

Behind the Scene of Side Channel Attacks

Victor Lomné, Emmanuel Prouff, and Thomas Roche

ANSSI

51, Bd de la Tour-Maubourg, 75700 Paris 07 SP, France

`firstname.lastname@ssi.gouv.fr`

Abstract. Since the introduction of side channel attacks in the nineties, a large amount of work has been devoted to their effectiveness and efficiency improvements. On the one side, general results and conclusions are drawn in theoretical frameworks, but the latter ones are often set in a too ideal context to capture the full complexity of an attack performed in real conditions. On the other side, practical improvements are proposed for specific contexts but the big picture is often put aside, which makes them difficult to adapt to different contexts. This paper tries to bridge the gap between both worlds. We specifically investigate which kind of issues is faced by a security evaluator when performing a state of the art attack. This analysis leads us to focus on the very common situation where the exact time of the sensitive processing is drown in a large number of leakage points. In this context we propose new ideas to improve the effectiveness and/or efficiency of the three considered attacks. In the particular case of stochastic attacks, we show that the existing literature, essentially developed under the assumption that the exact sensitive time is known, cannot be directly applied when the latter assumption is relaxed. To deal with this issue, we propose an improvement which makes stochastic attack a real alternative to the classical correlation power analysis. Our study is illustrated by various attack experiments performed on several copies of three micro-controllers with different CMOS technologies (respectively 350, 130 and 90 nanometers).

1 Introduction

Since the seminal differential power analysis of Kocher *et al.* [17], various side channel Attacks (SCA) have been proposed and improved (*e.g.* [8, 9, 11, 12, 31]). In order to compare and classify them, **theoretical** frameworks have then been introduced [11, 22, 35, 39]. Their main purpose is to identify the attacks similarities and differences, and to exhibit contexts where one is better than another. They have laid the foundation stones for a general comparison and evaluation framework. In parallel, several **practical** works have addressed issues arising when applying an SCA in the real world (*e.g.* in an industrial context) [2, 5, 16, 24, 37]. Those works essentially attempt to fill the gap between the theoretical analysis of the attacks and their application in non-idealized contexts. However, whereas the published theoretical analyses usually tend towards generic and formal statements (sometimes at the cost of too simple models), many of

the practical analyses only focus on a particular attack specificity and often put the big picture aside. The latter analyses are indeed usually dedicated to one specific attack running against a specific target device, which makes them hard to generalize. This paper tries to be at the intersection of both worlds: we study practice-driven issues while keeping a generic approach w.r.t. attacks mechanisms and targeted platforms. This approach and our final purpose are close to those in the works of Standaert *et al.* [33] and Renaud *et al.* [29].

The starting observation of our study is that side channel traces are never reduced to one point in practice, even when they rely on the manipulation of a single variable. In contrary, those traces are often composed of a large number of points (typically several thousands). In spite of the evidence of this observation, it is rarely taken into account when analysing the effectiveness of a side channel attack. Such an analysis is indeed frequently done under the assumption, sometimes implicit, that a small number of points of interest (POI) has already been extracted from the traces either by pattern matching [21], or by dimension reduction [1, 6, 7, 32] or thanks to a previous successful attack [10, 29]. However, the two first categories of techniques are not yet perfect and, after reduction, the traces are often still composed of several points in practice. And, what is more important, the risk of information loss during the reduction process leads most of attack practitioners to not apply them. The third technique (performing a first attack to identify the POI) allows for interesting analyses, but it does not correspond to a *real attack* context. Moreover, the best POI for one attack type may not be so good for another one. Eventually, we come to a situation where attacks are analysed in a (uni-dimensional) context which does not fit with the (multi-dimensional) reality faced by the attack practitioners.

We argue in this paper that the state-of-the-art uni-dimensional analyses cannot be straightforwardly adapted to multi-dimensional contexts, which raises new interesting issues. The selection of the most likely candidate among the results of several instantaneous attacks is one of them. Indeed when the leakage traces are composed of several points, a side channel attack against the targeted sensitive variable must be performed for each point (*a.k.a* time index) in the traces. Then, the adversary must apply a strategy to select the most likely candidate among the different *instantaneous attacks* results. A classical method is to select the one with the highest score (*e.g.* the highest correlation coefficient in a Correlation Power Analysis – CPA – [8]). Nevertheless we argue that this strategy can be ineffective for some attack categories, including the case of the Linear Regression Analysis (LRA) [11, 30, 31]. For the latter one, we propose a new strategy to select the most likely candidate and we demonstrate its effectiveness in practice.

Another interesting issue when dealing with a large number of high dimensional traces is the reduction of the computational complexity. Here again, some works have investigated the use of parallel computing to decrease the data processing time [4, 19] but their goal was not to diminish the algorithmic complexity of the attacks. This work studies the LRA¹ and the Template Attacks (TA) with

¹ In this paper we only consider the unprofiled version of LRA [11].

this goal in mind. A common structure in their algorithmic description is exhibited and then used to propose a new general *modus operandi* which enables to significantly reduce the computation time when the number of traces is non-negligible. The strategy can also be applied to other attack (*e.g.* the CPA).

Finally, to make sure that our analysis is consistent with the reality, we completed our investigations by several experiments performed on three micro-controllers based on different CMOS technologies (350, 130 and 90 nanometers process). We report here on these experiments results. We moreover use them to confirm and complete the interesting behaviours observed in [29]: (1) the leakage seems to diverge from the classical Hamming weight model as the CMOS technology tends to the nanometer scale, which makes LRA a promising tool for side channel evaluations of nano-scale devices and (2) TA is effective in practice, even when the *templates* are built on one copy of the device and the attack is done on another copy.

The paper is organized as follows. In Section 2, we introduce the theoretical background for our study and we present the outlines of our proposal. Then, two sections are dedicated to the application of our ideas to the LRA and TA attacks respectively².

2 SCA: Practical Issues

In this section, we introduce some basics and we get into the specifics of the problematic focussed in this paper.

2.1 Notations

Throughout this paper, random variables are denoted by large letters. A realization of a random variable, said X , is denoted by the corresponding lower-case letter, said x . A *sample* of several observations of X is denoted by $(x_i)_i$. It will sometimes be viewed as a vector defined over the definition set of X . The notation $(x_i)_i \leftarrow X$ denotes the instantiation of the set of observations $(x_i)_i$ from X . The *mean* of X is denoted $\mathbb{E}[X]$, its standard deviation by $\sigma[X]$ and its *variance* by $\text{var}[X]$. The latter equals $\mathbb{E}[(X - \mathbb{E}[X])^2]$. The *covariance* of two random variables X and Y is denoted by $\text{cov}(X, Y)$ and satisfies $\text{cov}(X, Y) = \mathbb{E}[(X - \mathbb{E}[X])(Y - \mathbb{E}[Y])]$. When we will need to specify the variable on which statistics are computed, we will write the variable in subscript (*e.g.* $\mathbb{E}_X[Y]$ instead of $\mathbb{E}[Y]$).

