

# ABBY: Automating leakage modeling for side-channels analysis

Omid Bazangani  
Radboud University  
Nijmegen, The Netherlands  
Omid.bazangani@ru.nl

Ileana Buhan  
Radboud University  
Nijmegen, The Netherlands  
Ileana.buhan@ru.nl

Alexandre IOOSS  
Inria Paris-Rocquencourt  
Rocquencourt, France  
alexandre.iooss@crans.org

Lejla Batina  
Radboud University  
Nijmegen, The Netherlands  
Lejla.Batina@ru.nl

## ABSTRACT

We introduce ABBY, an open-source side-channel leakage profiling framework that targets the microarchitectural layer. Existing solutions to characterize the microarchitectural layer are device-specific and require extensive manual effort. The main innovation of ABBY is the collection of data, which can automatically characterize the microarchitecture of a target device and has the additional benefit of being scalable.

Using ABBY, we create two sets of data that capture the interaction of instructions for the ARM CORTEX-M0/M3 architecture. These sets are the first to capture detailed information on the microarchitectural layer. They can be used to explore various leakage models suitable for creating side-channel leakage simulators. A preliminary evaluation of a leakage model produced with our dataset of real-world cryptographic implementations shows performance comparable to state-of-the-art leakage simulators.

## CCS CONCEPTS

• **Security and privacy** → **Embedded systems security; Side-channel analysis and countermeasures.**

## KEYWORDS

microarchitecture, side-channel simulator, deep learning, leakage model

## ACM Reference Format:

Omid Bazangani, Alexandre IOOSS, Ileana Buhan, and Lejla Batina. 2018. ABBY: Automating leakage modeling for side-channels analysis. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 14 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](mailto:permissions@acm.org).

Conference acronym 'XX, June 03–05, 2018, Woodstock, NY

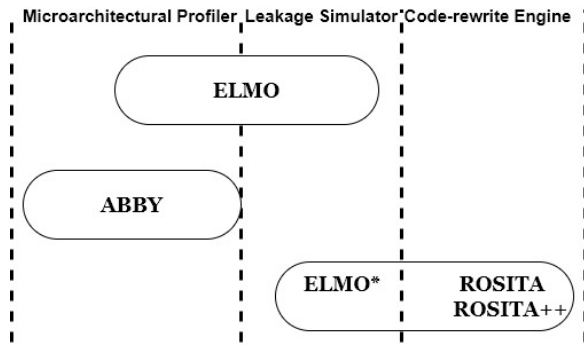
© 2018 Association for Computing Machinery.  
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00  
<https://doi.org/XXXXXXXX.XXXXXXX>

## 1 INTRODUCTION

when a code snippet runs on a chip, it interacts with the hardware environment. Kocher et al. [35] showed that this interaction can reveal the key during an encryption process. This interaction appears in physical side-channel(s) such as power [3, 11, 37], electromagnetic emanation [2, 49] or photonic emission [14]. An adversary can take advantage of these side channels and learn secret information during processing. Many studies show successful side-channel attacks leading to the recovery of secret keys or shares on various platforms [9, 22, 23, 27, 31, 43]. Technological advancements and twenty years of sustained effort by the cryptographic community significantly increased the workload required for successful key extraction. However, the problem of implementing a secure cryptographic algorithm on a given target device is not solved. A developer's challenge is balancing the presence of countermeasures against information leaks. As the product changes during development, it is important to understand whether the changes are beneficial or, in contrast, whether they compromise the security of the implementation.

The appeal of side-channel *leakage simulators*, which model the instantaneous power consumption of a device, is evident from the effort towards creating such tools [13]. A leakage simulator generates side-channel measurements from a sequence of instructions with the help of a *leakage model*, a function that describes how the target devices consume power. In the absence of tools such as leakage simulators, a security researcher tasked with hardening a cryptographic implementation, will measure traces, detect leakage, change the implementation, and reiterate until the implementation stops leaking. The process is slow, error-prone, and expensive. Moreover, the absence of leakage does not guarantee that there is no attack possible. A leakage simulator can automate the detection of side-channel leaks and, more importantly, can be used to explain the cause of a leak.

A leakage simulator transforms high-level code into traces similar to those collected from the target architecture. When adequate and informative, leakage simulators assist in the design of secure cryptographic implementation provided that the leakage model accurately reflects the reality of a key recovery attack. As in [45], we distinguish between *value* and *transition*-based leakage models. A leakage model is value-based if it takes the intermediate values of a cryptographic algorithm as arguments. Examples include the Hamming Weight (HW) or the Identity Model (ID). A leakage model is *transition-based* if it takes as parameters any pairwise combination



**Figure 1: Overview of ABBY, ELMO, and ROSITA in application level design**

of intermediate values [45] such as the Hamming distance (HD) model which captures events such as the update of a register.

A popular defense against side-channel attacks is masking, where each secret value uses multiple shares [15, 34] to break the dependency between the power consumption and the processing of a given variable. Capturing interactions between intermediate variables is essential for verifying the correctness of a masked cryptographic implementation. However, theoretically secure masking implementation often fails in practice due to unexpected interactions between variables [10, 47].

Creating a transition-based leakage model is specific to the target and requires intensive manual effort [38, 39, 52, 63]. Capturing the interaction of intermediate values requires profiling the microarchitectural layer, which may contain hidden storage elements where unexpected interactions between instructions can occur. Although the importance of transition-based leakage models has been established [8, 26, 38, 52], the details of microarchitectural implementation are often considered a trade secret and, therefore, not public information. By automating the creation of leakage models, more side-channel simulators, such as ELMO [39] which cover different architectures can be created with reasonable effort. *Automated microarchitectural profiling*, which captures the interaction between variables due to microarchitecture optimization is necessary for creating leakage simulators for different architectures.

Leakage simulators are the building block for rule-driven code rewriting engines such as ROSITA [52] and ROSITA++ [51] that patch the code automatically once the leakage is detected. Fig 1 shows the functional relationship between these tools, from an architectural profiler to a side-channel code rewriter engine. For example, ELMO, as a side-channel simulator, developed its microarchitectural profiler based on Thumbulator [60] (an Instruction Set Simulator). ROSITA developed a code rewrite engine based on the ELMO\* (upgraded ELMO) simulator combined with an assembly code modifier.

Although the methodology for building leakage simulators is known, the main limiting factor for their wide adoption is the limited number of supported target devices, a direct consequence of the effort required to reverse engineering the microarchitectural implementation. As the main barrier to overcome for the widespread deployment of leakage simulators is the characterization

of the target device, we ask the following question: *Can we automate microarchitectural profiling of a chip to speed up the process for side-channel simulator design?*

**Summary of contribution.** This paper makes the following contributions:

- We present ABBY, the first generic framework to automate the creation of training data used for creating a side-channel leakage simulator;
- Using ABBY, we create two dataset, ABBY-CM0 and ABBY-CM3, which capture the interaction of instructions for the ARM Cortex-M0 and Cortex-M3 boards.
- We applied a qualitative and quantitative analysis on ARM Cortex-M0 and Cortex-M3 leakage models conducted with statistical tests, dedicated leakage test vectors, and correlation power analysis.
- We demonstrate that deep learning can be used for realistic leakage models.

We propose ABBY, an open-source framework that *automatically* captures microarchitectural leakage. Furthermore, we offer the ABBY-CM0/CM3 dataset for ARM Cortex-M0/M3 chips (STM32F0/F1 families) produced with the ABBY framework. Based on our knowledge, this is the first open-source dataset to profile the microarchitectural layer, which can be used to study a target device's profiling further. In addition, we develop different transition-based leakage models, which allow us to create different simulators. We investigate the model's performance based on statistical parameters and learn how different microarchitectural features contribute to leakage. We compare the performance in detecting leakage of these simulators with ELMO and show that the performance is comparable to ELMO. We believe it is possible to improve further the leakage model.

**Paper organization:** Related works on microarchitectural leakage simulators are mentioned in Section 2. Section 3 briefly introduces the background on side-channel attacks and leakage detection and describes the hardware setup we used. Section 4, discusses building transition-based leakage models. We introduce our dataset in Section 5. In Section 6, we build several leakage simulators using the ABBY-CM0/CM3 dataset and discuss their performance. Section 7 concludes the article.

## 2 RELATED WORKS

Pinpas [21] is the first side-channel leakage simulator. Soon many more followed [20, 24, 46, 56]. However, the first open-source side-channel leakage simulator was SILK [58], which captures no specific hardware architecture and targets data-dependent power consumption. SAVRASCA [59] takes as input the compiled binary code and, using the tracing feature of the SimulAVR tool, will output simulated power traces for the AVR architecture. SAVRASCA was used to report a bug in the implementation code used for the DPAv4 trace set. ASCOLD [45] checks violations of the AVR architecture's independent leakage assumption (ILA). It takes as input the assembly file of the masked implementation and a configuration file of the system. The device shows the location of the leak (line number) and the rule that was violated. The physical causes of the ILA breaching effects are device-specific (cannot be generalized) and

counterintuitive when related to the assembly description of the target.

