

Plan your defense: A comparative analysis of leakage detection methods on RISC-V cores

Konstantina Miteloudi, Asmita Adhikary, Niels van Drueten, Lejla Batina, and Ileana Buhan

Digital Security Group, Radboud University, Nijmegen, The Netherlands
`firstname.lastname@ru.nl`, `niels@vandrueten.nl`

Abstract. Hardening microprocessors against side-channel attacks is a critical aspect of ensuring their security. A key step in this process is identifying and mitigating “leaky” hardware modules, which leak information during the execution of cryptographic algorithms. In this paper, we explore how different leakage detection methods, the Side-channel Vulnerability Factor (SVF) and the Test Vector Leakage Assessment (TVLA), contribute to hardening of microprocessors. We conduct experiments on two RISC-V cores, SHAKTI and Ibex, using two cryptographic algorithms, SHA-3 and AES. Our findings suggest that SVF and TVLA can provide valuable insights into identifying leaky modules. However, the effectiveness of these methods can vary depending on the specific core and cryptographic algorithm in use. We conclude that the choice of leakage detection method should be based not only on computational cost but also on the specific requirements of the system, the implementation of the algorithm examined and the nature of the potential threats.

Keywords: Pre-Silicon Design · Side-Channel Analysis · Leakage Detection Methods · Hardening Microprocessors.

1 Introduction

Microprocessors are at the heart of modern digital systems, from everyday consumer electronics to critical infrastructure. Ensuring their security against various forms of attacks is crucial. One such form of attack, side-channel attacks, exploits information leaked during the execution of cryptographic algorithms, potentially compromising the system’s security. As such, a key step in hardening microprocessors against side-channel attacks is identifying and mitigating these “leaky” hardware modules.

Several methods exist for detecting such leakages, each with strengths and weaknesses. Two methods, the Side-channel Vulnerability Factor (SVF) [5] and the Test Vector Leakage Assessment (TVLA) [9], have been commonly used in the field. SVF provides a measure of the potential exploitability of a leak, while TVLA offers a statistical framework for identifying whether a device under test is vulnerable to information leakage. However, the effectiveness of these methods can vary depending on the specific hardware and algorithm in use [20], [2]. In

this work, we applied these leakage assessment methods on two RISC-V cores: SHAKTI and Ibex. These cores, popular in various applications for their open-source nature, present a range of opportunities and challenges for security.

Our research question is: "How effectively can we harden microprocessors of varying sizes that run full cryptographic algorithms against side-channel attacks?". We aim to provide insights into how SVF and TVLA can be used to identify vulnerabilities in RISC-V cores of different sizes and how these findings can enforce the best strategies for hardening these cores. Through our experiments with different cryptographic algorithms, i.e. SHA-3 and AES, we seek to offer a comprehensive understanding of the ways these leakage detection methods can contribute to the hardening of RISC-V cores. Our findings will be of interest to security designers and architects, contributing to the development of more secure microprocessors/tools. [18].

2 Related work

Several approaches have been developed to identify and eliminate leaks during the pre-silicon phase [12]. These approaches can be roughly divided based on which device layer or development phase they aim to harden [4]. The level of leakage can be detected at top architectural choices [1] or/and can emerge from the micro-architectural behaviour [14].

De Mulder et al. [15] proposed a solution to protect an AES implementation against side channel leakage related to memory accesses on a RISC-V core. Gigerl et al. [8] introduced COCO, a tool that can detect gate-level leakage by simulating execution with Verilator. They annotate the registers and memory that hold secret data and trace their flow through the circuit to find possible sources of leakage. He et al. [10] estimate the power profile of a hardware design using functional simulation at the RTL level. Gao et al. [7] designed and implemented an ISE (Instruction Set Extension) called FENL that localizes and reduces microarchitectural leakage. The ISE acts as a leakage fence that prevents interaction between instructions. A similar approach is taken by Pham et al. [16], which combines a diversified ISE with hardware diversification through a co-processor to achieve leakage mitigation. Bloem et al. [3] extended the concept of hardware-software contracts to power-side channels and formally verified a wide range of instructions for implementing cryptographic algorithms for the RISC-V Ibex core. ACA [19] uses a gate-level model for a target design, typically available after logic synthesis and a side-channel leakage model. Kiaei and Schau-mont proposed Root Canal [11], a framework to help a designer with white-box access to the embedded CPU system uncover the origin of a side-channel leak. Root Canal can eliminate side-channel leaks before tape-out. After tape-out, changes to the hardware are no longer possible.

