PROACT - Physical Attack Resistance of Cryptographic Algorithms and Circuits with Reduced Time to Market

Asmita Adhikary[1], Abraham Basurto[1], Lejla Batina[1], Ileana Buhan[1], Joan Daemen[1], Silvia Mella[1], Nele Mentens[2(✉)], Stjepan Picek[1], Durga Lakshmi Ramachandran[3], Abolfazl Sajadi[2], Todor Stefanov[2], Dennis Vermoen[3], and Nusa Zidaric[2]

[1] Radboud University, Nijmegen, The Netherlands
[2] Leiden University, Leiden, The Netherlands
n.mentens@liacs.leidenuniv.nl
[3] Riscure, Delft, The Netherlands

Abstract. Electronic devices that populate the Internet of Things play increasingly important roles in our everyday lives. When these devices process, store, or communicate personal or company-critical data, digital security becomes a necessity. However, mechanisms to secure electronic systems have a significant influence on the cost of the system and come with an overhead in energy consumption, computational delay, and (silicon) chip area. Therefore, developing secure electronic systems is a balancing act between minimizing the overhead and maximizing the security. Moreover, in rapidly evolving markets, there is another parameter that can have a negative influence on the security strength of electronic devices, namely the time to market: it takes longer to bring a secure product to the market than to develop a product with no or little security measures in place.

In the PROACT project, we tackle the challenge of maximizing the security strength while minimizing the overhead w.r.t. energy consumption, computational delay, and hardware resources, as well as reducing the time to market of digital electronic systems. We specifically focus on the fast development of efficient cryptographic hardware with protection against physical attacks, i.e., attacks that exploit the physical implementation of cryptographic algorithms. Physical attacks are categorized into (1) side-channel analysis attacks that target the extraction of secret information by monitoring side-channels like the power consumption, the electromagnetic emanation or the timing of the device, and (2) fault analysis attacks that aim at introducing computational errors that lead to the leakage of secret information. Physical security is of vital importance when potential attackers can easily get in the vicinity of an electronic system. This is the case in, e.g., medical sensor devices, wearables and implants, which are typically constrained in energy budget, cost and form factor, and are therefore the perfect use case for the results of PROACT.

Keywords: side-channel analysis · fault analysis · cryptography · hardware security

I. Skliarova et al. (Eds.): ARC 2024, LNCS 14553, pp. 255–266, 2024.
https://doi.org/10.1007/978-3-031-55673-9_18

1 Introduction and Envisioned Contributions

As digital data are omnipresent in our daily lives, the need for digital security is growing rapidly. This is illustrated by popular media frequently reporting on attacks that expose the security flaws of real-life electronic systems. A very powerful type of attack is a physical attack, which exploits the physical implementation of a cryptographic algorithm, as shown in Fig. 1. The first category of physical attacks is side-channel analysis attacks [1], which analyze the information available through side channels, such as the power consumption, the electromagnetic (EM) emanation, or the timing behavior of an electronic system. Another type of physical attack is a fault analysis attack [2], which perturbs the system, e.g., through the injection of a laser beam, a clock glitch, or a power supply glitch, in order to retrieve secret information. Especially for Internet-of-Things (IoT) devices, physical attacks form an underestimated threat and must be dealt with through proper countermeasures.

To achieve the highest level of physical security, protection mechanisms must be foreseen throughout all steps in the knowledge value chain: in the design of cryptographic algorithms, the design of cryptographic circuits, and the physical implementation of cryptographic chips. Additionally, design choices made in one of these steps introduce constraints in other steps, such that interaction between the steps in the chain is indispensable. PROACT covers the entire knowledge value chain in the development of physically secure cryptographic hardware, from algorithms to fabricated chips.

