

# Hound: Locating Cryptographic Primitives in Desynchronized Side-Channel Traces Using Deep-Learning

Davide Galli<sup>✉</sup>  
 DEIB  
 Politecnico di Milano  
 Milan, Italy  
 davide.galli@polimi.it

Giuseppe Chiari<sup>✉</sup>  
 DEIB  
 Politecnico di Milano  
 Milan, Italy  
 giuseppe.chiari@polimi.it

Davide Zoni<sup>✉</sup>  
 DEIB  
 Politecnico di Milano  
 Milan, Italy  
 davide.zoni@polimi.it

**Abstract**—Side-channel attacks allow the extraction of sensitive information from cryptographic primitives by correlating the partially known computed data and the measured side-channel signal. Starting from the raw side-channel trace, the preprocessing of the side-channel trace to pinpoint the time at which each cryptographic primitive is executed, and, then, to re-align all the collected data to this specific time represent a critical step to setup a successful side-channel attack. The use of hiding techniques has been widely adopted as a low-cost solution to hinder the preprocessing of side-channel traces, thus limiting side-channel attacks in real scenarios. This work introduces Hound, a novel deep-learning-based pipeline to locate the execution of cryptographic primitives within the side-channel trace even in the presence of trace deformations introduced by the use of dynamic frequency scaling actuators. Hound has been validated through successful attacks on various cryptographic primitives executed on an FPGA-based system-on-chip incorporating a RISC-V CPU while dynamic frequency scaling is active. Experimental results demonstrate the possibility of identifying the cryptographic primitives in DFS-deformed side-channel traces.

**Index Terms**—Side-channel analysis, dynamic frequency scaling, deep-learning, locating of cryptographic primitives.

## I. INTRODUCTION

Side-channel attacks pose a significant threat to modern cryptographic implementations, even when the underlying algorithms boast mathematical soundness. By exploiting vulnerabilities in the physical implementation of the cryptographic primitives (CPs), side-channel attacks can exploit the unintended information leakage generated from electronic devices during the execution of CPs [1]. The leakage, often manifested in power consumption [2] or electromagnetic emissions [3], can be analyzed by attackers to potentially reveal confidential information by correlating it with partially known processed data.

Traditionally, properly implemented masking [4], [5] and hiding [6]–[8] countermeasures offer robust protection against side-channel attacks, leading to successful security evaluations. Masking countermeasures split sensitive intermediate values into independent shares that are processed separately to reduce significantly the correlation between each share and

the secret key. Hiding countermeasures introduce randomness into the targeted side-channel signal to degrade the signal-to-noise ratio.

Over the past two decades, researchers have proposed various methods to enhance the effectiveness of side-channel analysis, such as differential power analysis (DPA) [2], template attacks [9], correlation power analysis (CPA) [10], as well as deep-learning-based solutions [11], [12]. Notably, all the proposed techniques share two key requirements. First, they require a large number of executions of the same CP with different inputs. Second, the attacker needs to locate and align in time all the CP executions in the side-channel trace to feed the attack method of choice.

In a controlled laboratory environment, the attacker has full access to the target device, simplifying the security assessment of a CP implementation. Security evaluation boards, like SASEBO SAKURA-II [13] and NewAE CW305 [14], provide so-called trigger pins, facilitating the temporal alignment between CP executions and corresponding side-channel signals.

In contrast, attacking the implementation of a CP in the real world presents a more significant challenge. Indeed, the attacker needs to locate the CP within the side-channel trace. Specific settings still permit a rough alignment of the measured side-channel signal with CP execution even in real-world scenarios, e.g., by leveraging specific logic events happening in the computing platform or by using pattern-matching techniques applied to the side-channel trace to generate so-called virtual triggers [15], [16].

However, the importance of pinpointing the executions of the CPs in the side-channel trace even without a triggering infrastructure motivated several contributions, i.e., [17], [18]. In a similar manner, the continuous improvement in the attack techniques fueled the use of hiding techniques to hinder the task of locating a CP, thus preventing the localization of the CPs in the side-channel trace as an additional security countermeasure.

[19] is the seminal work that presented a deep-learning pipeline to locate the CPs within side-channel traces deformed by means of random delay countermeasures. However, random

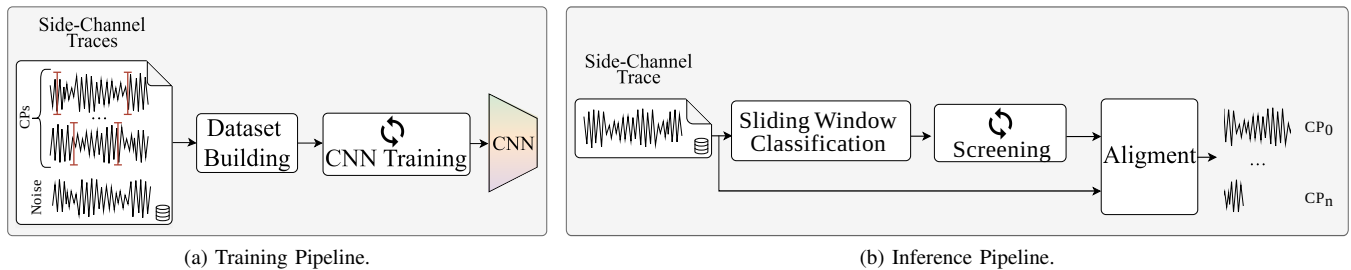


Fig. 1: Overview of the proposed Hound pipeline for locating cryptographic primitives in frequency-scaled side-channel traces, divided into *training* and *inference* pipelines.

delay introduces a limited trace deformation, and its implementation requires custom hardware or software implementations inducing non-negligible performance overheads [20]. In contrast, Dynamic Frequency Scaling (DFS) actuators offer a cheap and standard solution that has been widely investigated to implement hiding techniques. The possibility of employing hundreds of operating frequencies allows to effectively deform the side-channel trace at the point of preventing the correct pinpointing of the CP and, consequently, invalidating the possibility of attacking the target device even if the attacker can deploy an effective side-channel attack.

