Improving Non-Profiled Attacks on Exponentiations Based on Clustering and Extracting Leakage from Multi-Channel High-Resolution EM Measurements

Robert Specht[1], Johann Heyszl[1], Martin Kleinsteuber[2], and Georg Sigl[2]

[1] Fraunhofer Institute AISEC, Munich, Germany
robert.specht@aisec.fraunhofer.de, johann.heyszl@aisec.fraunhofer.de
[2] Technische Universität München, Munich, Germany
kleinsteuber@tum.de, sigl@tum.de

Abstract. The success probability of side-channel attacks depends on the used measurement techniques as well as the algorithmic processing to exploit available leakage. This is particularly critical in case of asymmetric cryptography, where attackers are only allowed single side-channel observations because secrets are either ephemeral or blinded by countermeasures. We focus on *non-profiled* attacks which require less attacker privileges and cannot be prevented easily. We *significantly improve* the algorithmic processing in *non-profiled* attacks *based on clustering* against exponentiation-based implementations compared to previous contributions. This improvement is mainly due to PCA and a strategy to select few mid-ranked components where exploitable, low-variance leakage is concentrated. As a result from a practical experiment using single-channel high-resolution magnetic field measurements, we report a significant improvement in the number of successful attacks. Further, we present the first practical results from using three such channels simultaneously. The combination of three channels leads to further improved results over the best individual channel when applying a *profiled* template attack. The *clustering-based* algorithmic approach for the *non-profiled* attack, however, does not show improvements from the combination.

1 Introduction

The side-channel information leakage about secret-dependent internal values is usually limited. Attackers who target implementations of symmetric ciphers may repeat measurements many times to collect sufficient leakage information while the secret remains unchanged. In case of asymmetric algorithms, however, the secret is either ephemeral or blinded through countermeasures and attackers are only allowed one side-channel observation. Hence, it is crucial to *record and exploit* as much leakage as possible. *Profiled* attacks, e.g. template attacks, are powerful in exploiting leakage efficiently, however, can be prevented by blinding or by preventing attackers from gaining full access for profiling. *Non-profiled*

attacks cannot be prevented in this way, because they do not require profiles and hence, are a much bigger threat to devices. Heyszl et al. [8] proposed to use *well-established* clustering algorithms for *non-profiled* attacks. They use k-means clustering after a simple sum-of-squares pre-processing of the measurement data in their practical experiments.

We follow their proposal and *significantly improve* the algorithmic approach. Principal Component Analysis (PCA) has been used for pre-processing and data reduction in other side-channel attacks [5, 2, 22, 3, 15]. Also, strategies to select only certain principal components have previously been mentioned [3]. We apply PCA to *clustering-based, non-profiled* attacks on exponentiation algorithms and performed practical experiments on an FPGA-based implementation of Elliptic Curve Cryptography (ECC) by using high-resolution electromagnetic measurements as side-channel. We find that PCA concentrates exploitable leakage with comparably low variance into few components which are not the highest-ranked ones. Hence, as an important step after transformation, we discard high-ranked as well as many low-ranked components during a parametrized selection. In our *non-profiled* setting, this requires some testing for the right selection parameters, hence, brute-force by the attacker. However, significantly improved attack results clearly justify this. For cluster classification, we use the expectation maximization algorithm instead of the k-means algorithm [8]. The resulting attack is successful with single-channel measurements in significantly more cases than if using the algorithmic approach by Heyszl et al. [8]. Most of the achieved algorithmic improvement can be attributed on using PCA and the component selection as pre-processing technique before clustering. Like expected, a *profiled* template attack still outperforms the improved *non-profiled* attack.

Another way to improve attacks in single-execution settings is to use *multiple simultaneous channels* and combine their leakage. Previous contributions have tested the combination of (low-resolution) magnetic field measurements and current consumption measurements [1, 22] using template attacks. High-resolution magnetic field measurements should generally provide better signal qualities [10] and allow to capture multiple independent channels because the signals highly depend on measurement locations [9]. We present the *first practical results from using three high-resolution magnetic field probes simultaneously* and combine them in the clustering-based non-profiled attack. However, we find that the combination of three channels does *not improve* the results using the *non-profiled PCA- and clustering-based* attack compared to the best individual channel. We conclude that in the *non-profiled* setting, our approach seems unsuitable for combining multi-channel data. The *profiled* template attack, however, leads to a *significant improvement* through the combination of channels. In profiled settings, attackers are able to find the best measurement positions for single channels. Hence, the additional cost for multi-channel equipment is only reasonable in *profiled* settings *and* if the available *leakage* is still *insufficient*.

We first explain the background and related work of (non-profiled) attacks against exponentiations in Sect. 2.1. In Sect. 2.2, we cover the background and related work of magnetic field side-channels and multi-channel measure-

ments. In our first main Sect. 3, we describe our algorithmic approach to improve clustering-based attacks on exponentiations and to handle multi-channel data. We back these considerations by practical experiments in Sect. 4 and discuss the results. We summarize our contribution and findings in Sect. 5.

2 Preliminaries

2.1 Non-Profiled Attacks Against Exponentiations

The main computation in public key cryptosystems is modular integer exponentiation with secret exponents (e.g. RSA, DSA) or elliptic curve scalar multiplication (e.g. ECDSA) with secret scalars. In this contribution, we use the generalized terms 'exponentiation algorithms' and 'secret exponents'. The secret exponent is usually either ephemeral by design (e.g. ECDSA) or blinded through countermeasures (e.g. exponent blinding in RSA, or in ECDSA to prevent profiling). Therefore, it is different for every execution and side-channel attackers may only exploit *single executions*. The first *single-execution* attack on exponentiations was presented by Kocher [13] who exploits data-dependent execution times of algorithms. To avoid this, improved algorithms like the square-and-multiply-always, double-and-add-always or the Montgomery ladder algorithm have constant operation sequences (e.g. side-channel atomic routines) to avoid such *simple side-channel attacks*. In all those algorithms, exponents are scanned bit- or digit-wise (depending on whether it is a binary, m-ary, or sliding window exponentiation) and the computation is performed in a loop iterating a constant sequence of operations. (We will continue to refer to the binary case in this contribution.) Nonetheless, some side-channel leakage about the processed exponent remains in many cases which can be referred to as single-execution leakage. Examples include data-dependent leakage from using pre-computed multiples in digit-wise multiplications [25], address-bit leakage [12], location-dependent leakage from accessing different storage locations [9], or operation-dependent leakage, e.g., when square and multiply operations can be distinguished [4].