Going one level lower are the simulators, which capture some of the microarchitectural effects of the target. These simulators capture a more descriptive target model at the cost of a larger engineering effort. MAPS [36] is a power simulator designed for the ARM Cortex M3, which takes as input the source code of the masked implementation and outputs a simulated power trace. To capture the microarchitectural details, the authors used an HDL file of the target architecture and mainly focused on the leakage caused by the pipeline. In most cases, however, the target’s HDL files are unavailable. We consider ELMO [39] to be the first transition-based leakage simulator for the ARM-Cortex-M0/M4 family. ELMO models power consumption as a linear combination of values and transitions. Most remarkably, the simulator was created without detailed information on the hardware description or the target microprocessor. ELMO\* [52] improves the leakage model of ELMO by capturing interactions that span multiple cycles. ROSITA [52] is a rule-driven code rewrite engine that automatically patches the code once a leakage is detected. ROSITA starts with a (masked) implementation of a cryptographic algorithm, cross-compiled to produce both the assembly and the binary executable. A very compelling feature of ROSITA is that it extends an existing leakage detection tool, ELMO [39], to report instructions that leak secret information. The new detection framework (ELMO\*) uses the binary file to detect leakage and identify the offending machine instruction; ROSITA then applies a set of rules that replace the leaky instruction with an equivalent one (functionally) that does not leak. ROSITA repeats the process until no more leakage is detected. While the importance of microarchitectural details in a security analysis has been established [26], [38] access to its implementation is typically not available. The authors of ELMO had to make a lot of manual efforts (instructions selection, clustering, pipeline effect coverage) for the microarchitectural implementation of the target ARM Cortex-M0 processor. The current state of the art allows reverse engineering a commercial ARM Cortex-M3 microprocessor [19, 26]. The authors note that the current methodology involves intensive manual effort. However, it is worth considering, as it shows the importance of capturing microarchitectural effects. A large body of works targets the creation of pre-silicon side-channel simulators [28, 30, 32, 44, 50, 53, 54, 62], which use design information of the target device to create the leakage model.

## 3 BACKGROUND

### 3.1 Side channel.

Power consumption and EM signals emitted from a device correlate with the processed data and the executed instructions. The amount of power required to maintain a signal’s value depends on the signal’s logical state. In CMOS technology, the predominant choice when manufacturing integrated circuits, changing the value of a bit requires a different power level than keeping a bit constant. Therefore, the power consumption of a circuit directly correlates with the data processed by the circuit. Monitoring the physical properties of devices can reveal information about the operations and data processed. To perform a side-channel attack, an attacker attempts to correlate the observed physical side channels with the

values processed by the device. Other effects, such as variations in signal propagation time or cross-capacitance effects, contribute to the device’s power consumption and correlate with the data processes.

**Leakage modeling.** Let the set  $X$  be the data that we wish to monitor. When using side-channel analysis,  $X$  is typically the set of intermediate values created when transforming plaintext into ciphertext. We denote by  $L(X)$  the leakage model of the variable  $X$ . An adversary collecting side-channel traces has access to variable  $y$ , defined using Equation 1. The measured power traces are conventionally considered noisy, and this Gaussian noise  $N(0, \sigma^2)$  is independent of leakage  $L(X)$  [25].

$$Y = L(X) + N(0, \sigma^2) \quad (1)$$

As  $L(X)$  depends on the architectural design of the target and originates from the interaction between software and hardware. To improve this estimation and get it close to the real leakage, we consider the most relevant microarchitectural features of the target, such as instruction interaction, pipeline effects on instructions, operand values, and memory interactions. In this study, we consider the target as a “gray box” model. Although we do not have access to the hardware description layer (white box), we have access to the instruction set architecture (ISA) and full control of firmware execution.

**Leakage assessment.** Test Vector Leakage Assessment (TVLA) [29] is one of the most popular methods for leakage assessment due to its simplicity and relative effectiveness. TVLA relies on statistical hypothesis tests and comes in two flavors: *specific* and *non-specific*. The ‘fixed-vs-random’ is the most common nonspecific test and compares a set of traces acquired with a fixed plaintext with another set of traces acquired with random plaintext. In the case of a specific test, the traces are divided according to a known intermediate value tested for leakage. In both cases, Welch’s two-sample t-test for equality of means is applied for all trace samples. A difference between two sets larger than a given threshold is evidence of a leak’s presence. Despite its simplicity, it is easy to misuse this test [61].

### 3.2 Regression model evaluation

The goal of most statistical models is to predict future events or to help explain reality[12]. In the former case, the quality of the model is defined by its predictive power. In contrast, the quality of the model in the latter is related to the number of relevant factors it can identify. In leakage simulators, predictive models are used to estimate the power consumption of an intermediate variable. We use *coefficient of determination* ( $R^2$ ) and *cross-validation* to judge the quality of the model we use, as these are popular choices to evaluate regression models [25]. For readability, we use the notation [25], [39].

$R^2$  measures how much of the variation in the dependent variable can be explained by the independent (explanatory) variable(s). An  $R^2$  value close to one shows a good fit between the predicted and measured values. To compute  $R^2$ , we need to compute two types of *sum of squares*. The first parameter is called Residual Sum of Square(RSS) which measures how much of the explanatory variables’ variation can not be explained by the model. It represents

the sum of the squared differences between the actual measurement  $y_i$  and the predicted value  $\tilde{L}(Z_i)$ (equation 2).

$$RSS = \sum_{i=1}^n (y_i - \tilde{L}(Z_i))^2 \quad (2)$$

where  $n$  is the number of samples. The second parameter Explained Sum of Square (ESS) measures the variation of the explanatory variables(equation 3).

$$ESS = \sum_{i=1}^n (\tilde{L}(Z_i) - \bar{y})^2 \quad (3)$$

The Total Sum of Square(TSS) is the sum of ESS and RSS(equation 4).

$$TSS = RSS + ESS = \sum_{i=1}^n (y_i - \bar{y})^2 \quad (4)$$

The coefficient of determination  $R^2$  can be calculated with equation 5 [25].

$$R^2 = \frac{ESS}{TSS} = 1 - \frac{RSS}{TSS} \quad (5)$$

The disadvantage of using the  $R^2$  metric is that its value increases with the increasing number of explanatory variables included in the model. To penalize additional explanatory variables added to the model and adjust this metric against the overfitting problem, we look at  $R^2$  adjusted denoted by  $R^2_{adj}$ .

$$R^2_{adj} = 1 - \frac{(1 - R^2)(n - 1)}{(n - p - 1)} \quad (6)$$

$n$  is the number of samples, and  $p$  represents the number of explanatory variables fed to the model. According to Equation 6, if the number of explanatory variables is negligible compared to the number of samples( $n \gg p$ ), then  $R^2 \approx R^2_{adj}$ .

**F-test.** To investigate the effect of adding different explanatory variables to the model, we used the F-test introduced in [39]. We explore the importance of explanatory variables based on their contribution to the model's performance. We check if a reduced model (fitted by a subset of explanatory variables) is missing a significant contribution compared to a full model, which consists of the full explanatory variables. Let us consider that B is a reduced model of model A. Therefore the number of explanatory variables of model A ( $p_A$ ) is larger compared to the number of explanatory variables of model B ( $p_B$ ), so we have  $p_A > p_B$ . In this example, the null hypothesis states that the extra parameters present in model A do not affect the model performance. The F-statistic is computed based on the residual sum of squares (RSS), while  $p_A - p_B$  and  $n - p_A$  are degrees of freedom as shown in equation 7.

$$F = \frac{\left( \frac{RSS_B - RSS_A}{p_A - p_B} \right)}{\left( \frac{RSS_A}{n - p_A} \right)} \quad (7)$$

For a specific significance level (usually  $\alpha = 5\%$ ), if the F value is more than the critical value under the  $F_{p_A - p_B, n - p_A}$  distribution, the null hypothesis is rejected, which means that the parameters of model A, which are not present in model B have a significant effect.

### 3.3 Cortex-M0 Vs Cortex-M3 architecture

Cortex-M0 is based on ARMv6-M architecture with three stages pipeline. This architecture fully supports the Thumb-1 instruction set while supporting some of the Thumb-2 instructions, a total of 56 instructions [6]. Cortex-M3 with three stages pipeline is developed based on ARMv7-m architecture, having full instruction coverage for Thumb-1 and Thumb-2 with total coverage for more than 90 instructions [7]. One of the key components used in this study for cycle counting is *Data Watchpoint and Trace unit (DWT)*. All Cortex-M3 microcontrollers support the DWT register. Although ARMv6-M supports this feature, not all the cortex-M0 microcontrollers support the DWT register. Especially the chip that ELMO developed based on.

