

ROSITA: Towards Automatic Elimination of Power-Analysis Leakage in Ciphers

Madura A. Shelton
University of Adelaide
madura.shelton@adelaide.edu.au

Niels Samwel
Radboud University
nsamwel@cs.ru.nl

Lejla Batina
Radboud University
lejla@cs.ru.nl

Francesco Regazzoni
ALaRI – USI
regazzoni@alari.ch

Markus Wagner
University of Adelaide
markus.wagner@adelaide.edu.au

Yuval Yarom
University of Adelaide and Data61
yval@cs.adelaide.edu.au

Abstract—Since their introduction over two decades ago, physical side-channel attacks have presented a serious security threat. While many ciphers’ implementations employ masking techniques to protect against such attacks, they often leak secret information due to unintended interactions in the hardware. We present ROSITA, a code rewrite engine that uses a leakage emulator which we amended to correctly emulate the microarchitecture of a target system. We use ROSITA to automatically protect a masked implementation of AES and show the absence of exploitable leakage at only a 11% penalty to the performance.

I. INTRODUCTION

The seminal work of Kocher [33] demonstrated that interactions of cryptographic implementations with their environment can result in *side channels*, which leak information on the internal state of ciphers. Since then multiple side channels have been demonstrated, exploiting various effects, such as timing [1, 11, 14, 15], power consumption [34], electromagnetic (EM) emanations [17, 27, 44], shared microarchitectural components [29, 50], and even acoustic and photonic emanations [4, 30, 35, 47]. These side channels pose a severe risk to the security of systems, and in particular to cryptographic implementations, and effective side channel attacks have been demonstrated against block and stream ciphers [31, 45], public key systems, both traditional [22, 41] and post quantum [43], cryptographic primitives implemented in real-world devices [5, 24], and even non-cryptographic algorithms [8].

Many approaches to protect devices have been suggested, in particular against power and EM attacks. These range from special logic styles that are designed to make leakage data-independent [19, 26, 49], through noise generation to hide the signal [39], to algorithmic changes designed to prevent certain leakage [38]. In particular, *masking* is a common algorithmic countermeasure in which all intermediate (secret-dependent) values in the ciphers are combined with random masks, so that the leakage of one or even a few values does not provide the attacker enough information to recover the secrets.

The resistance of implementations to side channel leakage is typically evaluated by assessing the information that a

hypothetical attacker can observe. The attacker is specified in terms of the technical capabilities of the measurement setup, such as the sampling rate and the accuracy of the equipment, and the attacker model, which specifies the scenario in which the attacker gets access to measure the device’s leakage. The assessment is then based on the number of observations that an attacker needs to collect to retrieve information from the side channel and on a, somewhat quantified, number of bits that could become available through this side channel. The more observations and capabilities are required, the more resistant the implementation is. Correspondingly, the less leakage is released through the system performing computation, the more secure the implementation is perceived.

Thus, to achieve a certain level of security, implementations of cryptographic primitives often go through an iterative process of evaluating the leakage and manually tweaking the code to remove the leakage [16]. Such a process is costly because it requires a high level of expertise and significant manual labor, especially taking into account state-of-the-art side channel adversaries.

Recently, several works have experimented with tools that provide a high resolution emulation of the power consumption by computing devices as they execute software [54]. The results of such emulations are combined with standard statistical tests [10] to perform leakage assessment of software without executing it on the actual hardware [42]. Observing that these tools eliminate the hardware from the leakage evaluation process, we ask the following question:

Can leakage emulators be used for automatic elimination of side channel leakage from software implementations?

In this work, we answer this question in the affirmative, albeit with some caveats. Specifically, we develop a semi-automatic tool, ROSITA that evaluates leakage from assembly code and applies fixes to remove the leakage.

On a masked implementation of AES¹, the tool successfully

¹https://github.com/Secure-Embedded-Systems/Masked-AES-Implementation/blob/master/Byte-Masked-AES/byte_mask_aes.c

eliminates leakage from up to 40,000 traces, at a performance cost of 11%.

To develop ROSITA, we first investigate the requirements that emulators should satisfy for the proposed use, including both the nature and the accuracy of the emulation output. Namely, cost-efficient emulators are typically built on simple models, which, by default, do not use large sample sets. Hence, this feature of emulators contradicts their desired ability to resist strong adversaries.

We then proceed to evaluate the ELMO leakage emulator [37]. We compare the leakage emulation with the leakage we measure directly from the physical device. We find that although the emulation is of a high quality, ELMO misses some leakage sources, requiring that the model be augmented to identify the missed leakage. We also modify ELMO, to produce the output required for leakage elimination.

Reproducing the work of McCann et al. [37], we observe that a masked implementation of AES² is leaking. We extend their work and find that the leakage is caused by unintended interactions between masked values. We identify two main sources for such unintended interactions. The first is register reuse, and the second is microarchitectural effects of the processor.

To eliminate leakage, we repeatedly execute ELMO to identify code locations that leak and then invoke ROSITA to rewrite the code. Finally, we test the code produced by ROSITA on the physical device, and show the absence of discernible leakage in up to 40,000 traces.

To sum up, the contributions of this work are:

- We propose a framework for generating leakage-resilient implementations of masked ciphers by iteratively rewriting the code at leakage points. (Section III)
- We identify multiple sources of leakage that ELMO leakage model fails to detect. We extend ELMO and augment it to identify these sources of leakage. We further augment ELMO to report instructions that leak secret information and the specific cause of leakage for each. (Section IV)
- We develop ROSITA, a code rewrite engine that uses the output of our augmented ELMO to rewrite leaking instructions and eliminate leakage. (Section V)
- We use ROSITA to rewrite a masked implementation of AES. We test the code ROSITA produces and show the absence of leakage in up to 40,000 traces. (Section VI)

II. BACKGROUND

A. Side-Channel Attacks

When software runs on some hardware, it affects the environment it executes in. This effect can be manifested as variations in the power consumption, electromagnetic emanations, temperature, and state of various CPU components. As these variations correlate with the operation of the algorithm, monitoring these variations discloses information about the internal state, and as such provide a ‘communication’ channel

that transfers information from the software being observed to the observer. These unintended communication channels are often known as *side channels*.

In 1996, Kocher [33] noted that the information acquired through a timing side channel may reveal secret information processed by a cryptographic algorithms. Since then a significant effort has been dedicated to analyzing and eliminating side-channel leakage, particularly in the context of cryptographic implementations [9, 27, 34, 36, 44].

Protection against side-channel attacks depends on both the channel and the attacker capabilities. The *constant-time* programming style [12, 32], which avoids secret-dependent branches and table lookups, has proven effective against attacks that depend on the cryptographic operation timing or on its effect on micro-architectural components [11, 14, 29]. However, when the leakage correlates with the data values being processed, constant-time programming is not sufficient to protect the implementation.

One of the main approaches to protect cryptographic implementations against side channels that leak information on data values is *masking* [13, 18]. In a nutshell, masking combines internal values processed by the cryptographic algorithm with random *masks*, typically by xoring a mask and the internal value. The masked internal values are uniformly distributed, thus the leakage of a masked value does not reveal any information to the attacker.

