Dissecting Leakage Resilient PRFs with Multivariate Localized EM Attacks

A Practical Security Evaluation on FPGA

Florian Unterstein[1], Johann Heyszl[1], Fabrizio De Santis[2], and Robert Specht[1]

[1] Fraunhofer Research Institution AISEC, Munich, Germany
name.surname@aisec.fraunhofer.de
[2] Technische Universität München, Munich, Germany
desantis@tum.de

Abstract. In leakage-resilient symmetric cryptography, two important concepts have been proposed in order to decrease the success rate of differential side-channel attacks. The first one is to limit the attacker's data complexity by restricting the number of observable inputs; the second one is to create correlated algorithmic noise by using parallel S-boxes with equal inputs. The latter hinders the typical divide and conquer approach of differential side-channel attacks and makes key recovery much more difficult in practice. The use of localized electromagnetic (EM) measurements has already been shown to limit the effectiveness of such measures in previous works based on PRESENT S-boxes and 90 nm FPGAs. However, it has been left for future investigation in recent publications based on AES S-boxes. We aim at providing helpful results and insights from LDA-preprocessed, multivariate, localized EM attacks against a 45 nm FPGA implementation using AES S-boxes. We show, that even in the case of densely placed S-boxes (with identical routing constraints), and even when limiting the data complexity to the minimum of only two inputs, the guessing entropy of the key is reduced to only 2^{48}, which remains well within the key enumeration capabilities of today's adversaries. Relaxing the S-box placement constraints further reduces the guessing entropy. Also, increasing the data complexity for efficiency, decreases it down to a direct key recovery. While our results are empirical and reflective of one device and implementation, they emphasize the threat of multivariate localized EM attacks to such AES-based leakage-resilient constructions, more than currently believed.

1 Introduction

Differential Power Analysis (DPA) is one of the most powerful classes of side-channel attacks against symmetric cryptographic implementations. It exploits multiple measurements obtained under the same key and different inputs to recover the secret key using statistical methods, and is particularly robust in the presence of noise. Conventional side-channel countermeasures like protected logic styles [13] and masking schemes [3] typically come with significant overhead in terms of implementation complexity, area and time resources. Leakage-resilient and re-keying techniques aim at bounding the side-channel leakage to a level which is not computationally exploitable for the adversary, while having less area overhead than conventional countermeasures. For instance, most leakage-resilient and re-keying schemes [14] reduce the number of observable computations by changing the secret key according to a predefined mechanism, e.g. at every execution. In such cases, the implementation needs to be protected only against single observation attacks. To generate session keys, possible approaches for re-keying schemes e.g. require to update a secret internal state (stateful devices), or use an internal random number generator to realize some form of key-update agreement protocol among the parties (stateless devices), as e.g. in [16] or CIPURSE from Infineon AG [9]. However, many embedded devices require stateless and non-interactive solutions, e.g. for encrypted software updates, or do not have secure random number generators. Leakage-resilient Pseudo-Random Functions (PRFs) [23] provide a stateless method to derive session keys based on a public input. Arguably, the PRF tree construction by Goldreich et al. [11] is one of the most influential leakage-resilient PRFs currently investigated in literature. It can easily be instantiated from block ciphers as shown in [17]. This allows to thwart differential side-channel attacks in two ways: (1) by reducing the data complexity (number of different observable inputs) by construction; (2) by adding correlated algorithmic noise to the measurements (which cannot be averaged out) by exploiting parallel S-boxes which are provided with the same plaintext inputs. Note that measurement complexity (number of measurements allowed), on the contrary, is generally not restricted.

In 2012, Medwed et al. [17] showed that limiting the data complexity alone is not sufficient to achieve protection against differential side-channel attacks, even if it is as low as 2[3]. Also, they concluded that the AES may not be a valid candidate for the construction of leakage-resilient PRFs (at least when the data complexity is > 2) due to the limited number of parallel S-boxes which can be instantiated, hence, leading to a remaining search complexity for enumeration of only $16! \approx 2^{44}$. Subsequently, Belaid et al. [2] investigated such a construction with 32 parallel PRESENT S-boxes and data complexity of 2^4, which led to a remaining search complexity of $32! \approx 2^{117}$ when faced with DPA. They also showed, however, that the security level can be reduced down to 2^{69} by employing univariate localized EM attacks [12] and breaking the equal leakages and correlated noise assumptions. Finally, in a recent contribution to ASIACRYPT 2016, Medwed et al. [18] used a Pseudo-Random Generator (PRG) (taken from Standaert et al. [22]) for the initialization of a novel unknown-inputs leakage-resilient PRF based on the AES block cipher. The security of the PRG part of the leakage-resilient construction is again based on minimal data complexity of two inputs and S-box parallelism to obtain correlated algorithmic noise. Their

[3] Reducing the data complexity to 1 would mean that only a single observation (with possibly unlimited measurement complexity) would be available to adversaries. This corresponds to a simple power analysis attack scenario which is not generally considered in most contributions in the field of leakage-resilient cryptography.

contribution explicitly mentions the lack of an empirical security evaluation using localized EM attacks, which is the main motivation for this work.