The notation \vec{X} will be used to denote column vectors and $\vec{X}[u]$ will denote its u^{th} coordinate. Calligraphic letters will be used to denote a matrix. The elements of a matrix \mathcal{M} will be denoted by $\mathcal{M}[i][j]$. Classical additions and multiplications (over real values, vectors or matrices) are denoted by $+$ and \times respectively. Scalar-vector operations are denoted by \cdot and $/$ (all the coordinates

² This work is completed in the extended version of this paper with a similar study on CPA.

of the vector are multiplied, respectively divided, by the scalar). When applied to vectors or matrices, the symbols \cdot^2 and $\sqrt{\cdot}$ denote the operation consisting in computing the square (resp. the square root) of all the vector/matrix coordinates. Eventually, a function from \mathbb{F}_2^n to \mathbb{F}_2^m will be called a (n, m) -function.

2.2 General Attacks Framework

In this paper, the attacks framework is described by considering that the adversary targets the manipulation of a single sensitive variable Z , but the study and results directly extend to contexts where several variables are targeted in parallel. The variable Z is supposed to functionally depend on a public variable X and a secret sub-part k such that $Z = F(X, k)$ where F is a $(n + n, m)$ -function (which implies $X, k \in \mathbb{F}_2^n$ and $Z \in \mathbb{F}_2^m$). The bit-lengths n and m depend on the cryptographic algorithm and the device architecture³.

The attacks are described under the assumption that the adversary owns N side channel traces $\vec{\ell}_0, \dots, \vec{\ell}_{N-1}$, each of them containing information about Z . Namely, the i^{th} leakage trace $\vec{\ell}_i \leftrightarrow \vec{\mathbf{L}}$ corresponds to the processing of a public value $x_i \leftrightarrow X$ and contains information on the value $z_i \leftrightarrow Z$ such that $z_i = F(x_i, k)$. The dimension of the traces (*i.e.* the number of different instantaneous leakage points) is denoted by d . By definition, we have $d \doteq \dim \vec{\mathbf{L}}$.

When little information is known about the implementation and the device (which is usually the case in practice), the exact manipulation time of z_i cannot be precisely determined *a priori*. Also, precision in the observation often comes at the cost of a high sampling rate⁴. As a consequence, the dimension of the traces is usually high (from several thousand of points up to millions) and the attack must be repeated on all of their coordinates independently (as *e.g.* in LRA) or must consider huge traces chunks globally (as *e.g.* for TA). Although bearing differences, most of side-channel attacks (including LRA and TA) follow a common process flow. Starting from this generic description, this paper studies, in Sections 3 and 4 respectively, the *effectiveness* and *efficiency* of the LRA and TA attacks. The core ideas of those analyses are presented in the two next subsections.

2.3 Effectiveness Discussions

A part of our study is dedicated to the *distinguisher value* definition and, more precisely its relevance when considering side channel traces with a large number of points. This study was motivated by the observation that the classical LRA distinguisher value for one leakage time is not comparable as such to that computed for another leakage time. Figure 1(a) illustrates this claim for an LRA targeting the device B described in Section 2.5: when directly applying the protocol given in [11, 31], the correct key candidate does not maximize the distinguisher value

³ An example of function F is the function that applies a so-called *sbox* transformation to the bitwise addition between k and X .

⁴ Especially in the case of Electro-Magnetic side channel measurements




Fig. 1. Instantaneous LRA scores computed over 10000 traces (scores for the correct key in black)

globally but only in a local area, which makes the attack unsuccessful unless this area is known by the adversary (which is not assumed here). This observation led us to study the handling of distinguishing values in SCA attacks. We for instance show that by normalizing the LRA distinguishing values, the correct key candidate becomes clearly distinguishable even when considering the full vector of instantaneous attack results (as depicted on Figure 1(b)).

More generally, our study relies on a well studied problem which is the comparison of the results of two different instantaneous attacks [11, 20, 33, 34, 36, 38]. For the LRA, it will lead to a modification of the candidate selection rule.

2.4 Algorithmic Complexity Improvements Proposals

The other important issue an evaluator faces when performing SCA, is the computational complexity of the attack when the number of traces N and their dimension grows to millions. Indeed, the execution time of naive attack implementations can easily reach several days of processing and this is not compatible with standard evaluation processes⁵.

We show in Sections 3 and 4 that the two considered attacks may be rewritten in a *partitioning* fashion that can be exploited to significantly decrease the algorithmic complexity. Roughly speaking, the basic idea is to lower the impact of the heavy computations so that its complexity does no longer depend on the traces number N but on the dimension n of the targeted data. To that purpose, we propose to modify the attack first step so that it processes separately the traces with respect to their input value x_i . As a result, the algorithmic complexity of the attacks is divided by $\frac{N}{2^n}$ making the algorithmic improvement interesting when $N \gg 2^n$ (which is often the case in practice).

⁵ In Common Criteria evaluations applied on hardware security devices, all penetration tests (including invasive and non-invasive attacks) have usually to be performed in 3 months, leaving only few weeks for the whole side channel evaluation.

Almost every SCA may be rewritten as a combination of tests on statistics estimated on leakage partitions. Some of them (*e.g.* the DPA [17] or the multi-bit DPA [23]) were actually originally written as such, whereas the other ones were developed in a partitioning way after their introduction (see *e.g.* [18] for the CPA, [10] for the LRA and [33] for the MIA). To the best of our knowledge, this property has however never been exploited to improve the attacks efficiency.

2.5 Experimental Setup

For each studied SCA, practical experiments were performed on three Micro-Controller Units (MCUs for short) with different CMOS technologies (350, 130 and 90 nanometers processes). The observed processing was that of an AES128 encryption handling one byte at a time. Each attack was performed against 4 sbx outputs of the first round. Furthermore, to measure the variability of our experiments, we used three different copies for each MCU (called copy 1, 2 and 3 in the sequel). This choice enabled us to perform the TA profiling step on one copy and to use the results to attack other ones. Also, it gives more credit to our experimental results as the templates consistency was checked on three different versions of the same MCU.

The side channel observations were obtained by measuring the electromagnetic (EM) radiations emitted by the device. To this aim, several sensors were used, all made of several coils of copper (the diameters of the coils were respectively of 1mm, 500 μ m and 250 μ m for the 350, 130 and 90nm MCUs), and were plugged into a low-noise amplifier. To sample measurements, a digital oscilloscope was used with a sampling rate of 1G samples per second for the 350nm MCU and 10G samples per second for the others, whereas the MCUs were running at few dozen of MHz.

We insist on the fact that the temporal acquisition window was set to record the first round of the AES only. This synchronization has been done thanks to *simple electromagnetic analysis* [26]. As the MCU clocks were not stable, we had to resynchronize the measurements. This process is out of the scope of this work, but we emphasize that it is always needed in a practical context and it impacts the measurements noise.

We sum-up the specificities of the three experimental campaigns hereafter:

- **Device A** (3 copies): 90nm CMOS technology with MCU based on a 8-bit 8051 architecture. EM traces composed of 12800 points each after resynchronization. Highest Signal to Noise Ratio (SNR) over the full traces equals to 0.09.
- **Device B** (3 copies): 130nm CMOS technology with MCU based on a 8-bit 8051 architecture. EM traces composed of 16800 points each after resynchronization. Highest SNR equals to 0.6.
- **Device C** (3 copies): 350nm CMOS technology with MCU based on a 8-bit AVR architecture. EM traces composed of 51600 points each after resynchronization. Highest SNR equals to 0.3.