Although theoretically secure, naive masking often fails to provide the required protection. The main cause is that side-channel leakage often correlates not only with the value being processed, but also with changes in the state of components within the processor. In particular, changing the value of a stored bit typically requires more power than keeping the value unchanged. Thus, the power required for changing the contents of a storage element correlates not only with the new value stored in the element, but also with the difference between the old and the new values.

For masked implementations, this unintended interaction between the old and the new values in storage elements can be disastrous. For example, let us assume that the software processes two values, $v_1 \oplus m$ and $v_2 \oplus m$, where v_1 and v_2 are two internal values that need to be kept secret, and m is a random mask. Leaking either $v_1 \oplus m$ and $v_2 \oplus m$ provides the attacker no information, because these values are uniformly distributed. However if the software writes $v_1 \oplus m$ to a register that contains $v_2 \oplus m$, the power consumption will correlate with $(v_1 \oplus m) \oplus (v_2 \oplus m) = v_1 \oplus v_2$, leaking information on the relationship between the two *unmasked* values.

One approach for fixing this issue is to go to higher order masking that can protect against leakage of multiple values [40]. Specifically, Balasch et al. [6] demonstrate that unintended interactions due to transitions half the masking order. However, algorithms that use high-order masking are significantly more complex than simple first-order masking, and require much more resources and randomness.

Thus, the common practice, from the practitioners’ point of view, is to use first order masking, and to combine it with

²https://github.com/bristol-sca/ELMO/blob/master/Examples/FixedvsRandom/MaskedAES_R1/MaskedAES_R1.s

ad-hoc countermeasures for unintended interactions. Then, the implementers of the cryptographic software typically perform leakage assessment to find code locations that leak information. If leakage is detected, they apply manual modifications to the code to eliminate the leakage, and they repeat this process until no further leakage is evident.

We now turn our attention to assessing leakage in cryptographic implementations.

B. Side Channel Leakage Assessment

The leakage assessment of a device is very important for both semiconductor and security evaluation industries, and it has accordingly received a lot of attention in the past years. Depending on the attacker model, there could be many attack vectors possible and exhaustive evaluation (by trying all possible attacks) is simply not feasible. As an alternative, a leakage evaluation methodology called Test Vector Leakage Assessment (TVLA) was proposed Cooper et al. [20], Tunstall and Goodwill [51]. The question that it answers relates to the presence of any sort of leakage (from side channel measurements) of the targeted implementation running on the device of interest. Clearly, only the negative answer suggests that the device is indeed secure and the positive one does not tell much on the exploitability of the leakage. Nevertheless, due to its simplicity and efficacy, the TVLA method has been found a useful first diagnostics in side-channel leakage assessment and has become the mainstream for security analysts. The main idea is as follows: the analysis aims at differentiating between two sets of measurements, one for fixed and one for random inputs, by means of the Welch's t -test [57]. Finding those two sets "sufficiently different" statistically, implies that the device leaks some data-dependent information through a side channel. The main limitation is in evaluating each point in time independently, so the leakage from combining multiple points is not detected. To overcome this limitation, Schneider and Moradi [48] extend the t -test to handle also multiple points. In addition, to address leakages distributed over multiple orders they propose the use of the χ^2 -test as a natural complement to the Welch's t -test. However, the leakage assessment of interest for this work requires just TVLA, as we are considering the first order masking scheme only.

C. Leakage Emulators

Since conducting real experiments for leakage detection is costly, leakage emulation has been adapted as an alternative. PINPAS [23] was the first of its kind that was used to detect leakage in Java-based smart cards. Since then, various other methods of emulating leakage have surfaced. SPICE-based methods such as the one presented by Aigner et al. [3] are the most accurate because they simulate internal circuitry of a CPU down to minute details. The drawback is that these emulators tend to be slower. Alternatively, researchers have looked at emulating at the source code level [52] and at machine instruction level [42, 53]. In source code level emulation, the emulator does not have any information about a specific CPU that will be used to run the compiled machine

code of a given source code. It emulates leakage having source code as its only input. In instruction level emulation, the emulation is based on the machine code that will be executed on a certain CPU or more generally a specific CPU kind. Recently, advanced instruction level emulators have been introduced that use power traces from real experiments to make better estimates [37]. Similarly, advanced characteristics of CPUs such as instruction pipelining have found their way into recent leakage emulators [21].

D. Automatic Approaches to Handling Side-Channel Leakage

Due to the numerous problems and pitfalls with countermeasures against side-channel attacks as previously discussed, researchers found several automated approaches for handling side-channel leakage. The approaches can be grouped into two categories, one of which simulates the leakage and bases their mitigations on the simulation results. The second category takes another approach where leakage is detected in the source code using several techniques.

Veshchikov [52] presents the SILK simulator, which simulates a high level abstraction of the source code of an algorithm that generates traces. Another simulator, MAPS [21] targets the Cortex-M3 and bases its leakage properties on the Hardware Description Language (HDL) source code. The simulator mainly focusses on leakage caused in the pipeline.

These two simulators only automate the generations of traces. Hence, they are basically assisting the leakage evaluation process in terms of speeding it up. In contrast to this, our work uses simulated traces to find leakage. Once the leakage is found, it determines the cause and applies a technique to the corresponding assembly code to counter the leakage.

There are also works approaching this problem via the formal verification of masking schemes. A success there would imply the absence of leakage. Barthe et al. [7] describe how to automatically verify higher-order masking schemes and present a new method based on program verification techniques. The work of Wang and Schaumont [55] explains how formal verification and program synthesis can be used to detect side-channel leakage, prove the absence of such leakage and modify software to prevent such leaks. However, both of these works remain limited in the ways they model the hardware and actual implementations.

Closer to ours are the works that, although loosing on their generality, address the problem of "fixing" the leakage from a specific device. For instance, Papagiannopoulos and Veshchikov [42] performed an in-depth investigation of device specific effects that violate the independent leakage assumption. They also provide an automated tool that is capable of detecting such violations in AVR assembly code.

Another method to eliminate timing side channels in software was proposed by Wu et al. [59]. Their method requires a list with secret variables as input and produces code that is functionally equivalent to the original code but without timing side channels. In a recently published work Wang et al. [56] describe a type-based method for detecting leaks in source code. They implemented their mitigations in a compiler

and evaluated their method. Eldib and Wang [25] propose a method to add countermeasures to source code that masks all intermediate computation results such that all intermediate results are statistically independent.

In Agosta et al. [2], the authors introduce a framework to automate the application of countermeasures against Differential Power Analysis (DPA). Their approach adds multiple versions of the code preventing an attacker from recognizing the exact point of leakage.

The previously mentioned works all automatically analyze the code to determine if there is leakage and modify the code accordingly. Our work also modifies the code to prevent leakage, but to find this leakage we also simulate traces and argue that they are adequately replacing real traces. Our approach is much more cost-efficient, albeit effective, as if when an actual chip would be tested.