Contributions The main contribution of this paper is a laboratory evaluation of a leakage-resilient implementation based on AES using localized EM measurements. We provide answers to questions left open by Medwed et al. [18], who re-proposed the use of AES with data complexity 2, and Belaid et al. [2], who analyzed *unconstrained* S-boxes on a 90 nm FPGA with a data complexity of 2^4 and *univariate* localized EM attacks. In particular, we (1) employ state of the art profiled *multivariate* localized EM attacks using linear discriminant analysis (LDA) preprocessing for the identification of the Points of Interests (PoIs); and (2), investigate a design with carefully *constrained* S-boxes and a data complexity of 2 on a 45 nm Xilinx Spartan 6 device. Our results show that even when the lowest data complexity of 2, full parallelism of 16 S-boxes, and constrained placement are used, the practical achieved security level[4] is only 2^{48}. This suggests that leakage-resilient constructions [18, 22, 24] will not provide a sufficient level of implementation security in face of multivariate localized EM attacks, when implemented on FPGA devices similar to the one used for this evaluation.

2 Background

Leakage-Resilient PRFs Leakage-resilient Pseudo-Random Functions (PRFs) have been introduced in [19, 20, 23] and essentially build on the tree construction of Goldreich, Goldwasser and Micali [11]. The input x to the PRF is split into parts of a small number of m bits, which are input to multiple subsequent block cipher operations using different keys, i.e. the result of every encryption iteration is used as the key for the next round. In each round, m bits are taken from x and are replicated for the plaintext input until all bits of x are processed, as depicted in Fig. 1. The replication of the input bits achieves what is referred to as carefully chosen plaintexts by Medwed et al. [17], i.e. the plaintext input to every S-box is the same. As a consequence, the data complexity in an attack on any intermediate key is restricted to 2^m possible plaintexts. The choice of m imposes a trade-off between data complexity and efficiency for designers, as the number of necessary block cipher iterations is $^{128}/_{m}$. Note that for $m = 1$, the data complexity equals 2 and the leakage of the PRF becomes equivalent to that of the PRG used in [18] and [22], where the data complexity is also limited to two. If the leakage of all S-boxes is assumed to be equal, then standard DPA attacks will recover all key bytes at once, but without any information about their correct order within the secret key. This leads to a search complexity of

[4] This corresponds to the ranking of the key after a practical laboratory evaluation using localized EM, where the order of the key bytes is discovered during the attacks, but not all correct subkeys are ranked first. In contrast, the previously mentioned 2^{44} corresponds to the remaining search complexity of global attacks, once all key bytes are assumed to be ranked first, despite the correlated algorithmic noise (theoretical best case).

$16! \approx 2^{44}$ in case of AES, if measurements and attack have led to perfect results (all key bytes are ranked first). Results from Belaid et al. [2], however, have already shown that the equal and concurrent leakage assumption does not hold when localized EM measurements are performed.

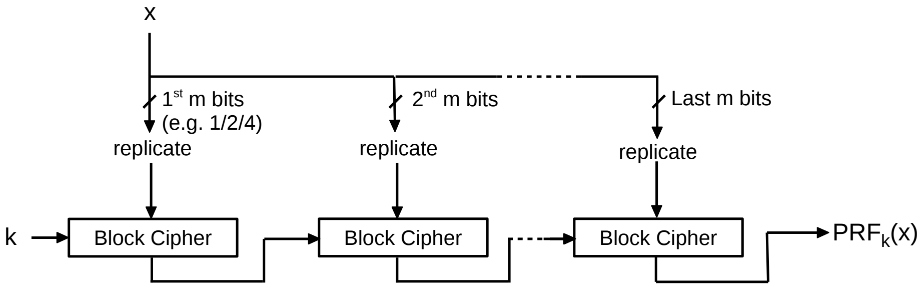

Fig. 1: Leakage resilient PRF.

Linear Discriminant Analysis While multivariate template attacks are amongst the most powerful differential side-channel attacks, they are also computationally intensive and can face numerical issues when the number of time-samples per trace is large. A common way to deal with this is to reduce the number of time-samples included in the calculation of the templates, i.e. the dimensionality of the trace. Fisher's Linear Discriminant Analysis (LDA) [8] has been proposed for template attacks by Archambeau et al. [1] and then specialized for EM attacks by Standaert et al. [21]. It has later been shown by Bruneau et al. [4] that this is in fact the optimal strategy to reduce the dimensionality of leakage traces. LDA also has the advantage that the transformation makes templates more robust against measurement campaign-dependent variations caused by temperature or environmental noise [7]. LDA stems from statistical classification and is a linear transformation of a dataset onto a lower-dimensional subspace with good class-separability. It calculates a transformation matrix $\mathbf{W}$, which maximizes the ratio of between-class to within-class scatter. In our case, we calculate one transformation matrix for each S-box and use the S-box input values as classes. Let $t_{i,j}$ be all traces with S-box input value i with $j \in [0, N_i - 1]$, $\mu_i = \frac{1}{N_i} \sum_{j=0}^{N_i-1} t_{i,j}$ the estimated class mean and $\mu = \frac{1}{256} \sum_{i=0}^{255} \mu_i$ the estimated overall mean. Then LDA calculates the within-class scatter matrix $\mathbf{S_w}$, between class scatter matrix $\mathbf{S_b}$ and $\mathbf{W}$, such that criterion J is maximized:[5]

$$\mathbf{S_w} = \sum_{i=0}^{255} \sum_{j=0}^{N_i-1} (t_{i,j} - \mu_i)(t_{i,j} - \mu_i)^T \tag{1}$$

$$\mathbf{S_b} = \sum_{i=0}^{255} N_i(\mu_i - \mu)(\mu_i - \mu)^T \tag{2}$$

$$J(\mathbf{W}) = \frac{\mathbf{W}^T\mathbf{S_b}\mathbf{W}}{\mathbf{W}^T\mathbf{S_w}\mathbf{W}} \tag{3}$$