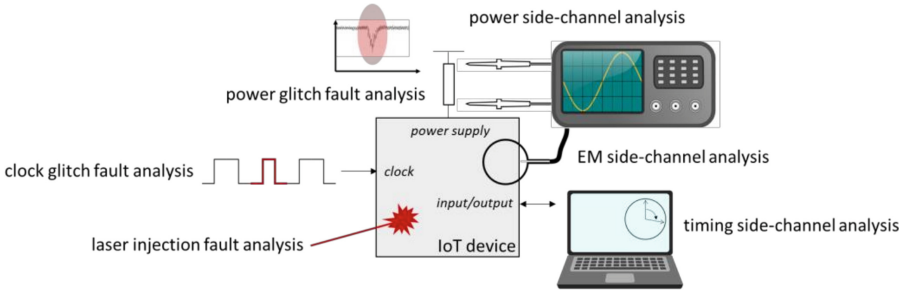

Fig. 1. Examples of physical attacks (i.e., fault analysis and side-channel analysis) on an IoT device.

Since many IoT devices are limited in energy/power consumption and computing resources, implementing countermeasures is challenging. Manual efforts by experienced designers can be effective but are prone to errors and do not lead to optimal results when the design space is large. Existing design automation tools can optimize towards low energy/power and low resources but do not take into account physical security. **PROACT develops design automation tools with low energy/power consumption, low computational resources, and high physical security strength as optimization goals.**

The IoT market consists of rapidly evolving applications. Therefore, minimizing the time to market of new products and services is crucial to survive for companies that operate in this market segment. However, evaluating the physical security of a cryptographic chip is typically something that is done after the (silicon) chip is fabricated. This makes the elapsed time between the design of a cryptographic algorithm/circuit and the physical security evaluation very long. A weakness detected in the evaluation phase leads to a re-spin of the chip, taking away the competitive advantage of the company that intends to be the first to bring a new IoT application to the market. **PROACT designs, implements, and validates a pre-silicon simulator for physical security to maximize the chances of first-time-right cryptographic chips.**

The aforementioned design automation tools and pre-silicon simulator facilitate our proposed modifications to the traditional design flow for cryptographic hardware as shown in Fig. 2. The physical security simulation becomes a new crucial step in the flow, just before the chip fabrication, because an early feedback on the security strengths of the chip could trigger important physical security related improvements on the cryptographic algorithm/hardware design, thereby significantly reducing the time to market of the chip.

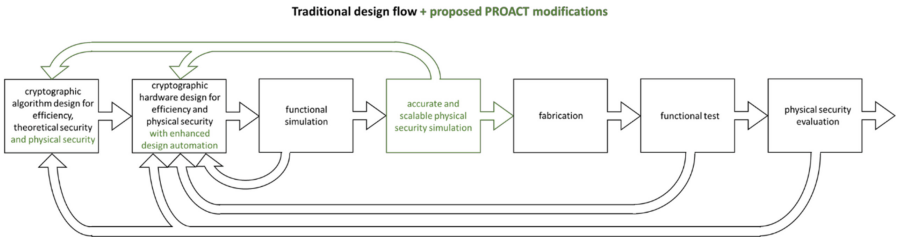

Fig. 2. Proposed modifications to the traditional design flow.

In PROACT we aim to answer the following research questions:

- What are the problems with respect to physical security in existing cryptographic algorithms, and how can we design algorithms that are resilient against physical attacks?
- Which cryptographic circuits have optimal physical security strength, energy consumption, resource occupation, and performance, or an optimal trade-off of these properties?
- How can we use design automation to improve the efficiency and the physical security of cryptographic circuits?
- How can we design a pre-silicon physical security simulator with optimal accuracy and simulation speed?
- How can we use artificial intelligence to improve the accuracy and simulation speed of the pre-silicon simulator?

– Which state-of-the-art and beyond-state-of-the-art analysis methods can be used to perform a systematic evaluation of the developed cryptographic chips and validation of the pre-silicon simulator?

2 Project Goals and Status of the Conducted Research

The research and development work in the PROACT project is divided across four work packages (WP1, WP2, WP3, and WP4), each having their own specific goals that jointly contribute to the main project objectives introduced in Sect. 1.