To our knowledge, none of the previous work investigates the impact of the leakage detection methods on identifying leaky modules. Arsath et al. [6] developed a framework, PLAN, that analyzes the RTL description of a processor and reports the information leakage in each of the processor modules. In PLAN, they

use a modified version of SVF as a leakage detection method, and they apply this method to a simulated RISC-V core running different cryptographic algorithms. The results of their analysis provide a ranking of the hardware modules based on their contribution to the overall leakage. In our work, we replicate the experimental setup of [6], applying both the SVF and the TVLA methods to compare their effectiveness in identifying leaky modules. This replication is the foundation for further exploration and comparison of these two leakage detection methods.

3 Preliminaries

3.1 Leakage detection methods

Let $\mathcal{X}(N, d)$ be a set of N traces. A *trace* is a time series with d samples recorded during the processing of an algorithm on a given device (e.g., an encryption operation) for a given input x .

Test Vector Leakage Assessment (TVLA) [9] is the most popular leakage detection method due to its simplicity and relative effectiveness. It comes in two flavours: *specific* and *non-specific*. The ‘fixed-vs-random’ is the most common non-specific 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. Welch’s two-sample t -value for equality of means applies to all trace samples in both cases. A difference between two sets larger than a given threshold is evidence of a leak’s presence.

Side-channel Vulnerability Factor (SVF) [5] measures side-channel information leakage by recognizing leaked execution patterns. SVF quantifies the similarity between patterns in the observations of the attackers or *side-channel traces* (\mathcal{S} , defined in (1), with the actual execution patterns of the victim or *oracle traces* (\mathcal{O} , defined in (2).

$$\mathcal{S} = \{(x_i; s_1^i, s_2^i, \dots, s_d^i), \text{ where } 1 \leq x_i \leq N\} \quad (1)$$

$$\mathcal{O} = \{(x_i; o_1^i, o_2^i, \dots, o_d^i), \text{ where } 1 \leq x_i \leq N\} \quad (2)$$

where s_j^i (and o_j^i) is sample j in side channel trace (and oracle trace, respectively) corresponding to input x_i .

The original version of the SVF algorithm proposed in [5], which we denote with SVF_{time} , quantifies *patterns in the time-domain* between an oracle and a side-channel trace. Arsath et al. [6] implemented a modified version of $\text{SVF}_{\text{input}}$ that is adapted to capture *patterns related to changes in the input data*, the typical cause of side-channel vulnerabilities. After data collection, both algorithms extract patterns in parallel for the oracle and side-channel trace. The difference between the implementation of SVF_{time} and $\text{SVF}_{\text{input}}$ is apparent in the construction of the similarity matrices as shown in equation (3), (4) and (5), (6) respectively.

$$M_{\text{time}}^{\mathcal{S}}(s_j^i, s_k^i) = \begin{cases} D(s_j^i, s_k^i), & \text{if } j < k \\ 0, & \text{if not.} \end{cases} \quad (3)$$

and

$$M_{\text{time}}^{\mathcal{O}}(o_j^i, o_k^i) = \begin{cases} D'(o_j^i, o_k^i), & \text{if } j < k \\ 0, & \text{if not.} \end{cases} \quad (4)$$

When computing SVF_{time} , the first step is to construct $M_{\text{time}}^{\mathcal{S}}$, the similarity matrix for a side-channel trace, using equation (3), where D is a distance (of choice) between samples s_j^i, s_k^i . The next step is to compute the similarity matrix, $M_{\text{time}}^{\mathcal{O}}$, for the oracle trace by computing the distance D' between the samples in the same power trace o_i, o_j , using equation (4). The correlation between $M_{\text{time}}^{\mathcal{S}}$ and $M_{\text{time}}^{\mathcal{O}}$ will give the SVF_{time} value.