The within-class scatter matrix is asymptotically equal to the *pooled* covariance matrix calculated over all traces. This assumes that all classes share the same covariance matrix (homoscedasticity), which is justified by the fact that the covariance values are determined by the influence of measurement noise, which should in most practical cases be independent of the inputs.

3 Hardware Design

The main building block of our design is a straightforward AES block cipher implementation with *full parallelism*, hence, 16 S-boxes (Canright S-boxes [5]) in the data path and 4 S-boxes in the key schedule as shown in Fig. 2. There is one state register, each AES round is computed in one clock cycle, and key scheduling is computed in parallel. As in the case of many other leakage-resilient constructions, this block cipher is used in two different stages, or modes: first the block cipher is used to generate a secret IV or session key from a public input, i.e. PRF mode, then, it is used for encryption (block cipher mode of operation). We emphasize that the block cipher remains the same (it shares the same hardware), while the input and key are different in those two modes. This is reasonable to avoid unnecessary overhead from duplication of the AES hardware block, but helps adversaries during profiling, since they may build profiles using either of the two modes.

The design is configured into a 45 nm Xilinx Spartan 6 XC6SLX9-3TQG144C FPGA. We synthesized the design in two ways, (1) without any routing constraints, and (2) with 16 hard-macro S-boxes placed as dense as possible. For the densely placed design, we first placed and routed one S-box (Fig. 6 in the appendix depicts the FPGA layout). Then we utilized Xilinx's relative location (RLOC) and area constraints to clone and place this 'hardmacro' as dense as possible (Fig. 7 in the appendix shows the placed S-boxes). This should help to fulfill the equal leakage assumption of the S-boxes as closely as possible because S-boxes are equal to a higher degree (apart from the routing to/from the S-box) and the area is generally smaller, which should make localized EM attacks more difficult. In addition, we constrained the placement of the rest of the AES to a

[5] The equations show the calculation of the transformation matrix for one S-box. We omitted an additional identifier for the S-box number for the sake of clarity.

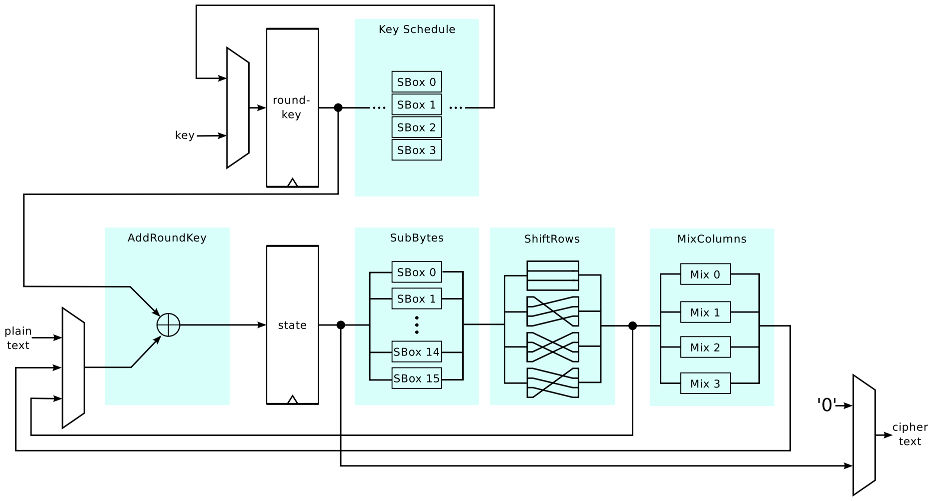

Fig. 2: AES hardware design.

confined area (black box in Fig. 7) in an attempt to make the routing, e.g. to the mix-columns logic, as short as possible. Based on the reports of the design tools, the estimated die area occupied by the AES is about 0.5 mm^2. Under these circumstances and for both placement options, we synthesized designs with $m = 1$ (data complexity of 2), and $m = 4$ (data complexity of 16).

