



HAL
open science

You only Get One-Shot: Eavesdropping Input Images to Neural Network by Spying SoC-FPGA Internal Bus

May Myat Thu, Maria Mendez Real, Maxime Pelcat, Philippe Besnier

► To cite this version:

May Myat Thu, Maria Mendez Real, Maxime Pelcat, Philippe Besnier. You only Get One-Shot: Eavesdropping Input Images to Neural Network by Spying SoC-FPGA Internal Bus. 18th International Conference on Availability, Reliability and Security, ARES 2023, Aug 2023, Benevento, Italy. pp.31, 10.1145/3600160.3600189 . hal-04226071

HAL Id: hal-04226071

<https://hal.science/hal-04226071v1>

Submitted on 16 Oct 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

You Only Get One-Shot: Eavesdropping Input Images to Neural Network by Spying SoC-FPGA Internal Buses

ABSTRACT

Deep learning is currently integrated into edge devices with strong energy consumption and real-time constraints. To fulfill such requirements, high hardware performances can be provided by hardware acceleration of heterogeneous integrated circuits (IC) such as System-on-Chip (SoC)-field programmable gate arrays (FPGAs). With the rising popularity of hardware accelerators for artificial intelligence (AI), more and more neural networks are employed in a variety of domains, involving computer vision applications. Autonomous driving, defence and medical domains are well-known examples from which the latter two in particular require processing sensitive and private data. Security issues of such systems should be addressed to prevent the breach of privacy and unauthorised exploitation of systems. In this paper, we demonstrate a confidentiality vulnerability in a SoC-based FPGA binarized neural network (BNN) accelerator implemented with a recent mainstream framework, FINN, and successfully extract the secret BNN input image by using an electromagnetic (EM) side-channel attack. Experiments demonstrate that with the help of a near-field magnetic probe, an attacker can, with only one inference, directly retrieve sensitive information from EM emanations produced by the internal bus of the SoC-FPGA. Our attack reconstructs SoC-FPGA internal images and recognizes a handwritten digit image with an average accuracy of 89% using a non-retrained MNIST classifier. Such vulnerability jeopardizes the confidentiality of SoC-FPGA embedded AI systems by exploiting side-channels that withstand the protection of chip I/Os through cryptographic methods.

CCS CONCEPTS

• **Hardware Security** → **Side-channel Attacks**; • **Embedded Systems** → **System-on-chip Field Programmable Gate Array**; • **Electromagnetic** → **Near-field**.

KEYWORDS

Electromagnetic Side-Channel Attack, Binarized Neural Network, Image Processing, System-on-Chip, Field Programmable Gate Array, On-chip Communication Buses

ACM Reference Format:

. 2018. You Only Get One-Shot: Eavesdropping Input Images to Neural Network by Spying SoC-FPGA Internal Buses. In *Proceedings of 18th International Conference on Availability, Reliability and Security (ARES 2023)*. ACM, New York, NY, USA, 7 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
ARES 2023, August 29- September-01, 2023, Benevento, Italy

© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/XXXXXXXX.XXXXXXX>

1 INTRODUCTION

Organizations and companies are storing most of their important data in digital form [1], hence, solutions to prevent exploitation of confidential internal information must be explored. The two main aspects of securing systems involve *software* and *hardware* security. While the main job of software security consists of protecting software applications and digital solutions from external viruses and malware, the importance of hardware security should also be equally considered. Among various types of hardware security, one of the most popular and common in the state-of-the-art is *side-channel attack*. A side-channel attack exploits unwanted signal emanations from a victim system to jeopardize its data confidentiality. Emitted signals can take many physical forms, from execution time variations to consumed power variations or EM emanations. This paper demonstrates a new form of electromagnetic side-channel attack on system internal buses. A security exploit of such EM emanations when focusing on AI algorithm can target different assets, e.g. sensitive data, secret algorithm parameters and details about AI models.