Attacks against an exponentiation are carried out by partitioning side-channel measurements into *trace-segments* with each segment corresponding to an independently processed bit of the secret exponent.The segmentation borders are either known a priori, or can often be derived from visual inspection or comparison of shifted trace parts. The trace for measuring n exponent bits consists of n trace-segments $\boldsymbol{t_d} = (t_{1+(d-1)\cdot l}, ..., t_{d\cdot l})$ with $d \in [1, n]$, each of which is of length l (time-samples) which is referred to as its dimensionality (of features). For analyzing and attacking the measurement data, a $n \times l$ matrix $\boldsymbol{M}$ is constructed by placing each segment in one row. The contained leakage is exploited to find a structure, or partitioning of the rows due to secret exponent values. Template attacks use a profiling step to create templates of the segments for different values. Profiling can be prevented in many cases by blinding countermeasures or not allowing attackers full access to devices for profiling.We concentrate on *non-profiled* attacks because they are more powerful and threatening.

There have been several published attacks on exponentiations which do *not* require profiling. Walter [25] was the first to describe an attack by using a *custom* algorithm (resembling a clustering algorithm) to partition the segments into buckets. Messerges et al. [16], Clavier et al. [6], and Witteman et al. [26] use cross-correlation in non-profiled single-execution attacks on exponentiations. We pursue the approach by Heyszl et al. [8] who promote the use of established clustering algorithms (such as e.g. k-means) for *non-profiled* attacks due to the generality of their approach and support for the combination of multiple channels. A correct classification of trace-segments equals the recovery of the secret exponent. (Later, Perin et al. [18] described a similar but heavily customized two-stage approach which seems tailored to their case and unreasonable for generalization.) *We extend and significantly improve previous work by using Principal Component Analysis (PCA) and expectation maximization clustering (instead of k-means and simple pre-processing).*

2.2 Multi-Probe Measurements of Magnetic Fields

Using multiple side-channels concurrently, and combining them in an attack is an important way of increasing the exploitable leakage in single-execution attacks. Agrawal et al. [1] first, and later Standaert et Archambeau [22], describe the combination of *current consumption* with *magnetic field* measurements in *profiled* attacks through *concatenation* of traces. Standaert and Archambeau [22] report better results from magnetic field than current measurements and report an improvement from the combination of both channels. Souissi et al. [21] first presented results from combining *two* simultaneous measurements of the *magnetic field*. They measure the field close to two different supply capacitors of an FPGA. In this way they measure the supply of two different parts of the FPGA.

We find that in many cases, side-channel measurements of the magnetic field are closely related to the consumption of an *entire* device because comparably large coil diameters ($> 500\,\text{um}$) are used at large distances to the integrated circuits ($> 300\,\text{um}$) [20, 17, 22, 7, 1]. Such measurements often capture the magnetic field of supply wires (bonding wires) which is directly proportional to the current consumption of the *entire* integrated circuit (including noise sources from within the device). In our opinion, it is *unreasonable* to simultaneously record more *than one magnetic field* channels in such cases due to this global character. Lately, high-resolution magnetic field measurements at close distances to an integrated circuit die have been investigated extensively by Heyszl et al. [9, 10]. Such high-resolution measurements require magnetic field probes with diameters of $\approx 150\,\mu\text{m}$ at close distances to an integrated circuit die ($< 100\,\mu\text{m}$). In our opinion, the capturing of *multiple simultaneous magnetic field* side-channels *only* makes sense in case of such *high-resolution* measurements which can be restricted to *parts* of integrated circuits because they will convey sufficiently different information (e.g. localized leakage [9]). Heyszl et al. [8] mention the combination of multiple high-resolution channels for *non-profiled single-execution* attacks, however, did not perform actual simultaneous measurements. *We extent their*

work and present first results from an extensive practical study using three high-resolution *micro-coil magnetic field channels.*

3 Improving Clustering-Based Attacks

In this section, we describe our algorithmic approach to clustering-based *non-profiled* attacks on exponentiations which improves previous work [8]. We explain how we use Principal Component Analysis (PCA) as a pre-processing step for dimensionality reduction and feature selection in Sect. 3.1. We continue and describe how expectation maximization clustering can be used to attack single- and multi-channel measurements in Sect. 3.2. Finally, we describe how classification errors can be handled and derive the brute-force complexity as a measure to assess attack outcomes in Sect. 3.3.

3.1 PCA for Dimensionality Reduction and Feature Selection

Side-channel measurements usually lead to big amounts of data, especially when high sampling rates for magnetic field measurements are required. This increases required computational power and memory consumption during subsequent data analysis. Only a small part of the data will contain exploitable leakage information. Hence, *feature selection* to discard other parts is desirable.

Simple trace compression [14] is commonly used and usually justified by electrical properties. This includes extracting the peak values or computing the sum-of-squares (such as Heyszl et al. [8]) during the time-period of one clock cycle. Another popular method is the selection of so-called points-of-interest. This subset is usually identified through profiling with known secrets.

We concentrate on powerful *non-profiled, unsupervised* methods, specifically, on PCA. PCA has been applied to side-channel analysis for data reduction in several contributions [5, 2, 22, 3, 15] for different attacks of which Archambeau et al. [2] were the first to describe the use of PCA in the context of template attacks. Standaert and Archambeau [22] later compare PCA and Linear Discriminant Analysis (LDA) in the context of template attacks and confirm that LDA leads to superior results. We disregard LDA because training data from profiling is used to achieve a representation which maximizes cluster separation.