4 Side-Channel Analysis

Our contribution evaluates the implementation security of the previously described techniques for leakage resilience. Since the parallelism of S-boxes and their ideally equal leakage characteristics are crucial to the idea of the construction, a high-precision EM measurement setup is especially relevant. Our assumption is that the localization capability thereof allows a spatial separation of the leakage of the individual S-boxes and the exploitation of even subtle differences in their characteristics. We use a state of the art high-end setup with a Langer ICR HH 100-27 100 µm diameter EM probe which is positioned about 10 µm over the decapsulated die surface. In addition to the built-in 30 dB amplifier of the probe, another Langer PA303 30 dB pre-amplifier is employed. A LeCroy WavePro 725Zi oscilloscope with 2.5 GHz bandwidth and a sampling rate of 5 GS/s records the measurements. The FPGA-based design is clocked at 20 MHz and this clock is synchronized to the oscilloscope. An X-Y-table is used to collect measurements on multiple locations over the die surface. The measurement positions are located within an area of about 2.8 mm by 2.8 mm, which should cover most, but not all of the floorplan (shown in Fig. 7 in the appendix) because the probe movement is limited by the bonding wires.

It is common practice to allow profiling for a meaningful implementation security analysis which is representative of the fact that adversaries may use their own devices where they could choose keys for profiling. In this profiled setting, the adversary is able to compute *all* internal states of the implementation. Based on this, we performed profiling using the block cipher mode of operation[6] of our implementation instead of the PRF mode. Analogously, an adversary would use stage 2 in the construction of Medwed et al. [18]. Our analysis is split into three tasks: (1) the localization of the measurement positions with the maximum leakage for each S-box, (2) the profiling phase on these positions, and (3), the attack phase. The first task is the most time consuming since it requires a full scan of the die surface. Considering that the measurement time grows quadratically when reducing the step size, we partitioned the measurement area in a grid of 20×20 for the unconstrained design and 40×40 for the dense design, which corresponds to a step size of 140 µm and 70 µm, respectively. We used a larger step size for the unconstrained design since we expect it to spread over a bigger area. On each position, we acquired $10,000$ traces. With our setup, the measurement takes roughly 1 day for the 20×20 grid and 4 days for the finer grid. For each S-box, we calculated the signal-to-noise ratio (SNR) by partitioning the traces according to the input values of the S-box [15]:

$$SNR^b = \frac{Var(Signal^b)}{Var(Noise^b)} = \frac{Var(\mu_0^b \ldots \mu_{255}^b)}{Mean(\sigma_0^{2^b} \ldots \sigma_{255}^{2^b})}, \tag{4}$$

with b being the index of the S-box and μ_i^b and $\sigma_i^{2^b}$ being the estimated mean and variance traces computed over all traces with input value i at this S-box. The result is a trace of many SNR values (SNR trace) which we evaluated within the timespan where the first AES round is computed. In our case one clock cycle corresponds to 250 samples, and the interesting part, i.e. the part where there is activity after the clock edge, is around 50 samples (10 ns) wide. There are several options how to chose positions from this part of the SNR trace. We found that in some cases, the positions with the highest peak SNR value gave the best results, and in others, the positions with the highest mean SNR (calculated over the 50 samples in the interesting region of the SNR trace) performed better. In cases where those metrics gave different positions or were ambiguous due to multiple peaks of similar amplitude, we conducted the rest of the analysis on all such positions for this S-box and kept the best result.

In two separate acquisition campaigns, we collected the profiling and attack traces. The attack traces were acquired with limited data complexity, i.e. 16 for $m = 4$, and 2 for $m = 1$. We cut the traces and limited our analysis to the time span where the first AES round is calculated. To reduce the number of samples included in the templates, we use LDA [8] as dimensionality reduction algorithm.

We compute full estimated Gaussian templates for each S-box and each of their S-box input values. As stated earlier, for LDA to be applicable, the traces belonging to one S-box are assumed to share a common covariance matrix, regardless of the input value. In this case, it suffices to calculate a single pooled

[6] We used OFB mode, but other modes would work as well.

covariance matrix for all templates belonging to one S-box. This gives a better estimate of the actual distribution and drastically reduces the computational effort for the template matching. Our experiments suggest that the assumption holds in our case and gave generally better results when using the pooled matrix when compared to separate covariance matrices. Thus, all our presented attacks were conducted using the pooled covariance matrix.

During the attack phase, the traces are matched against the templates in a template based DPA. Since we are using the pooled covariance matrix, we can make use of simplifications detailed by Choudary et al. [6] and calculate the logarithmic score. To combine the score of multiple attack traces, we sum the scores and calculate the average. This results in a list with scores for each subkey candidate. In order to calculate the overall key rank we used the key rank estimator proposed by Glowacz et al. [10]. The estimated key rank is, within its error boundaries, equivalent to the metric of guessing entropy used in other publications.

5 Results and Discussion

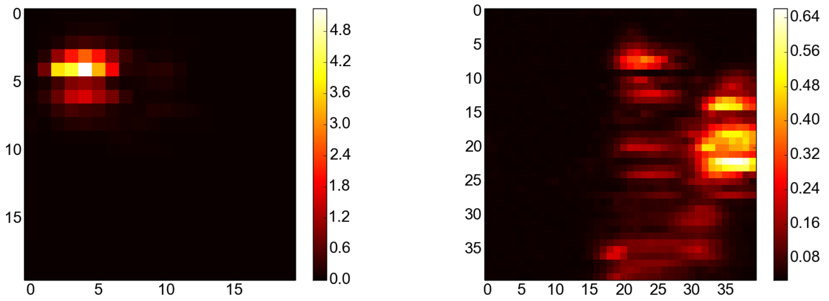

(a) Unconstrained placement.