SoC-FPGAs are gaining momentum as AI hardware accelerators thanks to their capacity to exploit the embarrassingly parallel nature of neural networks so as to provide energy-efficient and computationally sufficient solutions [2]. This is true especially for power hungry convolutional neural network (CNN) AI accelerators which are commonly deployed in computer vision solutions, and then applied to, for instance, medical and defence domains [3]. Jeopardizing the input image of a computer vision system gives advanced information to the attacker, and we demonstrate in this paper that the very nature of the accelerator creates EM side-channels that jeopardize image information. Our work aims to eavesdrop the input image sent to the neural network for inference process by making use of electromagnetic vulnerability in internal buses of SoC-FPGA. The implemented attack in this paper is called "*You Only Get One-Shot*", as data is retrieved from the observation of the emanations during a unique inference, and in the rest of the paper it will be referred to as YOGOS (abbrev. for You Only Get One-Shot). The main contributions of this work are the following:

- We demonstrate that on-chip buses on SoC-FPGAs are vulnerable to EM side-channel attacks. Electromagnetic activities of these communication buses can be spied on using near-field magnetic probes, which can lead to the extraction of private data regardless of the application on FPGA.
- By exploiting the electromagnetically induced signals emitted by Advanced eXtensible Interface (AXI) bus, we propose an attack against a neural network implemented on SoC-FPGA without the knowledge of the neural network architecture or implementation details. Our attack can extract the input images sent on the AXI bus to the FPGA device where the inference of these images takes place.
- Our method does not require any other form of interaction with the victim system, *i.e.*, no need to trigger CNN inference

of the victim. Synchronisation of each individual trace is possible just from analyzing the emanations in signal processing stage.

- We propose YOGOS, a one-shot attack where a single acquisition of one inference EM trace is sufficient to perform the attack.
- An experimental setup for studying the YOGOS attack is built and further proves that neural network (NN) input images retrieved by this attack can be recognized with 89% of average accuracy.

2 BACKGROUND

This section includes a brief explanation of the SoC-FPGAs and its on-chip bus protocol, EM emanations and FINN framework used in this experiment.

2.1 System-on-Chip and Bus Protocol

SoC-FPGAs combine a processing system (PS), usually embedding ARM cores, and programmable logic (PL), embedding FPGA computing resources, onto a single integrated device. PS and PL communicate through complex interconnects. A series of open-standard on-chip interconnect protocols for the connection and management of functional blocks in SoC designs is known as the advanced micro-controller bus architecture (AMBA) [4]. One of the specifications of AMBA is AXI-4 which is the established protocol between PS and PL of Zynq-7000 and usually comes with data widths of 32, 64 or 128 bits. Our attack focuses on AXI buses which provide data access from the PL to the on-chip static random access memory (SRAM) in the PS, and to the external dynamic random access memory (DRAM) connected to the PS side. They can be configurable as either 32-bit or 64-bit widths and use a FIFO controller to connect to the memory interconnect [5]. In this paper, we show that the AXI buses of Zynq-7000 SoC can compromise the security of the system.

2.2 Electromagnetic Emanations

EM emanation of ICs are mainly attributed to numerous and simultaneous commutations of logic gates transistors [6]. These commutations are responsible for voltage fluctuations between IC power and ground planes, and for current loops. In case of a bus line, current loops may be created in a complementary metal-oxide-semiconductor (CMOS) inverter gate when, according to clock rate, both p and n transistors are simultaneously conducting for a short period of time. Moreover, charge or discharge of the bus line may be responsible for current leakage due to parasitic capacitance. At high-speed and high clock rate data bus (100 MHz or above), these two phenomena are likely to occur giving more chance to pick up the reactive near-field with an adequate, preferably magnetic probe. EM emanations emitted from this behaviour of ICs may contain confidential information and can be maliciously exploited.

2.3 FINN

FINN, developed by Xilinx, is a well supported experimental framework to build fast and scalable BNN accelerators onto FPGAs [7]. It supports weights quantization and reduction of memory footprint.

FINN provides a series of prototypes [8] that accelerate BNN inference on standard datasets: MNIST [9], CIFAR-10 and SVHN [10]. Using high-level synthesis (HLS), FINN allows automated implementation of pre-trained BNN models, exploiting FPGA capabilities such as arbitrary data size manipulations [3]. FINN supports Pynq Z1, Pynq Z2, Ultra96, ZCU102 and ZCU104 development boards. The choice of the hardware platform relies on the size and complexity of the NN. For instance, ZCU104 is recommended for FINN dataflow-style accelerators trained with ImageNet-sized datasets while Pynq Z1 is suggested for smaller networks.