PCA is based on Singular Value Decomposition (SVD) and transforms the data into another coordinate system subspace with linearly uncorrelated coordinates by using the variance as score function, hence, maximizing the retained variance of the data. As described in Sect. 2.1, recorded side-channel measurements are cut into trace-segments corresponding to exponent bits. This leads to the real matrix $\boldsymbol{M}$ of measurement data, with the shape $n \times l$ for every probe (see Sect. 2.1). The SVD of $\boldsymbol{M} = U * \Sigma * V^*$ and the transformation into the orthogonal subspace of $\boldsymbol{M}$ equals $U * \Sigma$. This matrix $U * \Sigma$ consists of column vectors $(\boldsymbol{PC_1}, ..., \boldsymbol{PC_r})$ with r being the number of row-vectors and $\boldsymbol{PC_j}$ being a column-vector of shape $n \times 1$, which is called a principal component. The maximum number of components equals the number of trace-segments n of

the original data, $\max|PC| = \min(n, l)$, because the segment-length l is usually much larger than n. After applying PCA, the components are ordered by their variance which can be found in the diagonal matrix Σ. In our experiments, we normalize the variances of the principal components to one, i.e. we directly use $\boldsymbol{M}_{\mathrm{PCA}} = U$ instead of $U * \Sigma$. Before applying PCA, we removed the mean of every trace-segment as a standard measure.

Ideally, a transformation into a reduced subspace should maintain the 'useful' information while neglecting 'not useful' information, which is difficult without supervision. PCA combines correlating input dimensions into single principal components. Archambeau et al. [2] propose to only retain the first-ranked components assuming that the leakage is contained there, while discarding the remaining low-variance ones, assuming only noise is contained. Batina et al. [3] found in their practical experiments, that results of correlation-based Differential Power Analysis (DPA) improved when removing first-ranked components. There are several reasons for high variances of the trace segments, e.g. data-dependent signal influences and noise, which are irrelevant to the desired classification. We suspect that relevant and irrelevant signal parts will aggregate within separate components. Also, from our experience, the 'interesting' leakage signal parts are rather low-variance in the case of single-execution attacks.

Hence, we *propose a selection strategy* which discards several highest-ranked as well as many low-ranked components because they either contain noise or information which we are not interested in. We either select *single* principal components or a number of *consecutive* components (random choices of multiple components will lead most likely to an untestable amount of possibilities). Reduced trace-segments $\boldsymbol{M}_{\mathrm{PCA},k:k+i} = (\boldsymbol{PC_k}, ..., \boldsymbol{PC_{k+i}})$ are derived with k the first selected component and $i \geq 1$ the number of consecutive components retained. We trialled values of $k \in [1, 20]$ and $i \in \{1, 2, 4, 6, 9\}$ in our practical experiments and found that using only one single component $i = 1$ leads to the best results in our attack on average, and that the $k \leq 3$ first-ranked components should be discarded. This selection strategy reflects the approach of an attacker who is unable to perform profiling. An optimal selection of components can certainly *not* be determined *a priori* because it is highly device- and application-specific (general issue in machine learning [27]). Hence, without a priori-knowledge, an attacker has to trial different values for k and i. This, however, only requires an additional *brute-force complexity* of a few bits and improved attack outcomes clearly justify this.

3.2 Expectation-Maximization Clustering of Multi-Channel Data

Clustering algorithms can generally be split into supervised, semi-supervised and unsupervised algorithms. Our focus on *non-profiled* attacks restricts the choice to *unsupervised* algorithms. Heyszl et al. [8] first describe how *unsupervised* clustering algorithms can be used in a *non-profiled* attack to partition n trace-segments into classes according to their secret exponent values. An unsupervised cluster classification is equivalent to estimating the free parameters of

the classes' assumed distribution model. The choice of the algorithm and free parameters depends on the assumed probability distribution model, hence shape of the clusters. While Heyszl et al. use k-means clustering, we improve this by using the expectation maximization algorithm while keeping the Gaussian distribution assumption which both algorithms are based on.

Expectation-maximization clustering provides more free parameters which leads to a generally improved approximation of the cluster distributions, which usually leads to better classification results. The algorithm is based on repeated expectation and maximization steps. During these iterations the maximum likelihood means and covariances for the Gaussian distribution are derived. The result is a classification and a class-membership probability which indicates the reliability of correct classification for each segment (resp. secret exponent bit). The number of free parameters in the clustering algorithm can be chosen. We assume that the cluster shapes are mainly defined by Gaussian distributed noise. Additionally, we assume the noise being independent of the processed bit value. Hence, we chose to estimate two means and one joint full covariance matrix.

Multiple simultaneous measurements channels are combined by concatenating the trace-segments from different channels which correspond to the same exponent bits [1, 8]. PCA is applied to all side-channel measurement channels separately before concatenation. For example, segments $\boldsymbol{M}^1_{\mathrm{PCA},k:k+i}$ from measurement channel 1 are combined with segments $\boldsymbol{M}^2_{\mathrm{PCA},k:k+i}$ from measurement 2 leading to combined segments $\boldsymbol{M}^{\mathrm{combined}}_{\mathrm{PCA},k:k+i} = (\boldsymbol{M}^1_{\mathrm{PCA},k:k+i}, \boldsymbol{M}^2_{\mathrm{PCA},k:k+i})$. An attacker would rather use the *same values for k and i in all channels* because it significantly increases the attack complexity to test different k-s and i-s for every channel *without profiling* (e.g. repeat clustering process $(20 * 5)^3$ times).