$$M_{\text{input}}^{\mathcal{S}}(s_t^i, s_t^j) = \begin{cases} \bar{D}(s_t^i, s_t^j), & \text{if } i < j \\ 0, & \text{if not.} \end{cases} \quad (5)$$

and

$$M_{\text{input}}^{\mathcal{O}}(o_t^i, o_t^j) = \begin{cases} \tilde{D}(o_t^i, o_t^j), & \text{if } i < j \\ 0, & \text{if not.} \end{cases} \quad (6)$$

For computing SVF_{time} , one side-channel trace is sufficient however for computing $\text{SVF}_{\text{input}}$, multiple side-channel traces are required (to capture changes in the input). The procedure for computing $\text{SVF}_{\text{input}}$ is very similar to SVF_{time} . The difference is in the choice of samples for computing the similarity matrix is illustrated in equations (5) and (6). We first construct $M_{\text{input}}^{\mathcal{S}}$, the similarity matrix for the side-channel traces by computing the distance \bar{D} between the sample s_t corresponding to different input values x_i, x_j . In the same way, we calculate $M_{\text{input}}^{\mathcal{O}}$, the similarity matrix for the oracle traces, by computing the distance \tilde{D} between the sample o_t and the input x_i, x_j . All similarity matrices are triangular, as the main diagonal, which contains only zero values, is removed, and distance measures are commutative. In this work we implement $\text{SVF}_{\text{input}}$.

4 Experimental setup

Simulation setup. In our setup, we use two different 32-bit RISC-V cores, SHAKTI-C and Ibex, and two different algorithms, AES and SHA-3. SHAKTI-C [17] is a 5-stage pipeline in-order processor, while Ibex [13] is a 2-stage in-order processor. From each core, we selected specific modules to examine. We targeted the ones responsible for processing data and instructions and we excluded those that do other work, such as error checking. Specifically, for SHAKTI-C, we examine:

1. RF (Register File): integer and floating point registers.
2. CSR (Control and Status Register): raises interrupts on the processor.

3. ALU (Arithmetic Logic Unit): performs the arithmetic and logic operations.
4. FPU (Floating Point Unit): handles operations with floating point numbers.
5. Dcache: cache memory connected with the ALU.
6. MBOX: implements the multiplication and division operations.
7. BPU (Branch Prediction Unit): decides the next program counter.
8. ITLB (Instruction Translation-Look aside Buffer): keeps track of instructions recently used to avoid second access to memory.
9. DTLB (Data Translation Look aside Buffer): keeps track of data recently used to avoid second access to memory.

For Ibex, we examine modules with the same functionality:

1. RF (Register File): integer registers.
2. CSR (Control and Status Register): raises interrupts on the processor.
3. ALU (Arithmetic Logic Unit): performs the arithmetic and logic operations.
4. MULT/DIV (Multiplier/Divider Block): performs multiplication and division.
5. PF-BUF (Prefetch Buffer): fetches instructions from the memory.
6. LSU (Load-Store Unit): interfaces with the RF and the main memory to deal with load/store operations.

We ran simulations of every algorithm with Verilator simulator (Ver. 4.210) for 256 different inputs, randomly generated. For each simulation run, we took one Value Change Dump (VCD) file. VCD files show the value of every signal, of every module of the RISC-V core for every timestamp of an implementation. We processed the vcd files, as well as analyzed them, using Python. We parsed every file, and for each module, signals were concatenated for each timestamp, creating a composite signal. This signal represents the collective behaviour of all signals and the module's state at that specific point in time. Then, all the concatenated signals for all the selected timestamps, were processed differently, depending on the methodology selected.

SVF computation. For SVF computation, the concatenated signals are retained in their original form. These signals serve as a complex representation of the module's state. Each module contains N rows of signal values per timestamp s_t . Computing SVF requires the generation of oracle traces, which contain the intermediate values of the cryptographic algorithm during its execution. They are the expected values that the algorithm will produce at each time step given a particular input. We run simulations with the gcc compiler on a Linux system, using the same 256 inputs that we used in the simulations on RISC-V. We record different intermediate values in order to examine how the choice of the oracle trace affects the SVF. For the oracle set, we use the Hamming distance metric to compute $M_{\text{input}}^{\mathcal{O}}$, as described in 5 and 6. The oracle set will contain one sample for each input, so the size of \mathcal{O} is N .