3 RELATED WORK

Additionally to being used to steal an encryption key from a cryptographic algorithm, starting from the late 2010's, side-channel attacks (SCAs) are also being increasingly employed to expand the attack capability to NNs implementations on FPGA [11]. Several attempts have been made since 2018 to recover private data or important details about NN models (for *e.g.* input, model architecture) implemented on FPGAs by exploiting different side-channels.

3.1 Non Electromagnetic Side-Channel Attacks

Among various types of SCAs which have been demonstrated in the literature, power SCAs are the most common, observing variations on power consumption from the device. In [12], the authors focus on spying on the secret neural network image inputs by first detecting the background pixels in an image-generated signal from captured power traces. Then, as for the foreground pixels, a power template is generated so that the collected power traces can be compared to it in order to deduce each original pixel value of the input image. Similarly, authors of [13] developed the first remote power SCA to recover input images to the BNN accelerator, implemented on multi-tenant FPGAs by making use of the power-related traces obtained from time-to-digital converters (TDCs). It has been shown in [14] that weights values of an FPGA-implemented NN can also be retrieved by the attackers using chain correlation power analysis. In the case of [15], parameters specific to FINN implementation (folded layers and number of neurons) are retrieved by analyzing power side-channels from BNN accelerator with the help of TDCs.

3.2 Electromagnetic Side-Channel Attacks

As for the EM side-channel attacks, in [16], authors demonstrated that EM traces captured from the FPGA surface, can be analyzed to reveal active operating regions of the chip. Authors also propose methods to efficiently make EM imaging and cartography on the chip which serves as a fundamental step in most EM side-channel attacks. DeepEM [17] demonstrates that attackers can accurately recover NN model hyper-parameters by exploiting EM side-channel information from an FPGA-based SoC implementation. Moreover, [18] illustrates how EM traces emitted by a targeted FPGA device may be utilized by an attacker to fully recover the secret BNN weights used in the network.

However, to our knowledge, there is no work on eavesdropping input images through EM emanations of the internal bus to compromise SoC-FPGAs. We demonstrate in this work an EM side-channel attack that exploits the vulnerability of SoC on-chip buses and proposed a novel algorithm to retrieve the input image information.

4 THREAT MODEL

YOGOS can be applied to real-life scenarios in both public and private sectors. Deep learning-enabled surveillance cameras in public places capture the real-time video frames which are then transmitted to the accelerator part of the camera where the deep learning algorithm executes. The algorithm then processes the videos frames and classifies to determine what is being detected in the frame. In order to perform this attack, EM traces have to be collected through a near-field probe that can be put into the camera device by obtaining physical access to the system. Given the small size of the probe loop, spying components or hardware Trojans, as small as a grain of rice, could be implemented during the fabrication of chips [19] to allow the attacker to capture the EM signals at a distance, which opens up ways to apply YOGOS attack at scale. Besides, the attack can also be implemented by an insider of a device supply chain. A malicious individual who has access to a surveillance camera in a restricted zone (e.g., laboratory, warehouse) can tag the spying component to steal the confidential information. This is in particular threatening for medical and defense related companies in which the information of the production and new research data can be leaked.

Knowledge: The attacker requires a general idea of the encoding of image (raster pixel scan, color planes). The attacker has no knowledge of the system running on the targeted hardware, i.e. neither the architecture of BNN nor its parameters. Image size can be deduced from the signal shape over time by counting the number of observed peaks.

Capability: The attacker is required to collect EM traces of the targeted system while it is operating. This requirement can be met in internet of things (IoT) devices such as smart cameras, sensors, etc, where the device is physically available to the attacker. The near-field access can also be obtained through a hardware Trojan inserted in the package of the device during the manufacturing process as explained above. On the other hand, the attacker does not need to control the inputs, nor interact with the BNN, nor trigger the inference. Hence, the inference inputs, weights, layers, and neurons of the BNN are unknown to the attacker.

Goal: The objective of the proposed attack is to infer the secret inputs to the BNN-based SoC-FPGA accelerator without the need of any knowledge on the system architecture, parameters or hardware implementation design of the BNN.