**(Data Watchpoint and Trace) DWT** is a debug unit that provides watchpoints, system profiling, and data tracing. This unit contains a cycle counter called *Clock Cycles Counter (CYCCNT)* to count the CPU clock cycle. We configure DWT in the hardware watchpoint mode to read DWT\_CYCCNT register after each instruction processing to calculate the required execution cycle for each instruction. There are two problems with using hardware breakpoints. First, there are only four breakpoints available. Second, breaking the normal process would affect the pipeline. We need to reuse available breakpoints (only one breakpoint is used in ABBY) to solve the former problem simply. To explore the second problem in detail, we can check it with a simple assembly code snippet like three LDR in a row. If we run this code without halting the process, it would take four cycles for these three LDR to get executed, while halting the processor for each LDR takes 6 CPU cycles. The reason for this difference is related to the optimization happening in the pipeline in which every LDR separately takes two cycles. In contrast, for a series LDR in a row, only the first one takes two cycles, and the rest would take one to execute.

To solve the process halting effect on the pipeline and cycle counting, we are recording DTW\_CYCCNT register while moving the breakpoint instruction by instruction until we reach the end of the target assembly snippet code Fig.13. In this way, we can calculate the cycle count for the first instruction(LDR) by a simple subtraction of  $X2 - X1$  and so on for the next instruction. To ensure no other prior instructions affect the targeted assembly code snippet, we add 10 NOP instructions after and before the targeted assembly code for isolation.

### 3.4 Hardware Setup

We use two different Cortex-M0 chips based on Armv6-M architecture manufactured by ST Microelectronics with STM Discovery Boards [41, 42]. The target boards have the STM32F051R8T6 or STM32F030R8T6 chip [40] and an external crystal oscillator (8MHz). Although initially tested our framework on both boards, we continued on STM32F030R8T6 target to compare the result with the ELMO. We modified the boards to reduce the measurement noise by removing the power line capacitors. We measured the current through a current probe (Riscure CP271), which is used as a proxy for the target's power consumption. We used a PicoScope 3207B for data acquisition at a sampling rate of 500MS/s while the target runs at 8MHz. This oscilloscope can store up to 512Ms due to memory limitations. A physical 48MHz low-pass filter is used before the




Figure 2: Setup Block Diagram

measurement. We are using two acquisition channels, one for the power signal (connected to the current probe), and the other for the trigger signal. The trigger is fed by a GPIO of the target when the desirable segment of the code is running. The oscilloscope will be armed for signal recording as soon as the trigger signal is detected. Fig. 2 shows the configuration block diagram. The setup for Cortex-M3 is the same as Cortex-M0 except for the clock source of the chip, which is fed by an external clock. The Cortex-M3 microcontroller is based on armv7-M architecture and is manufactured by ST Microelectronics (STM32F107vct) [40].

#### 4 AUTOMATED MICROARCHITECTURAL PROFILER

This section explains the challenges and innovations behind ABBY methodology to create an automated microarchitectural profiler. We address important questions such as 1) Why do we need automation? 2) what are the design requirements? 3) what are the design challenges? and 4) How does the ABBY framework tackle it?

**Why do we need automation?** Designing a side-channel simulator like ELMO has different steps. These steps include selecting relevant instructions to decrease data space, profiling selected instructions, clustering instructions based on their profile, adding support for sequence dependency, and training a simulator model. Despite the amount of effort and time-consuming, each step needs an expert decision to continue to the next step. Moreover, the ELMO model only supports symmetric crypto algorithms to reduce the effort in different design steps. We need automation to put all aforementioned efforts into a toolchain to design more sophisticated side-channel simulators.

**Design requirements.** First, We need an algorithm-independent model that covers as many instructions as possible from the target instruction set, not to be limited to some specific algorithms. Second, we need to remove human expert decisions between each step to decrease the needed effort. Third, we require a method to be scalable for other ARM Cortex-M families or even more different architectures.

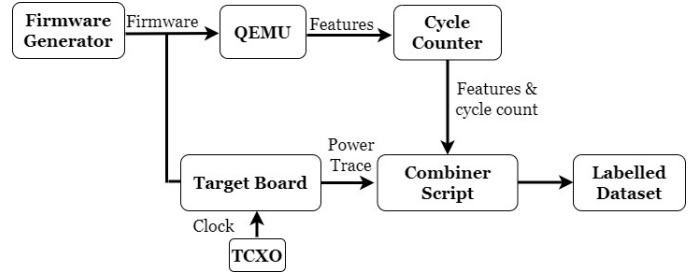


Figure 3: ABBY Cortex-M series block design

**Challenges.** To satisfy the aforementioned requirements, we have some challenges to tackle. First, to cover more assembly instructions, we must consider how the instructions combination grows dramatically due to the pipeline effect. For instance, for a Cortex-M0 chip with 56 supported Thumb instructions, there are  $56^3 \approx 175K$  combinations. The second challenge is to annotate the power trace with executing instruction in each clock cycle, we need to know how many cycles a specific instruction requires to be executed. The third challenge is that the clock drift might affect this annotation.

**ABBY solution.** ABBY Framework tackles these challenges by offering a methodology shown in Fig 3. ABBY framework handles the needed effort by adding more assembly instructions through an automated Random Assembly Generator block, automatically generating, programming, and capturing power traces for each firmware. To collect microarchitectural features, ABBY uses QEMU [55] and Genu Debugger(GDB) tools which are standard for not only different ARM cortex-M families but also other architectures. ABBY uses *Data Watchpoint and Trace (DWT)* register for instruction cycle counting while feeding the target a precise external clock using a TCXO component. These two techniques help us label our power trace precisely with the related instruction processed in a specific clock cycle. The annotation process occurs in the Combiner script (Fig 3).

**Engineering challenges.** Engineering challenges to fitting ABBY on a specific architecture are as follows.

- Supported by QEMU or support JTAG debugger
- Support GDB basic commands
- Support a cycle counter register like DWT

To extract microarchitectural features from a specific firmware, We can run it on a real target with a JTAG debugger or an emulation tool. ABBY uses a popular emulator called QEMU to speed up the process. QEMU supports many ARM Cortex-M families and many others under development. Even if QEMU does not support a specific target, using the JTAG debugger is an alternative option supported by all ARM Cortex-M families and even more architecture, as well as GDB.

All ARM Cortex-M families support the DWT register for cycle counting except Cortex-M0. To solve the cycle counting problem for Cortex-M0, we applied some modifications to the block design of ABBY. To annotate the power trace, we use the ELMO model that gives us a simulated trace with annotations. Next, we align the simulated trace by ELMO with the measured power trace using the




Figure 4: ABBY Cortex-M0 block design

DTW algorithm. As ELMO was developed based on Thumbulator, we replaced QEMU with Thumbulator for feature extraction.

## 5 DATASET CREATION

### 5.1 Feature Selection

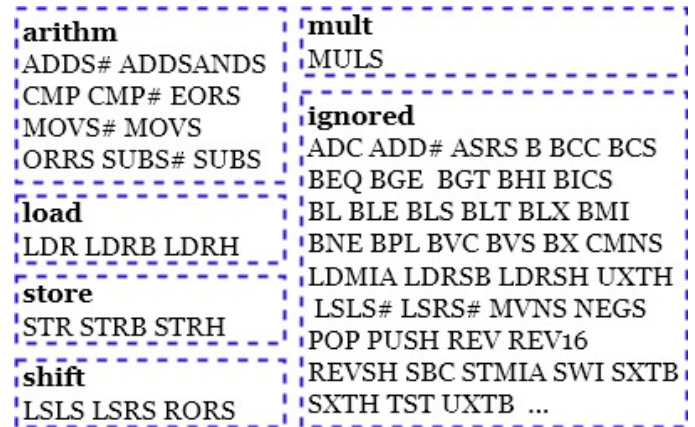
The critical observation made by McCann et al. [39] when building the ELMO leakage model is that the power consumption of the *current instruction*,  $I_c$  depends on *the preceding instruction*,  $I_p$  and *the subsequent instruction*  $I_s$  [57]. The reason behind this observation can be found in the three-stage design of the target pipeline. Not only the instructions but also the operand values of these instructions contribute to the power consumption of the chip [39].

**Instruction coverage of classical leakages model.** Although executing instructions is one of the main contributors to the chip’s power consumption, traditional leakage models only look at the HW or HD of each operand with its previous value. These models cover neither the pipeline effect nor instructions.

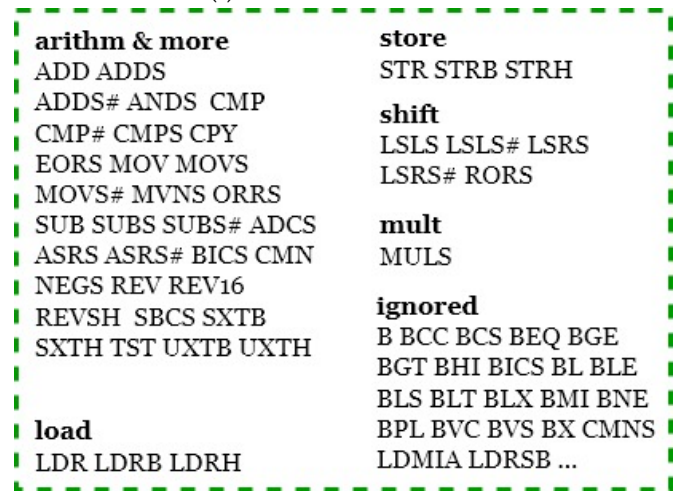
**ELMO instruction coverage.** ELMO is instruction-accurate, which has the advantage of allowing the quick identification of a leaky instruction. Following a cluster analysis to group “similar instructions” (that is, that leak information in the same way), the authors identify five groups, all of which include 21 instructions, see Fig 5.

