C. Giuffrida, and T. Dumitras, "Terminal Brain Damage: Exposing the Graceless Degradation in Deep Neural Networks Under Hardware Fault Attacks," in *28th USENIX Security Symposium, USENIX Security 2019, Santa Clara, CA, USA, August 14-16, 2019*, N. Heninger and P. Traynor, Eds., USENIX Association, 2019, pp. 497–514.

[88] M. C. Tol, S. Islam, B. Sunar, and Z. Zhang, "An optimization perspective on realizing backdoor injection attacks on deep neural networks in hardware," *CoRR*, vol. abs/2110.07683, 2021.

[89] M. M. Alam, S. Tajik, F. Ganji, M. M. Tehranipoor, and D. Forte, "RAM-Jam: Remote Temperature and Voltage Fault Attack on FPGAs using Memory Collisions," in *2019 Workshop on Fault Diagnosis and Tolerance in Cryptography, FDTC 2019, Atlanta, GA, USA, August 24, 2019*, IEEE, 2019, pp. 48–55.

[90] European Comission, *Cyber Resilience Act*, Brussels, Sep. 2022.