

ASCA, SASCA and DPA with Enumeration: Which One Beats the Other and When?

Vincent Grosso, François-Xavier Standaert.

ICTEAM/ELEN/Crypto Group, Université catholique de Louvain, Belgium.

Abstract. We describe three contributions regarding the Soft Analytical Side-Channel Attacks (SASCA) introduced at Asiacrypt 2014. First, we compare them with Algebraic Side-Channel Attacks (ASCA) in a noise-free simulated setting. We observe that SASCA allow more efficient key recoveries than ASCA, even in this context (favorable to the latter). Second, we describe the first working experiments of SASCA against an actual AES implementation. Doing so, we analyse their profiling requirements, put forward the significant gains they provide over profiled Differential Power Analysis (DPA) in terms of number of traces needed for key recoveries, and discuss the specificities of such concrete attacks compared to simulated ones. Third, we evaluate the distance between SASCA and DPA enhanced with computational power to perform enumeration, and show that the gap between both attacks can be quite reduced in this case. Therefore, our results bring interesting feedback for evaluation laboratories. They suggest that in several relevant scenarios (e.g. attacks exploiting many known plaintexts), taking a small margin over the security level indicated by standard DPA with enumeration should be sufficient to prevent more elaborate attacks such as SASCA. By contrast, SASCA may remain the only option in more extreme scenarios (e.g. attacks with unknown plaintexts/ciphertexts or against leakage-resilient primitives). We conclude by recalling the algorithmic dependency of the latter attacks, and therefore that our conclusions are specific to the AES.

1 Introduction

State-of-the-art. Strategies to exploit side-channel leakages can be classified as Divide and Conquer (DC) and analytical. In the first case, the adversary recovers information about different bytes of (e.g.) a block cipher key independently, and then combines this information, e.g. via enumeration [30]. In the second case, she rather tries to recover the full key at once, exploiting more algorithmic approaches to cryptanalysis with leakage. Rephrasing Banciu et al, one can see these different strategies are a tradeoff between pragmatism and elegance [1].

In brief, the “DC+enumeration” approach is pragmatic, i.e. it is easy to implement, requires little knowledge about the target implementation, and can take advantage of a variety of popular (profiled and non-profiled) distinguishers, such as Correlation Power Analysis (CPA) [5], Mutual Information Analysis (MIA) [10], Linear Regression (LR) [28] or Template Attacks (TA) [7]. We will use the term Differential Power Analysis (DPA) to denote them all [17].

By contrast, analytical approaches are (more) elegant, since they theoretically exploit all the information leaked by an implementation (vs. the leakages of the first rounds for DC attacks). As a result, these attacks can (theoretically) succeed in conditions where the number of measurements available to the adversary is very limited. But this elegance (and the power that comes with it) usually implies stronger assumptions on the target implementation (e.g. most of them require some type of profiling). The Algebraic Side-Channel Attacks (ASCA) described in [24] and further analyzed in [6, 26] are an extreme solution in this direction. In this case, the target block cipher and its leakages are represented as a set of equations that are then solved (e.g. with a SAT solver, or Groebner bases). This typically implies a weak resistance to the noise that is usually observed in side-channel measurements. As a result, various heuristics have been suggested to better deal with errors in the information leakages, such as [19, 33]. The Tolerant Algebraic Side-Channel Attacks (TASCA) proposed in [20, 21] made one additional step in this direction, by replacing the solvers used in ASCA by an optimizer. But they were limited by their high memory complexity (since they essentially deal with noise by exhaustively encoding the errors they may cause). More recently, two independent proposals suggested to design a dedicated solver specialized to byte-oriented ciphers such as the AES [12, 22]. The latter ones were more efficient and based on smart heuristics exploiting enumeration. Eventually, Soft Analytical Side-Channel Attacks (SASCA) were introduced at Asiacrypt 2014 as a conceptually different way to exploit side-channel leakages analytically [32]. Namely, rather than encoding them as equations, SASCA rather describe an implementation and its leakages as a code, that one can efficiently decode using the Belief Propagation (BP) algorithm. As a result, they are able to directly exploit the (soft) information provided by profiled side-channel attacks (such as LR or TA), in a very efficient manner, with very limited memory complexity, and for multiple plaintexts. Concretely, this implies that they provide a natural bridge between DC attacks and analytical ones.