3.3 Classification Errors and Required Brute-Force Complexity

If the recovered exponent is incorrect, faulty bits need to be identified, which is usually hard. As described by Heyszl et al [8], an attacker can use the bits' probabilities of correctness to judge which need to be trialled for correctness and follow a simple strategy to enumerate possible keys. This strategy leads to an estimated remaining *brute-force complexity* which we use to assess practical attack outcomes. Better, even optimal, key enumeration strategies [23, 24] will result in a lower amount of required brute-force *if* the attacker applies them. However, the typically large key sizes in asymmetric cryptography make the application of such algorithms challenging for attackers as well as evaluators. The said *brute-force complexity* which is used instead can be seen as an upper bound for the rank of the correct key as derived from an optimal enumeration.

We chose to use the silhouette index score [19] for the bits' error probability. It is based on the cumulative distance of each trace-segment to other trace-segments of each cluster. The silhouette index is calculated for every $\boldsymbol{m}_{\mathrm{PCA}}$, which corresponds to one row of $\boldsymbol{M}_{\mathrm{PCA},k:k+i}$, with $\mathcal{C}_1$ being the set of trace segments $\boldsymbol{t_d}$ of the same cluster like $\boldsymbol{m}_{\mathrm{PCA}}$ (determined by the expectation maximization algorithm) and $\mathcal{C}_2$ being the set of trace segments belonging to

other clusters. With the distance function $dist(a, b)$ (we use Euclidean distance due to the Gaussian noise assumption) the silhouette index s is computed as:

$$s(\boldsymbol{m}_{\mathrm{PCA}}, \mathcal{C}_1, \mathcal{C}_2) = \frac{f(\mathcal{C}_1, \boldsymbol{m}_{\mathrm{PCA}}) - f(\mathcal{C}_2, \boldsymbol{m}_{\mathrm{PCA}})}{max(f(\mathcal{C}_1, \boldsymbol{m}_{\mathrm{PCA}}), f(\mathcal{C}_2, \boldsymbol{m}_{\mathrm{PCA}}))} \tag{1}$$

$$f(\mathcal{C}, \boldsymbol{m}_{\mathrm{PCA}}) = \frac{1}{|\mathcal{C}|} \sum_{x \in \mathcal{C}} dist(x, \boldsymbol{m}_{\mathrm{PCA}}) \tag{2}$$

After calculating the score for all n segments, the ones with the lowest s are brute-forced in repetitions while including an increasing number of bits [8]. Let q be the last bit which is trialled until the correct exponent is found, then $2^{(q+1+1)}$ different exponents have to be tested at maximum which can be referred to as remaining *brute-force complexity* after the attack [8]. One additional bit is included for both possibilities to assign labels to the two classes. It equals $2^{(n+1+1)}$ at maximum and 2^1 at minimum.

4 Practical Evaluation

We present the first practical results from the *simultaneous* use of *three high-resolution* magnetic field probes. We chose a fixed geometric arrangement of the measurement probes close to the surface of an FPGA die and performed 400 measurements at different positions to gain conclusive insights from a high number of tests. We succeed in *demonstrating the algorithmic improvement* from our approach and derive conclusions about the *benefit from using multiple channels simultaneously*.

4.1 Design-Under-Test and Multi-Probe Setup

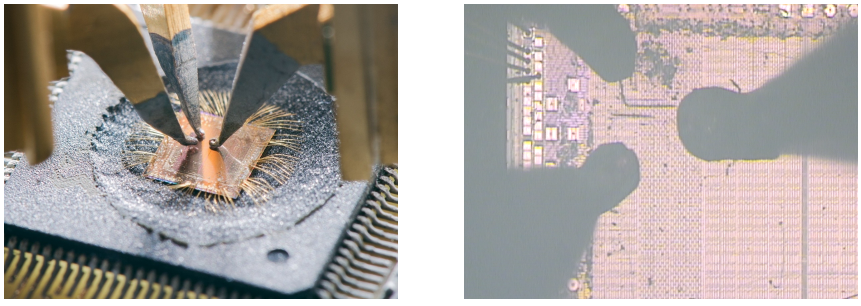

(a) View from side (b) Top-view

Fig. 1: Geometric arrangement of measurement-probes on FPGA die surface

As a device under test, we use a Xilinx Spartan 3A FPGA chip (see Fig. 1a) which is configured with an Elliptic Curve Cryptography (ECC) design and

performs an 163 bit elliptic curve scalar multiplication using a Montgomery ladder. This algorithm is a classical candidate for attacks against exponentiation algorithms since it processes the secret exponent bit-wise in n constant time segments. As a single-execution side-channel leakage about the consecutively processed exponent bits we exploit location-based leakage which is revealed by high-resolution measurements of the electromagnetic field [9].

After decapsulating the FPGA die (see Fig. 1a), we use an area of 1700 µm × 1700 µm on the surface of the die between bonding wires to place probes. We arrange three probes in a fixed formation, and place them on 400 (20 × 20) different positions within this area to able to evaluate 400 data sets by our analysis. Figure 1 depicts the geometric arrangement of the probes from the side and from the top. The distance of the probes to the die surface is approximately 100 µm. We used three near-H-field (magnetic) probes with coil diameters of 250 µm, 150 µm and 100 µm which we had available in our laboratory. The bandwidth of the probes is 6 GHz with a built-in 30 dB amplifier. The signal is sampled synchronously to the device's clock at 2.5 GS/s. Contrary to other contributions [8, 18] no simple compression or pre-proccessing (e.g. averaging, maximum extraction or sum-of-squares during clock cycles) is applied before Principal Component Analysis (PCA) and clustering. Such simple trace pre-processing techniques have been shown to have negative effects on results [10].