**Contributions** - This paper introduces Hound, a deep-learning methodology for locating cryptographic primitives within heavily randomized side-channel traces where the trace deformation is obtained by means of a randomized DFS actuator. Hound offers three contributions to the state of the art:

- We propose a deep-learning pipeline that automatically locates and aligns CPs within side-channel power traces, even in the presence of DFS countermeasures. Thus, the need for triggering infrastructure is eliminated, which is a significant hurdle in real-world side-channel analysis.
- We evaluate the effectiveness of our solution across various cryptographic primitives. Furthermore, we perform successful side-channel attacks after CP localization and alignment to validate the quality of the results and compare our solution against state-of-the-art techniques.
- To facilitate further research and ensure reproducibility, we have released Hound as an open-source tool along with a dataset of relevant side-channel traces. The tool and dataset are available on GitHub <sup>1</sup>.

The rest of the paper is organized into four sections. Section II discusses the academic and commercial tools to locate the CPs within side-channel traces. Section III presents the proposed Hound approach. Section IV details the experimental results. Finally, Section V presents the conclusions.

## II. RELATED WORKS

Existing methods for locating CPs within side-channel traces primarily rely on trigger signals from security evaluation boards (SASEBO SAKURA-II [13] and NewAE CW305 [14])

or pattern matching against pre-computed CP templates. Commercially available FPGA-based devices, like the Riscure icWaves [16] and NewAE ChipWhisperer Pro [15], offer virtual-triggering features, producing a trigger pulse upon real-time identification of a pattern in the side-channel trace being monitored. Barengi et al. [17] introduced a method utilizing matched filters for efficiently pinpointing the AES-128 cryptosystem within a power trace, achieving robustness in environments with interrupts and busy waits that could change the trace’s profile. Nonetheless, this approach fails when it comes to recognizing interrupted CPs, which are discarded by the matched filter. Becker et al. [21] suggested a waveform-matching trigger system designed to identify CPs by comparing them to a pre-calculated CP template.

All these methods need a pre-computed template of the CP to be localized to work correctly. To address this issue, Trautmann et al. explored in [18] a semi-automatic technique for locating CPs by building the template online, starting from some CP characteristics, such as the number of rounds.

However, using architectural-level techniques to morph the power trace represents an easy-to-implement and effective countermeasure to deceive pattern-matching-based solutions. For instance, employing time-sharing multithreading on a single-core microcontroller [22] can obstruct an attacker’s ability to accurately locate CPs, with interrupt service routines further modifying the side-channel trace’s form. Consequently, more sophisticated methods for detecting CPs in the side-channel trace have been developed. Chiari et al. [19] developed a deep-learning approach for identifying deformed CPs through random delay, employing a Convolutional Neural Network (CNN) trained to differentiate CPs from the remaining portions of the power trace. Even though random delay-based hiding techniques generate some degree of CP deformation, they are limited in their ability to introduce significant variability in the power trace [20].

Despite these advancements, current state-of-the-art techniques struggle with highly deformed traces, particularly when the computing platform employs effective hiding countermeasures like DFS [23]. These countermeasures profoundly alter the timing of operations and continuously morph the side-channel signal, making it challenging to isolate CPs using traditional methods. To address this limitation, we propose a

<sup>1</sup><https://github.com/hardware-fab/Hound>




Fig. 2: Focus on the proposed Hound training pipeline, divided into *Dataset Building* and *CNN Training*.  $Class_0$ ,  $class_1$  and  $class_2$  contain CP start part windows, CP spare parts windows, and noise windows, respectively.

deep-learning-based technique for CP localization in power traces collected from platforms implementing DFS as a countermeasure.

### III. METHODOLOGY

This section introduces Hound, a novel deep-learning method for CP localization within highly deformed side-channel traces, specifically designed to address scenarios where the target platform employs DFS as countermeasure to effectively obfuscate the measurements.

**Threat model** - Similarly to profiled attacks, we assume the attacker has access to an identical copy of the target device, i.e., a clone. However, we focus on a realistic scenario where the attacker can only run selected applications and monitor the resulting side channels without complete control over the cloned device. Specifically, the attacker cannot activate or deactivate the hiding countermeasure, namely DFS. Nonetheless, the attacker retains the ability to probe the clone device, enabling them to detect when the CP initiates. This probing facilitates precise labeling of the beginning of each CP on the clone device.