2.1 WP1: Algorithm Design

The goal of the first work package is to have a process allowing the feedback from the other work packages on the design of cryptographic algorithms and their implementation, in the form of requirements, restrictions and recommendations. Requirements we are thinking of are the following:

– Low latency: applications such as memory or pointer encryption [16] for protection against micro-architectural attacks require low latency: a very short time between the availability of specific inputs (plaintext, ciphertext, memory address, ...) and the output (plaintext or ciphertext). Moreover, the latency of a circuit is correlated to its energy consumption [2,10], so low latency is also a recommendation for lightweight use cases that run on battery power.
– Suitability for masking and threshold implementation: if side channel attacks such as power or electromagnetic analysis are a threat and countermeasures against these attacks at mode level are not an option or undesirable, masking or threshold implementations may be a cost-effective countermeasure. This requirement usually boils down to limiting non-linear operations in the round function to have low algebraic degree.
– Suitability for protection against sophisticated fault attacks like Statistical Ineffective Fault Attacks (SIFA) [13]. One very promising countermeasure is to implement the cipher using the combination of *toffoli* gates and masking as laid out in [8]. Clearly the cipher shall be suited to be implemented efficiently with those building blocks.
– Suitability for instruction/operation shuffling and having equivalent representations. When protecting against power or electromagnetic analysis but also faults, the ability to execute the operations in a round function in different orders, called *shuffling* provides an additional layer of protection on top of masking or threshold implementations. The suitability of a round function for shuffling is strongly determined by the amount of parallelism in the cipher. For example in AES [11] the 16 S-box computations can be done in any order. Equivalent representations are the result of symmetry and can be useful against fault attacks in the presence of redundant computations. The attack vector is then to force the same fault twice and the existence of equivalent representations allows randomizing the *location* where the computation takes place [17].

These are just a few examples and we expect a myriad of requirements, restrictions and recommendations to originate from the other work packages or more in general of the design flow.

2.2 WP2: Circuit Design and Design Automation

We are designing a circuit, nicknamed the PROACT chip, to obtain real measurements (e.g., circuit power consumption, timing performance, etc.) needed for the design of the pre-silicon physical security simulator, and for evaluation and validation of the simulator. Currently, we are designing and selecting benchmark circuits for the PROACT chip. The PROACT chip will allow side-channel analysis of cryptographic software executed on a general-purpose CPU as well as side-channel analysis of custom cryptographic hardware, i.e., cryptographic co-processors. We are also considering protected software and hardware implementations. This approach ensures a holistic evaluation of the security of a system on a chip.

The first version of the PROACT chip, shown in Fig. 3, is currently prototyped on a Field-Programmable Gate Array (FPGA) and will later be fabricated as a dedicated ASIC chip. This design leverages a RISC-V core (Ctrl-RV) and two 32-bit registers (Control Reg and Status Reg) to control the system. The main general-purpose CPU component is a second RISC-V core (SW-RV in Fig. 3) dedicated to executing cryptographic software. We selected the Ibex core for this project, an open-source 32-bit RISC-V CPU written in System Verilog [12]. Furthermore, the PROACT chip will include cryptographic co-processors, such as Ascon [7], Xoodyak [9], and AES [1].

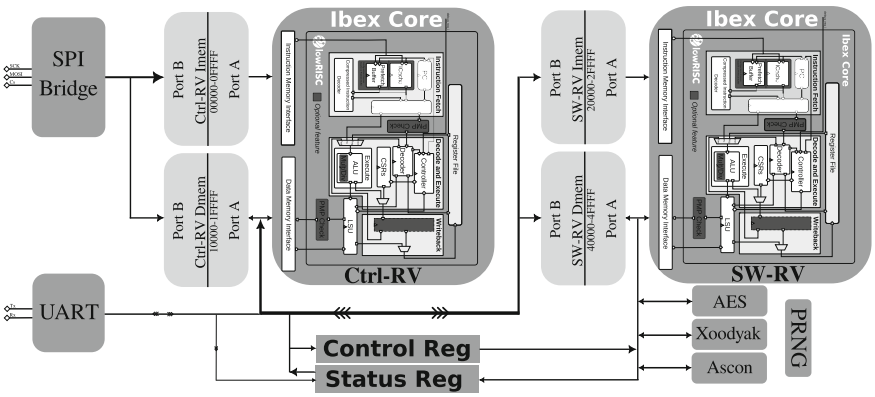

Fig. 3. The PROACT chip block diagram.