4.2 Quality of Principal Components

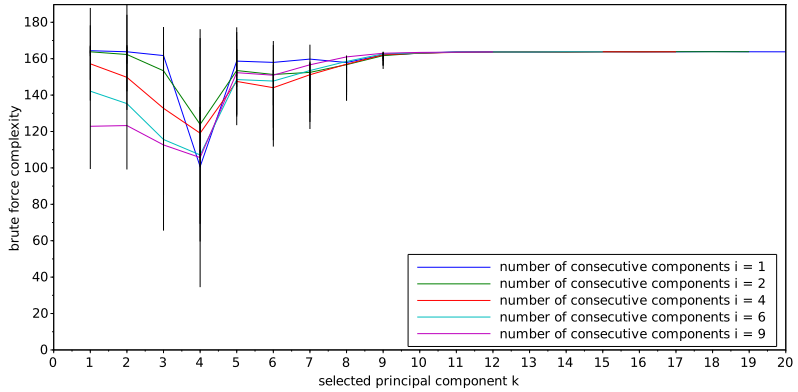

Fig. 2: Mean brute-force complexity for different selected principal components (k and i) over all measurement positions *including standard deviation as bars*

Our algorithmic approach includes the selection of principal components after PCA as a first step before clustering. The selection can be described by two parameters, k the first selected component, and, $i \geq 1$ the number of consecutively selected components after the k-th one as described in Sect. 3.1. In this section we investigate the quality of different parameter choices. We executed

the clustering-based attack on *every* single measurement from all 3 probes and 400 positions with choices of $k \in [1, 20]$ and $i \in \{1, 2, 4, 6, 9\}$ and assess the quality using the remaining brute-force complexity explained in Sect. 3.3.
We show the means over $3 * 400$ results for the resulting brute-force complexities for each combination of parameters k and i in Fig. 2. Hence we are able to equally compare the results of different probes and show some fundamental properties of our measurements. These high *mean* brute-force complexities of > 100 bits are certainly not within the range of realistic computing capabilities. They result from including many low-scoring results. The standard deviations are shown as vertical bars and indicate that there are multiple results with significantly lower brute-force complexities (the diagram does not include $+1$ bits for assigning labels to classes). As an important observation, low-ranked components ($k < 10$) seem preferable overall and *first-ranked principal components do not contain exploitable leakage* (see curve with $i = 1$ or $i = 2$ in Fig. 2). This confirms our assumptions from Sect. 3.1 as well as similar observations from Batina et al. [3]. Thus, we *discard first-ranked as well as low-ranked principal components* before further analysis and achieve significantly improved brute-force complexities.

In Fig. 2 it can also be noted that the component number $k = 4$ seems to contain the most leakage on average, reaching the lowest mean brute-force complexities. It seems that PCA *concentrates most of the exploitable leakage information into a single principal component.* This means that a choice of $i = 1$ for the number of selected consecutive principal components led to the best results in our circumstances. We used this choice in the practical evaluation in the next Sect. 4.3. As another observation, curves with $i > 2$ lead to low complexities as soon as component 4 is included in the consecutively selected components. For illustrative purposes, we show the resulting principal components after PCA transformation of an examples trace in the Appendix 6.1.

4.3 Analyzing Separate Channels

For every probe, we have 400 measurements from different positions. We analyze the data from the three available channels separately: Firstly we perform pre-processing by applying PCA, secondly we perform clustering using the expectation maximization algorithm and thirdly we compute the remaining brute-force complexity. For every probe separately, and for every selection of principal components (for every $k \in [1, 20]$ while $i = 1$), we summarize the results from 400 tests in Figures 3a, 3b, and 3c. Figure 3a shows results for the 250 µm coil probe, Fig. 3b for the 150 µm coil probe and Fig. 3c for the 100 µm coil probe. The figures show, *how many of the* 400 *measurements of each probe*, and for every selection of k, lead to which brute-force complexities. The occurrence rate is visually indicated by the size of the respective dots. Bigger dots mean that the corresponding brute-force complexity has occurred more often. For example, in Figure 3a, almost all of the 400 measurements lead to a maximum brute-force complexity of 163 for $k < 5$ and $k > 10$. For $k = 5$, however, many measurements lead to lower resulting brute-force complexities, some even of the minimum. The

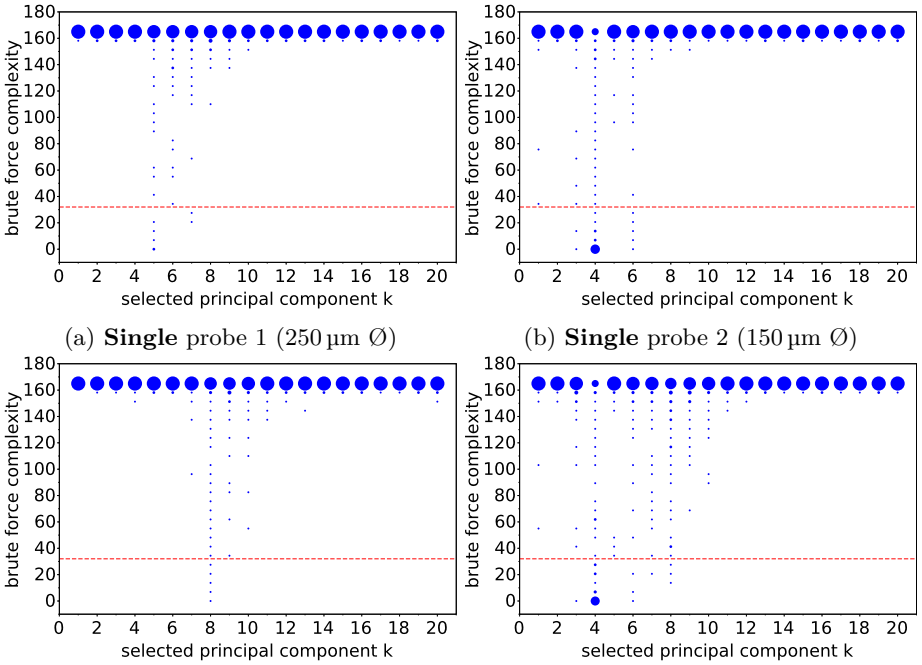

(a) **Single** probe 1 (250 µm Ø)

(b) **Single** probe 2 (150 µm Ø)

(c) **Single** probe 3 (100 µm Ø)

(d) **Combined** probes

Fig. 3: Brute force complexity occurrences over different principal components

red dashed line highlights the 32 bit complexity level up to which all outcomes are *easily manageable* for attackers through computation.