The groups correspond to the same processor component: ALU instructions in one group, shift instructions in another group, load, and stores that interact with the memory are in two different groups, and the MULS instruction with a distinct profile due to its fit in the implementation of the single cycle in a separate group. These groups also represent the internal structure we could expect from an ARM core since shift, multiplication, and arithmetic operations do not use the same CPU part.

**ABBY-CM0 instruction coverage** Similar to ELMO, ABBY captures the interaction between instructions in the pipeline registers and memory leakage. No assumptions about the operand interaction are made to simplify data collection and training. ABBY also adds memory read/write values for current and previous memory access to cover memory leaks, as these were shown to be important in [52]. As part of the effort to simplify training, we further remove the clustering of instructions. We keep the thumb instructions for crypto algorithms, representing 44 instructions for ARM Cortex-M0, including arithmetic, shift, store, load, and multiplication operations. We do not profile branching, stack operations PUSH, POP, LDR/STR sp, or operations changing the PC register (Fig. 5) as these operations are not used for implementing cryptographic algorithms. For the storage and loading operations, we reserved the r0 register to point to a memory address of an empty data section. Furthermore, ignoring instructions such as branching B, BEQ, BL



(a) ELMO instruction clusters



(b) ABBY instructions

Figure 5: ELMO instruction clusters vs ABBY instructions


makes sense as block ciphers avoid using them to be time constant and to prevent leaks from branch prediction.

For a 3-stage pipeline microprocessor, and considering 56 possible Thumb instructions for Cortex-M0, the entire instructions space is  $56^3 \approx 175k$ . Fig. 6 shows the entire instruction space vs. ABBY and ELMO coverage. We see that ABBY covers more instruction combinations compared to ELMO.

### ABBY-CM3 instruction coverage.

The data collection phase is the same as CM0, with these differences that we target only ALU component in a CM3 chipset. Based on our knowledge, there is no cycle-accurate instruction simulator for CM3 like Thumbulator, and it’s not our goal here to develop an instruction simulator. We reduced profiling space to the ALU component of CM3 as it guaranteed timing of one clock per each different supported instruction [4]. Also, this unit provides many popular instructions which are used in crypto algorithm [39]. We covered seven different instructions of ALU including ADD, AND, CMP, EOR, MOV, ORR, and SUB. Moreover, we considered different possible variants for these instructions, like the version with suffix S or with





**Figure 6: Possible instruction space in gray vs. ABBY and ELMO in green and blue respectively**

and without immediate value and the extended version with suffix .W, which ended up with 17 different variants for these seven instructions. We implement 50 different firmware, each including 80 triplets of randomly chosen instructions from the aforementioned instruction space. For each instruction, we are collecting 1000 different traces while each of them is fed with random operands through the serial port. We ended up with  $50 \times 80 \times 1000 = 4M$  data samples.

### 5.2 Automated Firmware creation

We automatically generate randomized assembly firmware, including triplet Thumb instructions, to cover the 3-stage pipeline of the target. We analyze the output to ensure the instructions are uniformly chosen and their operand value is uniformly distributed. Unlike ELMO, we do not have the clustering limitation (all instructions in one cluster), so we keep all Thumb instructions typical for cryptographic implementations. The result is a group of 44 instructions for the ARM Cortex-M0, including arithmetic, shift, store, load, and multiplication operations. For Cortex-M3, the result is 17 different instructions of the ALU component.

### 5.3 Dataset construction

**Cortex-M0.** Using the generated firmware, we collect 1,000 triplets<sup>1</sup> (or pipeline states) in a single acquisition. By generating and flashing<sup>2</sup> 50 000 different Random assembly firmware automatically, we collect for each triplet of instructions (with random operands)  $\frac{50\,000 \times 1\,000}{4^3} \approx 587$  data points.

Firmware with random instructions is loaded on the target device, and power consumption is measured while the target executes the firmware.

<sup>1</sup>A compromise between the oscilloscope memory and the duration of running Dynamic Time Warping for alignment ( $t \sim O(n^2)$ ).

<sup>2</sup>We could run from RAM to preserve flash endurance, but it influences leaks and the number of cycles per instruction. Running the firmware from the flash is closer to real-world scenarios.

**Cortex-M3.** It is the same as the Cortex-M0 process, except we have 50 different firmware for Cortex-M3, each including 80 triplets with  $\frac{4M}{17^3} \approx 815$  data points.

### 5.4 Dataset labeling.

**Cortex-M0.** To label the measured power samples, we need to identify the corresponding triplet of instructions. A challenge when annotating the measured traces with the executed instructions is that different instructions might take different cycles, depending on the optimizations made by the manufacturer.

Our solution, dictated by the simplicity of the ARM Cortex-M0 processor, is not perfect, but it is effective. We use ELMO as an initially estimated power consumption to compare with the feature annotation. Next, we replace the value generated by ELMO with the corresponding value extracted from the measured traces. ELMO is not cycle-accurate; the measured and estimated traces do not have the same number of samples. Our solution for aligning two-time series of varying sizes is *Dynamic Time Warping (DTW)*, a popular algorithm used for speech recognition. This technique solves the signal alignment problem and finds each instruction triplet’s corresponding power consumption value (Fig. 4). DTW is not perfect, and alignment errors are possible. As a result of the first attempt to align the data, 22% of instructions are dropped. We also expect errors in the remaining data, so we perform several passes. We observe that the alignment distance converges after four passes (there is no significant improvement with more iterations).

**Cortex-M3.** Dataset labeling for CM3 is a challenge due to the following limitations:

- (1) Lack of an open source ISS like Thumbulator
- (2) Lack of an available power model like ELMO

**Lack of an open source ISS.** Thumbulator works only on ARMV6-m architecture. Based on our knowledge, there is no open-source cycle-accurate simulator for ARMV7-m, which is the based architecture for CM3. To solve the feature extraction problem without an ISS tool, we take advantage of GNU Debugger (GDB). After loading firmware on the target using a script, another script would extract features using GDB-related commands through a J-link debugger. After we ran feature extraction using GDB on a real target, we observed that this process was time-consuming as we needed to stop on every assembly instruction to collect all operand values and related registers. We solved this problem using a well-known emulator called Quick Emulator (QEMU) [55]. Using QEMU would make the process at least ten times faster, and you would not need a physical setup for feature extraction.

**Lack of an available power model for CM3.** Another limitation to extending the previous model for CM3 chips is that there is no available power model to apply DTW based on it. To create our labeled dataset, we must annotate recorded power traces from an actual measurement with the extracted features from the same execution. The main challenge for this annotation is that even with a slight drift in the signal annotation, like 100 nanoseconds in our case, as our target is running on 10MHz, we are not annotating the desired instruction but the next or previous instruction in the subsequent based on a positive or negative drift. So, if drift happens in power annotation, it leads to false labeling in the dataset. A

corrupted dataset would not guarantee what the ML model would learn, even if it learns anything from this dataset. One solution for this problem would be a precise execution timing to calculate the execution timing of each instruction. For this accurate calculation, first, we need to calculate the cycle count for each instruction and ensure that the target is fed with a precise clock. For the former, we know that our target execution cycle is deterministic, and the ARM developer manual guarantees it [5]. Based on the developer manual, all the ALU instructions require one clock cycle to execute. We confirm it in practice with the following methods:

- (1) Using Data Watchpoint and Trace unit (DWT)
- (2) Clock signal recording

**Clock Signal Recording.** As we record the clock signal fed to the Processor beside the power signal, we can count the number of clock cycles needed for our target assembly code to execute. We can double-check the calculated clock with the ARM developer manual. We cannot rely on the chip's internal clock for clock precision, as it's prone to drifts and changes by voltage or temperature changes. Among the options for external clock sources, we used a TCXO (Temperature compensated crystal oscillator) based clock source equipped with temperature changes compensation component. This solution provides a precise and steady enough clock cycle for execution timing measurement.

## 5.5 The ABBY dataset specification

**ABBY-CM0 dataset.** The ABBY-CM0 dataset is a Pandas DataFrame [33] file, which includes  $\approx 35$  million samples with 12 columns representing all extracted features alongside the chip's power consumption while executing related features.