The next step is to calculate a similarity matrix for the oracle trace and a similarity matrix for the side channel trace as described in 3.1. This step is necessary for correlation because the two traces contain different information and cannot be compared directly. We get two lists of Hamming distance values, and for each timestamp, we compute the Pearson correlation value between

the oracle list and the side-channel list. This value shows whether there is a linear correlation between the two lists in our implementation. The SVF value of a module and an oracle is the maximum of all Pearson correlation values. A module’s final SVF value is the oracles’ maximum SVF value. Arsath et al. [6] use 4 categories to show how much a module leaks: (1) 0.0 - 0.1: No leakage, (2) 0.1 - 0.3: Mild leakage, (3) 0.3 - 0.6: Medium leakage, (4) 0.6 - 1.0: Severe leakage.

TVLA computation. Our TVLA computation is based on the nonspecific fixed versus random test as specified in [9]. Since we worked on simulated executions, we needed a hypothetical power consumption model. This model is implemented with the Hamming weight (HW), and every timestamp of the concatenated signals takes an HW value. To calculate the t -value per timestamp, we used the `ttest_ind` function from the SciPy Python library, that calculates the t -value for the means of two independent samples of values. The non-specific fixed versus random test executes the fixed set multiple times to eliminate noise during a run. In our case, the runs do not contain noise because the run is simulated, and we know exactly all the signal values at any given time. Once all data for TVLA had been collected, we computed a t -value per cycle. This t -value is calculated from the fixed set of size one and the random set of size $N = 128$ for SHAKTI and $N = 256$ for Ibex.

4.1 Target cryptographic implementations

We chose unprotected implementations without any countermeasures, as our goal was to examine how leaks impact the different hardware modules. We ran the full algorithms, then chose to zoom in on different parts of them. We wanted to find any vulnerabilities that could arise, regardless of their nature or the specific type of data that might be exposed.

AES: We used the Tiny-AES¹ implementation, written in C. There are options for 128-bit, 192-bit, or 256-bit key sizes and options for ECB, CTR, and CBC modes. We used a key size of 128 bits in the ECB mode, and we encrypted one block of data.

SHA-3: We used the `tiny_sha3`² implementation, written in C. SHA-3 is a sponge function with the KECCAK-f[1600] as permutation function. We provided 832 bits of data as input, so we did not need padding. The output of SHA-3 is 384 bits.

5 Experimental results

To determine whether the choice of the leakage detection function influences the decision about the leakiness of a module, we used the experimental setup described in 4. We selected nine and six modules for the SHAKTI core and

¹ <https://github.com/kokke/tiny-AES-c>

² https://github.com/mjosaarinen/tiny_sha3

the Ibex core, respectively, that target processing instructions and data. We determined how “leaky” a module is by recording the maximum SVF value. In addition, we ran TVLA in fixed versus random mode. As this is a nonspecific test, we did not explicitly target intermediate variables. When comparing the results of TVLA with the results of SVF, we expected that TVLA will reveal more leaky points since our SVF procedure did not target all the intermediates.

AES. For experiments, we chose the typical candidate intermediate variables as the target: the first byte of the S-box output, `sbox_out1`, fifth byte of the S-box output, `sbox_out5`, and the first byte of the S-box input, $p \oplus k$ (`sbox_in1`). For the round output, we have oracles for the full round output (`mc_out`) and for the first byte of the round output (`mc_out1`).

SHA3. SHA-3 is the other implementation we analyzed. From the first round of the Keccak permutation function, we targeted χ as the only nonlinear operation. We defined three oracles based on the χ step: `bc`, a SHA-3 implementation-specific operation $bc[i] = st[j + i]$ where $i = 0$ and $j = 0$, not operation in $x \leftarrow x \oplus (\neg y \& z)$ where $i = 0$ and $j = 0$ and xor operation in $x \leftarrow x \oplus (\neg y \& z)$ where $i = 0$ and $j = 0$.