III. ROSITA OVERVIEW

ROSITA aims to automate the process of producing leakage-resilient software. We assume that the underlying algorithm employs a protection technique, such as masking. However, unintended interactions between data, introduced in the execution of the software, e.g. when overwriting one value with another where both values are masked with the same mask, can break the *independent leakage assumption* (ILA) [46] and leak secret information through a physical side channel, such as the power channel. Thus, to fix such leaks, implementers typically go through a manual, iterative process whereby the software is installed on the target device, the leakage is measured, and fixes are applied to the machine code, until the leakage is reduced to an acceptable level for the target use case.

This process, naturally, requires a significant level of expertise both in setting up and conducting the experiment to assess the leakage and in fixing the software to reduce the leakage. Moreover, because the assessment requires a large number of encryption rounds on relatively low-performing devices, and a number of repetitions in repairing the leakage and evaluating, the process is time consuming.

ROSITA automates this process as shown in the diagram in Figure 1. To produce leakage-resilient cryptographic software, we start with a (masked) implementation of the cryptographic primitive. We use cross-compilation to produce both the assembly code and the binary executable for the target device. The binary executable is then passed to a leakage emulator, in our case ELMO [37], to perform leakage assessment. This assessment identifies the leakage and the machine instructions that cause it. ROSITA processes the output of ELMO, together with the assembly code. It applies a set of rules that replace leaky assembly instructions with functionally-equivalent sequences of instructions that do not leak. Afterwards, the produced assembly program is assembled and fed back to ELMO and the process repeats until no further fixes can be applied, at which time ROSITA produces a report indicating the remaining leakage, if any.

Note that our approach makes use of a leakage emulator. Prior static-analysis-based solutions, such as [21, 42, 52, 55,

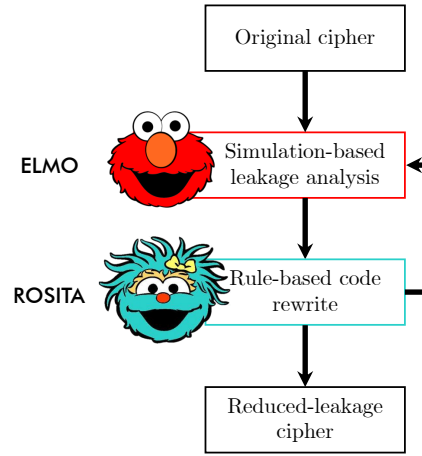


Figure 1: Automatic Elimination of Side-Channel Leaks with ROSITA.

59], rely on tags that identify the nature of values within the program. For example, in ASCOLD [42] the programmer needs to assign tags to values, e.g. identifying them as random or masked. The main downside of the manual approach is that any mistake the programmer makes in tagging values can be translated to missed leakage. In contrast, ROSITA applies de-facto industry-standard tools, such as TVLA, to the emulated power trace. As such, ROSITA depends neither on the programmer’s proficiency nor on specific properties of the masking scheme to detect leakage. Subject to the accuracy of the emulator and the strength of the statistical tools applied, ROSITA will detect leakage in the implementation (up to the level which the masking scheme used is meant to protect).

IV. THE ELMO LEAKAGE EMULATOR

Due to ROSITA’s reliance on a leakage emulator, care should be taken when selecting one. We now present ELMO [37], the leakage emulator we use with ROSITA. We first describe how ELMO models the device it emulates and the leakage. We then identify limitations for using ELMO with ROSITA and describe how we address these.

A. The ELMO Leakage Model

Emulating the hardware at the transistor level would produce the most accurate leakage estimate. However, this is often infeasible, both due to the complexity of such analysis and because the hardware implementation details are not available to the security evaluators and software developers.

Instead, leakage emulators use an abstract model of the device and of its power consumption. The abstract model is significantly simpler than emulating at the transistor level. At the same time, using an abstract model reduces accuracy and may result in missing some leakage. Thus, the *leakage model* presents a trade-off between modeling cost and accuracy.

ELMO’s model of the hardware considers bit values and changes in bit values over the ALU inputs and outputs and memory instructions. Specifically, each operand is compared

to the corresponding operand of the preceding instruction. Power consumption is modeled as linear combinations of bit values or bit changes.

ELMO models 21 instructions that its authors claim cover typical use in cryptography. These 21 instructions are divided into five groups, each modeled separately. To generate the model, power traces are collected while the processor executes sequences of three instructions. Each trace is processed to select a point-of-interest to be used as a representative of the trace. ELMO then performs a linear regression on the data collected in the traces to find the coefficients for the model.

The model itself consists of 19 main components, each modeling a specific part of the architecture. These cover:

- A linear combination of the bit flips between each operand of the current instruction and the corresponding operand of the previous and the subsequent instructions.
- A linear combination of the bit values of the operands of the current instruction.
- The instruction groups of the previous and subsequent instructions.

ELMO provides a pre-computed model of the STM32F030³ evaluation board which features an ARM Cortex-M0 based STM32F030R8T6 System-on-Chip (SoC).⁴

B. Evaluation Setup

To evaluate ELMO, we compare its output with leakage assessment of the code on the real hardware. Our evaluation setup is shown in Figure 2.

We evaluate ELMO with the same STM32F030 Discovery evaluation board used in McCann et al. [37]. Following the instructions of McCann et al., we disconnect one of the two power inputs of the System on Chip (SoC) and attach a 330 Ω shunt resistor to the second power input.

We use a PicoScope 6404D with a Pico Technology TA046 differential probe connected to the oscilloscope via a Langer PA 303 preamplifier, to measure the voltage drop across the shunt resistor as a proxy for the power consumption of the SoC. See circuit diagram in Figure 3.

We sample every 12.8 ns, which, with a clock rate of 8 MHz, is roughly 9.77 samples per clock cycle. The samples are 8-bit wide and our PicoScope can store up to 2 giga samples before running out of memory.

We use a control PC to orchestrate the experiments. The PC controls the oscilloscope and the STM32F030 Discovery evaluation board. It sends the software to be tested and the data to be used to the evaluation board, and collects the trace data from the oscilloscope.

Each experiment collects multiple power traces from running the software on the evaluation board. The execution of the software alternates between the fixed and the random cases. Thus, half of the collected traces are for the fixed case and the other half is for the random tests. To identify the start and end of the segment that we monitor, we use output pins of




Figure 2: Evaluation Setup.

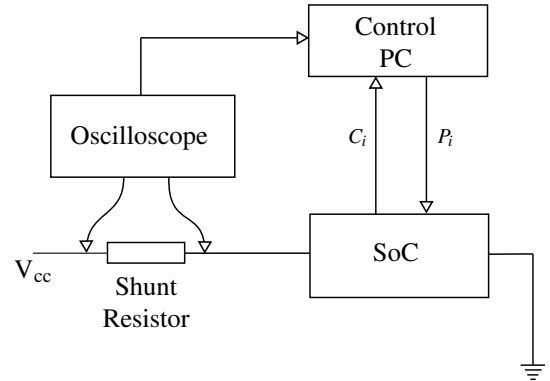


Figure 3: Evaluation Setup — Circuit Diagram

the device to trigger the trace collection and to mark the end of the points of interest. Later, we use these trigger points to filter out traces with clock jitter. In order to align the traces, we use the eShard SCARED library⁵.