As an important finding, it can be observed, that the *probe* with the 150 µm coil diameter depicted in Fig. 3b leads to the *best results* by far. For the principal component $k = 4$, *an astonishing percentage of* 56 % out of the 400 measurements led to a remaining brute-force complexity ≤ 32 bit (summing up all outcomes equal or lower the red dashed line). This high number was unexpected and means that *with the improved algorithmic, more than half of all measurement positions exhibited sufficient leakage* for a complete break. The 100 µm probe depicted in Fig. 3a leads to only 3 % ≤ 32 bit for $k = 5$. Also the 250 µm probe depicted in Fig. 3c only leads to 3 % ≤ 32 bit for $k = 8$. Hence, the 150 µm coil probe seems to work best under our circumstances. Since finding suitable measurement positions is rather easy (using the best probe), attackers should test different measurement positions instead of employing extensive computational brute-force, testing is comparably easy in case of single-execution attacks because only single measurements need to be analyzed at every position.

Without knowing $k = 4$ and $i = 1$ a priori, attackers could make minimal heuristic assumptions like $k \in [3, 10]$ and $i \in [1, 4, 9]$ which could fit similar circumstances. This would result in an additional brute-force complexity of +4 bits which is not included in Fig. 3 and justified by significantly improved results.

To demonstrate the improvement of our proposal, we performed the *original attack of Heyszl et al. [8]* on the *same measurements*. A remaining brute-force complexity ≤ 32 bit is reached in *none* (0 %) of the 400 measurement cases using the 150 µm coil probe. Compared to 56 % from the improved attack, this means that we achieve *astonishingly improved results* from applying PCA and expectation maximization clustering. (Only the 250 µm coil probe led to marginally better results using the previous method, i.e., 8 % instead of 3 % of the cases ≤ 32 bit, however, this does not invalidate the previous statement in our opinion.)

We compared the performance of the k-means versus the expectation maximization clustering algorithm in the context of single channels. Since we only select single components ($i = 1$) after PCA, channels only consist of single dimensions and there is not much benefit from more free parameters in the clustering algorithm. This is confirmed by the fact that expectation maximization and k-means clustering *lead to almost equal results*. This means that our reported *improvement is mainly due to the PCA transformation and the selection of components*. In the multi-channel case, however, more dimensions aggregate from separate channels making expectation maximization more eligible.

As *benchmark* for *high-resolution magnetic field measurements*, we tested the improved *non-profiled* attack on a *current consumption measurement*. We use a 1 Ohm measurement resistor and a differential probe at unchanged sampling rate. To cancel one-time effects such as disturbances or noise, we repeated this 12 times and averaged the results. The outcome is a *significantly high brute-force complexity of* 152 *bits*. Hence, it is completely impossible to exploit leakage from such current measurements. This underlines that high-resolution magnetic field measurements are clearly *superior* in leakage signal quality in our circumstances.

4.4 Combining Multiple Channels

After the individual analysis of the three measurement channels, we combined the channels for analysis as described in Sect. 3.2. The motivation for attackers to combine channels is to increase the exploitable leakage to improve attack outcomes, e.g. instead of trying to find better measurement positions.

Figure 3d shows the brute-force complexity results for the combined measurements in the same way as described in the previous Sect. 4.3. A *visual comparison* of the combined results in Fig. 3d to the individual results in Figures 3a, 3b, and 3c gives the impression,that the overall result is comparable to Fig. 3b. However, expressed quantitatively in the same way as before, the combined channels lead to a remaining brute-force complexity of ≤ 32 bit in only 52 % of the cases for $k = 4$. Hence, as an important result, *instead of an improvement*, we observe a *slight degradation* compared to the best individual case which led to 56 % of cases ≤ 32 bit. This means that the described clustering-based *non-profiled* attack is *unable to benefit* from a combination of channels (in our circumstances).

We suspect that this is due to the fact that our selection strategy selects equal values k and i to pre-process all three channels in the same way using PCA in case of combined attacks. This should be a significant disadvantage in our case where different k are best for different channels (see results in Sect. 4.3).

Unfortunately, it would significantly increase the complexity to test different k-s for every channel (e.g. repeat clustering 20^3 times). Increasing the number of selected components i to prevent this would include more noise, in our circumstances, which in turn would degrade classification results significantly (see how curves with $i > 1$ result in higher mean-values in Fig. 2).

We compared the improved *non-profiled* attack against a *profiled* template attack. This requires one additional trace for profiling at each position and for every probe. Templates consist of two means and a single full covariance matrix. To derive the remaining brute-force complexity as described in Sect. 3.3, we use bit-wise template matching results. For a fair comparison, we also apply PCA including the selection strategy for k and i. A *higher number of* 61 % of positions (compared to the 56 % from the non-profiled attack) lead to remaining brute-force complexities $\leq$ 32 bit for the 150 µm coil probe, with $i = 1$ and $k = 4$.The *profiled* template attack outperforms the *non-profiled* attack. As the most important observation, we find that the *combination of channels leads to an improved* 66 % of the cases with a remaining brute-force complexity of ≤ 32 bit, with $i = 9$ and $k = 3$. This clearly demonstrates *the gain of combining channels in the profiled setting.*

In a *profiled setting*, attackers are able to test and find the 'best' measurement positions. This means that, in our circumstances, the use of multiple channels is only reasonable if the leakage of such best single channels is insufficient which diminishes the good results to a certain extent.