Our contribution. In view of this state-of-the-art, we consider three open problems regarding DC and analytical strategies in side-channel analysis.

First, we observe that the recent work in [32] experimented SASCA in the context of noisy leakages. While this context allowed showing that SASCA are indeed applicable in (noisy) environments where ASCA would fail, it leaves the question whether this comes at the cost of a lower efficiency in a noise-free context open. Therefore, we launched various experiments with noise-free leakages to compare ASCA and SASCA. These experiments allowed us to confirm that also in this context, SASCA are equally (even slightly more) efficient.

Second, the experiments in [32] exploited simulations in order to exhibit the strong noise-resilience of SASCA (since the amount of noise can then be used as a parameter of such simulations). But this naturally eludes the question of the profiling of a concrete device, which can be a challenging task, and for which the leakage models of different target intermediate values may turn out to be quite different [9]. Therefore, we describe the first working experiments of SASCA against an actual AES implementation, for which a bivariate TA ex-

exploiting the S-box input/output leakages would typically be successful after more than 50 measurements. We further consider two cases for the adversary’s knowledge about the implementation. In the first one, she has a precise description in hand (i.e. the assembly code, typically). In the second one, she only knows AES is running, and therefore only exploits the generic operations that one can assume from the algorithm specification. Our experiments confirm that SASCA are applicable in a simple profiled scenario, and lead to successful key recoveries with less traces than a DC attack (by an approximate factor up to 5). They also allow us to discuss the profiling cost, and the consequences of the different leakage models in our target implementation. A relevant observation regarding these models is that weak leakages in the MixColumns operations are especially damaging, which can be explained by the (factor) graph describing an AES implementation: indeed, XORing two values with limited information significantly reduces the information propagation during the execution of the BP algorithm. This suggests interesting research directions for preventing such attacks, since protecting the linear parts of a block cipher is usually easier/cheaper.

Third, we note that SASCA are in general more computationally intensive than DC attacks. Therefore, a fair comparison should allow some enumeration power to the DC attacks as well. We complement our previous experimental attacks by considering this last scenario. That is, we compare the success rate of SASCA with the ones of DC attacks exploiting a computational power corresponding to up to 2^{30} encryptions (which corresponds to more than the execution time of SASCA on our computing platform). Our results put forward that SASCA remain the most powerful attack in this case, but with a lower gain.

Summary. These contributions allow answering the question of our title. First, SASCA are in general preferable to ASCA (i.e. both for noise-free and noisy scenarios). Second, the tradeoff between SASCA and DC attacks is more balanced. As previously mentioned, DC attacks are more pragmatic. So the interest of SASCA essentially depends on the success rate gains it provides, which itself depends on the scenarios. If multiple known plaintexts are available, our experiments suggest that the gain of SASCA over DPA with enumeration is somewhat limited, and may not justify such an elegant approach. This conclusion backs up the results in [1], but in a more general scenario, since we consider multiple-plaintext attacks rather than single-plaintext ones, together with more powerful analytical attacks. By contrast, if plaintexts/ciphertexts are unknown, or if the number of plaintexts one can observe is very limited (e.g. by design, due to a leakage-resilient primitive [2]), SASCA may be the best/only option.

Preliminary remark. Our focus in this paper is on a couple of extreme approaches to side-channel analysis, i.e. the most pragmatic DC attacks against 8-bit targets of the first block cipher round, and the most elegant ASCA/SASCA exploiting most/all such targets in the implementation. Quite naturally, the other analytical attacks mentioned in this introduction would provide various tradeoffs between these extremes. Besides, more computationally-intensive DPA attacks (based on larger key hypotheses) are also possible, as recently discussed by Mather et al. [18]. Such attacks are complementary and may further reduce

the gain of SASCA over DPA, possibly at the cost of increased computational requirements (e.g. the latter work exploited high-performance computing whereas all our experiments were carried out on a single desktop computer).

2 Background

In this section we first describe the measurement setup used in our experiments. Then, we describe two tools we used to identify and evaluate information leakages in the traces. Finally, we recall the basics of the different attacks we compare.