3 Practical Evaluation of Linear Regression Attacks

Linear regression attacks (a.k.a. stochastic attacks) have been introduced by Schindler *et al.* in 2005 [31]. Initially, they were presented with a profiling step and were viewed as an alternative to the template attacks [13]. In [11], the authors have shown how to express the linear regression attacks such that the profiling stage is no longer required. They also argued that the LRA can be applied in the same context as the CPA, but with weaker assumption on the device behavior. Subsequently, these results of Doget *et al.* have been extended in [10] to apply against masked implementations. In parallel, linear regression attacks have been used to analyse/model the deterministic part of the information leakage for complex circuits [14, 15]. As a matter of fact, all those analyses assume that the side-channel traces are composed of a single leakage point: the issue raised in Section 2.3 is thus put aside. Moreover, the question of the efficient processing of the attack, when applied against high dimensional leakage traces, is not tackled. The rest of this section aims at dealing with two issues.

3.1 Attack Description

In LRA, the adversary chooses a so-called *basis* of functions⁶ $(\mathbf{m}_p)_{1 \leq p \leq s}$ with the only condition that \mathbf{m}_1 is a constant function (usually $\mathbf{m}_1 = 1$). Then, for each x_i and each sub-key hypothesis \hat{k} , the prediction $\hat{z}_i = F(x_i, \hat{k})$ is calculated. The basis functions \mathbf{m}_p are then applied to the \hat{z}_i independently, leading to the construction of a $(N \times s)$ -matrix $\mathcal{M}_{\hat{k}} \doteq (\mathbf{m}_p(F(x_i, \hat{k})))_{i,p}$. The comparison of this matrix with the set of d -dimensional leakages $(\vec{\ell}_i)_{i \leq N} \leftarrow \vec{\mathbf{L}}$ is done by processing a linear regression of each coordinate of $\vec{\ell}_i$ in the basis formed by the row elements of $\mathcal{M}_{\hat{k}}$. Namely, a real-valued $(s \times d)$ -matrix $\mathcal{B}_{\hat{k}}$ with column vectors $\vec{\beta}_1, \dots, \vec{\beta}_d$ is estimated in order to minimize the error when approximating $\vec{\ell}_i^\top$ by $(\mathbf{m}_1(F(x_i, \hat{k})), \dots, \mathbf{m}_s(F(x_i, \hat{k}))) \times \mathcal{B}_{\hat{k}}$. The matrix $\mathcal{B}_{\hat{k}}$ is defined such that:

$$\mathcal{B}_{\hat{k}} = \underbrace{(\mathcal{M}_{\hat{k}}^\top \times \mathcal{M}_{\hat{k}})^{-1} \times \mathcal{M}_{\hat{k}}^\top}_{\mathcal{P}_{\hat{k}}} \times \mathcal{L}, \quad (1)$$

where \mathcal{L} denotes the $(N \times d)$ -matrix with the $\vec{\ell}_i^\top$ as row vectors. In the following, the u^{th} column vector of \mathcal{L} (composed of the u^{th} coordinate of all the $\vec{\ell}_i$) is denoted by $\vec{\mathcal{L}}[u]$. Moreover, the $(s \times N)$ -matrix $(\mathcal{M}_{\hat{k}}^\top \times \mathcal{M}_{\hat{k}})^{-1} \times \mathcal{M}_{\hat{k}}^\top$, which does not depend on the leakage values, is denoted by $\mathcal{P}_{\hat{k}}$.

To quantify the estimation error, the *goodness of fit model* is used and the *correlation coefficient of determination* \mathcal{R}^2 is computed for each u . The latter is defined by $\mathcal{R}^2 = 1 - \text{SSR}/\text{SST}$, where SSR and SST respectively denote the *residual sum of squares* (deduced from $\mathcal{B}_{\hat{k}}$) and the *total sum of squares*⁷ (deduced from \mathcal{L}). We give in Algorithm 1 the pseudo-code corresponding to a classical LRA attack processing.

⁶ The basis choice and its impact are not a trivial matter, see [10] for a detailed study.

⁷ For their exact definitions, see their construction in Alg. 1

Algorithm 1: LRA - Linear Regression Analysis

```

Input : a set of  $d$ -dimensional leakages  $(\vec{\ell}_i)_{i \leq N}$  and the corresponding plaintexts
           $(x_i)_{i \leq N}$ , a set of model functions  $(m_p)_{p \leq s}$ 
Output: A candidate sub-key  $\hat{k}$ 

/* Processing of the leakage Total Sum of Squares (SST) */
1 for  $i = 0$  to  $N - 1$  do */
2    $\mu_{\vec{L}} = \mu_{\vec{L}} + \vec{\ell}_i$ 
3    $\sigma_{\vec{L}} = \sigma_{\vec{L}} + \vec{\ell}_i^2$ 
4  $\overrightarrow{SST} = \sigma_{\vec{L}} - 1/N \cdot \mu_{\vec{L}}^2$ 

/* Processing of the  $2^n$  predictions matrices  $\mathcal{M}_{\hat{k}}$  and  $\mathcal{P}_{\hat{k}}$  */
5 for  $\hat{k} = 0$  to  $2^n - 1$  do */
6   /* Construct the matrix  $\mathcal{M}_{\hat{k}}$  and  $\mathcal{P}_{\hat{k}}$  */ */
7   for  $p = 1$  to  $s$  do */
8     for  $i = 0$  to  $N - 1$  do */
9        $\mathcal{M}_{\hat{k}}[i][p] \leftarrow m_p[F(x_i, \hat{k})]$ 
10       $\mathcal{P}_{\hat{k}} = (\mathcal{M}_{\hat{k}}^T \times \mathcal{M}_{\hat{k}})^{-1} \times \mathcal{M}_{\hat{k}}^T$ 

11 for  $\hat{k} = 0$  to  $2^n - 1$  do */
12   /* Test hyp.  $\hat{k}$  for all leakage coordinates */ */
13   for  $u = 0$  to  $d - 1$  do */
14     /* Instantaneous attack (at time  $u$ ) */ */
15      $\vec{\beta} = \mathcal{P}_{\hat{k}} \times \vec{\mathcal{L}}[u]$ 
16     /* Compute an estimator  $\vec{E}$  of  $\vec{\mathcal{L}}[u] = (\ell_{0[u]}, \dots, \ell_{N-1[u]})^T$  */ */
17      $\vec{E} = \mathcal{M}_{\hat{k}} \times \vec{\beta}$ 
18     /* Compute the estimation error (i.e. the SSR) */ */
19     SSR = 0
20     for  $i = 0$  to  $N - 1$  do
21       SSR = SSR +  $(\vec{E}[i] - \vec{\ell}_i[u])^2$ 
22     /* Compute the coefficient of determination */ */
23      $\mathcal{R}[\hat{k}][u] = 1 - \text{SSR} / \overrightarrow{SST}[u]$ 

/* Most likely candidate selections */
24 candidate =  $\text{argmax}_{\hat{k}} (\max_u \mathcal{R}[\hat{k}][u])$ 
25 return candidate