5 METHODOLOGY OF THE PROPOSED ATTACK

In this section, the different steps necessary for performing the attack are presented.

5.1 Leakage Detection

First, the attacker explores the EM emanations of the SoC-FPGA to locate the zones presenting notable EM leakage. In our case, we want to locate the internal bus. For this step, the attacker is required to have access to a similar SoC-FPGA to identify the location of the bus independently of what is implemented in the FPGA. By scanning the surface of the SoC with a near field magnetic probe, during data transfers, the attacker can identify the specific energy leakage pattern on the SoC. The latter can be easily inspected as it

presents high activity visible through significant peaks on the signal compared to noise floor. After careful observation of this energy leakage pattern, it can be expected that when the buses are *idle*, the EM signals are at the lowest and are very close to the noise floor. However, signals with substantial peaks appear whenever there is *data activity* on the AXI bus. The difference between these two EM behaviors can be visually spotted, as shown in Fig. 1. In other words, the detected energy pattern represents the data leakages from the embedded AXI internal buses during its communication between PS and PL. This process of finding the exact location of bus leakage through the real-time EM emanations can also be automated.

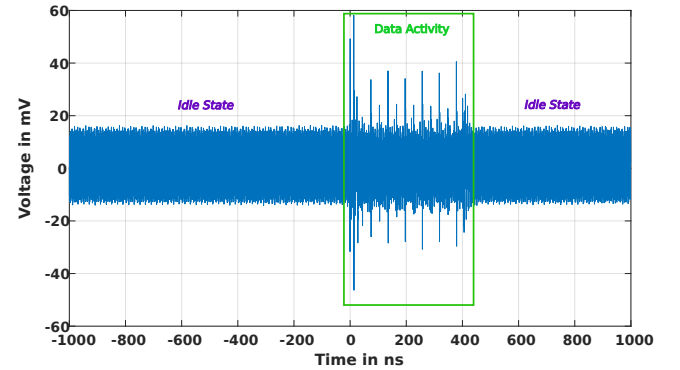


Figure 1: Idle vs. data activity state in the EM signal observed from the near-field probe.

5.2 Acquisition and Synchronization

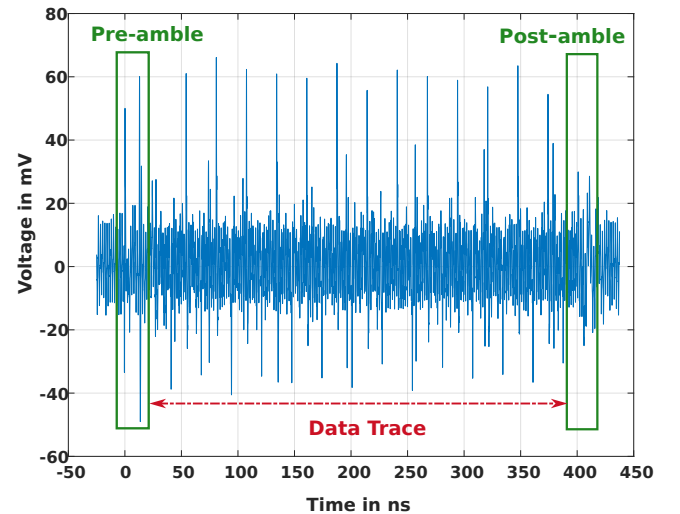


Figure 2: Zoom-in on observed traces when data activity is present on the AXI bus, i.e. when bus lines transitions from 0 to 1 and from 1 to 0 due to the AXI write data channel carrying information.

EM signal obtained from the leakage zone must be dealt with to ensure the correct reconstruction of the input image in the end. For

this single shot attack, the recorded trace must be synchronized to ensure identification of the exact part of the trace that is correlated to the input data. Unlike the majority of attacks, synchronization step in our work can be achieved without interacting with the victim system as it has been found that the acquired near field EM signal, seen in Fig. 2, can be divided into three parts: *Pre-amble*, *Data Trace* and *Post-amble*. The *pre-amble* and *post-amble* are two data independent signals that frame the *data trace*. *Pre-amble* signal defines the beginning of the bus data transmission, and hence, allowing the attacker to precisely synchronize the traces every single time. Similarly, *post-amble* signal marks the end of the data transfer on the bus which, by combining with *pre-amble* signals, can be interpreted as the total length of the data series. The *data trace* in between *pre-* and *post-amble* is utilized in the image reconstruction algorithm.