- **Assembly Instructions.** Concerning the three stages pipeline of the target, three different instruction columns exist in the ABBY-CM0 dataset. These columns represent the current executing instruction ( $I_c$ ), previous instructions that have been executed ( $I_p$ ), and subsequent instruction, which is in the decoding phase to get executed ( $I_s$ ).
- **Operand Values.** These columns are related to the data processing by current and previous instructions. `op1_value_current` and `op2_value_current` Represent registers value used for executing  $I_c$  while `op1_value_previous` and `op2_value_previous` Represent values of corresponding registers used by previous instruction.
- **Memory transactions.** These features represent the data values to store to or load from the memory processing by STR or LDR assembly instructions. The `readbus_value_current` and the `readbus_value_previous` present the loaded data on the memory bus while the current or previous load operation processing, respectively. The `writebus_value_current` and the `writebus_value_previous` present the same scenario for the store operation.
- **Power sample.** Shows a proxy value of target power consumption at the moment that all other features are processing.

We must pre-process features to fit any simulator model on the dataset. Categorical data, Assembly instructions ( $I_p$ ,  $I_c$ , and  $I_s$ ), are hot-encoded and numerical data are represented in a 32-bit

binary system ( $ID_{32}$ )<sup>3</sup> instead of the decimal system ( $ID_{10}$ ) to decompress information. Furthermore, we also add HW of operand and memory transaction values with their previous values. Moreover, the HD of each operand value is calculated and located in the dataset.

**ABBY-M3 dataset.** We have two versions of the CM3 dataset, ABBY-M3V1 and ABBY-M3V2.

**ABBY-M3V1.** It's a Pandas DataFrame [33] file with 4M samples and includes nine columns of data representing all extracted features alongside the target power consumption. All features are the same as ABBY-CM0 except, in ABBY-M3V1, there are no memory transaction columns as we only cover the ALU component.

**ABBY-M3V2.** The only difference between this dataset with the previous version is the power consumption column. In this dataset, instead of considering the maximum value of the power sample in each clock as a proxy for the power consumption of the target, We simply add all ten samples we are acquiring per each clock cycle. In this way, we are not losing any information because of the downsampling, and our dataset represents the transient power consumption of the chip during the execution of an instruction with specific features.

## 6 FITTING POWER SIMULATORS USING THE ABBY DATASET

In this section, we implement different power simulators based on the ABBY dataset and evaluate each model's performance. The leakage model is the heart of a simulator, and it's a bridge that connects the side-channel signal (power in this study) to the processing data (instructions and operands value) in the target. More precise leakage models can help detect more leakages. We analyze conventional linear regression and nonlinear deep learning models in this study.

### 6.1 Fitting Linear regression models on ABBY-CM0

We constructed different leakage models using linear regression to evaluate the ABBY-CM0 dataset and investigate microarchitectural leakage. In Equation 8,  $Y$  is the estimated power consumption of the target,  $\beta_0$  is a constant called intercept, while  $\beta_i$  is the coefficient of the explanatory variable  $X_i$ , and  $\epsilon$  is the error. A regression model aims to find the best coefficients for each explanatory variable.

$$Y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \dots + \beta_n X_n + \epsilon \quad (8)$$

In this study,  $Y$  would be our leakage model  $Y = L(X)$ . In this section, we fit different leakage models to investigate how the different microarchitectural features contribute to the leakage.

**HW \HD classic leakage model** As discussed in Section 4, classical leakage models do not consider instructions but operands. HW and HD are the most popular classical leakage models. We fit a leakage model based on the HW of the previous instruction and

<sup>3</sup>ID represent the identity model in side-channel



the current instruction operand values. Moreover, the HD between each operand’s current and previous values was added to the model. Furthermore, we add HW for memory interactions. Equation 9 shows the fitted leakage model.

$$L_{HW}(X) = [HW(OPs) | HD(OPs) | HW(MRs) | HW(MWs)]\beta + \epsilon \quad (9)$$

Where:

HW(OPs): a matrix that includes the HW of each operand for the previous and current instruction.

HD(OPs): a matrix that includes HD between current and previous operands.

HW(MRs): a matrix that includes the HW of the value of the read memory for the current and previous interaction.

HW(MWs): a matrix that includes the HW of the memory write value for the current and previous interaction.

**Identity leakage model.** The identity model is a straightforward classical model that uses the value of operands without any changes. Fitting a model based on the identity of the operand values requires normalization of the data. We chose a binary representation of the data instead of normalization because dividing the values to a 32-bit size makes the result less sensitive to small changes. Equation 11 shows the identity leakage model. Where  $ID_2$  represents the binary representation of values, all other parameters are the same as the HW\HD model, except that the values are represented as binary instead of HW or HD.

$$L_{ID}(X) = [ID_2(OPs) | ID_2(MRs) | ID_2(MWs)]\beta + \epsilon \quad (10)$$

**Instruction-based leakage model.** Based on section4, not only the processing data (operational values) but also the execution of operations (instructions) contribute to leakage. Concerning the three-stage pipeline of our target, equation 11 models the leakage related to the instructions  $L_{INS}(X)$ . To fit the mnemonic assembly instructions, we use a one-hot encoding technique. After one-hot encoding, we end up with a 44-bit value representing our 44 supported instructions in the framework for each instruction level in the pipeline.

$$L_{INS}(X) = [(Ip) | (Ic) | (Is)]\beta + \epsilon \quad (11)$$

**Comprehensive (CH) leakage model.** To consider the effect of operations and operand values, we build a leakage model based on operand values and instructions. For operand values, we use identity leakage model ( $L_{ID}(X)$ ) as it shows slightly better performance compared to HW\HD model<sup>4</sup>. Combining the ID with the instruction-based leakage model, we make a more comprehensive leakage model and call it ( $L_{CH}(X)$ ) that models not only the instruction and corresponding operand values but also the linear interaction between them.

$$L_{CH}(X) = L_{ID}(X) + L_{INS}(X) \quad (12)$$

<sup>4</sup>We didn’t consider the complexity of the model in this step and performance came first. From a complexity point of view, we have ten coefficients for the HW\HD model, while it is (8 \* 32) 256 for the ID. In addition, the identity model might interact better with the instructions

**Table 1: Evaluation of linear regression leakage models**

| Model                  | $L_{HW}(X)$ | $L_{ID}(X)$ | $L_{INS}(X)$ | $L_{CH}(X)$ |
|------------------------|-------------|-------------|--------------|-------------|
| OHE( $I_P, I_C, I_S$ ) |             |             | X            | X           |
| ID(Ops)                |             | X           |              | X           |
| ID(Mem_Bus)            |             | X           |              | X           |
| HW(Ops)                | X           |             |              |             |
| HW(Mem_Bus)            | X           |             |              |             |
| R <sup>2</sup> adj     | 0.30        | 0.32        | 0.57         | 0.58        |
| F-Stat                 | 6.92e + 5   | 9.39e + 5   | 1.67e + 5    | 1.04e + 5   |

Considering  $L_{ID}(X)$  and  $L_{INS}(X)$  as a reduced model of the  $L_{CH}(X)$  model, we applied F-test. The result shows that the combination of characteristics in  $L_{CH}(X)$  shows a significant effect with  $\alpha = 0.05$ .


**Model selection.** We fit all the aforementioned leakage models in the ABBY-CM0 dataset, and Table 1 summarizes the result. Each column represents the constructed leakage model, and rows represent features. The X mark shows that the corresponding feature is used in that specific leakage model. The  $R^2_{adj}$  and F-Stat are calculated separately for each leakage model.

Although the identity leakage model shows slightly better performance, the classical HW\HD is a powerful operand leakage model (compared to the identity model), considering its simplicity. The instruction-based leakage model performs better than the operand leakage model based on the parameter  $R^2_{adj}$ . Albeit the  $L_{CH}(X)$  leakage model performance increased slightly with combining instructions and operand values, this model covers the leakage related to both operands and instructions.


**Model evaluation** The goal of a side-channel simulator is leakage detection to help the developers during the design phase. To evaluate our model, we applied a correlation-based DPA attack on an AES Crypto algorithm to investigate the performance of leakage detection for our model versus measured power trace and ELMO simulated trace while all are running the same firmware. Based on the previous section, we chose the CH model as the best model. For the AES implementation, we choose a first-order protected implementation by Yao et al.[64], which we refer to as *Byte Masked AES*. Although this implementation should be secure against first-order leaks, correlation analysis indicates leaks on the measured and simulated power traces by ELMO and CH models. Fig. 7 shows a successful DPA attack on the Byte-Masked AES implementation.

## 6.2 Fitting Deep learning regression models on ABBY-CM0


In this section, we apply the deep learning method to our dataset to evaluate nonlinear model performance on the ABBY-CM0 dataset. We use the popular Multi-Layer Perceptron (MLP) model to train the ABBY-CM0 dataset. We use TensorFlow2 [1] with the Keras submodule to build a preprocessing pipeline and an MLP model with three hidden layers. The input layer consists of three different groups of data, mnemonic assembly instructions, operands value for these instructions, and memory bus values. Before delivering input data to an MLP model, the input must be normalized to a



(a) Correlation based on the real measurement



(b) Correlation based on the CH model



(c) Correlation based on the ELMO model


**Figure 7: DPA attack based on correlation on the first round SBox of Byte-Masked AES shows successful leakage detection on the real power trace vs. the ELMO and the CH model.**

number between zero and one. For mnemonic assembly instructions which are not ordinal categorical features, we use one-hot encoding [48]. For normalizing operands and memory read/write values, we extended each value to a 32-bits binary array. The input layer consists of  $(8 \times 32) + (3 \times 44) = 388$  neurons. Table 2 summarizes the model architecture. The chip power consumption is predicted by the output layer, which consists of one neuron. After training, our MLP model reached  $r^2 = 0.771$  on a training set of  $\approx 35$  Millions samples with a test set of  $\approx 15$  Millions while 20% of the training set is used as validation. Although the  $r^2$  metric evaluates the model’s performance, we apply standard side-channel evaluation metrics to determine its usefulness for leakage detection.