2.1 Measurement setup

Our measurements are based on the open source AES FURIOUS implementation (<http://point-at-infinity.org/avraes>) run by an 8-bit Atmel ATMEGA644p microcontroller at a 20 MHz clock frequency. We monitored the voltage variations across a 22Ω resistor. Acquisitions were performed using a Lecroy WaveRunner HRO 66 ZI providing 8-bit samples, running at 400 Msamples/second. For SASCA, we can exploit any intermediate values that appear during the AES computation. Hence, we measured the full encryption. Our traces are composed of 94,000 points, containing the key scheduling and encryption rounds. Our profiling is based on 256,000 traces corresponding to random plaintexts and keys. As a result, we expect around 1,000 traces for each value of each intermediate computation. We use $l_{n,x}^i$ for the value x of the n^{th} intermediate value in the i^{th} leakage trace, and $l_{n,x}^i(t)$ when we access at the t^{th} point (sample) of this trace.

2.2 Information detection tools

Since SASCA can exploit many target intermediate values, we need to identify the time samples that contain information about them in our traces, next referred to as Points Of Interest (POI). We recall two simple methods for this purpose, and denote the POI of the n^{th} intermediate value in our traces with t_n .

(a) Correlation Power Analysis (CPA) [5] is a standard side-channel distinguisher that estimates the correlation between the measured leakages and some key-dependent model for a target intermediate value. In its standard version, an a-priori (here, Hamming weight) model is used for this purpose. In practice, this estimation is performed by sampling (i.e. measuring) traces from a leakage variable L and a model variable M_k , using Pearson’s correlation coefficient:

$$\rho_k(L, M_k) = \frac{\hat{E}[(L - \hat{\mu}_L)(M_k - \hat{\mu}_{M_k})]}{\sqrt{\hat{\text{var}}(L)\hat{\text{var}}(M_k)}}.$$

In this equation, \hat{E} and $\hat{\text{var}}$ respectively denote the sample mean and variance operators, and $\hat{\mu}_L$ is the sample mean of the leakage distribution L . CPA is a univariate distinguisher and therefore launched sample by sample.

(b) **The Signal-to-Noise Ratio (SNR)** [16] of the n^{th} intermediate value at the time t can be defined according to Mangard’s formula [16]:

$$\text{SNR}_n(t) = \frac{\text{v}\hat{\text{a}}r_x\left(\hat{\text{E}}_i(\mathbf{l}_{n,x}^i(t))\right)}{\hat{\text{E}}_x\left(\text{v}\hat{\text{a}}r_i(\mathbf{l}_{n,x}^i(t))\right)}.$$

Despite connected (high SNRs imply efficient CPA if the right model is used), these metrics allow slightly different intuitions. In particular, the SNR cannot tell apart the input and output leakages of a bijective operation (such as an S-box), since both intermediate values will generate useful signal. This separation can be achieved by CPA thanks to its a-priori leakage predictions.

2.3 Gaussian templates attacks

Gaussian TA [7] are the most popular profiled distinguisher. They assume that the leakages can be interpreted as the realizations of a random variable which generates samples according a Gaussian distribution and work in two steps. In a profiling phase, the adversary estimates a mean $\hat{\mu}_{n,x}$ and variance $\hat{\sigma}_{n,x}^2$ for each value x of the n^{th} intermediate computation. In practice, this is done for the time sample t_n obtained thanks to the previously mentioned POI detection tool. Next, in the attack phase and for each trace \mathbf{l} , she can calculate the likelihood to observe this leakage at the time t_n for each x as:

$$\hat{\text{P}}r[\mathbf{l}(t_n)|x] \sim \mathcal{N}(\hat{\mu}_{n,x}, \hat{\sigma}_{n,x}^2).$$

In the context of standard DPA, we typically have $x = p \oplus k$, with p a known plaintext and k the target subkey. Therefore, the adversary can easily calculate $\hat{\text{P}}r[k^*|p, \mathbf{l}(t_n)]$ using Bayes theorem, for each subkey candidate k^* :

$$\hat{\text{P}}r[k^*] = \prod_i \hat{\text{P}}r[k^*|p, \mathbf{l}^i(t_n)].$$

To recover the full key, she can run a TA on each subkey independently.