5.3 Image Reconstruction

The reconstruction of images from the EM side-channel traces requires several steps.

5.3.1 Hamming Distance. The number of input pixels P , which can be sent concurrently at a single clock cycle, is determined by the data width of the AXI bus. Hamming Distance (HD) between two groups of P pixels, transferred between two consecutive clock cycles, is the number of bit positions at which the corresponding bit has flipped (from $0 \rightarrow 1$, or from $1 \rightarrow 0$). The HD provides crucial knowledge regarding the transmission of data, which can be exploited in this case. If the victim bus data width is 32 bits, a group of pixels comprises $P = 4$ pixels.

5.3.2 Threshold. In order to define significant EM leaks in the collected traces, a positive and a negative threshold is required to detect the peaks. In our experiments, these thresholds are determined from the noise floor level, keeping a security margin to avoid false positives.

5.3.3 Peak Analysis. The peaks, positive or negative, in EM trace represent the HD between two continuous groups of P pixels sent on the AXI bus. The higher the magnitude of the peak, the bigger the HD value. Additionally, the two signs of the peaks, positive or negative, indicate the nature of the bit transition between the pixels group ($0 \rightarrow 1$ for *ascending* and $1 \rightarrow 0$ for *descending*).

Therefore, if the bit values of the current pixel group sent in cycle t is higher than the bit values of the previous pixel group $t-1$, it will induce a positive peak which represents the *ascending* transition among the P bits sent. Likewise, if the bit values of the previous pixel group sent at $t-1$ is higher than the bit values of the current pixel group at t , a negative peak will be formed in the trace which can be translated as a *descending* transition.

Note that a positive and negative bit transition on a bus corresponds to the charging and discharging of output gate capacitance or the parasitic capacitance of the bus line. For grayscale images used in this work, this ascending or descending nature of the bit transition determines whether the upcoming group of P pixels is likely to become lighter, or darker, than the current group. EM emanations induced by these transitions, $0 \rightarrow 1$ and $1 \rightarrow 0$, are additive and signed, leading to a form of *signed HD*. This signed HD is sensitive to variations in between two successive pixel groups

and not to actual pixel values directly, thus, information is lost in the interception process. Moreover, transitions on least significant bits have the same physical effect than transitions on most significant bits. This method of analyzing the peak's nature of the EM footprint to infer the input pixels, which are sent through AXI for inference process inside the FPGA is called *Peak Analysis*.

5.3.4 Reconstruction Algorithm. We propose here a method to reconstruct a binarized version of the original image from the EM traces acquired. Since the transitions on least significant bits cannot be distinguished from transitions on most significant bits, binarization is hypothesized to be the best solution to extract a decent amount of information for limited complexity. To start reconstructing the binary version of the secret input image, an image matrix is set to the background color (e.g., black in our experiment). The dimension of this matrix is deduced by the attacker as the size of input images is a public knowledge. The regeneration of the image is sequentially done, i.e., pixel by pixel, until the whole image matrix is obtained. Let P be the number of pixels which can be transmitted simultaneously on AXI bus 5.3.1. For e.g., the value of P is 4 for 32 bit bus and 8 for 64 bit bus. Starting off, the voltage values of the trace at each individual time instance is steadily scanned. When the scanned voltage value reaches positive threshold, the pixel at that certain time instance, as well as the upcoming $P-1$ pixels in the scanning order, will be assigned to white. Now the current white pixel value is maintained for the rest of the following pixels until a negative threshold is triggered. Similarly, once the negative threshold is reached during the scan, the corresponding pixel at the detected time instance and its following $P-1$ pixels are assigned to black until the detection of another positive peak. This process is repeated continuously for the entire length of EM trace signal while every pixel value inside image matrix is being assigned to either black or white.