The controller (Ctrl-RV) is programmed using the Serial Peripheral Interface protocol (SPI Bridge in Fig. 3). Once initialized, this controller establishes a UART-based communication channel with a host computer for processing commands and controls. For instance, when loading the SW-RV instruction and data memories, the configuration commands need to specify the size and desired start addresses for the SW-RV instruction and data memory segments. When the SW-RV memories are loaded, the Ctrl-RV sets a flag in the Control Reg to activate SW-RV and waits for a specified flag in the Status Reg, indicating that SW-RV completed the execution. Similar control mechanism is implemented for the co-processors. The PROACT chip is designed as a cooperative target and sets appropriate triggers for precise power (or other) trace acquisition.

In parallel to the PROACT chip development, we are evaluating the accuracy and usefulness of power consumption estimation data, obtained from Cadence synthesis tools at different levels of abstraction, for the purpose of physical security evaluation of VLSI circuits. This includes selecting appropriate scope, e.g., circuit submodules and (interface) signals, and signals coverage of the circuit testbenches.

2.3 WP3: Pre-silicon Physical Security Simulator Design

Side-channel attacks that leak sensitive information through a computing device's interaction with its physical environment have proven to severely threaten the security of a device when adversaries have physical access to the device. Traditional approaches for leakage detection measure the physical properties of the device. Hence, they cannot be used during the design process and fail to provide root cause analysis. An alternative approach gaining traction is automating leakage detection by modeling the device. This work package aims to develop a simulator that takes a digital circuit design after synthesis, placement, and routing as input. The goal is to perform simulations of the physical security of the design, i.e., of the side-channel leakage and the effect of a fault injection. Artificial intelligence will enable the fast and accurate simulation of large circuit designs. The input for the training phase will consist of fabricated benchmark circuits and accompanying measurement data for physical security. These benchmark circuits will be selected based on a thorough analysis of the typical components of cryptographic circuits. The PROACT chip, discussed in Sect. 2.2, includes such benchmark circuits that will be measured. The goal is to determine which circuits are the most useful for training the simulator. At the end of the PROACT project, this will lead to guidelines for technology companies to generate the benchmark circuits and make the measurement data available. Following this approach, the technology can be characterized for physical security, and the simulator can be trained for this specific technology.

As a first step in the project, we investigated the state of the art for designing side channel simulators. In this study [6], we classify approaches to automated leakage detection based on the model's source of truth. We organize the existing tools on two main parameters: whether the model includes measurements from a concrete device and the abstraction level of the device specification used for

constructing the model. We survey the proposed tools to determine the current knowledge level across the domain and identify open problems. In particular, we highlight the absence of evaluation methodologies and metrics that compare proposals' effectiveness across the field. Our results help practitioners who want to use automated leakage detection and researchers interested in advancing the knowledge and improving automatic leakage detection. In this study, we made an inventory of available tools[1], and we determined the methods for automating leakage detection and summarized open problems. One of the conclusions is that no side channel simulator is available for RISC-V architectures like the one we made in this project. One of our main findings shows that creating a side channel simulator is primarily based on manual work and prohibitively effort-intensive. Capturing microarchitecture events characteristic to complex processors, such as pipeline stalls or misprediction, is an open problem. We found no evidence that AI techniques have been used to automate the manual work required to build a side-channel simulator.