To detect leakage, we employ non-specific TVLA. That is, we check the distribution of the values at each trace point (after alignment) and use the Welch t -test to check if the samples in fixed and in the random traces are drawn from the same distribution. Following the common practice in the domain, we use a t -test value above 4.5 or below -4.5 as an indication of leakage.

³<https://www.st.com/en/evaluation-tools/32f0308discovery.html>

⁴ELMO also provides a model for the Cortex-M4-based STM32F4 Discovery board, which we do not use in this work.

⁵<https://gitlab.com/eshard/scared>

C. Validating the Setup

First, to validate our setup, we reproduce the results of McCann et al. [37]. Specifically, we perform a fixed vs. random test on the code in Listing 1, which contains an implementation of one of the steps in the AES encryption known as the SHIFTRROWS operation. Specifically, register `r1` points to the 16 bytes that represent the state of the AES encryption. SHIFTRROWS performs a fixed permutation of these bytes. The implementation loads three four-byte words and uses the `rors` instruction to rotate the bytes, before storing them back to the state.

```
ldr r4, [ r1, #4 ]
rors r4, r5
str r4, [ r1, #4 ]
ldr r4, [ r1, #8 ]
rors r4, r6
str r4, [ r1, #8 ]
ldr r4, [ r1, #12 ]
rors r4, r3
str r4, [ r1, #12 ]
```

Listing 1: SHIFTRROWS from McCann et al. [37]

For the fixed vs. random test we collect 2500 traces where the state contains fixed data masked with the same mask value and 2500 traces where the state consists of random values masked with the same mask. The value of the mask is chosen in random for each trace. We compare the distribution of the power reading in each sample point between the fixed and the random traces, and calculate the Welch t -test to check the likelihood that the two distributions are the same. As mentioned before, following common practice in side-channel analysis, we consider the distributions different enough to indicate leakage if the absolute value of the t -test value is above 4.5.


Figure 4a shows the result of the fixed vs. random test. The horizontal axis shows the time and the vertical axis shows the t -test value. We indicate instruction boundaries with vertical bars, and the t -test threshold of ± 4.5 with horizontal red lines. Comparing the figure to the results of running ELMO on the same code, shown in Figure 4b, we see that ELMO produces a fairly accurate simulation of the leakage.

In particular, our figure resembles Figure 5 of McCann et al. [37], with only minor differences that reflect the different test environment.


D. Storage Elements and the ELMO Model

The ELMO model of the hardware only looks at interactions between arguments and outputs of successive instructions. However, it overlooks interactions due to storage of values in registers, memory, and other storage elements.

To evaluate interaction overlooked by ELMO, we design a systematic battery of small sequences of code that aim at highlighting interactions via storage elements between instructions. We use a fixed vs. random test with each of these sequences



(a) Real traces in STM32F030.



(b) Simulated traces from ELMO.

Figure 4: Fixed vs. random of the AES SHIFTRROWS operation.

to identify the existence (or absence) of a common storage element.

```
1 movs r7, r7
2 movs r7, r7
3 movs r4, r7
4 ldr r4, [r1, #4]
5 movs r7, r7
6 movs r7, r7
7 movs r7, r7
8 movs r7, r7
9 movs r5, r7
10 ldr r5, [r1, #8]
11 movs r7, r7
12 movs r7, r7
13 movs r7, r7
```

Listing 2: Evaluating the `ldr` instruction

Each of the tests focuses on two instructions that we check for interaction. Listing 2 shows an example of such a sequence which tests consecutive load instructions. In this example, the register `r7` is initialized to a random 32-bit value and the memory locations at `r1+4` and `r1+8` contain two “plaintext” values, both masked with the same mask. For all of the fixed tests, the values of the plaintexts remain the same. In the random tests each of these values is selected at random.

The two instructions we check for interactions are in Lines 4 and 10. Before each of these loads we set registers `r4`

and `r5`, respectively, to the random value in `r7` (Lines 3 and 9), to mask any interactions between previous values of these registers and the values we load from memory. Finally, the `movs r7, r7` instructions separate the instructions we monitor to avoid pipeline interactions already covered in the ELMO model and to better highlight each of the instructions in the captured traces.

As Figure 5a shows, performing the fixed vs. random test on the hardware shows a clear leakage at the second `ldr` instruction. This is in stark contrast to the ELMO results (Figure 5b) which show no sign of leakage. We hypothesize that the leakage is caused by a storage element in the memory bus that stores the most recent value loaded from memory. When the second load instruction overwrites the contents of the storage element, the transition effect correlates with the exclusive-or of the two plaintext. Hence the distributions at the fixed and the random tests differ, demonstrating a leakage.

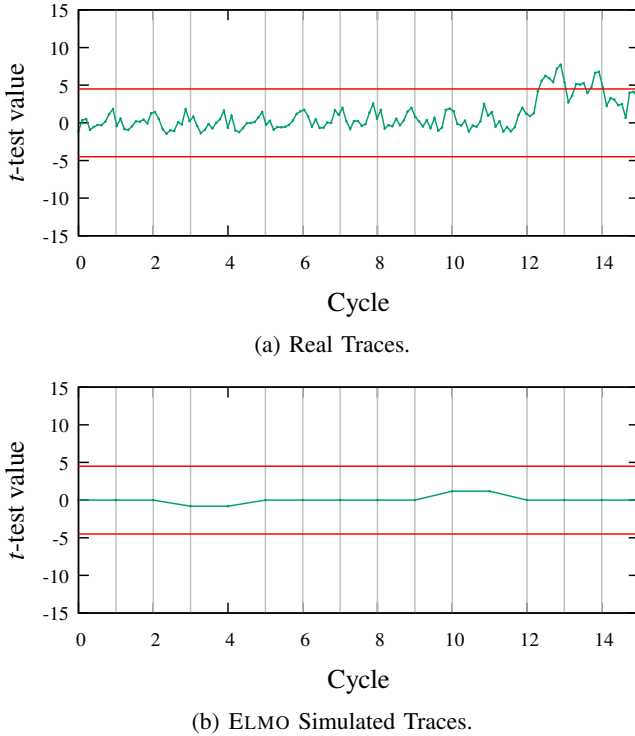


Figure 5: Evaluating consecutive `ldr` instruction.

E. Findings

We run a broad range of such tests, with (1) some focusing on architecturally known storage elements, such as registers and memory, and (2) others aiming to find microarchitectural storage elements by testing interactions between pairs of instructions. We find several sources of leakage that ELMO does not identify. We note that Gao [28] also identifies many of the issues we find; however, their identification was driven by the iterative tweaking of a cipher, while our systematic approach is cipher-agnostic.

1) *Registers*: We find that overwriting a register leaks the (weighted) Hamming distance between the previous value and the new value. This is a significant leakage source, because reusing a register that contains a masked value for another value with the same mask leaks secret information. Unlike Papagiannopoulos and Veshchikov [42], we do not find leakage across different registers.