By contrast, in the context of SASCA, we will directly insert the knowledge (i.e. probabilities) about any intermediate value x in the (factor) graph describing the implementation, and try to recover the full key at once. Note that our experiments consider univariate Gaussian TA because our focus is on exploiting multiple intermediate computations. Yet, they can be easily extended to multivariate ones, where we use more than one samples per intermediate value.

2.4 Key enumeration and rank estimation

At the end of a DC side-channel attack (as the previous TA), the attacker has probabilities on each subkey. If the master key is not the most probable one, she can perform enumeration up to some threshold thanks to enumeration algorithms, e.g. [30]. This threshold depends on the computational power of the

adversary, since enumerating all keys is computationally impossible. If the key is beyond the threshold of computationally feasible enumeration, and in order to gain intuition about the computational security remaining after an attack, key rank estimation algorithms can be used [11, 31]. A key rank estimation takes in input the list of probabilities of all subkeys and the probability of the correct key (which is only available in an evaluation context), and returns an estimation on the number of keys that are more likely than the actual key. Rank estimation allows to approximate d^{th} -order success rates (i.e. the probability that the correct key lies among the d first ones rated by the attack) efficiently and quite accurately. The security graphs introduced in [31] provide a visual representation of higher-order success rates in function of the number attack traces.

2.5 Algebraic side-channel attacks

ASCA were introduced in [24] as one of the (if not the) first method to efficiently exploit all the informative samples in a leakage trace. We briefly recall their three main steps and refer to previous publications for the details.

1. *Construction* consists in representing the cipher as an instance of an algebraic problem (e.g. Boolean satisfiability, Groebner bases). Because of their large memory (RAM) requirements, ASCA generally build a system corresponding to one (or a few) traces only. For example, the SAT representation of a single AES trace in [26] has approximately 18,000 equations in 10,000 variables.
2. *Information extraction* consists in getting exploitable leakages from the measurements. For ASCA, the main constraint is that actual solvers require hard information. Therefore, this phase usually translates the result of a TA into deterministic leakages such as the Hamming weight of the target intermediate values (the type of leakages does not matter, but we need them without errors).
3. *Solving*. Eventually, the side-channel information extracted in the second phase is added to the system of equations constructed in the first phase, and generic solvers are launched to solve the system and recover the key.

Summarizing, ASCA are powerful attacks since they can theoretically recover a key from very few leakage traces, but this comes at the cost of low noise-resilience, which motivated various heuristic improvements listed in introduction. The next SASCA are a more founded solution to get rid of this limitation.

2.6 Soft analytical side-channel attacks

SASCA [32] describe the target block cipher implementation and its leakages in a way similar to a Low-Density Parity Check codes (LDPC) [8]. Since the latter can be decoded using soft decoding algorithms, it implies that SASCA can directly use the posterior probabilities obtained during a TA. Similar to ASCA, they can also be described in three main steps.

1. *Construction*. The cipher is represented as a so-called “factor graph” with two types of nodes and bidirectional edges. First, variable nodes represent the

intermediate values. Second, function nodes represent the a-priori knowledge about the variables (e.g. the known plaintexts and leakages) and the operations connecting the different variables. Those nodes are connected with bidirectional edges that carry two types of messages (i.e. propagate the information) through the graph: the type q message are from variables to factors and the type r messages are from factors to variables (see [15] for more details).

2. *Information extraction.* The description of this phase is trivial. The probabilities provided by TA on any intermediate variable of the encryption process can be directly exploited, and added as a function node to the factor graph.

3. *Decoding.* Similar to LDPC codes, the factor graph is then decoded using the BP algorithm [23]. Intuitively, it essentially iterates the local propagation of the information about the variable nodes of the target implementation.

Since our work is mostly focused on concrete investigations of SASCA, we now describe the BP algorithm in more details. Our description is largely inspired by the description of [15, Chapter 26]. For this purpose, we denote by x_i the i^{th} intermediate value and by f_i the i^{th} function node. As just mentioned, the nodes will be connected by edges that carry two types of messages. The first ones go from a variable node to a function node, and are denoted as $q_{v_n \rightarrow f_m}$. The second ones go from a function node to a variable node, and are denoted as $r_{f_n \rightarrow v_m}$. In both cases, n is the index of the sending node and m the index of the recipient node. The messages carried correspond to the scores for the different values of the variable nodes. At the beginning of the algorithm execution, the messages from variable nodes to function nodes are initialized with no information on the variable. That is, for all n, m and for all x_n we have:

$$q_{v_n \rightarrow f_m}(x_n) = 1.$$

The scores are then updated according to two rules (one per type of messages):

$$r_{f_m \rightarrow v_n}(x_n) = \sum_{x_{n'}, n' \neq n} (f_m(x_{n'}, x_n) \prod_{n'} q_{v_{n'} \rightarrow f_m}(x_{n'})). \quad (1)$$

$$q_{v_n \rightarrow f_m}(x_n) = \prod_{m' \neq m} r_{f_{m'} \rightarrow v_n}(x_n). \quad (2)$$