**Hound** - Figure 1 depicts the proposed Hound methodology consisting of a training and inference pipeline. Figure 1a reports the *Training Pipeline*, which aims at creating a classifier that can sort a window of the side-channel trace into three categories: *start of a CP*, *spare part of a CP*, or just *noise* (see Figure 2). This framework locates where the CPs begin in the side-channel trace without removing the obfuscation countermeasure. We assume that the attacker leverages an exact copy of the target device to collect a noise trace and some cipher traces for the training. The noise trace is obtained by running general-purpose applications that are not the CP. Such information is necessary to train a neural network to distinguish the beginning of the CPs from the execution of other applications. Each cipher trace is gathered while a single CP is running, where the attacker decides what data and keys to use. All measurements are collected with the obfuscation mechanism active since the attacker cannot turn it off. Starting with measured raw traces, the *Dataset Building* stage creates a database of windows from the noise and cipher traces, each being  $N$  samples long. This database is then used to train a CNN classifier.

Figure 1b reports the *Inference Pipeline*. At inference time, the goal is to identify the CPs within a novel side-channel trace

from the target device by employing the trained CNN. This pipeline is divided into three key stages: *Sliding Window Classification*, *Screening*, and *Alignment*. The inference pipeline takes in a single side-channel trace and segments it to mark the start of each CP. Initially, the *Slicing* stage processes a side-channel trace and generates a series of windows to classify with the trained CNN. Starting from the classified windows, representing a segmentation of the input trace, the *Screening* stage produces a sequence of time instants marking the start of each CP in the trace. Finally, the *Alignment* stage chunks the original side-channel trace based on these outputs, arranging the identified CPs accordingly.

The details related to the dataset creation and the CNN architecture are discussed in Section III-A, while the sliding window classification and the screening procedure are detailed in Section III-B.

#### A. Training Pipeline

The Training Pipeline, detailed in Figure 2, is divided into *Dataset Building* and *CNN Training* stages. Starting from a raw collection of side-channel traces, the pipeline generates a dataset and trains a CNN classifier to distinguish the starting of a CP.

**Dataset Building** - The first stage in the training pipeline is *Dataset Building*, which receives a collection of side-channel traces as input and generates a dataset suitable for training the CNN. The attacker uses a clone of the target computing platform to create a noise trace and a set of cipher traces. The noise trace captures the execution of various general-purpose applications, excluding the target CP. Each cipher trace is collected while executing a single CP, where the attacker controls the plaintext and secret key. Notably, the attacker cannot disable the obfuscation mechanism, so the DFS remains active during all trace collection. Considering the proposed threat model, the attacker can only execute a chosen application and measure the corresponding side channel on the clone device. However, the attacker can probe the clone device to detect the start of each CP. In contrast, probing is not required for the noise trace. Once trained, the CNN works on the target architecture that implements DFS without the need for probing. For each cipher trace  $i$  of length  $L_i$  samples, the starting  $N$  samples are labeled as *start of the CP* (see  $class_0$  in Figure 2). The remaining  $L_i - N$  samples are equally split into consecutive windows of width  $N$  and labeled as *spare*




Fig. 3: Focus on the proposed Hound inference pipeline, divided into *Sliding Window Classification*, *Screening*, and *Alignment*.

part of the CP (see  $class_1$  in Figure 2). Moreover, we extract a random set of  $N$ -sample windows from the noise trace and we label each of them as *noise* (see  $class_2$  in Figure 2). Although the methodology aims to identify only the beginning of each CP, the division into three classes allows CNN to maximize the distinction over the three cases.

**Convolutional Neural Network** - The architecture of the 1D CNN is adapted from the CNN proposed in [19]. The CNN takes a window of  $N$  samples from the side-channel trace as input and outputs a classification score vector. The CNN’s structure is organized into six sequential components: a convolutional block, a pair of residual blocks, a global average pooling layer, a fully-connected block, and a softmax layer. Each convolutional block features a 1D convolutional layer, a batch normalization layer [24], and a ReLU activation function. The residual blocks [25], each containing two convolutional blocks, are augmented with shortcut connections to perform element-wise feature summation. All the convolutional layers implement a kernel size of 64, a stride of 1, and zero padding to maintain the sample count to  $N$ . The first convolutional layer and the one in the first residual block implement 16 filters, while the subsequent residual block doubles the filter count to 32. The global average pooling layer averages the obtained features over the temporal dimension  $N$ . The feature vector is then processed through two fully-connected layers with a ReLU activation function. The softmax layer completes the architecture, generating a 3-dimensional vector with the classification scores for each class.

Notably, DFS introduces a significant amount of randomization in the side-channel trace. To address this, the CNN proposed in [19] has been adapted to work with three classes, i.e., *start of the CP*, *spare part of the CP*, and *noise*. This modification allows the CNN to learn the pattern of the highly obfuscated side-channel trace better, maximizing the discrimination between the *start of the CP* and the other two classes.

### B. Inference Pipeline

The Inference Pipeline, detailed in Figure 3, is divided into *Sliding Window Classification*, *Screening*, and *Alignment* stages. Leveraging the trained CNN classifier, the pipeline locates the starting sample of each CP present in a new side-channel trace. **Sliding Window Classification** - The first stage

in the inference pipeline is *Sliding Window Classification*. It takes a new side-channel trace, slices it into  $N$ -sample windows, and uses the trained CNN to output a classification score to label each window as *start of the CP*, *spare part of the CP*, or *noise* (see *Sliding Window Classification* in Figure 3). The *Slicing* stage implements a sliding window procedure to slice the side-channel trace in input. It takes three inputs, i.e., the side-channel trace, the size of the sliding window ( $N$ ), and the stride ( $s$ ), and outputs an ordered set of  $N$ -sample windows to feed the CNN.