(b) Dense hard-macro placement.

Fig. 3: SNR heat maps for S-box #0 with different placements.

Using the SNR analysis, we were able to localize useful measurement positions for all S-boxes on both tested designs. Figure 3 shows one example SNR heat map of S-box #0 on the two different designs. All other heat maps can be found in Figures 8 and 9 in the appendix. Each colored pixel represents the peak SNR value of the SNR trace at that measurement position for this S-box. In both maps, regions with the highest SNR are clearly distinguishable and most likely correspond to the actual physical location of the logic of S-box #0. An important observation is that the SNR values of the design with the densely placed hard-macro S-boxes are - on average - by a factor of 2 smaller than the ones from the unconstrained placement. The average peak SNR of the S-boxes on the dense placement is 0.87, compared to 1.61 on the unconstrained placement. In the case of dense placement, where SNR values are generally smaller, there are multiple positions which exhibit a relatively high SNR. As described, we simply evaluated all such locations for the corresponding S-box in the attack instead of choosing just one, which increased the measurement time of the attack.

S-Box Placement	**Data Complexity**	**Est. Key Rank**
Unconstrained	2	2^{20}
Dense	2	2^{48}
Unconstrained	16	1
Dense	16	1

Table 1: Estimated key ranks after the attacks.

For the profiling phase, we used a maximum of $65,000$ traces per position for the unconstrained design and $650,000$ traces for the dense design in an effort to compensate for the lower SNR. During the attack, up to $100,000$ traces were used per S-box. Table 1 summarizes the results of the attacks using all available traces. With a data complexity of 2^4 during the attack, security is completely broken and all key bytes are successfully recovered, regardless of the placement. This is a result which is similar to the findings of Belaid et al. [2].

As expected, a data complexity of 2 leads to better results. Several subkeys are not ranked first and consequentially, a higher key rank of 2^{20} remains for the unconstrained placement case with data complexity 2. However, as an important result, this is an obvious insufficient level of security.

The dense design improves the security significantly and provides a higher security level of 2^{48} compared to 2^{20}. In both cases, the achieved security level is insufficient, which is the main contribution of our investigations. This means that a minimum data complexity of 2 together with parallel S-box inputs is not suited to achieve meaningful leakage-resilient constructions, at least under the present circumstances of a 45 nm feature size FPGA implementation.

While the security level is established to be insufficient, an interesting question is, whether more profiling traces would further improve the attack, or whether the lower bound is reached. We repeated the attack with different numbers of profiling traces while using all available attack traces. The results for both designs are shown in Fig. 4. It can be noted that the gain of using more traces for profiling diminishes and the key ranks seem to approach a lower bound at about 2^{20} for the unconstrained, and about 2^{48} for the dense design. We conclude that increasing the number of profiling traces even further seems useless and that the efficiency of the attack is in fact limited by the leakage-resilience, and not by insufficient profiling due to the lower SNR. In other words, we expect that other uncorrelated noise sources are averaged out sufficiently.

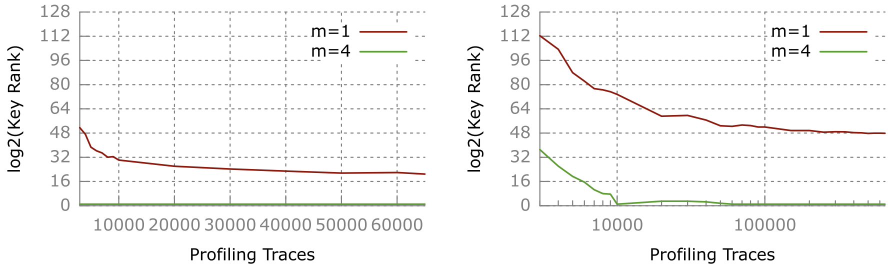

(a) Unconstrained design.

(b) Dense design.

Fig. 4: Key rank evolution with varying number of profiling traces and maximum number of attack traces.

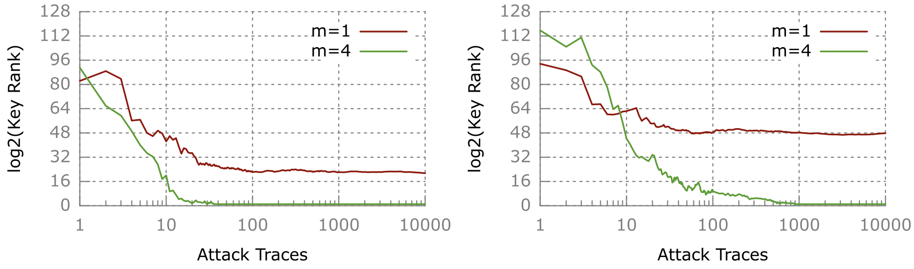

(a) Unconstrained design.

(b) Dense design.

Fig. 5: Key rank evolution with varying number of attack traces and maximum number of profiling traces.

In a similar manner, we investigated the number of traces required for the attack. In a real-world scenario, adversaries may have full access to one device for profiling, but limited access to the attacked device. Figure 5 shows the influence of the number of attack traces on the key rank when using templates built from the maximum number of available profiling traces. As an interesting observation, we report that the key rank seems to reach its lower bound after only about 100 attack traces, which is a surprisingly low number.