2) *Memory*: Writing data to memory interacts with data already stored in the same location. Hence, overwriting one masked value with another may remove the mask, leaking the values.

3) *Memory Bus*: The memory bus seems to have a storage element that stores the most recent value stored to or loaded from the memory. When loading from or storing to memory, the value of the storage element is overwritten, leaking the Hamming distance between the previous and the new value. This leakage differs from the two described above, and happens irrespective of the registers and the memory addresses used. Consequently, when writing to or reading from memory, care should be taken to only access non-secret values or values masked with different masks.

It is important to note that the storage element always stores a 32-bit word. Thus, when loading or storing a byte, the whole 4-byte aligned 32-bit word that contains the byte is moved to the storage element. This may create memory interaction between memory operations that seem completely unrelated. For example, consider the code in Listing 3. In this example we assume that memory locations `0x300` and `0x400` both contain one secret byte each, both masked with the same mask. The code in this example performs two memory operations, the first stores a byte into address `0x303` and the second reads a byte from location `0x402`. We note that none of these locations contains secret data, and the data stored is also not secret. However, the store operation loads the 32-bit word in memory locations `0x300–0x303` into the memory bus, and the following load operation replaces the contents with the 32-bit word in memory location `0x400–0x403`. This causes an interaction between the data in memory locations `0x300` and `0x400`, leaking the Hamming distance between the data stored in these locations.

```

1  movs r3, 0x303
2  movs r4, 0x402
3  movs r7, r7
4  movs r7, r7
5  movs r7, r7
6  strb r5, [r3]
7  movs r7, r7
8  movs r7, r7
9  movs r7, r7
10 ldrb r6, [r4]
11 movs r7, r7
12 movs r7, r7
13 movs r7, r7

```

Listing 3: Example of word interaction

A further issue in the memory bus is an interaction between the bytes of words loaded from or stored to memory. Specifically, our analysis shows that when memory data is accessed, consecutive bytes in the word interact with each other. Thus, if a word contains multiple bytes that are all masked with the same mask, loading it from or storing it to memory will leak the Hamming distance between consecutive bytes. We note that due to the memory bus storage element described above, the leakage occurs even if the memory access operations access a single byte of a 32-bit word.

F. Extending the ELMO Model

Recall (Section IV-A) that ELMO builds its model using a linear regression from traces collected from sequences of three instructions. To account for the effects of storage elements we identified, we update the model to include a few more components.

Whereas the ELMO model treats each operand separately, we also look at combinations of bits across the two operands of the instruction. Because the first operand is typically the destination register, correlating the two operands captures the effect of calculating the result of the operation and overwriting the destination register.

To capture interactions in the memory bus and the memory, we run tests that set the state (memory and memory bus contents) to a known value before measuring the leakage from a sequence of three instructions. We then correlate this state with the leakage.


In total, our model has 21 components. We validate our model by repeating the test cases used for identifying the storage elements. For example, Figure 6 shows the real and the emulated leakage from running the code in Listing 2 on the real hardware (top) and in ELMO updated with our extended model (bottom). Comparing to Figure 5b we see that our model identifies the leakage that the original ELMO model misses. We note that ELMO associates the leakage with the cycle in which the instruction enters the pipeline, even if the leakage occurs in later cycles. Consequently, the leakage in the emulated trace appears earlier than in the real trace.

V. CODE REWRITE IN ROSITA


As Section III describes, the core of ROSITA is a rewrite engine that uses the output of the ELMO emulator to drive code fixes for leakage. We assume that the original code is masked, i.e. it does not leak at the algorithm level. However, the translation of the algorithm into machine code and the execution of this machine code can result in unintended and unexpected leakage. In this section, we review the causes of leakage we identify, and describe the fixes the ROSITA applies for each. We begin with a high-level description of ROSITA and proceed with details on the rewrite rules it applies.

A. ROSITA Design

ROSITA is a rewrite engine that takes the code and the output of ELMO, and rewrites the code to avoid leakage. To decide which rewrite rule to apply, ROSITA needs more information



(a) Real Traces.



(b) ELMO Simulated Traces with Updated Model.

Figure 6: Evaluating consecutive `ldr` instruction.

than the original ELMO provides. Specifically, when leakage is identified, ROSITA needs to know the cause of the leakage.

To find the cause of the leakage, we modify ELMO to not only emulate the leaked signal, but also to find which of the components of the model is causing the leak. We, therefore calculate the t -test value not only for the combined emulated signal, but also to separate components of the signal.⁶ Thus, for example, we keep track of the t -test value of the part of the signal contributed by each of the instruction operands, by interactions between the instruction result and its operands and by interactions between the instruction results or operands with values previously stored to or loaded from memory. Using this information, when ELMO reports that an instruction leaks, we can inspect the components and identify the leakage cause.

The main strategy ROSITA uses to fix the leakage is to wipe stored state with a random mask. For that, ROSITA dedicates a *mask register* (ROSITA uses register `r7`), which is initialized with a random 32-bit mask. When compiling the software, we use the flag `-ffixed-r7` to direct the compiler not to use the mask register, ensuring that its contents are not modified except by ROSITA.

B. Operand Interaction

One of the common forms of unintended interaction is between the operands of successive instructions. Technically, as McCann et al. [37] note, loading an operand to the bus leaks

⁶Implementation note: to reduce memory usage, we calculate the t -test values incrementally, using Welford's algorithm [58].

the Hamming distance between the value previously held in the bus and the new value. If both values use the same mask, the Hamming distance between the masked values is the same as that between the original values.

ROSITA identifies such leakage by checking the various t -test values calculated for the operands and their relationship with those of prior instructions. In the case that the leakage is caused by such interaction, ROSITA inserts an access to the mask register, using `movs r7, r7`. The instruction moves the contents of the mask register into the mask register, and is therefore functionally a no-op. However, because the value of the mask register goes through the bus, the previous contents of the bus is forgotten, removing the interaction between two masked values.

C. Register Reuse

Due to the limited number of registers, compilers and programmers often reuse those, e.g. when the old contents is either consumed or stored in memory. Reusing a register rarely removes the old contents from it. Consequently, when new data is loaded into a register, it interacts with algorithmically unrelated data that remains from prior uses of the register.

Papagiannopoulos and Veshchikov [42] note that if the old contents and the new contents are both masked using the same mask value, the difference, between the masked contents, i.e. their exclusive or, is the same as the difference between the unmasked contents. Consequently, when a register is used consecutively for two values with the same mask, it leaks the difference between the values.

To identify this form of leakage, ROSITA checks the t -test value of overwriting register value. Once identified, ROSITA wipes the old contents of the register by copying the contents of the mask register to the destination register of the leaking instruction, as Papagiannopoulos and Veshchikov [42] suggest. For example, suppose that the instruction `movs r3, r4` leaks because both `r3` and `r4` contain values masked with the same mask. To eliminate the leak, ROSITA inserts `movs r3, r7` before the leaking instruction.