6 EVALUATION AND RESULTS




Figure 3: Hardware setup of the experiment

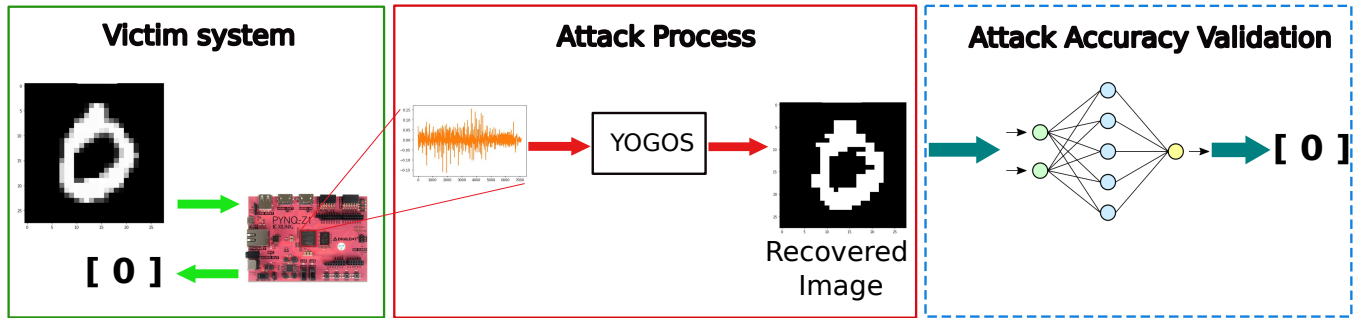


Figure 4: Victim system, YOGOS attack and validation process of recovered images with a CNN classification

6.1 Hardware Setup

We implement our proposed methodology using Pynq Z1 as it owns the advantage of having Zynq-7000, which is a SoC-FPGA with on-chip AXI buses. Pynq Z1 is equipped with FINN generated LeNet-5 BNN, pretrained on MNIST dataset. The targeted board is placed on a positioning surface with LANGER RF-R 3-2 attached to its arm for stability. A low noise amplifier with 100 MHz - 12 GHz bandwidth and 30 dB gain is attached to amplify the input signal. The oscilloscope with 20 GS/s and 6 GHz bandwidth, is used to observe the signal. The described setup can be found in Fig. 3.

6.2 Attack Implementation

According to 5.1, we sniff around the whole surface area of the Zynq-7000 SoC to find out where the EM leakages take place. After a promising leakage zone has been detected, we fix the probe to that certain position in order to capture EM side-channel traces. The exploration of real-time EM traces are done in two orthogonal orientation of the probe, 0° and 90° with a spatial resolution of 0.5 mm. It is found that the probe is more sensitive to magnetic flux on the horizontal, 0° direction where the peaks are significantly larger than the noise floor. Advanced details about identification of EM leakage zones are provided in [20].


It is important to notice that only one acquisition of EM trace, produced by the transmission of input image through AXI buses is sufficient to successfully exploit the channel and reconstruct the input image of the BNN. This is possible thanks to the data independent *pre-amble* pattern which lets us know the beginning of data transfer on the buses as explained in Section 5.2. Due to this phenomenon, we are able to detect the data related EM trace in every single acquisition. The collected one-shot EM signal will be used to reconstruct the original input image.

According to Section 5.3.2, a positive and a negative threshold must be defined initially to distinguish the nature of the peaks present in EM trace signal. In this experiment, the thresholds are set as 25% of the respective peak values. These thresholds decisions empirically ensure accurate discrimination of data signal from noise and reduce the risk of incorrect detection. After the thresholds have been determined, the images are extracted out of the EM traces by employing *peak analysis* method and reconstruction algorithm, mentioned in Sections 5.3.3 and 5.3.

The proposed method, *peak analysis*, is enacted on 100 different input images from MNIST database, 10 per 0 to 9 digits. Fig. 5 shows the original input images and the reconstructed image resulting from YOGOS.