To verify the efficiency of the leakage-resilient construction against regular power attacks, we also conducted a template attack where we measured the global power consumption over a resistor in the power line with a differential probe. For increased SNR, all capacities were removed from the board. Despite using $1,000,000$ profiling traces, the attack fails to result in any significant key rank reduction. Interestingly, the correct subkeys were not even ranked highly but instead were distributed evenly across the subkey list. This is far from optimal, where correct subkeys would be ranked in the first 16 positions in all subkey lists and leave only the permutation complexity for the enumeration of the whole key. For the case of unlimited data complexity, we report that an univariate CPA using the Hamming distance leakage model already succeeds with 20.000 traces. Even though this aspect was not the focus of our research, this discrepancy is an encouraging result when adversaries are limited to global (power) attacks.

Given that our analysis is reflective of one technology, namely 45 nm FPGAs, it remains unclear, how our results affect other and smaller technologies such as ASICs or upcoming 16 nm FPGA devices. In our case study, the die area occupied by the AES is about 0.5 mm^2 and relatively large compared to the probe diameter of 100 µm. For a rough comparison to an ASIC design, we synthesized our AES core for UMC's 55 nm process using Synopsys Design Compiler. The resulting design uses about 10.000 gate equivalents with an estimated die area of less than 0.02 mm^2 when place and route overhead is taken into account. This is significantly smaller than our FPGA design and comes close to the size of the probe itself.

6 Conclusion

We demonstrated that the achieved security level of AES-based leakage resilient implementations employing minimum data complexity and S-box parallelism is insufficient in the localized EM scenario, at least in cases similar to our FPGA with 45 nm feature size. In particular, we were able to isolate the leakage of individual S-boxes and attack them separately using LDA-based, profiled, multivariate attacks, thus, circumventing the "equally leaking" and "correlated algorithmic noise" assumptions. We were able to completely recover the correct key for all designs with data complexity 2^4. A data complexity of 2 proved to be more resilient, but we were still able to reduce the key rank to 2^{20} and 2^{48} for the unconstrained and dense placement, respectively. Finally, it remains as an open question whether a denser placement and smaller feature sizes on ASIC will suffice to reach acceptable security levels against localized EM attacks. In this regard, we advise further analysis.

Acknowledgements. The work presented in this contribution was supported by the German Federal Ministry of Education and Research in the project *ALESSIO* through grant number 16KIS0629.

References

1. Archambeau, C., Peeters, E., Standaert, F., Quisquater, J.: Template attacks in principal subspaces. In: Cryptographic Hardware and Embedded Systems - CHES 2006, 8th International Workshop, Yokohama, Japan, October 10-13, 2006, Proceedings. pp. 1–14 (2006)
2. Belaïd, S., De Santis, F., Heyszl, J., Mangard, S., Medwed, M., Schmidt, J.M., Standaert, F.X., Tillich, S.: Towards fresh re-keying with leakage-resilient PRFs: cipher design principles and analysis. Journal of Cryptographic Engineering 4(3), 157–171 (2014)

3. Belaïd, S., Grosso, V., Standaert, F.X.: Masking and leakage-resilient primitives: One, the other (s) or both? Cryptography and Communications 7(1), 163–184 (2015)
4. Bruneau, N., Guilley, S., Heuser, A., Marion, D., Rioul, O.: Less is more - dimensionality reduction from a theoretical perspective. In: Cryptographic Hardware and Embedded Systems - CHES 2015 - 17th International Workshop, Saint-Malo, France, September 13-16, 2015, Proceedings. pp. 22–41 (2015)
5. Canright, D.: A very compact s-box for AES. In: Cryptographic Hardware and Embedded Systems - CHES 2005, 7th International Workshop, Edinburgh, UK, August 29 - September 1, 2005, Proceedings. pp. 441–455 (2005)
6. Choudary, O., Kuhn, M.: Efficient template attacks. In: Francillon, A., Rohatgi, P. (eds.) Smart Card Research and Advanced Applications, Lecture Notes in Computer Science, vol. 8419, pp. 253–270. Springer International Publishing (2014)
7. Choudary, O., Kuhn, M.: Template attacks on different devices. In: Prouff, E. (ed.) Constructive Side-Channel Analysis and Secure Design, Lecture Notes in Computer Science, vol. 8622, pp. 179–198. Springer International Publishing (2014)
8. Fisher, R.A.: The use of multiple measurements in taxonomic problems. Annals of Eugenics 7(7), 179–188 (1936)
9. Gammel, B., Fischer, W., Mangard, S.: Generating a session key for authentication and secure data transfer (Nov 7 2013), US Patent 2014016955
10. Glowacz, C., Grosso, V., Poussier, R., Schüth, J., Standaert, F.: Simpler and more efficient rank estimation for side-channel security assessment. In: Fast Software Encryption - 22nd International Workshop, FSE 2015, Istanbul, Turkey, March 8-11, 2015, Revised Selected Papers. pp. 117–129 (2015)
11. Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions. Journal of the ACM (JACM) 33(4), 792–807 (1986)
12. Heyszl, J., Mangard, S., Heinz, B., Stumpf, F., Sigl, G.: Localized electromagnetic analysis of cryptographic implementations. In: Dunkelman, O. (ed.) Topics in Cryptology - CT-RSA 2012. Lecture Notes in Computer Science, vol. 7178, pp. 231–244. Springer Berlin / Heidelberg (2012)
13. Kirschbaum, M.: Power Analysis Resistant Logic Styles – Design, Implementation, and Evaluation. Ph.D. thesis (2011)
14. Kocher, P.C.: Leak-resistant cryptographic indexed key update (Mar 25 2003), US Patent 6,539,092
15. Mangard, S., Oswald, E., Popp, T.: Power Analysis Attacks. Springer Science & Business Media (2008)
16. Medwed, M., Standaert, F.X., Großschädl, J., Regazzoni, F.: Fresh re-keying: Security against side-channel and fault attacks for low-cost devices. In: AFRICACRYPT. pp. 279–296 (2010)
17. Medwed, M., Standaert, F., Joux, A.: Towards super-exponential side-channel security with efficient leakage-resilient PRFs. In: Prouff, E., Schaumont, P. (eds.) Cryptographic Hardware and Embedded Systems - CHES 2012 - 14th International Workshop, Leuven, Belgium, September 9-12, 2012. Proceedings. Lecture Notes in Computer Science, vol. 7428, pp. 193–212. Springer (2012)
18. Medwed, M., Standaert, F.X., Nikov, V., Feldhofer, M.: Unknown-input attacks in the parallel setting: Improving the security of the CHES 2012 leakage-resilient PRF. In: Advances in Cryptology–ASIACRYPT 2016: 22nd International Conference on the Theory and Application of Cryptology and Information Security, Hanoi, Vietnam, December 4-8, 2016, Proceedings, Part I. pp. 602–623. Springer (2016)