**Model evaluation.** TVLA is one of the side-channel domain’s most popular leakage assessment methods. See Section 3 for a brief introduction. We applied the TVLA test on two cryptographic algorithms, AES [18] and Xoodoo [16]. We chose the AES implementation the same as we did for the correlation attack [64]. Although this implementation should be secure against first-order leaks, TVLA indicates leaks on the measured power trace as well as on the simulated power traces by ELMO and ABBY models (Fig. 8). As proof

**Table 2: MLP hyperparameter description**

| Layer Types            | Details            |
|------------------------|--------------------|
| Input Fully-connected  | #neurons 388, Relu |
| Hidden Fully-connected | #neurons 388, Relu |
| Hidden Fully-connected | #neurons 16, Relu  |
| Hidden Fully-connected | #neurons 16, Relu  |
| Output                 | #neurons 1, Linear |



(a) Measured power trace




(b) ELMO model vs ABBY model


**Figure 8: Byte-Masked-AES TVLA result real vs. simulation**

that ABBY has learned the same leaks as ELMO, we notice how similar the  $t$ -trace scores produced by ELMO and ABBY are, with ABBY showing more leakage points compared to ELMO. Notice the difference in the  $x$ -axis between the measured traces (cycles) compared to the labels of the simulated traces (instructions).

To confirm that the positive results obtained for Masked AES are not just a lucky coincidence, we compare the  $t$ -test results of ELMO and ABBY on a different cryptographic algorithm, namely Xoodoo [16]. Xoodoo is the underlying permutation used in Xoodoo [17], one of the finalists in the NIST Lightweight Cryptography Standardization process. As shown in Fig. 9, we confirm that also, in this case, the output of ABBY is comparable to the ELMO model and the measured power traces.



(a) Measured power trace



(b) ELMO model vs ABBY model

Figure 9: Xoodoo TVLA result real vs. simulation

### 6.3 Fitting deep learning regression models on ABBY-M3

As the deep learning model showed a better performance compared to linear regression, we continued the CM3 model fitting based on an MLP regression model. The differences with the CM0 model are related to the input layer and model architecture. We have 17 mnemonic assembly instructions, ALU related, for the input layer. After normalization of mnemonic assembly instructions using one-hot encoding and a 32-bit binary array extension for operand values, The input layer consists of  $(4 \times 32) + (3 \times 17) + 1 = 180$  neurons. Table 3 summarizes the model architecture. The predicted ALU power consumption can be collected from the output layer, which consists of one neuron.

After training, our MLP model reached  $r^2 = 0.977$  on a training set of  $\approx 21$  Millions samples with a test set of  $\approx 14$  Millions while 20% of the training set used as validation. Although the  $r^2$  metric evaluates the model’s performance, we apply standard side-channel evaluation metrics to determine its usefulness for leakage detection.

**Model evaluation.** Fig. 10 shows an estimated power trace by our model vs. measured power trace while both are running the same code snippet (random assembly). Although the visual inspection of the estimated power shows a good similarity with the measured power, we need a metric to compare them.

One of the widespread side-channel attacks is the Differential Power Analysis attack. This attack is one of the strong attacks that can

Table 3: Extended MLP model hyperparameter description

| Layer Types            | Details            |
|------------------------|--------------------|
| Input Fully-connected  | #neurons 180, Relu |
| Hidden Fully-connected | #neurons 32, Relu  |
| Hidden Fully-connected | #neurons 32, Relu  |
| Output                 | #neurons 1, Linear |




Figure 10: ABBY estimated power trace Vs. Measured trace





Figure 11: Correlation result based on the ABBY model prediction

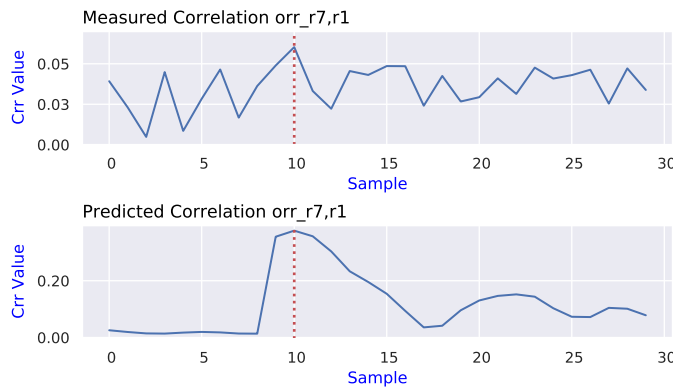
reveal secret information due to the microarchitectural leakages of a chip. We apply correlation analysis for a random assembly firmware designed to target ALU instructions on the measured power trace as same as the estimated power trace, which ABBY-CM3 deep learning model produces. We feed the designed firmware with 10K different values and use these values as an operand using the register r7 for different instructions. Fig 11 shows the correlation for the ABBY model. Fig 12 compares the correlation peaks found by the ABBY model with those found in the measured traces. With this experiment, we could confirm that the model correlation result is similar to the measured power traces from the real target. Although the figure shows the result for some ALU instructions, for all covered ALU instructions, we observed that maximum correlation was founded, and it was happening in the same position that is happening for the actual measurement.

## 7 CONCLUSIONS AND FUTURE WORK

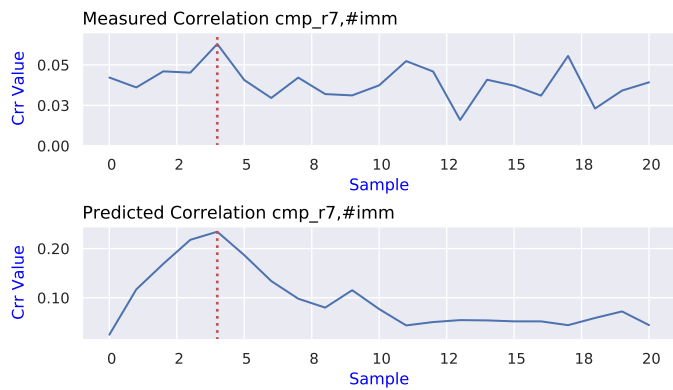
We propose ABBY, the first framework to automate the profiling of the architectural layer. As a result, ABBY significantly reduces the



(a) Correlation result for add instruction estimation Vs measured



(b) Correlation result for orr instruction estimation Vs measured



(c) Correlation result for cmp instruction estimation Vs measured

**Figure 12: Correlation result of 1K estimated traces Vs. measured traces on ALU instructions**

human effort necessary to create leakage models. ABBY is scalable and can be transferred to different architectures. The most challenging aspect of porting ABBY to different architectures is the creation of the labeled dataset. We used standard tools for developing the ABBY framework. Using the ABBY-CM0 dataset, we explored several leakage models ranging from primitive to transition-based, which include pipeline effects and instruction interaction via hidden registers. We evaluate and compare the performance of these leakage models using statistics metrics such as  $R^2$  and  $F$ -test and side-channel attacks. When considering side-channel attacks, transition-based leakage models such as ELMO and our CH model are superior to primitive ones. When comparing the performance of ELMO with our deep learning leakage model, the results are very close, demonstrating the quality of the ABBY-CM0 dataset and, ultimately, the effectiveness of the ABBY framework. We constructed the ABBY-CM3 dataset to investigate the scalability of the ABBY framework while we developed a side-channel power simulator for the ALU component based on this dataset. Despite statistical metrics, correlation results for simulation are close to the measured trace.

In future work, our objective is to develop a simulator for Cortex-M3 based on the ABBY framework that covers all the components besides the ALU. Moreover, we will investigate how targeted microarchitectural benchmarks such as [38] and optimizations of the model architecture can further enhance the performance of ABBY.

## 8 AVAILABILITY

Our framework and dataset will be published upon the paper's acceptance and can be downloaded from GitHub: ABBY-Framework while documentation is accessible on ABBY-Documentation.