(a) Example of original MNIST images (first line) with highly accurate recovered images through YOGOS attack (second line)



(b) Example of original MNIST images (first line) with less accurate recovered images through YOGOS attack (second line)

Figure 5: Some examples of retrieved MNIST images with YOGOS attack

6.3 Recognition Accuracy

The recovered images from YOGOS can be visually distinguished as seen in Fig. 5. However, depending on the clarity of each input MNIST image, the accuracy might change accordingly. Therefore, for the purpose of evaluating the fair precision of our attack, we used a third-party neural network classifier trained on MNIST, to compute a recognition accuracy as the process shown in Fig 4. A total of 100 randomly-chosen different input images, 10 per MNIST digit, are retrieved by applying YOGOS attack. These 100 recovered images resulting from the attack are fed into the classifier to determine the image recognition accuracy. The result is shown in Fig. 6 to compare the recognition accuracy between the original input images from standard MNIST dataset, and the recovered images by our YOGOS attack.




Figure 6: White column: Recognition accuracy (in %) per digit of the original MNIST images sent to the BNN inference. Purple column: Recognition accuracy (in %) per digit of the recovered images using YOGOS. The average accuracy for YOGOS recovered images is 89% (vs. 94% with the original images)

7 DISCUSSION

Since the concept of the YOGOS attack is based on exploiting the signed HD between two 32-bit AXI bus transitions, and not the original pixel values themselves, loss of information occurs during the process. Therefore, it is also expected that information loss will increase with the increase of bus width (*e.g.*, to 64-bits). However, even with this loss, our experiments on 32-bit bus proved to reach decent image reconstruction in terms of retrieval quality.

Moreover, it has been found that some specific handwritten number images have lower accuracy compared to the others. For instance, in the case of 4, the classifier can misclassify as 9 depending on the shape it is written. Another case would be the false detection of peaks in the acquired EM trace. As our experiments are done in open space and not in the noise free Faraday’s cage, noises occur. Nevertheless, YOGOS successfully recovers the input digit with 89% average accuracy.

8 CONCLUSION

This work has demonstrated that internal buses on SoC-FPGAs are exposed to EM side-channel attacks and leak information that can be maliciously exploited. Using a state-of-the-art AI frameworks, specifically designed for easy and efficient implementation of BNN accelerators on SoC-FPGAs, the buses, transferring clear data between PL and PS, are vulnerable. By taking advantage of this vulnerability, we have implemented an attack against FINN-implemented BNN on a Pynq Z1 target. Experiments show that we are able to successfully retrieve the secret input images in the form of MNIST handwritten characters with only one-shot inference observation. Indeed, character recognition of the BNN when fed by eavesdropped images efficiently classifies the images with 89% overall accuracy. Since our focus is on AXI on-chip communication buses, YOGOS attack is likely not to be limited to Zynq-7000 and may be applicable to other SoC platforms, as long as they contain AXI or internal buses for data transmission. Countermeasures to tackle the presented attack are not straightforward. Even if input data is encrypted when sent to the system, *e.g.*, through the network, to reinforce the system security, once the data flows inside the chip, it is decrypted and circulates as clear data in PS, before being sent to PL through internal buses for AI inference process. This means that the highlighted vulnerability and YOGOS proposed in this paper still cause a threat when system I/Os are encrypted.

On the other hand, inference on encrypted data is being studied and is a promising solution. However, it is still very costly as of now and it not yet being used [21]. Future work will include implementing the attack on different SoC hardware platforms to prove that the vulnerability exists in a broad set of systems.