19. Petit, C., Standaert, F.X., Pereira, O., Malkin, T.G., Yung, M.: A block cipher based pseudo random number generator secure against side-channel key recovery. In: Proceedings of the 2008 ACM symposium on Information, computer and communications security. pp. 56–65. ACM (2008)
20. Pietrzak, K.: A leakage-resilient mode of operation. In: Proceedings of the 28th Annual International Conference on Advances in Cryptology: The Theory and Applications of Cryptographic Techniques. pp. 462–482. EUROCRYPT '09, Springer-Verlag, Berlin, Heidelberg (2009)
21. Standaert, F., Archambeau, C.: Using subspace-based template attacks to compare and combine power and electromagnetic information leakages. In: Cryptographic Hardware and Embedded Systems - CHES 2008, 10th International Workshop, Washington, D.C., USA, August 10-13, 2008. Proceedings. pp. 411–425 (2008)
22. Standaert, F.X., Pereira, O., Yu, Y.: Leakage-resilient symmetric cryptography under empirically verifiable assumptions. In: Advances in Cryptology–CRYPTO 2013, pp. 335–352. Springer Berlin Heidelberg (2013)
23. Standaert, F.X., Pereira, O., Yu, Y., Quisquater, J.J., Yung, M., Oswald, E.: Leakage resilient cryptography in practice. In: Towards Hardware-Intrinsic Security, pp. 99–134. Springer (2010)
24. Taha, M.M.I., Schaumont, P.: Key updating for leakage resiliency with application to AES modes of operation. IEEE Trans. Information Forensics and Security 10(3), 519–528 (2015)

A Floorplanning

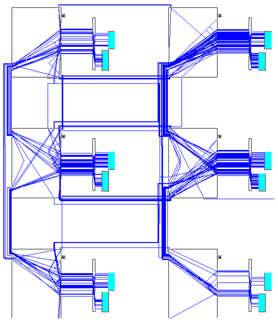

Fig. 6: Layout of one S-box in the Xilinx IDE.

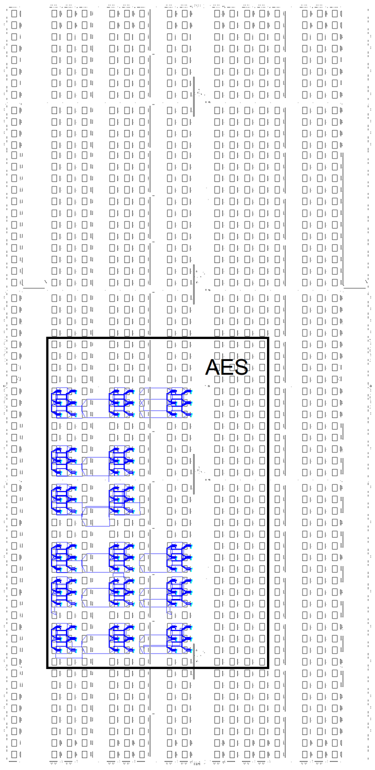

Fig. 7: Position of 16 S-boxes on the floorplan of the Xilinx Spartan 6 FPGA. The entire AES is placed within the black box.