One challenge for using AI techniques is the availability of labeled datasets. The conventional side-channel analysis demands substantial manual effort for setup preparation and trace recording, rendering it more intricate during the dynamic design phase, where software alterations occur frequently. Additionally, limited hardware descriptions and restricted access to process technology information have hindered identifying the specific instruction(s) responsible for leakage. We introduce ABBY [5], an open-source side-channel leakage profiling framework that targets the microarchitectural layer. Existing solutions to characterize the microarchitectural layer are device-specific and require extensive manual effort. ABBY's main innovation is data collection, which can then automatically describe the microarchitectural behavior of a target device and has the additional benefit of being extendable to similar devices. Using ABBY, we created two datasets that capture the interaction of instructions for the ARM CORTEX-M0/M3 architecture. These sets are the first to capture detailed information on the microarchitectural layer. They can be used to explore various leakage models suitable for creating side-channel leakage simulators. These attributes encompassed instruction interactions, operand interactions, pipeline effects, and memory transaction interactions. We delved into linear and nonlinear (predominantly deep learning-based) leakage models with the datasets acquired. The effectiveness of these leakage models was subsequently evaluated and compared using evaluative metrics such as the adjusted R^2, F-test, and actual side-channel attack outcomes. A preliminary evaluation of a leakage model produced with our dataset of real-world cryptographic implementations shows performance comparable to state-of-the-art leakage simulators. To showcase the effectiveness of the ABBY framework and assess the quality of the dataset it produces, we constructed a leakage model based on the ABBY-CM0 dataset. Our comparisons between this model and ELMO yielded strikingly similar results, underscoring the high caliber of the ABBY-CM0 dataset. To examine ABBY's scalability, we curated the ABBY-CM3 dataset. Concurrently, we designed a side-channel power

[1] https://ileanabuhan.github.io/Tools/.

simulator targeting the ALU component, drawing insights from the ABBY-CM3 dataset. Compared with the actual board, the simulator's performance was on par, further attesting to the impeccable quality of the ABBY-CM3 dataset.

We performed a similar study considering tools for protecting implementations against fault injection attacks [3]. Fault injection attacks have caused implementations to behave unexpectedly, resulting in a spectacular bypass of security features and even the extraction of cryptographic keys. Developers want to ensure the robustness of the software against faults and eliminate production weaknesses that could lead to exploitation. Several fault simulators that promise cost-effective evaluations against fault attacks have been released. In [3], we set out to discover how suitable such tools are for a developer who wishes to create robust software against fault attacks. We found four open-source fault simulators that employ different techniques to navigate faults, which we objectively compare and discuss their benefits and drawbacks. Unfortunately, none of the four open-source fault simulators use artificial intelligence (AI) techniques. However, AI was successfully applied to improve the fault simulation of cryptographic algorithms, though none of these tools is open source. We suggest improvements to open-source fault simulators inspired by the AI techniques used by cryptographic fault simulators.

2.4 WP4: Evaluation and Validation

Millions of products undergo rigorous security evaluations every day in evaluation laboratories around the world [20]. Testing for side-channel resistance is a key aspect of security evaluations for implementations featuring cryptography. The effort involved in testing is considerable, and the stakes for companies are high [4]. Moreover, cryptographic implementations often go through multiple cycles of leakage evaluation, e.g., as specified in ISO/IEC 17825:2016. Such a process is costly because it requires a high level of expertise and significant manual labor, especially when considering resourceful adversaries [23].

In recent years, developments in deep learning-based side-channel analysis (DLSCA) have made it one of the obvious choices when evaluating/validating the security of devices. While this trend is mostly academic for now, we expect the industry will soon follow with developments of various standards; for instance, there is a new standard draft for minimal requirements for evaluating machine learning-based SCA resistance.[2] In the process of the DLSCA evaluation, there are a number of questions to consider. What AI technique to use? How to tune it? Do we need trace pre-processing or feature engineering? What are the appropriate metrics to evaluate the security? What threat model to assume? How many side-channel measurements are necessary? How does a neural network defeat an implementation with a countermeasure?

Hyperparameter tuning represents one of the central points to achieving powerful deep learning performance, and SCA is no exception, with several directions to follow:

[2] Minimum Requirements for Evaluating Machine-Learning based Side-Channel Att ack Resistance.