5.1 Case study: the SHAKTI Core

Figure 1 shows the leakage in ALU from the SHAKTI core, for both AES and SHA-3. The first plot shows the evolution in time of the SVF value, during the first round of AES. We represented the different target intermediates with different unique symbols. Horizontal dotted lines are drawn to indicate leakage thresholds. The yellow dotted line at 0.3 shows the minimum threshold for what we consider to be medium leakage and the red dotted line at 0.6 is for the severe leakage. Also, the oracles are highlighted with yellow where the SVF value is ≥ 0.3 and red when $\text{SVF} \geq 0.6$. The second plot shows the combined SVF with the t-value for the same module, i.e., ALU. In TVLA, we see leaks in almost all cycles and for most of the cycles, it shows leaks unrelated to the SVF oracles. For example, during the execution of the SubBytes operation (cycle 23.000 - 25.000), the SVF only finds leaks when the first or fifth S-box is computed. TVLA looks at all S-box operations and shows more leaks in cycles 23.000 - 25.000. This observation also holds for the cycles after SubBytes. TVLA indicates leaks in cycles where SVF does not show leakage. There are two possible explanations for this behaviour. The first is that TVLA shows false positives. The second is that the Hamming distance between the power traces and the oracles, that were used to calculate the SVF values, does not capture all the relations between the samples. The third plot shows the evolution in time of the SVF value, during the first round of the first execution of KECCAK-f, the SHA-3 permutation, for ALU. The fourth plot shows the combined SVF with the t-value for the same module and implementation, i.e., ALU and SHA3. We observe that SHA-3 is extremely leaky, according to TVLA. While SHA-3 executes the KECCAK-f function multiple times, and one execution of KECCAK-f takes multiple rounds, SVF will only find leaks in the intermediate values we target. TVLA finds multiple leaks during the whole execution of SHA3.

5.2 Case study: the Ibex Core

Figure 2 shows the leakage in ALU from the Ibex core for both AES and SHA-3. Similar to SHAKTI, the first plot shows the evolution in time of the SVF value, during the first round of AES and the second plot shows combined the SVF with the t-value. Also, the third plot shows the evolution in time of the SVF value during the first execution of SHA3 and the fourth plot shows the combined SVF with the t-value for the same module, ALU. As we observed in SHAKTI, for ALU, TVLA shows leakage in almost all cycles, while SVF shows only at some. We also observe that both methods find the same pattern of leaky cycles. This might indicate that the leakage is not caused only by one instruction but by a sequence of instructions as they are processed over time.

Figure 3 shows the leakage in the Register File from Ibex core for both AES and SHA-3. The second plot shows the SVF combined with the t-value for the same module. The Register File and ALU identify the same sequence of leaky operations. Again, TVLA shows more leaky points than SVF. The third plot shows the leakage value for SHA3, and the fourth plot, the combined SVF-TVLA. We observe that SVF shows severe leakage in the second half of this execution timing window. Specifically, the oracle *not* shows the same leaky points as TVLA, while in the first half of the execution, it identifies only a couple of leaky instructions. We also observe that oracle *xor* identifies leaky instructions, while oracle *bc* does not.

Figure 4 shows the combined SVF-TVLA on module MBOX from SHAKTI and Mult/Div from Ibex. Both modules are responsible for multiplication division. The first two plots show AES and SHA-3, respectively, for MBOX, while the last two show AES and SHA3, respectively, for Mult/Div. If we compare these plots with the plots from ALU, we observe a similar pattern of leakage, unexpected given the absence of multiplications or divisions in our code. Our examination of Ibex’s RTL code revealed that the Register File outputs are directly connected to both the ALU and Mult/Div unit, keeping both components active within the processor, regardless of their usage. Similar behavior was observed between the Register File and LSU, with leakage patterns on LSU mirroring those of the Register File, even while not only load/store instructions were executed. These observations confirmed that even non-utilized components within the processor can become sources of information leakage.

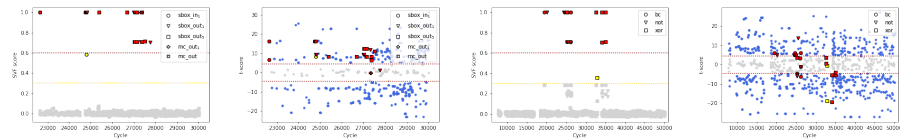


Fig. 1. SVF AES, combined SVF-TVLA AES, SVF SHA-3, combined SVF-TVLA SHA-3. For ALU module on SHAKTI core (left to right).