```

3.2 On the LRA Effectiveness

Let us focus on the best candidate selection step in a classical LRA. Each sub-key hypothesis \hat{k} is first associated with a score which is the greatest *instantaneous coefficient of determination* when testing it for all temporal coordinates u . It is denoted by $\max_u \mathcal{R}[\hat{k}][u]$ in Alg. 1. The second phase of the selection consists in the processing of the maximum $\text{argmax}_{\hat{k}} (\max_u \mathcal{R}[\hat{k}][u])$. The purpose of the latter step is to identify the candidate that maximises the greatest instantaneous coefficient. Implicitly, such a classical approach by *total maximisation of the distinguisher value* assumes that the most likely candidate corresponds to the greatest value taken by the distinguisher not only over all sub-key hypotheses

but also over all the leakage times. This assumption relies on another one, often done in the embedded security community, which states that the value of a distinguisher computed between wrong hypotheses (*i.e.* computed for a wrong sub-key value or a wrong time) and the leakage values tends toward its minimum value (often 0) when the sample size N increases (see *e.g.* [20]). However, as already noticed in several papers (*e.g.* by Messerges in [23], Brier *et al.* in [8] or by Whitnal *et al.* in [40]), both assumptions are often not verified in practice, where the adversary must for instance deal with the *ghost peaks* phenomenon. The situation is even worst for the LRA attacks since the vector of coefficients $\vec{\beta}$ (and thus the set of predictions) depends not only on \hat{k} but also on the attack time u . The strength of the LRA, namely its ability to adapt to the instantaneous leakage, is also its weakness as it makes it difficult to compare the different instantaneous attacks results.

To illustrate the issue raised in the previous paragraph, we experimented a LRA against an AES sbox processing running on Device B (see Section 2.5). The full leakage traces were composed of 16800 points. We performed the attack on the full trace length and, for each time coordinate, we recorded the scores of all the 256 key-candidates after $N = 1000$ observations. For clarity reasons, we present in Figure 2 the results only for a temporal window of size 250 points where the targeted variable was known to occur. In the top of the figure, the rank of the correct key is plotted and it can yet be observed that it is 0 for few times. In the second trace of Figure 2, the instantaneous maximum scores comprised in $[0.9982, 0.999]$ are plotted⁸: it may be checked that the maximum among those scores corresponds to a time ($t = 238$) when the correct key is not ranked first. This explains why the total maximisation approach fails in returning the correct key candidate in this case.

To build a better rule than the total maximisation test, we respectively plotted in the third and fourth traces of Figure 2 the mean (plain green trace) and the variance (plain red trace) of the instantaneous scores (*i.e.* the values $\mu(u) = 2^{-8} \sum_{\hat{k}} \mathcal{R}[\hat{k}][u]$ and $\sigma(u) = 2^{-8} \sum_{\hat{k}} (\mathcal{R}[\hat{k}][u] - \mu(u))^2$ with u denoting the time coordinate in abscissa). For each time, we also plotted in black dashed line, the maximum score $\max(u) = \max_{\hat{k}} (\mathcal{R}[\hat{k}][u])$. It may be observed that the correct key is ranked first at the time u when the distance $\max(u) - \mu(u)$ is large and $\sigma(u)$ is small. The third (red) trace and the fourth (gray) trace aim at supporting this claim. Eventually, they suggest us the following pre-processing before comparing the instantaneous attack results: for each leakage coordinate, center the maximum of the coefficients of determination and divide it by their standard deviation. The resulting scoring is plotted in the fifth (magenta) trace, where it can be checked that the maximum is indeed achieved for the correct key.

⁸ For visibility purpose, we chose to not plot the scores lower than 0.9982.




Fig. 2. LRA on Device B (over 1000 traces): Scores Statistics

As a conclusion, and in the light of our analysis, we propose to replace the candidate selection step of the LRA by the following ones⁹:

```

18 for  $u = 0$  to  $d - 1$  do
19    $\left[ \text{attackRes}[u] = \left\{ \text{argmax}_{\hat{k}} (\mathcal{R}[\hat{k}][u]), \frac{\max_{\hat{k}} (\mathcal{R}[\hat{k}][u]) - \mathbb{E}_{\hat{k}} [\mathcal{R}[\hat{k}][u]]}{\sigma_{\hat{k}} [\mathcal{R}[\hat{k}][u] ]} \right\} \right.$ 
20   candidate =  $\text{arg1max2}(\text{attackRes})$ 

```

In Section 3.4, our scores pre-processing technique is applied to attack samples of Device A and Device C in order to test whether our observations, about (1) the ineffectiveness of the classical LRA and (2) the soundness of the new pre-processing, stay valid for other devices than Device B.

3.3 On the LRA Efficiency

The construction of the prediction matrices in Alg. 1 implies, for each \hat{k} , the processing of 3 products of matrices with one dimension equal to s (number of basis functions) and the second dimension equal to N (number of leakage traces). The processing of the instantaneous attacks also requires two such matrix products for each pair (\hat{k}, u) with $u \leq d$. This makes the application of a linear regression attack as depicted in Alg. 1 difficult to perform (and even impossible) when the number N of leakage traces and/or the number d of attack times are

⁹ Where arg1max2 is a function returning the first coordinate of the maximum of an array of 2-dimensional elements, the maximisation being computed with respect to the second coordinate of the array elements.

large. Fortunately this complexity can be significantly reduced. It can indeed be easily shown (see [10]) that the processing of the vectors $\vec{\beta}$ is unchanged if performed for the set of *averaged leakages* $(\frac{1}{\#\{i:x_i=x\}} \sum_{i,x_i=x} \vec{\ell}_i)_{x \in \mathbb{F}_2^n}$ instead of $(\vec{\ell}_i)_i$. Actually, this amounts to change the definition of the matrices $\mathcal{M}_{\hat{k}}$ and \mathcal{L} in (1) such that $\mathcal{M}_{\hat{k}} \doteq (\mathfrak{m}_p(F(x, \hat{k})))_{x \in \mathbb{F}_2^n, p \leq s}$ and \mathcal{L} is a $(2^n \times d)$ -matrix whose x^{th} row vector $\vec{\mathcal{L}}^\top[x]$ equals $\frac{1}{\#\{x_i=x\}} \sum_{i,x_i=x} \vec{\ell}_i$. This improvement essentially lets the first 9 steps of Alg. 1 unchanged except the loop 7-8 which is now computed over $x \in \mathbb{F}_2^n$ instead of over $i \in [0; N - 1]$. Then, before Step 10, the following processing is done to compute the elements of the matrix \mathcal{L} :

```

for  $i = 0$  to  $N - 1$  do
   $\vec{\mathcal{L}}^\top[x_i] = \vec{\mathcal{L}}^\top[x_i] + \vec{\ell}_i$ 
   $\text{count}[x_i] = \text{count}[x_i] + 1$ 
for  $x = 0$  to  $2^n - 1$  do
   $\vec{\mathcal{L}}^\top[x] = \vec{\mathcal{L}}^\top[x] / \text{count}[x]$ 