1. Random/grid search. While random/grid search is easy to mount and can give excellent results (like in the current state-of-the-art), one still needs to define appropriate hyperparameter ranges that should be sufficiently small. Additionally, one commonly needs to evaluate many random models to improve the chances of obtaining good models.
2. Advanced tuning techniques. DLSCA investigated techniques like Bayesian optimization [26] and reinforcement learning [21], which exhibit excellent attack performance. Still, such techniques can be computationally expensive and have additional parameters to tune (shifting the problem from tuning the neural network hyperparameters to tuning the search technique parameters).
3. Methodologies. Methodologies can provide a systematic way to build neural networks that perform well in DLSCA. Unfortunately, it is difficult to design a methodology that is easy to follow and works for diverse targets/leakage models/neural network architectures.

It is also necessary to consider what features of side-channel traces will be inputted into neural networks. The first works on machine learning (and template attack) required a precise selection of features, making the effort in the feature engineering phase often much more significant than running the attack itself. Moving to the deep learning techniques brought promise that we require less feature engineering, allowing more time for hyperparameter tuning. As such, the common approach was to consider an interval of features that leaks the most [20]. Still, recent results showed that there is a further benefit when providing raw traces to neural networks, as it is possible to mount even optimal attacks (those that require only a single attack trace), but at the cost of more effort in hyperparameter tuning [18]. Thus, we reached a trade-off between the effort in feature engineering and hyperparameter tuning. Interestingly, the latest works showcased that it is possible to reach optimal attacks even if we provide "only" an interval of features, but then, the neural network architectures must be more complex, even using language models [15]. That being said, using all features is not possible for any profiling attack, as already discussed. Thus, it was shown possible to make a more powerful feature engineering phase based on a novel distance metric customized for SCAs, allowing the template attack to compete or even outperform state-of-the-art DLSCA [25]. Extending this concept further, it is possible to design custom loss functions for DLSCA that consider the most relevant features (e.g., features processed by the deep layers of neural networks) [30]. Finally, an important part of making the attacks more powerful is also understanding why the attacks work, as such knowledge can improve not only the attack perspective but also the future design of countermeasures [19].

Running supervised deep learning-based SCA is not ideal for any attack scenario. Often, it becomes necessary to relax the assumptions on the attacker power and not assume anymore that there is a clone device or that the leakage model is known. One example in that direction is when the adversary possesses a similar implementation that can be used as a white-box reference design [14]. Moreover, recent works show it is possible to move toward non-profiled DLSCA by using the bijective relationship between the plaintexts and a fixed key [27]. Then, by

following this, it is possible to mount attacks that even rival profiling DLSCA. Still, this approach can be considered non-profiled but not unsupervised because we still build labels. Luckily, it is possible to move toward non-profiled DLSCA by using, e.g., the multi-regression output approach [22]. Such an approach can be further improved by using techniques like ensembles and data augmentation. On the other hand, it is possible to consider different SCA paradigms, like collision attacks in the DLSCA setting [28]. In settings where we cannot assume the knowledge of leakage models, it is possible to run the attacks in model-free settings [24]. Finally, for settings where the evaluator has only a limited number of profiling traces, it is possible to make such attacks more powerful by examining the relationship between all possible key candidates, which leads to a novel metric describing the generalization power of a profiling model [29].

3 Conclusion and Next Steps

PROACT aims at adding the physical security dimension to the design flow of ASICs, which typically only focuses on the optimization of energy consumption, computational delay and hardware resources. The achieved results so far include (1) the design of a system on chip that is prototyped on an FPGA and almost ready for tape-out, (2) a suitability analysis and comparison of existing physical security simulators, and (3) the improvement of physical attack strategies. Two tape-outs are planned within the project in order to validate the envisioned physical security simulator that will be built within the project.