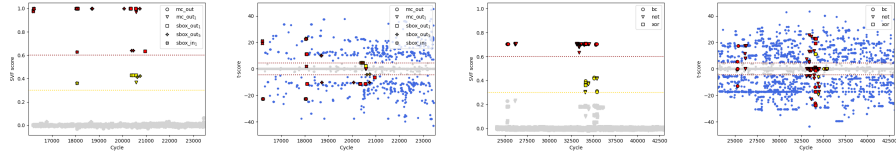


Fig. 2. SVF AES, combined SVF-TVLA AES, SVF SHA-3, combined SVF-TVLA SHA-3. For ALU module on Ibex core (left to right).

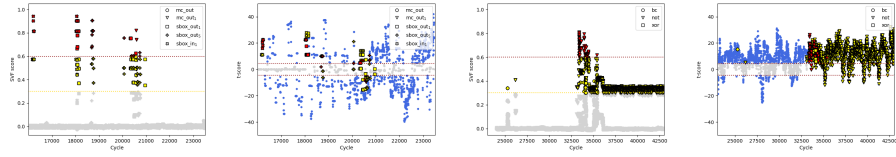


Fig. 3. SVF AES, combined SVF-TVLA AES, SVF SHA-3, combined SVF-TVLA SHA-3. For Register File on Ibex core(left to right).

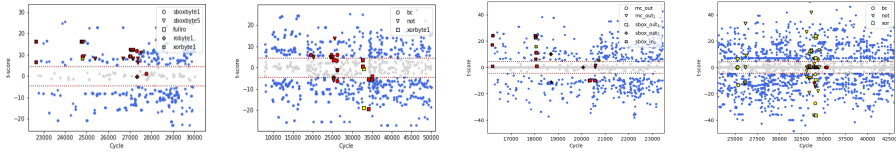


Fig. 4. SVF-TVLA for: AES MBOX(SHAKTI), SHA-3 MBOX(SHAKTI), AES Mult/Div(Ibex) and SHA-3 Mult/Div(Ibex) (left to right).

Table 1. Results for AES and SHA-3 on SHAKTI core

Type	AES					TVLA	SHA-3		
	max(SVF)						max(SVF)		
	sbbox_in1	sbbox_out1	sbbox_out5	mc_out1	mc_out		bc	not	xor
Dcache	1	1	1	1	0.05	✓	0.99	1	1
RF	1	1	1	1	0.05	✓	1	1	1
CSR	0.98	0.98	0.98	0.98	0.05	✓	0.97	0.97	0.97
ALU	1	1	1	1	0.05	✓	1	1	1
FPU	1	1	1	1	0.05	✓	1	1	1
MBox	1	1	1	1	0.05	✓	1	1	1
BPU	0	0	0	0	0	-	0	0	0
ITLB	0	0	0	0	0	-	0	0	0
DTLB	0.87	0.87	0.87	0.87	0.06	✓	0.3	0.3	0.31

5.3 Performance results

All experiments, for both the SHAKTI and the Ibex core, were done on an AMD Ryzen THREADRIPPER 3990X 4.3GHz CPU with 128 cores and 256GB

Table 2. Results for AES and SHA-3 on Ibex core

Type	AES					TVLA	SHA-3			TVLA
	max(SVF)						max(SVF)			
	sbox_in ₁	sbox_out ₁	sbox_out ₅	mc_out ₁	mc_out		bc	not	xor	
RF	0.94	0.9	0.9	0.82	0.05	✓	0.59	0.81	0.58	✓
CSR	0.97	0.97	0.97	0.95	0.05	✓	0.69	0.70	0.70	✓
ALU	1	0.99	1	0.99	0.06	✓	0.70	0.70	0.70	✓
MULT/DIV	0.99	0.99	1	0.84	0.04	✓	0.63	0.63	0.64	✓
PF-BUF	0	0	0	0	0	-	0	0	0	-
LSU	1	0.99	1	0.98	0.06	✓	0.96	0.96	0.96	✓

RAM. On this PC, the run-time of the two methods is significantly different. For the SHAKTI core, the SVF computation of all the modules for the AES case study took around 30k cycles, which lasted approximately eight hours. The same experiment for the SHA-3 case study took around one day of computation. On the other hand, TVLA computation lasted about an hour for all modules for all case studies. For the Ibex core, the SVF computation of the algorithms took about twice the time compared to the SHAKTI core. The TVLA computation lasted less than an hour for all the modules of AES and SHA3. We can easily observe the contrast in run-time efficiency between the SVF and TVLA, which is crucial when choosing the appropriate method.