5 Conclusion

We significantly improved the algorithmic approach for *non-profiled* attacks against exponentiation by applying Principal Component Analysis (PCA) and disregarding high- as well as low-ranked ones following a simple strategy. This selection strategy requires some trying-out (additional brute-force), but this is highly rewarded by improved attack outcomes in terms of low brute-force complexities. With this approach, the unsupervised attack using a single-channel high-resolution magnetic field measurement is remarkably threatening and leads to manageable brute-force levels in over half of the tested measurement positions. This emphasizes the need to prevent *all possible cause for exploitable single-execution leakage.* Regarding our results from three simultaneous channels, we find that the combination of channels only significantly improves the attack results, if a *profiled* attack is used. In case of the *clustering-based, non profiled* attack, the results from the combination are only comparable to the best individual one. In profiled settings attackers are also able to look for the 'best' measurement positions. Hence, multi-channel attacks are only reasonable if the exploitable leakage is insufficient at such best positions.

Acknowledgements This work was partly funded by the German Federal Ministry of Education and Research in the project SIBASE through grant number 01IS13020.

References

1. Agrawal, D., Rao, J., Rohatgi, P.: Multi-channel attacks. In: Walter, C., Koç, Ç., Paar, C. (eds.) Cryptographic Hardware and Embedded Systems - CHES 2003. Lecture Notes in Computer Science, vol. 2779, pp. 2–16. Springer Berlin / Heidelberg (2003)
2. Archambeau, C., Peeters, E., Standaert, F.X., Quisquater, J.J.: Template attacks in principal subspaces. In: Cryptographic Hardware and Embedded Systems-CHES 2006, pp. 1–14. Springer (2006)
3. Batina, L., Hogenboom, J., Woudenberg, J.: Getting more from PCA: First results of using principal component analysis for extensive power analysis. In: Dunkelman, O. (ed.) Topics in Cryptology - CT-RSA 2012, Lecture Notes in Computer Science, vol. 7178, pp. 383–397. Springer Berlin Heidelberg (2012)
4. Bauer, S.: Attacking exponent blinding in RSA without CRT. In: Schindler, W., Huss, S. (eds.) Constructive Side-Channel Analysis and Secure Design, Lecture Notes in Computer Science, vol. 7275, pp. 82–88. Springer Berlin / Heidelberg (2012)
5. Bohy, L., Neve, M., Samyde, D., Quisquater, J.J.: Principal and independent component analysis for crypto-systems with hardware unmasked units. In: Proceedings of e-Smart (2003)
6. Clavier, C., Feix, B., Gagnerot, G., Roussellet, M., Verneuil, V.: Horizontal correlation analysis on exponentiation. In: Soriano, M., Qing, S., López, J. (eds.) Information and Communications Security, Lecture Notes in Computer Science, vol. 6476, pp. 46–61. Springer Berlin Heidelberg (2010)
7. De Mulder, E., Örs, S.B., Preneel, B., Verbauwhede, I.: Differential power and electromagnetic attacks on a FPGA implementation of elliptic curve cryptosystems. Comput. Electr. Eng. 33, 367–382 (Sep 2007)
8. Heyszl, J., Ibing, A., Mangard, S., De Santis, F., Sigl, G.: Clustering algorithms for non-profiled single-execution attacks on exponentiations. In: Francillon, A., Rohatgi, P. (eds.) Smart Card Research and Advanced Applications, pp. 79–93. Lecture Notes in Computer Science, Springer International Publishing (2014)
9. 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)
10. Heyszl, J., Merli, D., Heinz, B., De Santis, F., Sigl, G.: Strengths and limitations of high-resolution electromagnetic field measurements for side-channel analysis. In: Mangard, S. (ed.) Smart Card Research and Advanced Applications. Lecture Notes in Computer Science, Springer Berlin Heidelberg (2012)
11. Homma, N., Hayashi, Y.i., Miura, N., Fujimoto, D., Tanaka, D., Nagata, M., Aoki, T.: Em attack is non-invasive?-design methodology and validity verification of em attack sensor. In: Cryptographic Hardware and Embedded Systems–CHES 2014, pp. 1–16. Springer (2014)
12. Itoh, K., Izu, T., Takenaka, M.: Address-bit differential power analysis of cryptographic schemes OK-ECDH and OK-ECDSA. In: Cryptographic Hardware and Embedded Systems - CHES 2002. Lecture Notes in Computer Science, vol. 2523, pp. 399–412. Springer Berlin / Heidelberg (2003)
13. Kocher, P.C.: Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In: Proceedings of the 16th Annual International Cryptology Conference on Advances in Cryptology. pp. 104–113. CRYPTO '96, Springer-Verlag, London, UK (1996)