The CNN’s softmax output is a 3-dimensional probability distribution of classes, classifying each  $N$ -sample window as *start of the CP*, *spare part of the CP* or *noise*. The inference process considers the class with the highest probability outcome among the three as the CNN’s output.

Since the CNN’s output might be noisy and not always suitable for directly determining the exact location of the CPs, the inference pipeline incorporates a subsequent *Screening* stage to refine the CNN’s output (see *Screening* in Figure 3).

**Screening** - The *Screening* stage refines the segmentation output from the *Sliding Window Classification* stage to identify the beginning of the CPs in the side-channel trace. Algorithm 1 highlights the screening procedure. The screening algorithm (see Algorithm 1) takes the segmentation output from the *Sliding Window Classification* stage, denoted as  $seg$ , and some parameters, i.e., the stride  $s$  used for the sliding window, the kernel size  $k$  for the polish procedure (see line 10 of Algorithm 1), and the average CP length  $avgCP$ . In the end, it returns the list  $starts$  of samples corresponding to the beginning of each CP instance in the processed side-channel trace.

The algorithm is an iterative process that involves three main steps: POLISH, EXTRACT, and REFINE. The core idea behind Algorithm 1 is to start with an aggressive polish and incrementally refine it by attempting to uncover new CPs. The polishing step (see line 10 of Algorithm 1) employs a majority filter across the whole segmentation. This filter traverses the input signal for a given kernel size  $k$ , replacing each point with the most frequently occurring value within a  $k$ -sized window. The extraction step (see line 14 of Algorithm 1) identifies the initial sample of each CP. It pinpoints the indices of the falling edge down to class 0, meaning those samples with a value of 0 where the preceding value is not 0. These samples are then

---

**Algorithm 1** Screening

---

**Input:**  $seg, s, k, avgCP$   
**Output:**  $starts$

```
1:  $starts \leftarrow []$ 
2: do
3:    $polishedSeg \leftarrow POLISH(seg, k)$ 
4:    $newStarts \leftarrow EXTRACT(polishedSeg, s)$ 
5:    $k, seg \leftarrow REFINE(k, avgCP, newStarts, seg)$ 
6:    $starts.update(newStarts)$ 
7: while  $seg$  not empty &&  $k \geq 1$ 
8: return  $starts$ 
9:
10: procedure  $POLISH(seg, k)$ 
11:    $polishedSeg \leftarrow majorityFilter(seg, k)$ 
12:   return  $polishedSeg$ 
13: end procedure
14: procedure  $EXTRACT(seg, s)$ 
15:    $newStarts \leftarrow []$ 
16:   for  $i, sample$  in  $enumerate(seg)$  do
17:     if  $sample$  is falling edge to class 0 then
18:        $newStarts.update(i \times s)$ 
19:     end if
20:   end for
21:   return  $newStarts$ 
22: end procedure
23: procedure  $REFINE(k, avgCP, starts, seg)$ 
24:    $k \leftarrow refineK(k)$ 
25:    $minCP \leftarrow refineMin(avgCP, starts)$ 
26:    $refineSegs \leftarrow []$ 
27:   // Look if can be still two consecutive CPs
28:   for  $s, s_{next}$  in  $starts$  do
29:     if  $s_{next} - s > 2 \times minCP$  then
30:        $subSeg \leftarrow seg[s : s_{next}]$ 
31:        $refineSegs.update(subSeg)$ 
32:     end if
33:   end for
34:   return  $k, refineSegs$ 
35: end procedure
```

---

scaled by  $s$ , i.e., the stride in the sliding window classification. Finally, the refining step (see line 23 of Algorithm 1) fine-tunes the algorithm’s parameters by reducing kernel size  $k$  and generating a list of sub-segmentation for subsequent iterations. These sub-segments encompass those regions of the input segmentation for which multiple hidden CPs might still exist. The process begins by determining the minimum length of a CP  $minCP$  (see line 25 of Algorithm 1), which is derived by comparing the input’s average CP length and the CP’s starting points identified so far. Next, the algorithm examines pairs of consecutive CP starting points (see line 28 of Algorithm 1). If the distance between them exceeds twice the  $minCP$ , it indicates the potential presence of additional CPs within that interval. Consequently, the corresponding sub-segmentation is queued for further processing.

The algorithm persists until no further refinement is possible, i.e., when  $k$  drops below one or the refining step yields an empty list of sub-segmentations  $seg$  (see line 7 of Algorithm 1).

**Alignment** - The last stage of the inference pipeline is *Alignment*. It leverages the CP starting points identified by the previous stages to chunk the input side-channel trace into temporally aligned segment.

#### IV. EXPERIMENTAL EVALUATION

This section presents the experimental results of Hound in four parts. Section IV-A outlines the experimental setup. Experimental results targeting the CNN training and inference pipelines are discussed in Section IV-B and Section IV-C, respectively. A comprehensive example workflow, including a successful side-channel attack and a comparison with state-of-the-art proposals, is provided in Section IV-D.