6 Conclusions and Future work

In this study, we compared the performance of two different leakage detection methods, SVF and TVLA, in detecting leaky modules on microprocessors. We have conducted experiments on two RISC-V cores, SHAKTI and Ibex, using the cryptographic algorithms AES and SHA3. Our investigation has revealed interesting differences in how these methods identify potential leaks. Specifically, we observe peaks in cycles where SVF does not find any correlation between the oracle trace and the side-channel trace. In contrast, TVLA identifies a difference in the probability density functions of our fixed set and random set. This distinction can be explained by the fact that SVF targets an intermediate computation when determining leaks. Consequently, leakage points of any other intermediate data value are not included in a single SVF test. TVLA, on the other hand, adopts a much broader approach, evaluating the overall difference in probability density functions between the fixed and random sets, which allows it to detect potential leaks that SVF might miss.

Our results, presented in Table 1 and Table 2 for SHAKTI and Ibex cores respectively, for all examined modules, reveal no leaks in the BPU and ITLB modules on SHAKTI or the Prefetch buffer on Ibex, with both algorithms. This finding shows the importance of selecting the appropriate method based on the desired level of analysis granularity. For more coarse-grained analysis where a broad overview is beneficial, TVLA is recommended due to its lower compu-

tational effort and quicker results. On the other hand, SVF is better suited for fine-grained analysis that requires examining micro-architectural behavior in detail.

Moreover, our analysis highlights that TVLA and SVF show leaks in different cycles, opening questions about which metric more indicates the “correct” leak. This aspect, alongside the critical role of selecting the appropriate oracles, emphasizes the complexity of accurately detecting leakage. Not all oracles show leakage, potentially leading to incorrect conclusions about a module’s security. Also, SVF uses the Pearson correlation coefficient to detect linear correlations. This means it might not find any non-linear relationships between the oracle trace and the side-channel trace. The generality of TVLA allows for broader applications, while the specificity of SVF, although powerful, necessitates careful oracle selection to ensure comprehensive leakage detection.

Moving forward, we plan to investigate further the differences between the SVF and TVLA methods, aiming to refine the process of hardening microprocessors against side-channel attacks. We also aim to explore other leakage detection methods and the impact of different implementations of cryptographic algorithms on the leakage profile of microprocessors. Ultimately, our goal is to contribute to developing more secure microprocessors that are robust against side-channel attacks.

Acknowledgments. This work was partially funded by the Dutch Research Council (NWO) through the PROACT project (NWA.1215.18.014).

References

1. Althoff, A., McMahan, J., Vega, L., Davidson, S., Sherwood, T., Taylor, M.B., Kastner, R.: Hiding intermittent information leakage with architectural support for blinking. In: 45th ACM/IEEE Annual International Symposium on Computer Architecture, ISCA, Los Angeles, CA, USA. pp. 638–649 (2018). <https://doi.org/10.1109/ISCA.2018.00059>
2. Arora, V., Buhan, I., Perin, G., Picek, S.: 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. pp. 80–96. LNCS, Springer (2021). https://doi.org/10.1007/978-3-030-97348-3_5
3. Bloem, R., Gigerl, B., Gourjon, M., Hadzic, V., Mangard, S., Primas, R.: Power contracts: Provably complete power leakage models for processors. In: Conference on Computer and Communications Security, CCS, Los Angeles, CA, USA. pp. 381–395. ACM (2022). <https://doi.org/10.1145/3548606.3560600>
4. Buhan, I., Batina, L., Yarom, Y., Schaumont, P.: Sok: Design tools for side-channel-aware implementations. In: ASIA CCS ’22: ACM Asia Conference on Computer and Communications Security, Nagasaki, Japan. pp. 756–770. ACM (2022). <https://doi.org/10.1145/3488932.3517415>
5. Demme, J., Martin, R., Waksman, A., Sethumadhavan, S.: Side-channel vulnerability factor: A metric for measuring information leakage. In: 39th International Symposium on Computer Architecture (ISCA. pp. 106–117 (2012). <https://doi.org/10.1109/ISCA.2012.6237010>