```

Eventually, Steps 10-17 are replaced by the following ones where we recall that $\vec{\mathcal{L}}[u]$ denotes the u^{th} column vector of \mathcal{L} .

```

for  $\hat{k} = 0$  to  $2^n - 1$  do
  /* Test hypothesis  $\hat{k}$  for all leakage coordinates */
  for  $u = 0$  to  $d - 1$  do
    /* Instantaneous attack (at time  $u$ ) */
     $\vec{\beta} = \mathcal{P}_{\hat{k}} \times \vec{\mathcal{L}}[u]$ 
    /* Compute an estimator  $\vec{E}$  of  $\vec{\mathcal{L}}[u] = (\mathcal{L}[0][u], \dots, \mathcal{L}[2^n - 1][u])^\top$  */
     $\vec{E} = \mathcal{M}_{\hat{k}} \times \vec{\beta}$ 
    /* Compute the estimation error (i.e. the residual sum of squares) */
     $\text{SSR} = 0$ 
    for  $x = 0$  to  $2^n - 1$  do
       $\text{SSR} = \text{SSR} + (\vec{E}[x] - \mathcal{L}[x][u])^2$ 
    /* Compute the coefficient of determination */
     $\mathcal{R}[\hat{k}][u] = 1 - \text{SSR} / \text{SST}[u]$ 

```

The efficiency improvements proposed here for the LRA attack allows for a significant time/memory gain. First, it replaces the $(N \times s)$ -matrix products at Step 13 by $(2^n \times s)$ -matrix products. More globally, the complexity is reduced from $O(s \times d \times N)$ to $O(s \times d \times 2^n)$. If the $\vec{\beta}$ values are not needed (*i.e.* the weights of the linear regression is of no interest to the attacker), the matrix products $\mathcal{M}_{\hat{k}} \times \mathcal{P}_{\hat{k}}$ can also be pre-processed. This enables to save one matrices product per loop iteration (over \hat{k} and u).

3.4 Experiments

We experimented the classical and improved LRA against against three copies of Devices A, B and C (see Section 2.5). The attacks target four bytes of the AES state after the first `SubBytes` operation and they are applied on the full side channel traces. Each attack has been performed 10 times against each of the three copies. The average rank over the four correct sub-keys is plotted in Figure 3 for each device. We recall that the rank of a sub-key k is here defined as the position of $\max_u \mathcal{R}[k][u]$ in the vector $(\max_u \mathcal{R}[\hat{k}][u])_{\hat{k}}$ after sorting (see Section 2.2 for a discussion about this choice). The experiments reported in Figure 3(a)-(c) are done with a *linear basis* (*i.e.* the functions m_i were chosen such that m_0 is constant equal to 1 and m_i , with $i \leq 8$, returns the i^{th} bit of its inputs). It may be observed that the classical attack always failed whereas the improved one succeeded with less than 2500 observations (and even less than 800 for Device B).




Fig. 3. LRA campaign – Rank evolution *versus* number of observations

4 Practical Evaluation of Template Attacks

Template attacks have been introduced in 2002 by Chari *et al.* [9]. Subsequent works have then been published which either show how to apply them against particular implementations (*e.g.* AES, RSA or ECDSA) or propose efficiency/effectiveness improvements [1, 3, 27, 29]. In [27], the authors reduce the complexity of template attacks by first applying a pre-processing on the measurements (to go from time domain to frequency domain) and then by applying dimension reduction techniques (*e.g.* PCA). The latter idea is also followed in [1] and [3]. In all those papers, the improvement of the template attacks efficiency is not studied at the algorithmic level. Moreover, the reported template attack experiments involve the same device for the profiling and matching phases of the attacks, which strongly reduces the practical significance of the argumentations. Indeed, as the profiling phase requires a full access to the device (and in particular the ability to chose the secret parameter), the latter experiments do not fit with the large majority of real attack/evaluation contexts where the adversary has no (or very few) control on the target device. In a more realistic attacker model the profiling phase is conducted on a different device. For such a model,

we have the following well known question¹⁰ about the efficiency of template attack: how sound/relevant is a profiling done on a device A when attacking another device B ? The first work, and to the best of our knowledge, the single one reporting on template attacks in such context is due to Renauld *et al.* [29]. On the latter article, the two devices used for the experiments are test chips implementing an AES s-box and made in 65-nanometer CMOS technology.

The results presented in the rest of this section improve the state-of-the-art recalled previously on two points. First, the efficiency improvement is done at the algorithmic level. It can hence be combined with the previous improvements which essentially correspond to measurements traces pre-processing. Secondly, the reported experiments concern a full AES implementation running on 3 different samples of 3 different technologies. This allowed us to complete the analyses done in [29] and to draw, for the first time, conclusions about the template attack efficiency for realistic scenarios.

4.1 Attack Description

A template attack (TA for short) assumes that a preliminary profiling step has been performed on an open copy of the targeted device. During this phase, the adversary has measured N' leakage traces $\vec{\ell}'_i$ for which he knows exactly the values taken by the corresponding sensitive value Z (which also implies that he knows the corresponding sub-key k). Those leakages have then been used to compute estimations $f_z(\cdot)$ of the *probability density function* of $(\vec{\mathbf{L}} \mid Z = z)$ for all possible z (which imposes $N' \gg 2^m$). The pdf estimations $f_z(\cdot)$ will play in a template attack, a similar role as the model functions in a CPA or LRA.

Once the adversary has the set of pdf estimations $(f_z(\cdot))_{z \in \mathbb{F}_2^m}$ in hand, a TA against the set of traces $(\vec{\ell}'_i)_{i \leq N}$ (for which the secrets are unknown) follows essentially the same outlines as the LRA: the hypothesis \hat{k} is tested by first computing the predictions $\hat{z}_i = F(x_i, \hat{k})$ and then by calculating the product $\prod_{i \leq N} f_{\hat{z}_i}(\vec{\ell}'_i)$. Usually, the pdf of the variables $(\vec{\mathbf{L}} \mid Z = z)$ is estimated by a multivariate normal law, which implies that f_z can be developed s.t.:

$$f_z(\vec{\ell}'_i) = \frac{1}{(2\pi)^d \det(\Sigma_z)} \exp\left(-\frac{1}{2}(\vec{\ell}'_i - \vec{\mu}_z)^\top \Sigma_z^{-1}(\vec{\ell}'_i - \vec{\mu}_z)\right), \quad (2)$$

where Σ_z denotes the $(d \times d)$ -matrix of covariances of $\vec{\mathbf{L}} \mid Z = z$ and where the (d) -dimensional vector $\vec{\mu}_z$ denotes its mean.

To minimize approximation errors induced by the processing of the product of exponential values, one usually prefers, in practice, a log-maximum likelihood processing to the classical maximum likelihood¹¹. Together with (2), this leads to the following computation to test the hypothesis \hat{k} :

¹⁰ This question is sometimes also related to the statistical problem of pdf estimations robustness [25].