D. Rotation Operations

Rotation operations show interaction between the value pre and post rotation. When a single masked value is rotated, this interaction is unlikely to leak secret data because the mask hides the contents. However, when rotating a word comprised of multiple masked values that all use the same mask, the result of the rotation may align the masked values, effectively nullifying the mask, leaking the difference of the unmasked values.

We propose two approaches to remove this leakage. For the discussion we assume the case from the AES implementation we investigate, where the AES Shift Rows operation is implemented by rotating 32 bit words that contain four masked bytes each. However, these approaches would work for similar cases, where the masked values are not bytes.

For the discussion, we assume that we would like to rotate the register `r2`, whose value is a concatenation of four masked

bytes: $(b_1 \oplus m) || (b_2 \oplus m) || (b_3 \oplus m) || (b_4 \oplus m)$. Rotation of `r2` by a multiple of 8 bits would result in leakage of information on the value of b_i . For example, assuming `r3` contains the value 8, the instruction `ror r2, r3` would set the value of `r2` to $(b_2 \oplus m) || (b_3 \oplus m) || (b_4 \oplus m) || (b_1 \oplus m)$, and the interaction between the original and the rotated values of `r2` would leak the Hamming weight of $(b_1 \oplus b_2) || (b_2 \oplus b_3) || (b_3 \oplus b_4) || (b_4 \oplus b_1)$.

Word Mask. A straightforward approach for preventing leakage when rotating a word is to mask the word with our mask register (`r7`), rotate both the word and the mask register and then use the rotated mask to unmask the word. Thus, instead of rotating `r2`, we rotate `r2 \oplus r7`. As an example, Figure 7 shows how ROSITA fixes a `rors r2, r3` instruction that ELMO indicates is leaking.

```
rors r2, r3          eors r2, r7
                    rors r2, r3
                    rors r7, r3
                    eors r2, r7
```

Figure 7: Masking rotation operations. The leaking `ror` operation on the left is replaced with a masking sequence on the right.

We note that this sequence modifies the contents of our mask register. However, this has no effect on the functionality because the mask register is assumed to be random and there is no long-term dependency on its exact contents.

Partial Rotations. An alternative approach is to combine multiple shifts to avoid rotations of multiples of the data size. For example, a rotation by 8 bits can be replaced with a rotation by 3 bits followed by a rotation by 5 bits.

ROSITA employs the word mask approach both because it is more general, i.e. does not depend on the size of the rotation, and because it already has the mask register, which it uses for the other fixes.

E. Memory Operations

As discussed in Section IV-E, there are several effects that can cause interactions between values used in memory operations. These include a storage element in the memory bus that remembers recently accessed memory value and consequently leaks the Hamming distance between the remembered value and the current one on memory access operations, interaction between loaded and stored values and the previous contents they overwrite, and an interaction between bytes in stored words.

When ELMO indicates that a load instruction leaks due to interaction with the memory bus storage element, ROSITA wipes the contents of the bus by pushing the mask register to the stack and popping from the stack to the destination register of the load instruction. Figure 8 shows an example of an `ldr` instruction (left) that leaks through interaction of the loaded value with a previously loaded value. To fix this, ROSITA inserts a `push` and a `pop` instructions before the load, yielding the code fragment in the right. Popping the

mask to the destination of the load instruction also protects against leakage through interaction with the previous value of the destination register.

```
ldr r2, [r3]      push r7
                  pop r2
                  ldr r2, [r3]
```

Figure 8: A leaking load instruction (left) and the fixed sequence (right).

Due to the more intricate potential interactions, the picture with store instructions is a bit more complex. To overcome interactions with the previous value used on the memory bus and to address possible interactions with the previous contents of memory, ROSITA first stores the mask register into the destination location and then performs the required store (See Figure 9).

```
str r2, [r3]      str r7, [r3]
                  str r2, [r3]
```

Figure 9: A leaking store instruction (left) and the fixed sequence (right).

When byte interaction within the stored data leaks, ROSITA stores one byte at a time. In such a case, care should be taken to ensure that these bytes and the operations required for their storage do not create unintended interactions, leading to a relatively long code segment in Figure 10. The code uses two registers chosen to not conflict with the store, `r0` and `r6` in the example in Figure 10. The first is used for selecting the byte to store, while the second is used for the byte. ROSITA uses two stores for each byte to avoid interactions on the memory bus or in the DRAM. While this rewrite rule eliminates the leakage, as we see in Section VI the performance cost of using it is significant. As such, it may be better to avoid stores of words that contain multiple values masked with the same mask. Changing the logic of the cipher is outside the scope of ROSITA.

VI. EVALUATION

We evaluate ROSITA with two examples of cryptographic software. For both implementations we use the methodology of Section IV-B to measure leakage before and after rewriting the code using ROSITA. We now describe the results.

A. AES SHIFTRROWS

McCann et al. [37] demonstrate the accuracy of ELMO using a short sequence of code from a trivially masked implementation of the AES SHIFTRROWS operation, see Listing 1. In Section IV-C we reproduce their results, showing leakage both in the ELMO simulation and in the actual hardware.

As apparent in the code, the SHIFTRROWS operation is implemented by loading a word consisting of four (masked) bytes of the cipher state, using a rotation operation to permute these bytes, and storing the results back. Because the same

```
str r2, [r3]      push {r6}
                  push {r0}
                  movs r0, #0xff
                  movs r6, r2
                  ands r7, r7
                  ands r6, r0
                  lsls r0, #0
                  strb r0, [r3, #0]
                  strb r6, [r3, #0]
                  movs r6, r7
                  movs r6, r2
                  movs r0, #0xff
                  lsls r0, #8
                  ands r7, r7
                  ands r6, r0
                  lsrs r0, #8
                  lsrs r6, #8
                  strb r0, [r3, #1]
                  strb r6, [r3, #1]
                  .
                  .
                  .
                  pop {r0}
                  pop {r6}
```

Figure 10: Addressing byte interaction in stores. A leaking store instruction (left) and the fixed sequence (right).

mask is used for all four bytes, all of the operations used in this implementation leak through interactions within the word.


Figure 11 shows the leakage of the original code in Listing 1, and the leakage of the code produced by ROSITA. As we can see, the rewritten code does not leak. However, due to the complexity of fixing the internal leakage in memory operations, ROSITA introduces a significant performance overhead, slowing the operation by a factor of 15. However, because the SHIFTRROWS operation takes only a small part of the full AES implementation, the impact of the high overhead for SHIFTRROWS is likely not to be so significant.

B. AES First Round


To evaluate ROSITA with a more complete implementation, we use the byte-masked implementation by the Secure Embedded Systems Lab, Virginia Tech⁷. The implementation follows the approach of Mangard et al. [36, Section 9.2.1]. Unlike the implementation of McCann et al. [37], this implementation uses byte loads and stores for the SHIFTRROWS operation. Following the suggestion of Gao [28], we use different masks for each row to avoid leakage through interactions between bytes in memory words.