14. Mangard, S., Oswald, E., Popp, T.: Power Analysis Attacks: Revealing the Secrets of Smart Cards (Advances in Information Security). Springer-Verlag New York, Inc., Secaucus, NJ, USA (2007)
15. Mavroeidis, D., Batina, L., Van Laarhoven, T., Marchiori, E.: Pca, eigenvector localization and clustering for side-channel attacks on cryptographic hardware devices. In: Machine Learning and Knowledge Discovery in Databases, pp. 253–268. Springer (2012)
16. Messerges, T., Dabbish, E., Sloan, R.: Power analysis attacks of modular exponentiation in smartcards. In: Cryptographic Hardware and Embedded Systems. Lecture Notes in Computer Science, vol. 1717, pp. 724–724. Springer Berlin / Heidelberg (1999)
17. Peeters, E., Standaert, F.X., Quisquater, J.J.: Power and electromagnetic analysis: improved model, consequences and comparisons. Integr. VLSI J. 40(1), 52–60 (Jan 2007)
18. Perin, G., Imbert, L., Torres, L., Maurine, P.: Attacking randomized exponentiations using unsupervised learning. In: Constructive Side-Channel Analysis and Secure Design - 5th International Workshop, COSADE 2014, Paris, France, April 13-15, 2014. Revised Selected Papers. pp. 144–160 (2014)
19. Rousseeuw, P.J.: Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. Journal of computational and applied mathematics 20, 53–65 (1987)
20. Sauvage, L., Guilley, S., Mathieu, Y.: Electromagnetic radiations of FPGAs: High spatial resolution cartography and attack on a cryptographic module. ACM Trans. Reconfigurable Technol. Syst. 2, 4:1–4:24 (Mar 2009)
21. Souissi, Y., Bhasin, S., Guilley, S., Nassar, M., Danger, J.L.: Towards different flavors of combined side channel attacks. In: Dunkelman, O. (ed.) Topics in Cryptology - CT-RSA 2012. Lecture Notes in Computer Science, vol. 7178, pp. 245–259. Springer Berlin / Heidelberg (2012)
22. Standaert, F.X., Archambeau, C.: Using subspace-based template attacks to compare and combine power and electromagnetic information leakages. In: Oswald, E., Rohatgi, P. (eds.) Cryptographic Hardware and Embedded Systems - CHES 2008. Lecture Notes in Computer Science, vol. 5154, pp. 411–425. Springer Berlin / Heidelberg (2008)
23. Veyrat-Charvillon, N., Gérard, B., Renauld, M., Standaert, F.X.: An optimal key enumeration algorithm and its application to side-channel attacks. In: Selected Areas in Cryptography. pp. 390–406. Springer (2013)
24. Veyrat-Charvillon, N., Gérard, B., Standaert, F.X.: Security evaluations beyond computing power. In: EUROCRYPT. vol. 7881, pp. 126–141. Springer (2013)
25. Walter, C.: Sliding windows succumbs to big mac attack. In: Koç, Ç., Naccache, D., Paar, C. (eds.) Cryptographic Hardware and Embedded Systems - CHES 2001. Lecture Notes in Computer Science, vol. 2162, pp. 286–299. Springer Berlin / Heidelberg (2001)
26. Witteman, M., van Woudenberg, J., Menarini, F.: Defeating RSA multiply-always and message blinding countermeasures. In: Kiayias, A. (ed.) Topics in Cryptology - CT-RSA 2011. Lecture Notes in Computer Science, vol. 6558, pp. 77–88. Springer Berlin / Heidelberg (2011)
27. Wolpert, D.H., Macready, W.G.: No free lunch theorems for optimization. Evolutionary Computation, IEEE Transactions on 1(1), 67–82 (1997)

6 Appendix

6.1 Illustration of Principal Components after Transformation

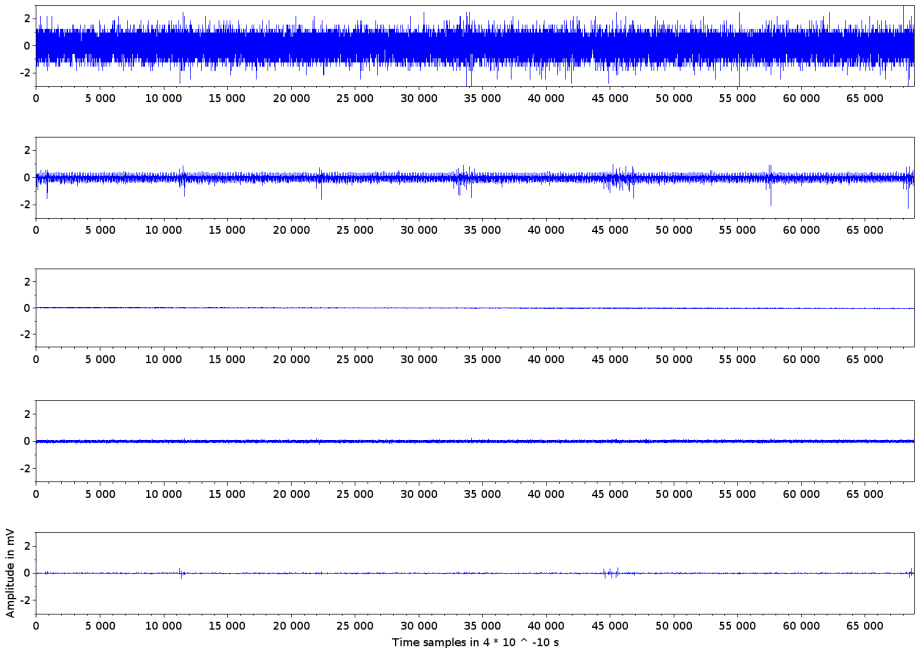

Fig. 4: Example of an original trace-segment (topmost) and its high-ranked principal components below. The 4-th (bottom) component contains signal leakage

Figure 4 depicts principal components after Principal Component Analysis (PCA) transformation for illustrative purposes. We used an example measurement where the side-channel leakage is sufficient for the attack to succeed without false classifications when selecting the $k = 4$-th component for expectation maximization clustering. The topmost diagram depicts one trace-segment in its original form. Below this, the four highest-ranked principal components of this segment are depicted. From the previous analysis we know that the exploitable leakage seems concentrated in component $k = 4$ which is depicted in the bottom diagram. The time-samples with higher values represent the times of exploitable leakage information in this component. The sparse occurrence fits to the description of data-dependent register accesses as source of this leakage [9]. A comparison to the other components in Fig. 4 clearly shows that the *leakage is small compared to the remaining signal parts.*

6.2 Countermeasures

As previously described by Heyszl et al. [8], countermeasures such as exponent blinding do not protect against *non-profiled* attacks. Many countermeasures address individual single-execution leakage sources of implementation (e.g. address-bit, or localized leakage).

As a conclusion from this contribution, we must emphasize the necessity to reduce all possible single-execution leakage sources as much as possible.

Homma et al. [11] present a general countermeasure against high-resolution magnetic field measurements. They describe an on-chip sensor which detects magnetic field probes in close distance to die surfaces. However, in our opinion this will not help since measurement probes are typically placed close to an integrated circuit before power-up. Hence, necessary calibration routines of the sensor will likely not be able to distinguish the static probes from other environmental influences.