¹¹ The two processes discriminate equivalently.

$$\mathcal{ML}[\hat{k}] = - \sum_{i \leq N} (\vec{\ell}_i - \vec{\mu}_{\hat{z}_i})^\top \Sigma_{\hat{z}_i}^{-1} (\vec{\ell}_i - \vec{\mu}_{\hat{z}_i}) - \sum_{i \leq N} \log((2\pi)^{d+1} \det(\Sigma_{\hat{z}_i})) . \quad (3)$$

We give in Alg. 2 the pseudo-code corresponding to the TA attack discussed previously.

Algorithm 2: TA - Template Attacks

```

Input : a set of  $d$ -dimensional leakages  $(\vec{\ell}_i)_{i \leq N}$  and the corresponding plaintexts  $(x_i)_{i \leq N}$ , a set of pdf estimations  $(\vec{\mu}_z, \Sigma_z)_{z \in \mathbb{F}_2^m}$ 
Output: A candidate sub-key  $\hat{k}$ 
/* Pre-Processing of the  $2^m$  log-determinants  $\log(2\pi^{d+1} \Sigma_z)$  and inverse-matrices  $\Sigma_z^{-1}$  */
1 for  $z = 0$  to  $2^m - 1$  do */
2    $\log \text{Det}_z = \log(2\pi^{d+1} \Sigma_z)$ 
3    $\text{invCov}_z = \Sigma_z^{-1}$ 

/* Instantaneous TA attacks Processing */
4 for  $\hat{k} = 0$  to  $2^m - 1$  do */
5   /* Test hyp.  $\hat{k}$  */ */
6    $\mathcal{ML}[\hat{k}] = 0$ 
7   for  $i = 0$  to  $N - 1$  do */
8      $\hat{z} = F(x_i, \hat{k})$ 
9      $\mathcal{ML}[\hat{k}] = \mathcal{ML}[\hat{k}] - (\vec{\ell}_i - \vec{\mu}_{\hat{z}})^\top \times \text{invCov}_{\hat{z}} \times (\vec{\ell}_i - \vec{\mu}_{\hat{z}}) - \log \text{Det}_{\hat{z}}$ 

/* Most likely candidate selections */
9 candidate =  $\text{argmax}_{\hat{k}} (\max \mathcal{ML}[\hat{k}])$ 
10 return candidate

```

4.2 On the TA Effectiveness

The idea developed in previous sections to improve the selection of the best candidate among the results of several instantaneous attacks is not relevant here. Indeed, for both the profiling and attack phases, a template attack exploits, by nature, all the leakage coordinates of the $\vec{\ell}_i$ simultaneously. There is consequently no need to compare the results of several (different) instantaneous attacks.

4.3 On the TA Efficiency

Applying the same idea as for the LRA, we propose hereafter an alternative writing of $\mathcal{ML}[\hat{k}]$ that leads to a much faster attack processing. For such a purpose, we focus on the term $(\vec{\ell}_i - \vec{\mu}_{\hat{z}_i})^\top \Sigma_{\hat{z}_i}^{-1} (\vec{\ell}_i - \vec{\mu}_{\hat{z}_i})$ in (3).

After denoting by \mathcal{L}_i each $(d \times d)$ -matrix $(\vec{\ell}_i[u] \vec{\ell}_i[u'])_{u,u'}$, we get the following rewriting of the latter term:

$$\sum_{u,u'} (\mathcal{L}_i[u][u'] \times \Sigma_{\hat{z}_i}^{-1}[u][u']) - \vec{\mu}_{\hat{z}_i}^\top \times (\Sigma_{\hat{z}_i}^{-1} + \Sigma_{\hat{z}_i}^{-1\top}) \times \vec{\ell}_i + \vec{\mu}_{\hat{z}_i}^\top \times \Sigma_{\hat{z}_i}^{-1} \times \vec{\mu}_{\hat{z}_i} .$$

After recalling that \hat{z}_i equals $F(x_i, \hat{k})$ and after denoting $F(x, \hat{k})$ by \hat{z} and $\#\{i, x_i = x\}$ by N_x , we deduce that the sum $\sum_i (\vec{\ell}_i - \vec{\mu}_{\hat{z}_i})^\top \Sigma_{\hat{z}_i}^{-1} (\vec{\ell}_i - \vec{\mu}_{\hat{z}_i})$ may be rewritten:

$$\sum_{x \in \mathbb{F}_2^n} \left(\sum_{u, u'} \left(\sum_{i, x_i = x} \mathcal{L}_i \right) [u][u'] \times \Sigma_{\hat{z}}^{-1} [u][u'] - \vec{\mu}_{\hat{z}}^\top \times (\Sigma_{\hat{z}}^{-1} + \Sigma_{\hat{z}}^{-1\top}) \times \left(\sum_{i, x_i = x} \vec{\ell}_i \right) + N_x \times \vec{\mu}_z^\top \times \Sigma_z^{-1} \times \vec{\mu}_z \right).$$

As a consequence, if the 2^n possible sums $\sum_{i, x_i = x} \mathcal{L}_i$ and $\sum_{i, x_i = x} \vec{\ell}_i$ have been precomputed, then the complexity of evaluating (3) for each \hat{k} goes from $O(Nd^2)$ to $O(2^n d^2)$. Algorithm 3 describes the improved TA attack.

Algorithm 3: TA - Template Attacks (Improved Version)

Input : a set of N leakages $(\vec{\ell}_i)_i$ and the corresponding plaintexts $(x_i)_i$, a set of pdf estimations $(\vec{\mu}_z, \Sigma_z)_{z \in \mathbb{F}_2^m}$

Output: A candidate subkey \hat{k}

```

/* Pre-Processing of the predictions data */
1 for z = 0 to 2^m - 1 do
2   logDet_z = log(2π^{d+1} Σ_z); invCov_z = Σ_z^{-1}; meanCov_z =
   [ μ_z^T × invCov_z × μ_z; sumMeanCov_z = μ_z^T × (Σ_z^{-1} + Σ_z^{-1\top})
/* Pre-Processing of the leakage data12 l_x = Σ_{i, x_i = x} l_i, L_x = Σ_{i, x_i = x} L_i and
   N[x] = # {i; x_i = x}. */
3 for i = 0 to N - 1 do
4   [ x = x_i; N[x] = N[x] + 1; L_x = L_x + l_i l_i^T; l_x = l_x + l_i
/* Instantaneous TA attacks Processing */
5 for k-hat = 0 to 2^n - 1 do
6   /* Test hyp. k-hat */
7   ML[k-hat] = 0
8   for x = 0 to 2^n - 1 do
9     z-hat = F_j(x, k-hat)
10    for u = 0 to d - 1 do
11      for u' = 0 to d - 1 do
12        [ ML[k-hat] = ML[k-hat] - L_x[u][u'] × invCov_z[u][u']
        ML[k-hat] = ML[k-hat] + sumMeanCov_z × l_x - N[x] × (meanCov_z + logDet_z)
/* Most likely candidate selections */
13 candidate = argmax_k-hat (max ML[k-hat])
14 return candidate

```

4.4 Experiments

To study the effectiveness of TA attacks in practice (and to confirm the observations reported in [29]) we experimented them against the families of devices A, B and C for three different scenarios. In the first scenario (referred to as "copy 1 \rightarrow copy 1"), the profiling and the attacks are performed on the same device copy. In the second and third scenarios (respectively referred to as "copy 1 \rightarrow copy 2" and "copy 1 \rightarrow copy 3"), the profiling made for copy 1 is used to attack the second and third copies. For each of these 9 attacks frameworks, we plot in Figure 4 the average rank of the correct sub-key (in color) with respect to both the number of traces used for the profiling (in ordinate) and the number of traces used for the attack (in abscissa). The rank averaging has been done over 10 attacks.