## REFERENCES

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. <https://www.tensorflow.org/> Software available from tensorflow.org.
- [2] Dakshi Agrawal, Bruce Archambeault, Josyula R. Rao, and Pankaj Rohatgi. 2003. The EM Side-Channel(s). In *Cryptographic Hardware and Embedded Systems - CHES 2002*, Burton S. Kaliski, çetin K. Koç, and Christof Paar (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 29–45.
- [3] Jude Angelo Ambrose, Naeill Aldon, Aleksandar Ignjatovic, and Sri Parameswaran. 2008. Anatomy of differential power analysis for AES. In *2008 10th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing*. IEEE, 459–466.
- [4] ARM. 2006. *Cortex™-M3 Revision r1p1 Technical Reference Manual*.
- [5] ARM. 2010. *Cortex™-M3 Revision r2p0 Technical Reference Manual*.
- [6] ARM. 2018. ARM® v6-M Architecture Reference Manual. <https://developer.arm.com/documentation/ddi0419/e/?lang=en>.
- [7] ARM. 2021. ARM®v7-M Architecture Reference Manual. <https://developer.arm.com/documentation/ddi0403/ee/?lang=en>.
- [8] Vipul Arora, Ileana Buhan, Guilherme Perin, and Stjepan Picek. 2021. A Tale of Two Boards: On the Influence of Microarchitecture on Side-Channel Leakage. In *Smart Card Research and Advanced Applications - 20th International Conference, CARDIS 2021, Lübeck, Germany, November 11-12, 2021, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 13173)*, Vincent Grosso and Thomas Pöppelmann (Eds.). Springer, 80–96. [https://doi.org/10.1007/978-3-030-97348-3\\_5](https://doi.org/10.1007/978-3-030-97348-3_5)
- [9] Josep Balasch, Benedikt Gierlichs, Roel Verdult, Lejla Batina, and Ingrid Verbauwhede. [n. d.]. Power analysis of Atmel CryptoMemory - recovering keys

- from secure EEPROMs, booktitle = Topics in Cryptology - CT-RSA 2012, The Cryptographers' Track at the RSA Conference, editor = O. Dunkelman, series = Lecture Notes in Computer Science, volume = 7178, publisher = Springer-Verlag, year = 2012, pages = 9–34.
- [10] Gilles Barthe, Marc Gourjon, Benjamin Grégoire, Maximilian Ortl, Clara Paglialonga, and Lars Porth. 2021. Masking in Fine-Grained Leakage Models: Construction, Implementation and Verification. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* 2021, 2 (2021), 189–228.
  - [11] Guido Bertoni, Vittorio Zaccaria, Luca Breviglieri, Matteo Monchiero, and Gianluca Palermo. 2005. AES power attack based on induced cache miss and countermeasure. In *International Conference on Information Technology: Coding and Computing (ITCC'05)-Volume II*. Vol. 1. IEEE, 586–591.
  - [12] Olivier Bronchain, Julien M. Hendrickx, Clément Massart, Alex Olshevsky, and François-Xavier Standaert. 2019. Leakage Certification Revisited: Bounding Model Errors in Side-Channel Security Evaluations. In *Advances in Cryptology - CRYPTO 2019*, Alexandra Boldyreva and Daniele Micciancio (Eds.). Springer International Publishing, Cham, 713–737.
  - [13] Ileana Buhari, Lejla Batina, Yuval Yarom, and Patrick Schaumont. 2022. SoK: Design Tools for Side-Channel-Aware Implementations. In *ASIA CCS '22: ACM Asia Conference on Computer and Communications Security, Nagasaki, Japan, 30 May 2022 - 3 June 2022*, Yuji Suga, Kouichi Sakurai, Xuhua Ding, and Kazuo Sako (Eds.). ACM, 756–770. <https://doi.org/10.1145/3488932.3517415>
  - [14] Elad Carmon, Jean-Pierre Seifert, and Avishai Wool. 2017. Photonic side channel attacks against RSA. In *2017 IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*. IEEE, 74–78.
  - [15] Suresh Chari, Josyula R. Rao, and Pankaj Rohatgi. 2002. Template Attacks. In *Cryptographic Hardware and Embedded Systems - CHES 2002, 4th International Workshop, Redwood Shores, CA, USA, August 13-15, 2002, Revised Papers (Lecture Notes in Computer Science, Vol. 2523)*, Burton S. Kaliski Jr., Çetin Kaya Koç, and Christof Paar (Eds.). Springer, 13–28. [https://doi.org/10.1007/3-540-36400-5\\_3](https://doi.org/10.1007/3-540-36400-5_3)
  - [16] Joan Daemen, Seth Hoffert, Gilles Van Assche, and Ronny Van Keer. 2018. The design of Xoodoo and Xooff. *IACR Trans. Symmetric Cryptol.* 2018 (2018), 1–38.
  - [17] Joan Daemen, Seth Hoffert, Michaël Peeters, Gilles Van Assche, and Ronny Van Keer. 2020. Xoodyak, a lightweight cryptographic scheme. *IACR Trans. Symmetric Cryptol.* 2020, S1 (2020), 60–87. <https://doi.org/10.13154/tosc.v2020.iS1.60-87>
  - [18] J. Daemen and Vincent Rijmen. 2002. AES the Advanced Encryption Standard. *The Design of Rijndael* 26 (01 2002).
  - [19] Arnaud de Grandmaison, Karine Heydemann, and Quentin L. Meunier. 2022. ARMISTICE: Microarchitectural Leakage Modeling for Masked Software Formal Verification. *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.* 41, 11 (2022), 3733–3744. <https://doi.org/10.1109/TCAD.2022.3197507>
  - [20] Nicolas Debande, Maël Berthier, Yves Bocktaels, and Thanh-Ha Le. 2012. Profiled Model Based Power Simulator for Side Channel Evaluation. *Cryptology ePrint Archive*, Report 2012/703. <https://eprint.iacr.org/2012/703>.
  - [21] Jerry den Hartog, Jan Verschuren, Erik P. de Vink, Jaap de Vos, and W. Wiersma. 2003. PINPAS: A Tool for Power Analysis of Smartcards. In *SEC*. 453–457.
  - [22] Margaux Dugardin, Louiza Papachristodoulou, Zakaria Najm, Lejla Batina, Jean-Luc Danger, and Sylvain Guilley. 2016. Dismantling Real-World ECC with Horizontal and Vertical Template Attacks. In *Constructive Side-Channel Analysis and Secure Design - 7th International Workshop, COSADE 2016, Graz, Austria, April 14-15, 2016, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 9689)*, François-Xavier Standaert and Elisabeth Oswald (Eds.). Springer, 88–108.
  - [23] Thomas Eisenbarth, Timo Kasper, Amir Moradi, Christof Paar, Mahmoud Salmasizadeh, and Mohammad T. Manzuri Shalmani. 2008. On the Power of Power Analysis in the Real World: A Complete Break of the KeeLoq Code Hopping Scheme. In *Advances in Cryptology - CRYPTO 2008 (Lecture Notes in Computer Science, Vol. 5157)*, David Wagner (Ed.). Springer-Verlag Berlin Heidelberg, 203–220.
  - [24] Georges Gagnerot. 2013. *Étude des attaques et des contre-mesures associées sur composants embarqués*. Ph. D. Dissertation. Université de Limoges.
  - [25] Si Gao and Elisabeth Oswald. 2022. A Novel Completeness Test for Leakage Models and Its Application to Side Channel Attacks and Responsibly Engineered Simulators. In *Advances in Cryptology - EUROCRYPT 2022*, Orr Dunkelman and Stefan Dziembowski (Eds.). Springer International Publishing, Cham, 254–283.
  - [26] Si Gao, Elisabeth Oswald, and Dan Page. 2021. Reverse Engineering the Micro-Architectural Leakage Features of a Commercial Processor. *IACR Cryptol. ePrint Arch.* (2021), 794. <https://eprint.iacr.org/2021/794>
  - [27] Aymeric Genêt, Natacha Linard de Guertechin, and Novak Kaludjerović. 2021. Full Key Recovery Side-Channel Attack Against Ephemeral SIKE on the Cortex-M4. In *Constructive Side-Channel Analysis and Secure Design*, Shivam Bhasin and Fabrizio De Santis (Eds.). Springer International Publishing, Cham, 228–254.
  - [28] Barbara Gigerl, Vedad Hadzic, Robert Primas, Stefan Mangard, and Roderick Bloem. 2020. Coco: Co-Design and Co-Verification of Masked Software Implementations on CPUs. *Cryptology ePrint Archive*, Report 2020/1294. <https://eprint.iacr.org/2020/1294>.
  - [29] G. Goodwill and J.J.B. Jun and P. Rohatgi. 2018. A testing methodology for side channel resistance validation. *NIST non-invasive attack testing workshop* (2018).
  - [30] Miao Tony He, Jungmin Park, Adib Nahiyan, Apostol Vassilev, Yier Jin, and Mark Mohammad Tehranipoor. 2019. RTL-PSC: Automated Power Side-Channel Leakage Assessment at Register-Transfer Level. In *VTS*. 1–6.
  - [31] Benjamin Hettwer, Stefan Gehrer, and Tim Güneysu. 2020. Deep Neural Network Attribution Methods for Leakage Analysis and Symmetric Key Recovery. In *Selected Areas in Cryptography - SAC 2019*, Kenneth G. Paterson and Douglas Stebila (Eds.). Springer International Publishing, Cham, 645–666.
  - [32] Sorin A. Huss, Marc Stöttinger, and Michael Zohner. 2013. *AMASIVE: An Adaptable and Modular Autonomous Side-Channel Vulnerability Evaluation Framework*. Lecture Notes in Computer Science, Vol. 8260. Springer, 151–165.
  - [33] NumFOCUS Inc. 2009. .
  - [34] Yuval Ishai, Amit Sahai, and David A. Wagner. 2003. Private Circuits: Securing Hardware against Probing Attacks. In *Advances in Cryptology - CRYPTO 2003, 23rd Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 2003, Proceedings (Lecture Notes in Computer Science, Vol. 2729)*, Dan Boneh (Ed.). Springer, 463–481. [https://doi.org/10.1007/978-3-540-45146-4\\_27](https://doi.org/10.1007/978-3-540-45146-4_27)
  - [35] Paul C. Kocher. 1996. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In *Advances in Cryptology - CRYPTO '96*, Neal Koblitz (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 104–113.
  - [36] Yann Le Corre, Johann Großschädl, and Daniel Dinu. 2018. Micro-architectural Power Simulator for Leakage Assessment of Cryptographic Software on ARM Cortex-M3 Processors. In *COSADE*. 82–98.
  - [37] Moritz Lipp, Andreas Kogler, David Oswald, Michael Schwarz, Catherine Eason, Claudio Canella, and Daniel Gruss. 2021. PLATYPUS: Software-based power side-channel attacks on x86. In *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 355–371.
  - [38] Ben Marshall, Dan Page, and James Webb. 2021. MIRACLE: MiCRo-Architectural Leakage Evaluation. *Cryptology ePrint Archive*, Report 2021/261. <https://ia.cr/2021/261>.
  - [39] David McCann, Elisabeth Oswald, and Carolyn Whitnall. 2017. Towards Practical Tools for Side Channel Aware Software Engineering: 'Grey Box' Modelling for Instruction Leakeages. In *USENIX Security Symposium*. 199–216.
  - [40] ST Microelectronic. [n. d.]. *STM32 Mainstream MCUs*.
  - [41] ST Microelectronic. 2015. *Discovery kit with STM32F030R8 MCU*. <https://www.st.com/en/evaluation-tools/32f0308discovery.html>
  - [42] ST Microelectronic. 2015. *Discovery kit with STM32F051R8 MCU*. <https://www.st.com/en/evaluation-tools/stm32f051r8discovery.html>
  - [43] Amir Moradi, Alessandro Barenghi, Timo Kasper, and Christof Paar. 2011. 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*. 111–124.
  - [44] Adib Nahiyan, Jungmin Park, Miao Tony He, Yousef Iskander, Farimah Farahmandi, Domenic Forte, and Mark Mohammad Tehranipoor. 2020. SCRIPT: A CAD Framework for Power Side-channel Vulnerability Assessment Using Information Flow Tracking and Pattern Generation. *ACM Trans. Design Autom. Electr. Syst.* 25, 3 (2020), 26:1–26:27.
  - [45] Kostas Papagiannopoulos and Nikita Veshchikov. 2017. Mind the Gap: Towards Secure 1st-Order Masking in Software. In *COSADE (Lecture Notes in Computer Science, Vol. 10348)*. Springer, 282–297.
  - [46] Riscure. 2002. Inspector-SCA. <https://www.riscure.com/security-tools/inspector-sca/>. [Online; accessed 7-Apr-2021].
  - [47] Matthieu Rivain and Emmanuel Prouff. 2010. Provably Secure Higher-Order Masking of AES. In *Cryptographic Hardware and Embedded Systems, CHES 2010, 12th International Workshop, Santa Barbara, CA, USA, August 17-20, 2010. Proceedings (Lecture Notes in Computer Science, Vol. 6225)*, Stefan Mangard and François-Xavier Standaert (Eds.). Springer, 413–427. [https://doi.org/10.1007/978-3-642-15031-9\\_28](https://doi.org/10.1007/978-3-642-15031-9_28)
  - [48] Pau Rodríguez, Miguel A. Bautista, Jordi González, and Sergio Escalera. 2018. Beyond one-hot encoding: Lower dimensional target embedding. *Image and Vision Computing* 75 (2018), 21–31. <https://doi.org/10.1016/j.imavis.2018.04.004>
  - [49] Asanka Sayakkara, Nhien-An Le-Khac, and Mark Scanlon. 2019. A survey of electromagnetic side-channel attacks and discussion on their case-progressing potential for digital forensics. *Digital Investigation* 29 (2019), 43–54.
  - [50] Nader Sehatbakhsh, Baki Berkay Yilmaz, Alenka G. Zajic, and Milos Prvulovic. 2020. EMSim: A Microarchitecture-Level Simulation Tool for Modeling Electromagnetic Side-Channel Signals. In *HPCA*. 71–85.
  - [51] Madura A Shelton, Lukasz Chmielewski, Niels Samwel, Markus Wagner, Lejla Batina, and Yuval Yarom. 2021. Rosita++: Automatic higher-order leakage elimination from cryptographic code. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*. 685–699.
  - [52] Madura A. Shelton, Niels Samwel, Lejla Batina, Francesco Regazzoni, Markus Wagner, and Yuval Yarom. 2021. Rosita: Towards Automatic Elimination of Power-Analysis Leakage in Ciphers. In *NDSS*.
  - [53] Danilo Sijacic, Josep Balasch, Bohan Yang, Santosh Ghosh, and Ingrid Verbauwhede. 2020. Towards efficient and automated side-channel evaluations at design time. *J. Cryptogr. Eng.* 10, 4 (2020), 305–319.
  - [54] Patanjali SLPSC, Prasanna Karthik Vairam, Chester Rebeiro, and V. Kamakoti. 2019. Karna: A Gate-Sizing based Security Aware EDA Flow for Improved Power