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

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. FIPS 197: Advanced Encryption Standard (AES) (2021). https://doi.org/10.6028/NIST.FIPS.197-upd1
2. Aagaard, M.D., Zidaric, N.: ASIC benchmarking of round 2 candidates in the NIST lightweight cryptography standardization process: (preliminary results). IACR Cryptol. ePrint Arch., p. 49 (2021). https://eprint.iacr.org/2021/049
3. Adhikary, A., Buhan, I.: SoK: assisted fault simulation - existing challenges and opportunities offered by AI. In: Zhou, J., et al. (eds.) ACNS 2023. LNCS, vol. 13907, pp. 178–195. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-41181-6_10
4. Azouaoui, M., et al.: A systematic appraisal of side channel evaluation strategies. In: van der Merwe, T., Mitchell, C., Mehrnezhad, M. (eds.) SSR 2020. LNCS, vol. 12529, pp. 46–66. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64357-7_3

5. Bazangani, O., Iooss, A., Buhan, I., Batina, L.: ABBY: automating leakage modeling for side-channels analysis. In: Proceedings of the 2024 ACM Asia Conference on Computer and Communications Security (2024, to appear)
6. Buhan, I., Batina, L., Yarom, Y., Schaumont, P.: SoK: design tools for side-channel-aware implementations. In: Suga, Y., Sakurai, K., Ding, X., Sako, K. (eds.) ASIA CCS 2022: ACM Asia Conference on Computer and Communications Security, Nagasaki, Japan, 30 May 2022–3 June 2022, pp. 756–770. ACM (2022). https://doi.org/10.1145/3488932.3517415
7. Christoph Dobraunig, I.: Ascon-a submission to CAESAR. In: 15th Central European Conference on Cryptology, p. 23 (2015)
8. Daemen, J., Dobraunig, C., Eichlseder, M., Groß, H., Mendel, F., Primas, R.: Protecting against statistical ineffective fault attacks. IACR Trans. Cryptogr. Hardw. Embed. Syst. **2020**(3), 508–543 (2020). https://doi.org/10.13154/TCHES.V2020.I3.508-543
9. Daemen, J., Hoffert, S., Peeters, M., Assche, G.V., Keer, R.V.: Xoodyak, a lightweight cryptographic scheme (2020)
10. Daemen, J., Massolino, P.M.C., Mehrdad, A., Rotella, Y.: The subterranean 2.0 cipher suite. IACR Trans. Symmetric Cryptol. **2020**(S1), 262–294 (2020). https://doi.org/10.13154/TOSC.V2020.IS1.262-294
11. Daemen, J., Rijmen, V.: The Design of Rijndael - The Advanced Encryption Standard (AES). Information Security and Cryptography, 2nd edn. Springer, Heidelberg (2020). https://doi.org/10.1007/978-3-662-60769-5
12. Davide Schiavone, P., et al.: Slow and steady wins the race? A comparison of ultra-low-power RISC-V cores for internet-of-things applications. In: 2017 27th International Symposium on Power and Timing Modeling, Optimization and Simulation (PATMOS), pp. 1–8 (2017). https://doi.org/10.1109/PATMOS.2017.8106976
13. Dobraunig, C., Eichlseder, M., Korak, T., Mangard, S., Mendel, F., Primas, R.: SIFA: exploiting ineffective fault inductions on symmetric cryptography. IACR Trans. Cryptogr. Hardw. Embed. Syst. **2018**(3), 547–572 (2018)
14. Karayalcin, S., Krcek, M., Wu, L., Picek, S., Perin, G.: It's a kind of magic: a novel conditional GAN framework for efficient profiling side-channel analysis. Cryptology ePrint Archive, Paper 2023/1108 (2023). https://eprint.iacr.org/2023/1108
15. Kulkarni, P., Verneuil, V., Picek, S., Batina, L.: Order vs. chaos: a language model approach for side-channel attacks. Cryptology ePrint Archive, Paper 2023/1615 (2023). https://eprint.iacr.org/2023/1615
16. LeMay, M., et al.: Cryptographic capability computing. In: MICRO 2021, pp. 253–267. ACM (2021)
17. Miteloudi, K., Batina, L., Daemen, J., Mentens, N.: ROCKY: rotation countermeasure for the protection of keys and other sensitive data. In: Orailoglu, A., Jung, M., Reichenbach, M. (eds.) SAMOS 2021. LNCS, vol. 13227, pp. 288–299. Springer, Cham (2021). https://doi.org/10.1007/978-3-031-04580-6_19
18. Perin, G., Wu, L., Picek, S.: Exploring feature selection scenarios for deep learning-based side-channel analysis. IACR Trans. Cryptogr. Hardw. Embed. Syst. **2022**(4), 828–861 (2022). https://doi.org/10.46586/tches.v2022.i4.828-861. https://tches.iacr.org/index.php/TCHES/article/view/9842
19. Perin, G., Wu, L., Picek, S.: I know what your layers did: layer-wise explainability of deep learning side-channel analysis. Cryptology ePrint Archive, Paper 2022/1087 (2022). https://eprint.iacr.org/2022/1087
20. Picek, S., Perin, G., Mariot, L., Wu, L., Batina, L.: SoK: deep learning-based physical side-channel analysis. ACM Comput. Surv. **55**(11), 1–35 (2023). https://doi.org/10.1145/3569577