As shown in Figure 12a, the implementation shows significant leakage at 10,000 traces. We use our improved ELMO to emulate the leakage at 10,000 traces, and use ROSITA

⁷<https://github.com/Secure-Embedded-Systems/Masked-AES-Implementation/tree/master/Byte-Masked-AES>



(a) Before applying code rewrites



(b) After applying code rewrites


Figure 11: Masked AES SHIFTRows by McCann et al. [37] (10000 traces), note the increase in code size from 15 cycles to 228 after rewrites are applied.

to fix identified leakage. As Figure 12b shows, the process clears all leakage at 10,000 traces. In this case, ROSITA only results in a moderate performance overhead of less than 11%. (1293 cycles for the original implementation vs. 1430 for the rewritten code.)


To determine the trend of leakage, we perform the fixed vs. random test on the hardware with a varying number of traces. Figure 13 shows the results for both the original and the fixed implementations. The horizontal axis shows the number of traces used for the fixed vs. random test, and the vertical is the maximum absolute value of the t -test for each of the implementations. As we can see, the original implementation shows increasing leakage, crossing the significance threshold at as little as 1000 traces. In contrast to this, our fixed implementation does not show any significant leakage at up to 40,000 traces.

VII. CONCLUSIONS

Since their introduction over two decades ago, physical side-channel attacks have presented a serious security threat, particularly to small computational devices that need to maintain secrets under the physical control of the adversary. To protect against such attacks, many ciphers' implementations employ masking techniques that combine intermediate values with randomly selected masks. As a consequence, due to the mask being uniformly distributed, leakage of a masked value does not reveal information to the adversary. While proven secure



(a) Before applying code rewrites (1293 cycles)



(b) After applying code rewrites (1430 cycles)

Figure 12: Masked AES Implementation from Virginia Tech (10000 traces)




Figure 13: AES t -test value trend

against certain attacks, in practice masked implementations often leak secret information due to unintended interactions between masked values involving hardware they are loaded and stored to. To fix these leaks, the common practice is to repeatedly “tweak the code until it stops leaking”.

In this work, we have set out to explore if leakage emulators can be used for the *automatic* elimination of side channel leakage from software implementations. To achieve this, we have created a code rewrite engine called ROSITA and combined it with the leakage emulator ELMO:

- ROSITA incorporates rules to mitigate leakage arising from operand interactions, register reuse, rotation operations, and memory operations.
- ELMO has undergone a major upgrade for two reasons:

firstly, it had to be able to tell ROSITA the cause of the leakage, and secondly, we have added support by including the values that would be in the memory bus and in the barrel shifter, as both can hold state that can leak information.

In our proof-of-concept, we used ROSITA with our version of ELMO to automatically protect a masked implementation of AES. Our experiments using the actual hardware show the absence of exploitable leakage at only a 11% penalty to the performance.

ACKNOWLEDGEMENTS

Francesco Regazzoni received support from the European Union Horizon 2020 research and innovation program under CERBERO project (grant agreement number 732105).

This research was partially supported by a gift from Intel.