In Equation 2, the variable node v_n sends the product of the messages about x_n received from the others function nodes ($m' \neq m$) to the function node f_m , for each value of x_n . And in Equation 1, the function node f_m sends a sum over all the possible input values of f_m of the value of f_m evaluated on the vector of $(x_{n'}, n' \neq n)$'s, multiplied by the product of the messages received by f_m for the considered values of $x_{n'}$. The BP algorithm essentially works by iteratively applying these rules on all nodes. If the factor graph is a tree (i.e. if it has no loop), a convergence should occur after a number of iterations at most equal to the diameter of the graph. In case the graph includes loops (e.g. as in our AES implementation case), convergence is not guaranteed, but usually occurs after a number of iterations slightly larger than the graph diameter. The main

parameters influencing the time and memory complexity of the BP algorithm are the number of possible values for each variable (i.e. 2^8 in our 8-bit example) and the number of edges. The time complexity additionally depends on the number of inputs of the function nodes representing the block cipher operations (since the first rule sums over all the possible input combinations of these operations).

3 Comparison with ASCA

ASCA and SASCA are both analytical attacks with very similar descriptions. As previously shown in [32], SASCA have a clear advantage when only noisy information is available. But when the information is noise-free, the advantage of one over the other has not been studied yet. In this section, we therefore tackle the question “which analytical attack is most efficient in noise-free scenario?”. To this end, we compare the results of SASCA and ASCA against a simulated AES implementation with noise-free (Hamming weight) leakages. We first describe the AES representation we used in our SASCA (which will also be used in the following sections), then describe the different settings we considered for our simulated attacks, and finally provide the results of our experiments.

3.1 Our representation for SASCA

As usual in analytical attacks, our description of the AES is based on its target implementation. This allows us to easily integrate the information obtained during its execution. For readability purposes, we start by illustrating the graph representation for the first round of one column of the AES in Figure 1. To build this graph for one plaintext, we start with 32 variable nodes (circles), 16 for the 8-bit subplaintexts (p_i), and 16 for the 8-bit subkeys (k_i). We first add a new variable node in the graph representation each time a new intermediate value is computed in the AES FURIOUS implementation¹, together with the corresponding function nodes (rectangles). There are three different operations that create intermediate values. First, the Boolean XOR takes two variables as inputs and outputs a new variable that is equal to the bitwise XOR of the two inputs. Next, two memory accesses to look-up tables are used for the S-box and Xtimes operations, which take one variable as input, and create a new variable as output. We finally add two types of leaf nodes to these three function nodes. The P’s reflect the knowledge of the plaintext used, and the \mathcal{L} ’s give the posterior probability of the value observed using Gaussian templates. A summary of the different function nodes used in our AES factor graph is given in Table 1.

The graph in Figure 1 naturally extends to a full AES execution. And when using several traces, we just keep a single description of the key scheduling, that links different subgraphs representing the different plaintext encryptions. Our description of the key scheduling requires 226 variable nodes and 210 function nodes. Our description of the rounds requires 1036 variable nodes and 1020

¹ Excluding memory copies which only increase the graph diameter.