##### A. Experimental Setup: Hardware and Software

We utilized the NewAE CW305 development board [14] as our validation platform, designed specifically to facilitate high-precision, low-noise side-channel measurements. This board incorporates an AMD Artix7-100 FPGA. Power traces were captured using a Picoscope 5244d digital sampling oscilloscope (DSO), operating at a sampling rate of 125 Msamples/s and a resolution of 12 bits. For our reference computing platform, we employed a 32-bit RISC-V System-on-Chip [26] implemented on the FPGA. Taking steps from [27], we implemented the Dynamic Frequency Scaling mechanism coupled with a true random number generator (TRNG) [28] to deliver a novel actuator that can randomize the operating frequency, thus introducing non-negligible deformations in the collected side-channel traces. The TRNG generates a random number used to select a frequency from a pool of 760 available frequencies ranging from 5 MHz to 100 MHz, with a step of 125 kHz. Notably, DFS remained enabled throughout the entire experimental phase. Consequently, a new frequency is requested as soon as a configuration is locked. As a result, each cryptographic primitive experiences a variable number of clock frequency reconfigurations, e.g., an average of 41 reconfigurations per AES encryption, leading to significant trace deformation. As the cryptographic primitives of choice, we selected the constant-time, unprotected version of three ciphers, i.e., AES-128, Clefia-128, and Camellia-128, from the OpenSSL software codebase [29], and a masked version of Tiny-AES-128 [30].

##### B. Training Evaluation

This section details the evaluation of the performance of the CNN by elaborating on the dataset-building process and the training metrics.

The training of the CNN leverages a *NVIDIA 1080 Ti* employing the PyTorch software programming framework. The training datasets were built following the methodology outlined in Section III-A, starting from a collection of 262 144 power traces. The three classes, namely *start of a CP*, *spare*

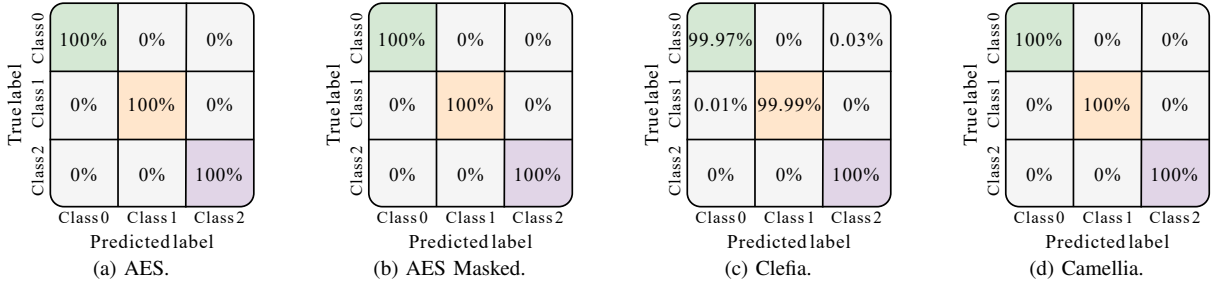


Fig. 4: Test confusion matrices for the different cryptosystems affected by DFS.

TABLE I: Parameters for each pipeline stage and metrics score over all the tested ciphers.

Cipher	General-Purpose Applications	Inference parameters				CNN hyperparameters			Metrics		
		avgCP	k	N	s	batch size	lr	dropout	Hits	Mean IoU	Std IoU
AES	✓	145k	150	10k	62	256	0.01	0.2	100%	97.01%	1.65%
	✗	120k	150	10k	62				100%	93.62%	2.75%
AES Masked	✓	50k	10	5k	50	256	0.007	0.35	100%	97.13%	1.85%
	✗	50k	150	5k	50				100%	95.05%	1.43%
Clefia	✓	80k	150	3k	80	256	0.007	0.3	100%	97.90%	2.88%
	✗	80k	150	3k	80				100%	98.46%	0.76%
Camellia	✓	4.4k	80	1.1k	50	128	0.007	0.4	100%	91.92%	4.79%
	✗	4.1k	63	1.1k	50				100%	93.09%	3.63%

part of a CP, and noise, were balanced with equal representation of 33% components each. Table I reports the window sizes  $N$  for each cipher. As a general guideline, the window size  $N$  is adjusted to approximate one round of the targeted CP. For fast encryption algorithms, like Camellia, the window size is increased to capture a more significant portion of the CP without compromising the quality of results. The windows corresponding to the ciphers, i.e., *start of a CP* and *spare part of a CP*, are evenly distributed across the key bytes. Following standard deep-learning practices, the collected datasets were divided into training, validation, and testing sets, constituting 80%, 10%, and 10% of the total, respectively.

Each neural network underwent training for 25 epochs, where an epoch encompasses a complete iteration over the training set. The Adam optimizer [31] was employed to minimize the cross-entropy loss, paired with the one-cycle learning rate scheduler. The validation error was assessed after each epoch, and the network with the lowest error was chosen for further evaluation. Subsequently, this network was employed to assess the performance of unseen traces during the inference phase. Since DFS introduces significant trace deformation with high variability, we noticed how CNNs that did not reach an accuracy of at least 99% were not always able to classify the CPs in the sliding windows procedure correctly. Hyperparameters such as training batch size, learning rate (lr), and dropout were carefully chosen to address this issue (see the *CNN hyperparameters* column in Table I).