In the first scenario, a profiling done on 15000 (resp. 47000) traces on Device B (resp. Device C) allows for a very efficient attack phase (the correct sub-key ranked first with less than ten traces). Moreover, it may be observed that a profiling on 8000 traces for Devices B and C is sufficient to have a successful attack in less than 23 (resp. 90) traces for device B (resp. C). For Device A, the TA attack in Scenario 1 is one order of magnitude less efficient (roughly speaking the values are multiplied per ten w.r.t. the traces for devices B and C).

Attacks on Devices B (resp. C) perform quite similarly in Scenarios 2 and 3. For Device A, a profiling performed on copy 1 for 18000 traces is sufficient to successfully attack copies 2 and 3 with less than 10 traces. Moreover, a profiling on 8000 traces enables successful attacks for less than 30 traces. For Device C, it may be observed that, even for a profiling performed on 50000 traces, the attacks on copies 2 and 3 require at least 80 traces to succeed. However, a profiling on 9000 traces is sufficient to have the TA succeeding in less than 130 traces.

As expected, we may observe a significant variability for the attack results in Scenarios 2 and 3 for Device A: templates done on copy 1 are almost as efficient to attack copy 2 than they were to attack copy 1 itself. They are however much less informative on the copy 3 behaviour since the profiling on copy 1 must be performed on at least 130000 traces to see the attack working on copy 3 with less than 700 traces. This observation is in-line with those done in [28] about the high variability of nano-scale technologies (we recall that Device A is made in a 90nm CMOS technology).

In the full version of this paper, we report on similar experiments results when only the leakage means (and not the covariance matrices) are involved in the templates. This approach indeed seems to be a natural alternative to the attacks described here since the traces contain instantaneous leakages. Our results actually confirm this feeling and it can even be noticed that it leads to improve the TA efficiency for Scenario 3 on Device A¹³. Another general remark on these simplified templates is that they perform much better than the full ones when the number of traces used for the profiling is small (around 4000).

¹³ This could be explained by the fact that the technology variability has more impact on the electromagnetic leakage variances than it has on the means.




Fig. 4. TA campaigns – Rank evolution *vs.* nb. of traces for the attack phase (x -axis) and the profiling (y -axis)

5 Conclusion

In this paper, we have studied the effectiveness and efficiency of the LRA and the TA attacks when performed in a context where the exact time of the sensitive computations is not known. In this situation, and even after the application of pattern matching or resynchronization techniques, the exploited leakage traces may be composed of several thousands of points and the same attack must be processed for each of those points. We noticed that the study of the side channel attacks effectiveness and efficiency in this multivariate context is an over-estimated problem. Most of the time, it is indeed assumed that the adversary succeeded in significantly reducing the traces size (*e.g.* by priorly processing a SNR, or a test attack, or even a dimension reduction). However, as argued in this paper, those techniques are either unrealistic or may lead to a significant

loss of useful information (a dimension reduction technique like the PCA may be sound for one attack – *e.g.* the CPA – and not for another one – *e.g.* the MIA or the LRA –). As a consequence, there was no work discussing about the rule to apply in order to select a candidate among all of those returned by a same attack performed against several time coordinates. To the best of our knowledge, the *de facto* rule was hence to simply choose the key candidate maximising all the attacks scores. In this paper, we have shown that this rule does not work for a LRA attack and we have conducted a statistical analysis to deduce a new selection rule that renders it effective in practice, even when the traces are composed of huge number of points. In this paper, we have also tackled out the efficiency problem for the multivariate LRA and TA attacks. For each of them, we have followed a similar approach which led us to significantly reduce their complexity when the number of traces and their dimension are high. It may be noticed that the approach could also be applied (almost straightforwardly) to improve the efficiency of the correlation power attack and of the mutual information attack (with histogram pdf estimation). Eventually, all our results and analyses have been illustrated by several attack experiments on three different copies of three different technologies. In particular, the latter experiments have enabled us to confirm the practicability of template attacks when the profiling phase and the attack are performed on different copies of the same device.

References

1. Archambeau, C., Peeters, E., Standaert, F.-X., Quisquater, J.-J.: Template Attacks in Principal Subspaces. In: Goubin, L., Matsui, M. (eds.) CHES 2006. LNCS, vol. 4249, pp. 1–14. Springer, Heidelberg (2006)
2. Balasch, J., Gierlichs, B., Verdult, R., Batina, L., Verbauwhede, I.: Power Analysis of Atmel CryptoMemory – Recovering Keys from Secure EEPROMs. In: Dunkelman, O. (ed.) CT-RSA 2012. LNCS, vol. 7178, pp. 19–34. Springer, Heidelberg (2012)
3. Bär, M., Drexler, H., Pulkus, J.: Improved Template Attacks. In: Proceedings of the Constructive Side-Channel Analysis and Secure Design - First International Workshop, COSADE 2010, Darmstadt, Germany, February 4-5 (2010)
4. Bartkewitz, T., Lemke-Rust, K.: A high-performance implementation of differential power analysis on graphics cards. In: Prouff, E. (ed.) CARDIS 2011. LNCS, vol. 7079, pp. 252–265. Springer, Heidelberg (2011)
5. Batina, L., Gierlichs, B., Lemke-Rust, K.: Comparative Evaluation of Rank Correlation Based DPA on an AES Prototype Chip. In: Wu, T.-C., Lei, C.-L., Rijmen, V., Lee, D.-T. (eds.) ISC 2008. LNCS, vol. 5222, pp. 341–354. Springer, Heidelberg (2008)
6. Batina, L., Hogenboom, J., van Woudenberg, J.G.J.: Getting More from PCA: First Results of Using Principal Component Analysis for Extensive Power Analysis. In: Dunkelman, O. (ed.) CT-RSA 2012. LNCS, vol. 7178, pp. 383–397. Springer, Heidelberg (2012)
7. 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)