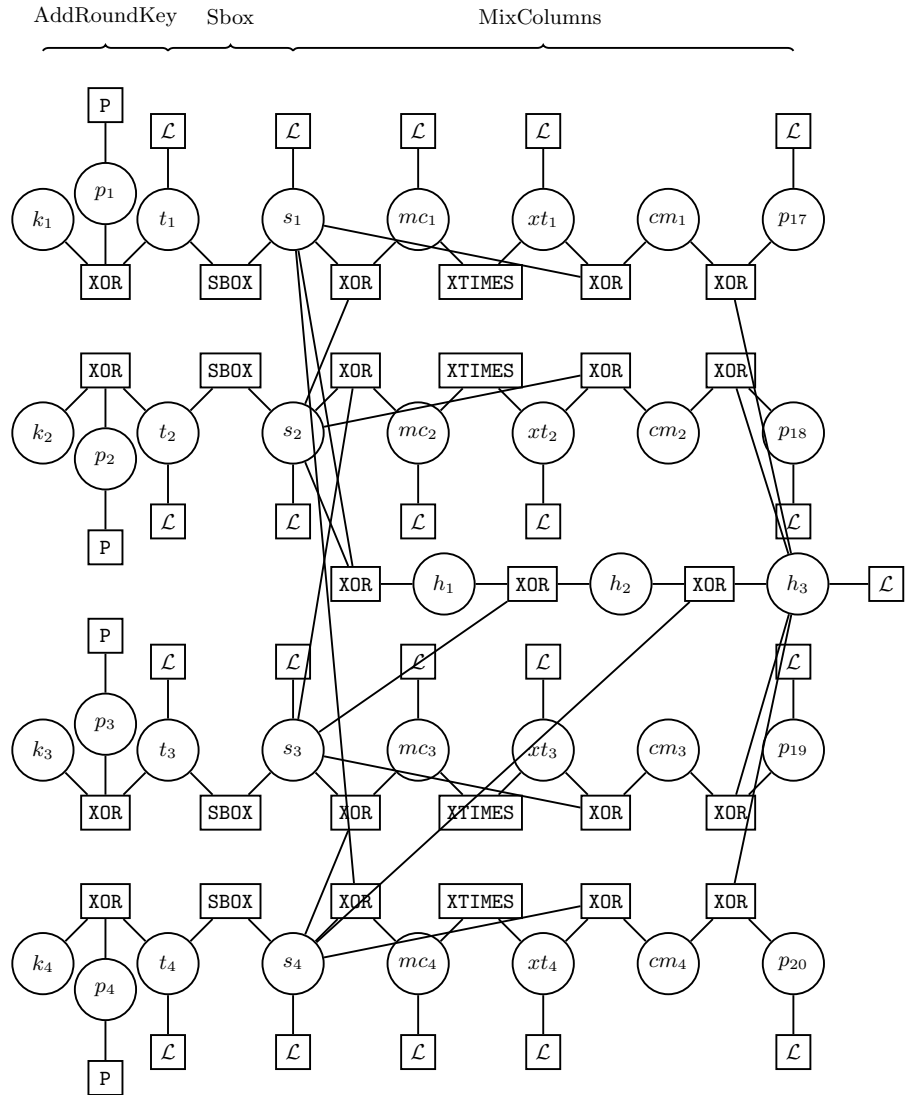


Fig. 1. Graph representation of first round of one column of the AES.

$$\begin{aligned} \text{XOR}(a, b, c) &= \begin{cases} 1 & \text{if } a = b \oplus c, \\ 0 & \text{otherwise.} \end{cases} & \text{SBOX}(a, b) &= \begin{cases} 1 & \text{if } a = \text{sbox}(b), \\ 0 & \text{otherwise.} \end{cases} \\ \text{XTIMES}(a, b) &= \begin{cases} 1 & \text{if } a = \text{xtimes}(b), \\ 0 & \text{otherwise.} \end{cases} & \text{P}(x_n) &= \begin{cases} 1 & \text{if } x_n = p, \\ 0 & \text{otherwise.} \end{cases} \\ \mathcal{L}(x_n) &= \Pr[x_n | \mathbf{l}(t_n)]. \end{aligned}$$

Table 1. Summary of the function nodes used in our AES factor graph.

function nodes. The key scheduling nodes are connected by 580 edges, and each round of the encryption contains 292 edges. As a result and overall, the factor graph for one plaintext contains 1262 variable nodes, 1230 function nodes and 3628 edges. On the top of that we finally add the leakage function nodes which account for up to 1262 edges (if all leakages are exploited). Concretely, each variable node represents an intermediate value that can take 2^8 different values. Hence, if we represent each edge by two tables in single precision of size 256, the memory required is: $256 \times (3628 \times 2 + 1262) \times 4 \text{ bytes} \approx 8\text{MB}$.²

3.2 Comparison setup

Our noise-free evaluations of ASCA and SASCA are based on single-plaintext attacks, which is due to the high memory requirements of ASCA (that hardly extend to more plaintexts). In order to stay comparable with the previous work in [26], we consider a Hamming weight (W_H) leakage function and specify the location of the leakages as follows:

- 16 W_H ’s for AddRoundKey,
- 16 W_H ’s for the output of SubBytes and ShiftRows,
- 36 W_H ’s for the XORs and 16 W_H for the look-up tables in MixColumns.