A distinct CNN model has been trained using a custom dataset for each cipher under examination. As an indicator of the effectiveness of the trained CNNs, their confusion matrices are presented in Figure 4. The column indices represent the true classes, while the row indices represent the predicted ones.

Notably, the trained classifiers can discriminate excellently between the three classes, as highlighted by the high percentages on the main diagonal of each matrix.

### C. Inference Evaluation

A brief experimental campaign was conducted to evaluate the best pipeline parameters. Table I displays the values of the average cipher length *avgCP*, the strides  $s$ , and the initial kernel size  $k$  for each cipher. We evaluate the performance of the *Sliding Window Classification* and *Screening* stages by computing two quality metrics: *hits* and *Intersection over Union (IoU)*. *Hits* represent the proportion of correctly identified CPs relative to the total number of actual CPs within the trace. The *IoU*, on the other hand, measures the normalized overlap between the predicted CP starting point and the ground truth, taking into account the CP’s overall length. A higher *IoU* score signifies better agreement between the predicted and actual CP starting points. It is formally defined in Equation 1, where  $P_i$  and  $GT_i$  are the segments corresponding to the CP  $i$  in the predicted and ground-truth segmentation, respectively.

$$IoU_i = \frac{|P_i \cap GT_i|}{|P_i \cup GT_i|} \quad (1)$$

Since Hound focuses on pinpointing the CP’s beginning within the trace, the *IoU* metric disregards the CP’s ending point. The predicted ending point is always assumed to be accurate. Table I shows the mean and standard deviation of *IoU*, where a higher score corresponds to a better segmentation.

We conducted tests on the inference pipeline for each cipher using consecutive cipher executions and encryptions interleaved with random general-purpose applications. The segmentation *hits* score achieved 100% for every cryptographic

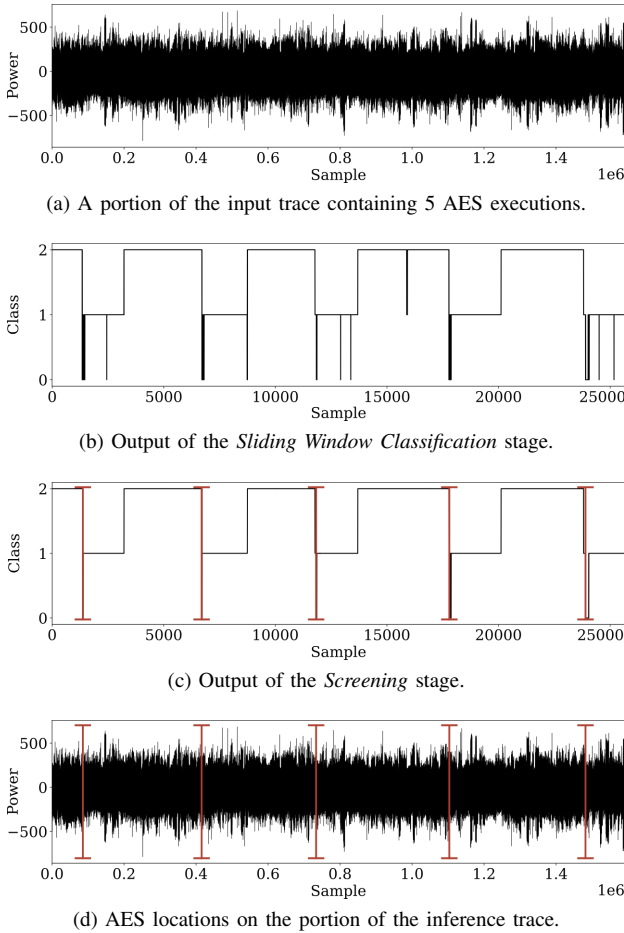


Fig. 5: Example of Hound inference pipeline applied to a side-channel trace affected by DFS that contains 5 AES executions mixed with general-purpose applications.

TABLE II: Segmentation and side-channel attack (SCA) results targeting AES-128. Reported results consider the presence (or not) of general-purpose applications interleaved with the CPs in the side-channel trace.

	Hiding Mechanism	General-Purpose Applications	Hits (%)	SCA (N. CPs)
[17]	None	✓	0%	✗
		✗	0%	✗
[18]	None	✓	0%	✗
		✗	0%	✗
[19]	Random Delay Insertion	✓	1637%	✗
		✗	879%	✗
Hound	DFS	✓	100%	571
		✗	100%	74

algorithm in both scenarios, i.e., consecutive encryption and interleaved with noisy general-purpose applications, consistently managing to identify all executions. The *IoU* score, on the other hand, ranges from 91.92% to 98.46% for the different ciphers. The results demonstrate the robustness of the proposed methodology in locating CPs with high accuracy and precision, even in the presence of DFS-induced trace deformation.

#### D. The Complete Attack Flow

This section showcases the efficacy of Hound in locating cryptographic primitives obfuscated by DFS. We present a comprehensive attack flow that takes an unknown side-channel trace and retrieves the secret key. The key is retrieved using a deep-learning-based attack as an effective side-channel analysis technique.