REFERENCES

- [1] O. Aciğmez, W. Schindler, and Ç. K. Koç, "Improving Brumley and Boneh timing attack on unprotected SSL implementations," in *CCS*, 2005.
- [2] G. Agosta, A. Barenghi, and G. Pelosi, "A code morphing methodology to automate power analysis countermeasures," in *DAC*, 2012, pp. 77–82.
- [3] M. J. Aigner, S. Mangard, F. Menichelli, R. Menicocci, M. Olivieri, T. Popp, G. Scotti, and A. Trifiletti, "Side channel analysis resistant design flow," in *ISCAS*, 2006.
- [4] M. Backes, M. Dürmuth, S. Gerling, M. Pinkal, and C. Spöck, "Acoustic side-channel attacks on printers," in *USENIX Security*, 2010, pp. 307–322.
- [5] J. Balasch, B. Gierlichs, R. Verdult, L. Batina, and I. Verbauwhede, "Power analysis of Atmel CryptoMemory - recovering keys from secure EEPROMs," in *CT-RSA*, 2012, pp. 9–34.
- [6] J. Balasch, B. Gierlichs, V. Grosso, O. Reparaz, and F. Standaert, "On the cost of lazy engineering for masked software implementations," in *CARDIS*, 2014, pp. 64–81.
- [7] G. Barthe, S. Belaïd, F. Dupressoir, P. Fouque, B. Grégoire, and P. Strub, "Verified proofs of higher-order masking," in *EUROCRYPT (1)*, 2015, pp. 457–485.
- [8] L. Batina, S. Bhasin, D. Jap, and S. Picek, "CSI NN: reverse engineering of neural network architectures through electromagnetic side channel," in *USENIX Security*, 2019, pp. 515–532.
- [9] A. G. Bayrak, F. Regazzoni, D. Novo, P. Brisk, F.-X. Standaert, and P. Jenne, "Automatic application of power analysis countermeasures," *IEEE Trans. Computers*, vol. 64, no. 2, pp. 329–341, 2015.
- [10] G. Becker, J. Cooper, E. DeMulder, G. Goodwill, J. Jaffe, G. Kenworthy, T. Kouzminov, A. Leiserson, M. Marson, P. Rohatgi, and S. Saab, "Test vector leakage assessment (TVLA) methodology in practice," http://icmc-2013.org/wp/wp-content/uploads/2013/09/Rohatgi_Test-Vector-Leakage-Assessment.pdf, 2013.
- [11] D. J. Bernstein, "Cache-timing attacks on AES," 2005, Preprint available at <http://cr.yp.to/papers.html#cachetiming>.
- [12] D. J. Bernstein, T. Lange, and P. Schwabe, "The security impact of a new cryptographic library," in *LATINCRYPT*, 2012, pp. 159–176.
- [13] J. Blömer, J. Guajardo, and V. Krummel, "Provably secure masking of AES," in *SAC*, 2004, pp. 69–83.
- [14] B. B. Brumley and N. Tuveri, "Remote timing attacks are still practical," in *ESORICS*, 2011, pp. 355–371.
- [15] D. Brumley and D. Boneh, "Remote timing attacks are practical," in *USENIX Security*, 2003, pp. 1–14.
- [16] C. Carlet, J. Danger, S. Guilley, and H. Maghrebi, "Leakage squeezing: Optimal implementation and security evaluation," *J. Mathematical Cryptology*, vol. 8, no. 3, pp. 249–295, 2014.
- [17] V. Carlier, H. Chabanne, E. Dottax, and H. Pelletier, "Electromagnetic side channels of an FPGA implementation of AES," *IACR Cryptology ePrint Archive* report 2004/145, 2004.
- [18] S. Chari, C. S. Jutla, J. R. Rao, and P. Rohatgi, "Towards sound approaches to counteract power-analysis attacks," in *CRYPTO*, 1999, pp. 398–412.
- [19] Z. Chen and Y. Zhou, "Dual-rail random switching logic: A countermeasure to reduce side channel leakage," in *CHES*, 2006, pp. 242–254.
- [20] J. Cooper, E. DeMulder, G. Goodwill, J. Jaffe, and G. Kenworthy, "Test vector leakage assessment methodology in practice," *International Cryptographic Module Conference*, 2013.
- [21] Y. L. Corre, J. Großschädl, and D. Dinu, "Micro-architectural power simulator for leakage assessment of cryptographic software on ARM Cortex-M3 processors," in *COSADE*, 2018, pp. 82–98.
- [22] E. De Mulder, S. B. Örs, B. Preneel, and I. Verbauwhede, "Differential power and electromagnetic attacks on a FPGA implementation of elliptic curve cryptosystems," *Comput. Electr. Eng.*, vol. 33, no. 5-6, pp. 367–382, 2007.
- [23] J. den Hartog, J. Verschuren, E. P. de Vink, J. de Vos, and W. Wiersma, "PINPAS: A tool for power analysis of smartcards," in *SEC*, 2003, pp. 453–457.
- [24] T. Eisenbarth, T. Kasper, A. Moradi, C. Paar, M. Salmasizadeh, and M. T. M. Shalmani, "On the power of power analysis in the real world: A complete break of the KeeLoq code hopping scheme," in *CRYPTO*, 2008, pp. 203–220.
- [25] H. Eldib and C. Wang, "Synthesis of masking countermeasures against side channel attacks," in *CAV*, 2014, pp. 114–130.
- [26] W. Fischer and B. M. Gammel, "Masking at gate level in the presence of glitches," in *CHES*, 2005, pp. 187–200.
- [27] K. Gandolfi, C. Moutel, and F. Olivier, "Electromagnetic analysis: Concrete results," in *CHES*, 2001, pp. 251–261.
- [28] S. Gao, "A Thumb assembly based byte-wise masked AES implementation," https://github.com/bristol-sca/ASM_MaskedAES, 2019.
- [29] Q. Ge, Y. Yarom, D. Cock, and G. Heiser, "A survey of microarchitectural timing attacks and countermeasures on contemporary hardware," *J. Cryptographic Engineering*, vol. 8, no. 1, pp. 1–27, 2018.
- [30] D. Genkin, A. Shamir, and E. Tromer, "RSA key extraction via low-bandwidth acoustic cryptanalysis," in *CRYPTO (1)*, 2014, pp. 444–461.
- [31] L. Goubin and J. Patarin, "DES and differential power analysis: the 'duplication' method," in *CHES*, Aug 1999, pp. 158–172.
- [32] E. Käsper and P. Schwabe, "Faster and timing-attack resistant AES-GCM," in *CHES*, 2009, pp. 1–17.
- [33] P. C. Kocher, "Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems," in *CRYPTO*, 1996, pp. 104–113.
- [34] P. C. Kocher, J. Jaffe, and B. Jun, "Differential power analysis," in *CRYPTO*, 1999, pp. 388–397.
- [35] J. Krämer, D. Nedospasov, A. Schlösser, and J.-P. Seifert, "Differential photonic emission analysis," in *COSADE*, 2013, pp. 1–16.
- [36] S. Mangard, E. Oswald, and T. Popp, *Power Analysis Attacks: Revealing the Secrets of Smart Cards*. Springer, 2007.
- [37] D. McCann, E. Oswald, and C. Whittall, "Towards practical tools for side channel aware software engineering: 'grey box' modelling for instruction leakages," in *USENIX Security*, 2017, pp. 199–216.
- [38] T. S. Messerges, "Securing the AES finalists against power analysis attacks," in *FSE*, April 2000, pp. 150–164.
- [39] —, "Power analysis attacks and countermeasures for cryptographic algorithms," Ph.D. dissertation, 2000.
- [40] —, "Using second-order power analysis to attack DPA resistant software," in *CHES*, 2000, pp. 238–251.
- [41] T. S. Messerges, E. A. Dabbish, and R. H. Sloan, "Investigations of power analysis attacks on smartcards," in *USENIX — Smartcard '99*, 1999, pp. 151–162.
- [42] K. Papagiannopoulos and N. Veshchikov, "Mind the gap: Towards secure 1st-order masking in software," in *COSADE*, 2017, pp. 282–297.
- [43] A. Park, K. Shim, N. Koo, and D. Han, "Side-channel attacks on post-quantum signature schemes based on multivariate quadratic equations," *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, vol. 2018, no. 3, pp. 500–523, 2018.
- [44] J.-J. Quisquater and D. Samyde, "Electromagnetic analysis (EMA): Measures and counter-measures for smart cards," in *Smart Card Programming and Security*, 2001, pp. 200–210.
- [45] C. Rechberger and E. Oswald, "Practical template attacks," in *WISA*, 2004, pp. 443–457.
- [46] M. Renaud, F. Standaert, N. Veyrat-Charvillon, D. Kamel, and D. Flan-dre, "A formal study of power variability issues and side-channel attacks for nanoscale devices," in *EUROCRYPT*, 2011, pp. 109–128.
- [47] A. Schlösser, D. Nedospasov, J. Krämer, S. Orlic, and J.-P. Seifert, "Simple photonic emission analysis of AES," *J. Cryptographic Engineering*, vol. 3, no. 1, pp. 3–15, 2013.
- [48] T. Schneider and A. Moradi, "Leakage assessment methodology - A

- clear roadmap for side-channel evaluations,” in *CHES*, 2015, pp. 495–513.
- [49] K. Tiri and I. Verbauwhede, “A digital design flow for secure integrated circuits,” *IEEE Trans. on CAD of Integrated Circuits and Systems*, vol. 25, no. 7, pp. 1197–1208, 2006.
 - [50] E. Tromer, D. A. Osvik, and A. Shamir, “Efficient cache attacks on AES, and countermeasures,” *J. Cryptology*, vol. 23, no. 1, pp. 37–71, 2010.
 - [51] M. Tunstall and G. Goodwill, “Applying TVLA to public key cryptographic algorithms,” IACR Cryptology ePrint Archive report 2016/513, 2016.
 - [52] N. Veshchikov, “SILK: high level of abstraction leakage simulator for side channel analysis,” in *PPREW@ACISAC*, 2014, pp. 3:1–3:11.
 - [53] N. Veshchikov and S. Guilley, “Use of simulators for side-channel analysis,” in *EuroS&P*, 2017, pp. 51–59.
 - [54] —, “Use of simulators for side-channel analysis,” in *Euro S&P*, 2017, pp. 104–112.
 - [55] C. Wang and P. Schaumont, “Security by compilation: An automated approach to comprehensive side-channel resistance,” *SIGLOG News*, vol. 4, no. 2, pp. 76–89, 2017.
 - [56] J. Wang, C. Sung, and C. Wang, “Mitigating power side channels during compilation,” in *ESEC/SIGSOFT FSE*, 2019, pp. 590–601.
 - [57] B. L. Welch, “The generalization of ‘Student’s problem when several different population variances are involved,” *Biometrika*, vol. 34, no. 1–2, pp. 28–35, Jan. 1947.
 - [58] B. P. Welford, “Note on a method for calculating corrected sums of squares and products,” *Technometrics*, vol. 4, no. 3, pp. 419–420, 1962.
 - [59] M. Wu, S. Guo, P. Schaumont, and C. Wang, “Eliminating timing side-channel leaks using program repair,” in *ISSTA*, 2018, pp. 15–26.