Side-Channel Attack Protection. In *ICCAD*. 1–8.

[55] QEMU team. [n. d.]. *QUICK EMULATOR tool*. Available at <https://www.qemu.org/>, version.

[56] C. Thuillet, P. Andouard, and O. Ly. 2009. A Smart Card Power Analysis Simulator. In *2009 International Conference on Computational Science and Engineering*, Vol. 2. 847–852. <https://doi.org/10.1109/CSE.2009.119>

[57] Vivek Tiwari, Sharad Malik, Andrew Wolfe, and Mike Tien-Chien Lee. 1996. Instruction Level Power Analysis and Optimization of Software. In *VLSI Design*. 326–328.

[58] Nikita Veshchikov. 2014. SILK: high level of abstraction leakage simulator for side channel analysis. In *PPREW@ACSAC*. 3:1–3:11.

[59] Nikita Veshchikov and Sylvain Guilley. 2017. Use of Simulators for Side-Channel Analysis. In *EuroS&P Workshops*. 104–112.

[60] David Welch. [n. d.]. *Thumbulator-Tool*. Available at <https://github.com/dwelch67/thumbulator>.

[61] Carolyn Whitnall and Elisabeth Oswald. 2019. A Critical Analysis of ISO 17825 ('Testing Methods for the Mitigation of Non-invasive Attack Classes Against Cryptographic Modules'). In *ASLACRYPT (3)*. 256–284.

[62] Yuan Yao, Tarun Kathuria, Baris Ege, and Patrick Schaumont. 2020. Architecture Correlation Analysis (ACA): Identifying the Source of Side-channel Leakage at Gate-level. In *HOST*. 188–196.

[63] Yuan Yao, Patrick Schaumont, Jasper Van Woudenberg, Cees-Bart Breunese, Edgar Mateos Santillan, and Steve Stecyk. 2020. Verification of Power-based Side-channel Leakage through Simulation. In *MWSCAS*. 1112–1115.

[64] Yuan Yao, Mo Yang, Conor Patrick, Bilgiday Yuce, and Patrick Schaumont. 2018. Fault-assisted side-channel analysis of masked implementations. In *2018 IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*. 57–64. <https://doi.org/10.1109/HST.2018.8383891>

## 9 APPENDICES

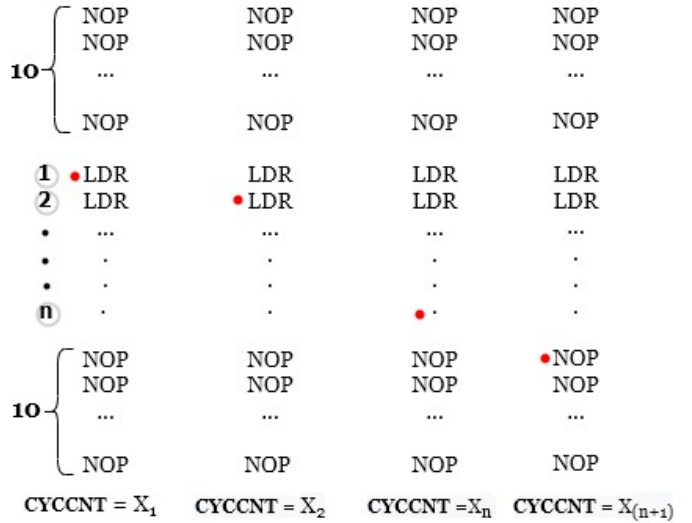


Figure 13: Solution for hardware breakpoints effect