REFERENCES

- [1] Sabrina De Capitani di Vimercati, Pierangela Samarati, and Sushil Jajodia. Hardware and software data security. *Computer Science and Engineering*, 15:270, 2009.
- [2] Kamel Abdelouahab, Maxime Pelcat, Jocelyn Serot, and François Berry. Accelerating cnn inference on fpgas: A survey. *arXiv preprint arXiv:1806.01683*, 2018.
- [3] Michaela Blott, Thomas B. Preußner, Nicholas J. Fraser, Giulio Gambardella, Kenneth O'Brien, Yaman Umuroglu, Miriam Leeser, and Kees Vissers. Finn-r: An end-to-end deep-learning framework for fast exploration of quantized neural networks. *ACM Trans. Reconfigurable Technol. Syst.*, 11(3), dec 2018.
- [4] Farhad Fallah. Demystifying axi interconnection for zynq soc fpga, 2020. <https://www.aldec.com/en/company/blog/145--demystifying-axi-interconnection-for-zynq-soc-fpga>, Last accessed on 2022-08-30.
- [5] Xilinx. Zynq-7000 soc data sheet: Overview, 2018. <https://docs.xilinx.com/v/u/en-US/ds190-Zynq-7000-Overview>, Last accessed on 2022-09-12.
- [6] Sonia Ben Dhia, Mohamed Ramdani, and Etienne Sicard. *Case Studies*, pages 311–394. Springer US, Boston, MA, 2006.
- [7] Yaman Umuroglu, Nicholas J. Fraser, Giulio Gambardella, Michaela Blott, Philip Heng Wai Leong, Magnus Jahre, and Kees A. Vissers. FINN: A framework for fast, scalable binarized neural network inference. *CoRR*, abs/1612.07119, 2016.
- [8] Yaman Umuroglu. Xilinx finn, 2020. <https://github.com/Xilinx/finn>.
- [9] Li Deng. The mnist database of handwritten digit images for machine learning research [best of the web]. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.
- [10] Ian Goodfellow, Yaroslav Bulatov, Julian Ibarz, Sacha Arnoud, and Vinay Shet. Multi-digit number recognition from street view imagery using deep convolutional neural networks. *arXiv preprint arXiv:1312.6082*, 12 2013.
- [11] Maria Méndez Real and Ruben Salvador. Physical side-channel attacks on embedded neural networks: A survey. *Applied Sciences*, 11(15):6790, 2021.
- [12] Lingxiao Wei, Bo Luo, Yu Li, Yannan Liu, and Qiang Xu. I know what you see: Power side-channel attack on convolutional neural network accelerators. In *Proceedings of the 34th Annual Computer Security Applications Conference, ACSAC '18*, page 393–406. Association for Computing Machinery, 2018.
- [13] Shayan Moini, Shanquan Tian, Daniel Holcomb, Jakub Szefer, and Russell Tessier. Power side-channel attacks on bnn accelerators in remote fpgas. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 11(2):357–370, 2021.
- [14] Kota Yoshida, Takaya Kubota, Shunsuke Okura, Mitsuru Shiozaki, and Takeshi Fujino. Model reverse-engineering attack using correlation power analysis against systolic array based neural network accelerator. In *2020 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–5, 2020.
- [15] Vincent Meyers, Dennis Gnad, and Mehdi Tahoori. Reverse engineering neural network folding with remote fpga power analysis. In *2022 IEEE 30th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, pages 1–10, 2022.
- [16] Laurent Sauvage, Sylvain Guilley, and Yves Mathieu. Electromagnetic radiations of fpgas: High spatial resolution cartography and attack on a cryptographic module. *ACM Trans. Reconfigurable Technol. Syst.*, 2(1), mar 2009.
- [17] Honggang Yu, Haocheng Ma, Kaichen Yang, Yiqiang Zhao, and Yier Jin. Deepem: Deep neural networks model recovery through em side-channel information leakage. In *2020 IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*, pages 209–218, 2020.
- [18] Ville Yli-Mäyry, Akira Ito, Naofumi Homma, Shivam Bhasin, and Dirmanto Jap. Extraction of binarized neural network architecture and secret parameters using side-channel information. In *2021 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–5, 2021.
- [19] DigiNews Desk. A tiny spy chip the size of a rice grain could have compromised servers of amazon, apple, cia and many other us federal agencies, 5 October 2018. <https://www.digit.in/news/general/a-tiny-spy-chip-the-size-of-a-rice-grain-could-have-compromised-servers-of-amazon-apple-cia-and-many-43945.html>.
- [20] Anonymous. Bus electrocardiogram: Vulnerability of soc-fpga internal axi buses to electromagnetic side-channel analysis. In *International Symposium and Exhibition on Electromagnetic Compatibility (EMC Europe 2023)*, 2023.
- [21] Salonik Resch, Zamshed I Chowdhury, Husrev Cilasun, Masoud Zabihi, Zhengyang Zhao, Jian-Ping Wang, Sachin Sapatnekar, and Ulya R Karpuzcu. Towards homomorphic inference beyond the edge. *arXiv preprint arXiv:2112.08943*, 2021.