6. F, M.A.K., Ganesan, V., Bodduna, R., Rebeiro, C.: PARAM: A microprocessor hardened for power side-channel attack resistance. In: International Symposium on Hardware Oriented Security and Trust, HOST, San Jose, CA, USA. pp. 23–34. IEEE (2020). <https://doi.org/10.1109/HOST45689.2020.9300263>
7. Gao, S., Großschädl, J., Marshall, B., Page, D., Pham, T.H., Regazzoni, F.: An instruction set extension to support software-based masking. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* pp. 283–325 (2021). <https://doi.org/10.46586/tches.v2021.i4.283-325>
8. Gigerl, B., Hadzic, V., Primas, R., Mangard, S., Bloem, R.: Coco: Co-design and co-verification of masked software implementations on cpus. In: 30th USENIX Security Symposium, USENIX. pp. 1469–1468 (2021)
9. Gilbert Goodwill, B.J., Jaffe, J., Rohatgi, P., et al.: A testing methodology for side-channel resistance validation. In: NIST non-invasive attack testing workshop. vol. 7, pp. 115–136 (2011)
10. He, M.T., Park, J., Nahiyan, A., Vassilev, A., Jin, Y., Tehranipoor, M.M.: RTL-PSC: automated power side-channel leakage assessment at register-transfer level. In: 37th IEEE VLSI Test Symposium, VTS Monterey, CA, USA. pp. 1–6 (2019). <https://doi.org/10.1109/VTS.2019.8758600>
11. Kiaei, P., Schaumont, P.: Soc root canal! root cause analysis of power side-channel leakage in system-on-chip designs. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* pp. 751–773 (2022). <https://doi.org/10.46586/tches.v2022.i4.751-773>
12. Lakshmy, A.V., Rebeiro, C., Bhunia, S.: FORTIFY: analytical pre-silicon side-channel characterization of digital designs. In: 27th Asia and South Pacific Design Automation Conference, ASP-DAC. pp. 660–665. IEEE (2022). <https://doi.org/10.1109/ASP-DAC52403.2022.9712551>
13. lowRISC: Lowrisc/ibex-demo-system: A demo system for ibex, <https://github.com/lowRISC/ibex-demo-system>
14. Marshall, B., Page, D., Webb, J.: MIRACLE: micro-architectural leakage evaluation A study of micro-architectural power leakage across many devices. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* pp. 175–220 (2022). <https://doi.org/10.46586/TCHES.V2022.I1.175-220>
15. Mulder, E.D., Gummalla, S., Hutter, M.: Protecting RISC-V against side-channel attacks. In: Proceedings of the 56th Annual Design Automation Conference 2019, DAC, Las Vegas, NV, USA. p. 45 (2019). <https://doi.org/10.1145/3316781.3323485>
16. Pham, T.H., Marshall, B., Fell, A., Lam, S., Page, D.: XDIVINSA: extended diversifying instruction agent to mitigate power side-channel leakage. In: 32nd IEEE International Conference on Application-specific Systems, Architectures and Processors, ASAP. pp. 179–186 (2021). <https://doi.org/10.1109/ASAP52443.2021.00034>
17. SHAKTI: Family of processors. <https://shakti.org.in/processors.html> (2021)
18. SLPSK, P., Vairam, P.K., Rebeiro, C., Kamakoti, V.: Karna: A gate-sizing based security aware EDA flow for improved power side-channel attack protection. In: Proceedings of the International Conference on Computer-Aided Design. pp. 1–8 (2019). <https://doi.org/10.1109/ICCAD45719.2019.8942173>
19. Yao, Y., Kathuria, T., Ege, B., Schaumont, P.: Architecture correlation analysis (ACA): identifying the source of side-channel leakage at gate-level. In: International Symposium on Hardware Oriented Security and Trust, HOST, San Jose, CA, USA. pp. 188–196. IEEE (2020). <https://doi.org/10.1109/HOST45689.2020.9300271>
20. Zhang, T., Liu, F., Chen, S., Lee, R.B.: Side channel vulnerability metrics: the promise and the pitfalls. In: HASP 2013, The Second Workshop on Hardware and Architectural Support for Security and Privacy, Tel-Aviv, Israel. p. 2. ACM (2013). <https://doi.org/10.1145/2487726.2487728>