8. Brier, E., Clavier, C., Olivier, F.: Correlation Power Analysis with a Leakage Model. In: Joye, M., Quisquater, J.-J. (eds.) CHES 2004. LNCS, vol. 3156, pp. 16–29. Springer, Heidelberg (2004)
9. Chari, S., Rao, J., Rohatgi, P.: Template attacks. In: Kaliski Jr., B.S., Koc, C.K., Paar, C. (eds.) CHES 2002. LNCS, vol. 2523, pp. 13–28. Springer, Heidelberg (2003)
10. Dabosville, G., Doget, J., Prouff, E.: A New Second Order Side Channel Attack Based on Linear Regression. *IEEE Transactions on Computers* (2012) (to appear)
11. Doget, J., Prouff, E., Rivain, M., Standaert, F.-X.: Univariate Side Channel Attacks and Leakage Modeling. *Journal of Cryptographic Engineering* 1(2), 123–144 (2011)
12. Gierlichs, B., Batina, L., Tuyls, P., Preneel, B.: Mutual Information Analysis. In: Oswald, E., Rohatgi, P. (eds.) CHES 2008. LNCS, vol. 5154, pp. 426–442. Springer, Heidelberg (2008)
13. Gierlichs, B., Lemke-Rust, K., Paar, C.: Templates vs. Stochastic Methods. In: Goubin, L., Matsui, M. (eds.) CHES 2006. LNCS, vol. 4249, pp. 15–29. Springer, Heidelberg (2006)
14. Heuser, A., Schindler, W., Stöttinger, M.: Revealing Side-Channel Issues of Complex Circuits by Enhanced Leakage Models. In: Rosenstiel, W., Thiele, L. (eds.) DATE, pp. 1179–1184. IEEE (2012)
15. Kasper, M., Schindler, W., Stöttinger, M.: A Stochastic Method for Security Evaluation of Cryptographic FPGA Implementations. In: IEEE International Conference on Field-Programmable Technology (FPT 2010). IEEE Press (December 2010)
16. Kizhvatov, I.: Side Channel Analysis of AVR XMEGA Crypto Engine. In: Serpanos, D.N., Wolf, W. (eds.) WESS. ACM (2009)
17. Kocher, P.C., Jaffe, J., Jun, B.: Differential Power Analysis. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 388–397. Springer, Heidelberg (1999)
18. Le, T.-H., Clédière, J., Canovas, C., Robisson, B., Servière, C., Lacoume, J.-L.: A Proposition for Correlation Power Analysis Enhancement. In: Goubin, L., Matsui, M. (eds.) CHES 2006. LNCS, vol. 4249, pp. 174–186. Springer, Heidelberg (2006)
19. Lee, S.J., Seo, S.C., Han, D.-G., Hong, S., Lee, S.: Acceleration of Differential Power Analysis through the Parallel Use of GPU and CPU. *IEICE Transactions* 93-A(9), 1688–1692 (2010)
20. Mangard, S.: Hardware Countermeasures against DPA – A Statistical Analysis of Their Effectiveness. In: Okamoto, T. (ed.) CT-RSA 2004. LNCS, vol. 2964, pp. 222–235. Springer, Heidelberg (2004)
21. Mangard, S., Oswald, E., Popp, T.: Power Analysis Attacks – Revealing the Secrets of Smartcards. Springer (2007)
22. Mangard, S., Oswald, E., Standaert, F.-X.: One for All - All for One: Unifying Standard DPA Attacks. *IET Information Security* (2011)
23. Messerges, T.S.: Using Second-Order Power Analysis to Attack DPA Resistant Software. In: Paar, C., Koc, C.K. (eds.) CHES 2000. LNCS, vol. 1965, pp. 238–251. Springer, Heidelberg (2000)
24. Oswald, D., Paar, C.: Breaking Mifare DESFire MF3ICD40: Power Analysis and Templates in the Real World. In: Preneel, B., Takagi, T. (eds.) CHES 2011. LNCS, vol. 6917, pp. 207–222. Springer, Heidelberg (2011)
25. Press, W.: Numerical Recipes in Fortran 77: The Art of Scientific Computing. *Fortran Numerical Recipes*. Cambridge University Press (1992)
26. Quisquater, J.-J., Samyde, D.: ElectroMagnetic Analysis (EMA): Measures and Counter-Measures for Smart Cards. In: Attali, S., Jensen, T. (eds.) E-smart 2001. LNCS, vol. 2140, pp. 200–210. Springer, Heidelberg (2001)

27. Rechberger, C., Oswald, E.: Practical Template Attacks. In: Lim, C.H., Yung, M. (eds.) WISA 2004. LNCS, vol. 3325, pp. 440–456. Springer, Heidelberg (2005)
28. Renauld, M., Kamel, D., Standaert, F.-X., Flandre, D.: Information Theoretic and Security Analysis of a 65-Nanometer DDSLL AES S-Box. In: Preneel, B., Takagi, T. (eds.) CHES 2011. LNCS, vol. 6917, pp. 223–239. Springer, Heidelberg (2011)
29. Renauld, M., Standaert, F.-X., Veyrat-Charvillon, N., Kamel, D., Flandre, D.: A Formal Study of Power Variability Issues and Side-Channel Attacks for Nanoscale Devices. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 109–128. Springer, Heidelberg (2011)
30. Schindler, W.: Advanced Stochastic Methods in Side Channel Analysis on Block Ciphers in the Presence of Masking. *Journal of Mathematical Cryptology* 2, 291–310 (2008)
31. Schindler, W., Lemke, K., Paar, C.: A Stochastic Model for Differential Side Channel Cryptanalysis. In: Rao, J.R., Sunar, B. (eds.) CHES 2005. LNCS, vol. 3659, pp. 30–46. Springer, Heidelberg (2005)
32. 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.) CHES 2008. LNCS, vol. 5154, pp. 411–425. Springer, Heidelberg (2008)
33. Standaert, F.-X., Gierlichs, B., Verbauwhede, I.: Partition *vs.* Comparison Side-Channel Distinguishers: An Empirical Evaluation of Statistical Tests for Univariate Side-Channel Attacks against Two Unprotected CMOS Devices. In: Lee, P.J., Cheon, J.H. (eds.) ICISC 2008. LNCS, vol. 5461, pp. 253–267. Springer, Heidelberg (2009)
34. Standaert, F.-X., Koeune, F., Schindler, W.: How to Compare Profiled Side-Channel Attacks? In: Abdalla, M., Pointcheval, D., Fouque, P.-A., Vergnaud, D. (eds.) ACNS 2009. LNCS, vol. 5536, pp. 485–498. Springer, Heidelberg (2009)
35. Standaert, F.-X., Malkin, T.G., Yung, M.: A Unified Framework for the Analysis of Side-Channel Key Recovery Attacks. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 443–461. Springer, Heidelberg (2009)
36. Standaert, F.-X., Veyrat-Charvillon, N., Oswald, E., Gierlichs, B., Medwed, M., Kasper, M., Mangard, S.: The World is not Enough: Another Look on Second-Order DPA. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 112–129. Springer, Heidelberg (2010)
37. Télécom ParisTech. DPA Contest v1 and v2, <http://www.dpacontest.org/> (retrieved on August 1, 2012)
38. Whitnall, C., Oswald, E.: A Comprehensive Evaluation of Mutual Information Analysis Using a Fair Evaluation Framework. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 316–334. Springer, Heidelberg (2011)
39. Whitnall, C., Oswald, E.: A Fair Evaluation Framework for Comparing Side-Channel Distinguishers. *J. Cryptographic Engineering* 1(2), 145–160 (2011)
40. Whitnall, C., Oswald, E., Mather, L.: An Exploration of the Kolmogorov-Smirnov Test as Competitor to Mutual Information Analysis. In: Prouff, E. (ed.) CARDIS 2011. LNCS, vol. 7079, pp. 234–251. Springer, Heidelberg (2011)