21. Rijsdijk, J., Wu, L., Perin, G., Picek, S.: Reinforcement learning for hyperparameter tuning in deep learning-based side-channel analysis. IACR Trans. Cryptogr. Hardw. Embed. Syst. **2021**(3), 677–707 (2021). https://doi.org/10.46586/tches.v2021.i3.677-707. https://tches.iacr.org/index.php/TCHES/article/view/8989
22. Savu, I., Krček, M., Perin, G., Wu, L., Picek, S.: The need for more: unsupervised side-channel analysis with single network training and multi-output regression. Cryptology ePrint Archive, Paper 2023/1681 (2023). https://eprint.iacr.org/2023/1681
23. Shelton, M.A., Samwel, N., Batina, L., Regazzoni, F., Wagner, M., Yarom, Y.: Rosita: towards automatic elimination of power-analysis leakage in ciphers. In: Proceedings 2021 Network and Distributed System Security Symposium. Internet Society, Virtual (2021). https://doi.org/10.14722/ndss.2021.23137. https://www.ndss-symposium.org/wp-content/uploads/ndss2021_4B-3_23137_paper.pdf
24. Wu, L., Ali-pour, A., Rezaeezade, A., Perin, G., Picek, S.: Breaking free: leakage model-free deep learning-based side-channel analysis. Cryptology ePrint Archive, Paper 2023/1110 (2023). https://eprint.iacr.org/2023/1110
25. Wu, L., Perin, G., Picek, S.: The best of two worlds: deep learning-assisted template attack. IACR Trans. Cryptogr. Hardw. Embed. Syst. **2022**(3), 413–437 (2022). https://doi.org/10.46586/tches.v2022.i3.413-437. https://tches.iacr.org/index.php/TCHES/article/view/9707
26. Wu, L., Perin, G., Picek, S.: I choose you: automated hyperparameter tuning for deep learning-based side-channel analysis. IEEE Trans. Emerg. Top. Comput., 1–12 (2022). https://doi.org/10.1109/TETC.2022.3218372
27. Wu, L., Perin, G., Picek, S.: Hiding in plain sight: non-profiling deep learning-based side-channel analysis with plaintext/ciphertext. Cryptology ePrint Archive, Paper 2023/209 (2023). https://eprint.iacr.org/2023/209
28. Wu, L., Tiran, S., Perin, G., Picek, S.: An end-to-end plaintext-based side-channel collision attack without trace segmentation. Cryptology ePrint Archive, Paper 2023/1109 (2023). https://eprint.iacr.org/2023/1109
29. Wu, L., et al.: Label correlation in deep learning-based side-channel analysis. IEEE Trans. Inf. Forensics Secur. **18**, 3849–3861 (2023). https://doi.org/10.1109/TIFS.2023.3287728
30. Yap, T., Picek, S., Bhasin, S.: Beyond the last layer: deep feature loss functions in side-channel analysis. In: Proceedings of the 2023 Workshop on Attacks and Solutions in Hardware Security, ASHES 2023, pp. 73–82. Association for Computing Machinery, New York (2023). https://doi.org/10.1145/3605769.3623996