Figure 5 shows the results of applying each step of the proposed inference pipeline to locate the CPs into a portion of a trace containing five executions of AES-128, i.e., the CP. The entire side-channel trace contains 367 CPs collected from a platform implementing a randomized dynamic frequency scaling mechanism. Starting from the portion of the side-channel trace in input, Figure 5b depicts the segmentation output of the *Sliding Window Classification* stage that correctly highlights the presence of five CPs interleaved with random general-purpose applications. Notably, it is easy to detect also the end of the AES encryptions by looking at *Class 1*, as the classifier is trained to recognize the spare part of the CPs. The number of samples in the x-axis of the plot is reduced by a factor of 62, equal to the value of stride  $s$  during the classification. The *Screening* stage cleans the segmentation output by removing spurious activations of the various classes (see Figure 5c). The time instants locating the beginning of the five CPs are applied to the original side-channel trace, thus allowing an easy time realignment (see Figure 5d).

As an effective side-channel attack, we employed a CNN developed using the tool proposed by [12]. The attacking CNN targets the *sub-byte* intermediate and has been trained starting from the same side-channel collection described in Section III-A. A minor sub-sampling is used to ease the training process of the attacking CNN, reducing the input size and expediting the process. Column SCA in Table II enumerates the number of CPs needed to reach a guessing entropy of 1, i.e., guess the secret key correctly.

**Comparison against state of the art** - Building on the prior evaluation, Table II further stresses the effectiveness of the proposed methodology, detailing the attack results under various scenarios. The target CP is AES-128. We compare our approach against three state-of-the-art proposals, i.e., [17], [18], and [19]. For each state-of-the-art approach, we report the hiding mechanism for which the methodologies were designed by their authors (see the *Hiding Mechanism* column in Table II). We considered the random CPs' execution within a set of noisy general-purpose applications and their execution without any noisy application (see the *General-Purpose Applications* column in Table II). For each analyzed scenario, Hound can correctly identify the beginning of all the CPs in the side-channel trace, leading to a successful side-channel attack. In contrast, the three state-of-the-art methodologies fail to locate the CPs in the side-channel trace due to the high DFS obfuscation, and thus, the subsequent side-channel attack is unsuccessful. More specifically, [17] and [18] achieve a 0% hit rate, indicating they cannot locate any CPs in the trace. On the

other hand, [19] exhibits anomalously high hit rates of 879% and 1637% for the two scenarios. These excessively high rate values highlight a significant number of false positives determining the subsequent failure of the side-channel attacks.

## V. CONCLUSIONS

We presented Hound, a novel deep-learning technique for pinpointing cryptographic primitives within highly desynchronized side-channel traces. Unlike existing methods, Hound successfully locates the CPs in side-channel traces collected from computing platforms implementing dynamic frequency scaling mechanism as an effective hiding countermeasure.

We conducted a comprehensive evaluation to locate both sequences of consecutive cryptographic primitives and those interleaved with other applications. The evaluation considered different cryptographic primitives. The experimental results, obtained by running a diverse set of applications on an FPGA-based RISC-V processor, validate the effectiveness of our approach in enabling successful side-channel attacks. These results also emphasize the shortcomings of current state-of-the-art solutions.

To promote reproducibility and future research, we have made Hound publicly available as open-source software, along with a collection of test traces.