B SNR Heat Maps For All S-Boxes

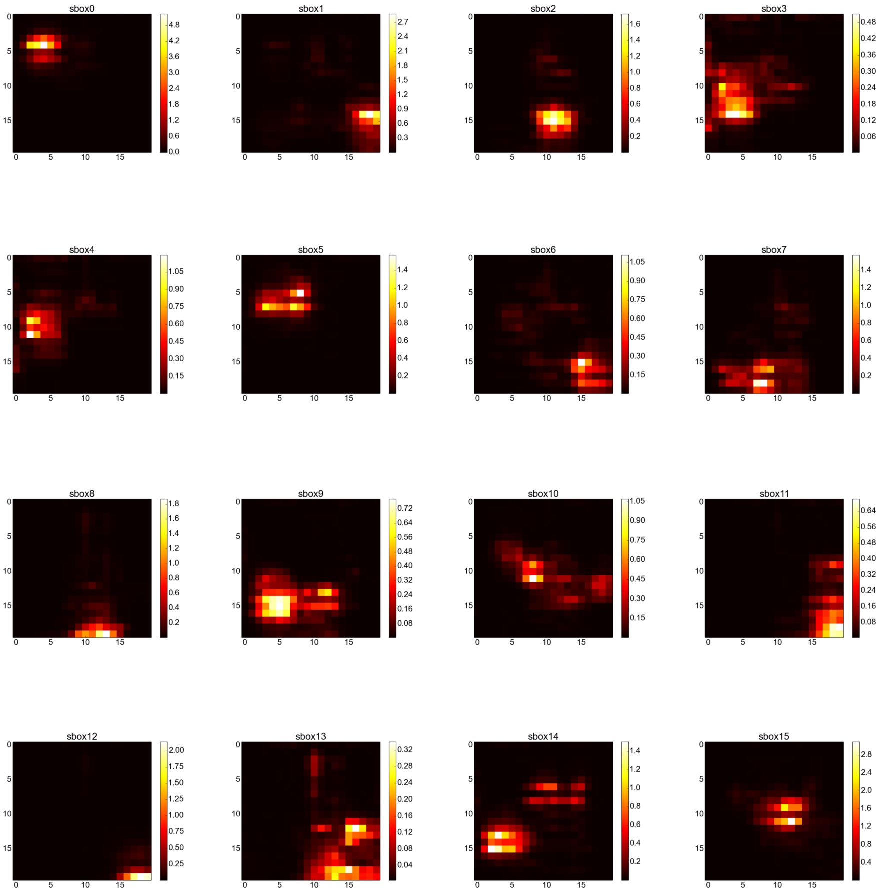

Fig. 8: SNR heat maps of unconstrained placement.

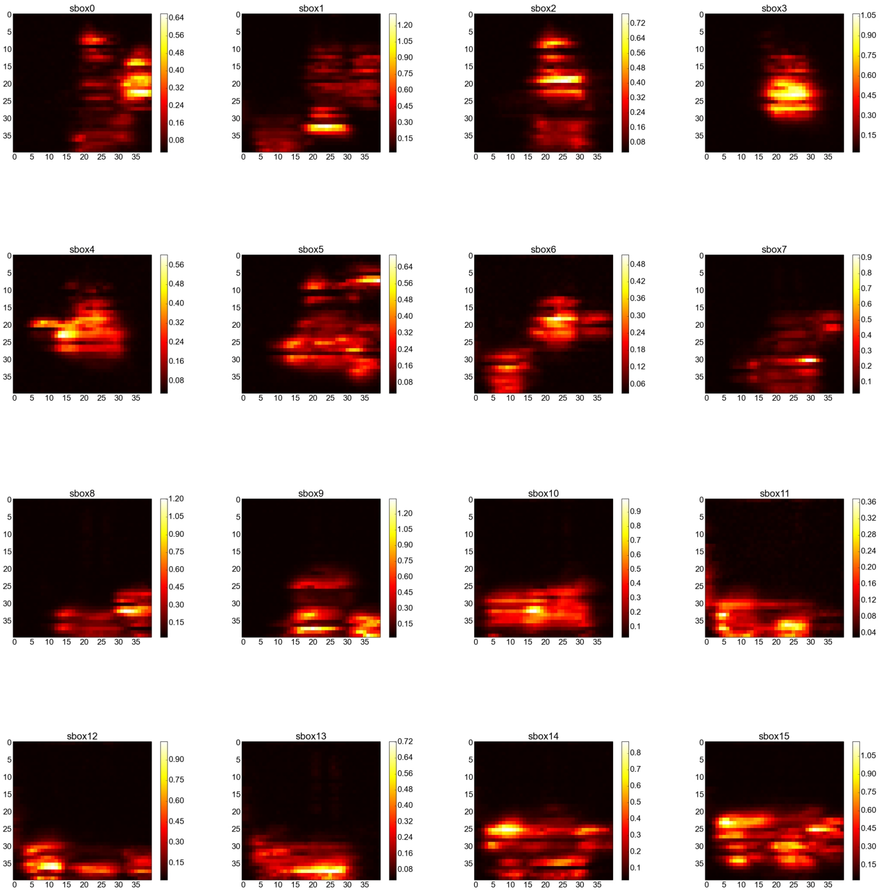

Fig. 9: SNR heat maps of dense hard-macro placement.