## REFERENCES

- [1] S. Mangard, E. Oswald, and T. Popp, *Power analysis attacks: Revealing the secrets of smart cards*. Springer Science & Business Media, 2008, vol. 31.
- [2] P. Kocher, J. Jaffe, and B. Jun, “Differential power analysis,” in *Annual international cryptology conference*. Springer, 1999, pp. 388–397.
- [3] D. Agrawal, B. Archambeault, J. R. Rao, and P. Rohatgi, “The em side—channel (s),” in *Cryptographic Hardware and Embedded Systems—CHES 2002: 4th International Workshop Redwood Shores, CA, USA, August 13–15, 2002 Revised Papers 4*. Springer, 2003, pp. 29–45.
- [4] L. Goubin and J. Patarin, “Des and differential power analysis the “duplication” method,” in *Cryptographic Hardware and Embedded Systems: First International Workshop, CHES’99 Worcester, MA, USA, August 12–13, 1999 Proceedings 1*. Springer, 1999, pp. 158–172.
- [5] H. Gross, S. Mangard, and T. Korak, “Domain-oriented masking: Compact masked hardware implementations with arbitrary protection order,” *Cryptology ePrint Archive*, Paper 2016/486, 2016, <https://eprint.iacr.org/2016/486>. [Online]. Available: <https://eprint.iacr.org/2016/486>
- [6] C. Clavier, J.-S. Coron, and N. Dabbous, “Differential power analysis in the presence of hardware countermeasures,” in *Cryptographic Hardware and Embedded Systems—CHES 2000: Second International Workshop Worcester, MA, USA, August 17–18, 2000 Proceedings 2*. Springer, 2000, pp. 252–263.
- [7] S. Yang, W. Wolf, N. Vijaykrishnan, D. N. Serpanos, and Y. Xie, “Power attack resistant cryptosystem design: A dynamic voltage and frequency switching approach,” in *Design, Automation and Test in Europe*. IEEE, 2005, pp. 64–69.
- [8] J.-S. Coron and I. Kizhvatov, “An efficient method for random delay generation in embedded software,” in *International Workshop on Cryptographic Hardware and Embedded Systems*. Springer, 2009, pp. 156–170.
- [9] S. Chari, J. R. Rao, and P. Rohatgi, “Template attacks,” in *Cryptographic Hardware and Embedded Systems - CHES 2002, 4th International Workshop, Redwood Shores, CA, USA, August 13-15, 2002, Revised Papers*, ser. Lecture Notes in Computer Science, vol. 2523. Springer, 2002, pp. 13–28.
- [10] E. Brier, C. Clavier, and F. Olivier, “Correlation power analysis with a leakage model,” in *International workshop on cryptographic hardware and embedded systems*. Springer, 2004, pp. 16–29.
- [11] R. Benadjila, E. Prouff, R. Strullu, E. Cagli, and C. Dumas, “Deep learning for side-channel analysis and introduction to ascad database,” *Journal of Cryptographic Engineering*, vol. 10, no. 2, pp. 163–188, 2020.
- [12] J. Rijdsdijk, L. Wu, G. Perin, and S. Picek, “Reinforcement learning for hyperparameter tuning in deep learning-based side-channel analysis,” *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pp. 677–707, 2021.
- [13] Inrevium Inc., “Sasebo-gii-32.” [Online]. Available: [https://digilent.com/reference/sasebo\\_g\\_ii\\_32/sasebo\\_g\\_ii\\_32](https://digilent.com/reference/sasebo_g_ii_32/sasebo_g_ii_32)
- [14] NewAE Technology Inc., “Cw305 artix fpga target,” 2018. [Online]. Available: <https://rtfm.newae.com/Targets/CW305%20Artix%20FPGA>
- [15] NewAE, “Chipwhisperer pro,” [https://wiki.newae.com/Tutorial\\_P1\\_Using\\_a\\_Custom\\_Trigger](https://wiki.newae.com/Tutorial_P1_Using_a_Custom_Trigger), 2020.
- [16] Riscure, “icwaves,” <https://www.riscure.com/security-tools/inspector-hardware>, 2020.
- [17] A. Barengi, G. Falcetti, and G. Pelosi, “Locating side channel leakage in time through matched filters,” *Cryptography*, vol. 6, no. 2, 2022. [Online]. Available: <https://www.mdpi.com/2410-387X/6/2/26>
- [18] J. Trautmann, A. Beckers, L. Wouters, S. Wildermann, I. Verbauwhede, and J. Teich, “Semi-automatic locating of cryptographic operations in side-channel traces,” *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2022, no. 1, p. 345–366, Nov. 2021. [Online]. Available: <https://tches.iacr.org/index.php/TCHES/article/view/9300>
- [19] G. Chiari, D. Galli, F. Lattari, M. Matteucci, and D. Zoni, “A deep-learning technique to locate cryptographic operations in side-channel traces,” in *2024 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2024, pp. 1–6.
- [20] F. Durvaux, M. Renaud, F.-X. Standaert, L. van Oldeneel tot Oldenzeel, and N. Veyrat-Charvillon, “Cryptanalysis of the ches 2009/2010 random delay countermeasure,” 2012, <https://eprint.iacr.org/2012/038>. [Online]. Available: <https://eprint.iacr.org/2012/038>
- [21] A. Beckers, J. Balasch, B. Gierlichs, I. Verbauwhede, F. Standaert, and E. Oswald, “Design and implementation of a waveform-matching based triggering system,” pp. 184 – 198, 2016-01-01.
- [22] I. Frieslaar and B. Irwin, “Investigating multi-thread utilization as a software defence mechanism against side channel attacks,” in *Proceedings of the 8th International Conference on Signal Processing Systems*, ser. ICSPS 2016. New York, NY, USA: Association for Computing Machinery, 2016, p. 189–193. [Online]. Available: <https://doi.org/10.1145/3015166.3015176>
- [23] B. Hettwer, K. Das, S. Leger, S. Gehrler, and T. Güneysu, “Lightweight side-channel protection using dynamic clock randomization,” in *2020 30th International Conference on Field-Programmable Logic and Applications (FPL)*. IEEE, 2020, pp. 200–207.
- [24] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” *CoRR*, vol. abs/1502.03167, 2015. [Online]. Available: <http://arxiv.org/abs/1502.03167>
- [25] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” 06 2016, pp. 770–778.
- [26] L. Denisov, A. Galimberti, D. Cattaneo, G. Agosta, and D. Zoni, “Design-time methodology for optimizing mixed-precision cpu architectures on fpga,” *Journal of Systems Architecture*, vol. 155, p. 103257, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1383762124001942>
- [27] D. Galli, A. Guarisco, W. Fornaciari, M. Matteucci, and D. Zoni, “The impact of run-time variability on side-channel attacks targeting fpgas,” 2024. [Online]. Available: <https://arxiv.org/abs/2409.01881>
- [28] D. Galli, A. Galimberti, W. Fornaciari, and D. Zoni, “On the effectiveness of true random number generators implemented on fpgas,” in *International Conference on Embedded Computer Systems*. Springer, 2022, pp. 315–326.
- [29] OpenSSL, “Tls/ssl and crypto library,” <https://github.com/openssl/openssl>, 2023.
- [30] MEALITY, “Masked aes implementation,” <https://github.com/CENSUS/masked-aes-c>, 2020.
- [31] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” 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. [Online]. Available: <http://arxiv.org/abs/1412.6980>