As previously mentioned, these leakages are represented by \mathcal{L} boxes in Figure 1. We also consider two different contexts for the information extraction:

- Consecutive weights (cw), i.e. the W_H ’s are obtained for consecutive rounds.
- Random weights (rw), i.e. we assume the knowledge of W_H ’s for randomly distributed intermediate values among the 804 possible ones.

Eventually, we analyzed attacks in a Known Plaintext (KP) and Unknown Plaintext (UP) scenario. And in all cases, we excluded the key scheduling leakages, as in [26]. Based on these settings, we evaluated the success rate in function of the quantity of information collected, counted in terms of “rounds of information”, where one round corresponds to 84 W_H ’s of 8-bit values.

3.3 Experimental results

The results of our SASCA with noise-free leakages are reported in Figure 2, and compared to the similar ASCA experiments provided in Reference [26].

² For the leakage nodes, messages from variable to function ($q_{v_n \rightarrow f_m}$) are not necessary.

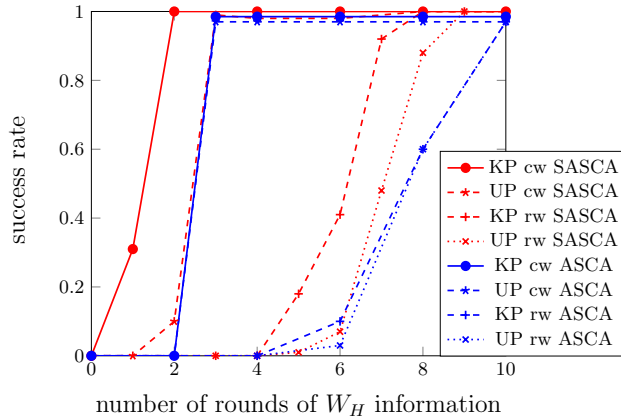


Fig. 2. Experimental results of comparison of ASCA and SASCA.

We first observe that 2 consecutive rounds of W_H 's are enough to recover the key for SASCA with the knowledge of plaintext and when the leakages are located in the first rounds.³ Next, if we do not have access to the plaintext, SASCA requires 3 consecutive rounds of leakage, as for ASCA. By contrast, and as previously underlined, the solving/decoding phase is significantly more challenging in case the leakage information is randomly distributed among the intermediate variables. This is intuitively connected to the fact that the solver and decoder both require to propagate information through the rounds, and that this information can rapidly vanish in case some intermediate variables are unknown. The simplest example is a XOR operation within MixColumns, as mentioned in introduction. So accumulating information on closely connected intermediate computations is always the best approach in such analytical attacks. This effect is of course amplified if the leakages are located in the middle rounds and the plaintext/ciphertext are unknown, as clear from Figure 2.

Overall, and since both SAT-solvers and the BP algorithm with loops in the factor graph are highly heuristic, it is of course difficult to make strong statements about their respective leakage requirements. However, these experiments confirm that at least in the relevant case-study of noise-free Hamming weight leakages, the better noise-resilience of SASCA does not imply weaker performances in a noise-free setting. Besides, and in terms of time complexity, the attacks also differ. Namely, the resolution time for ASCA depends of the quantity of information, whereas it is independent of this quantity in SASCA, and approximately 20 times lower than the fastest resolution times for ASCA.

³ We considered leakages for the two first rounds in this case, which seems the more realistic, and is the only minor differences with the experiments in [26], which considered middle rounds. However, we note that by considering middle round leakages with known plaintext, we then require three rounds of W_H 's, as for ASCA.

Note finally that moving to a noisy scenario can only be detrimental to ASCA. Indeed, and as discussed in [21], ASCA requires correct hard information for the key recovery to succeed. In case of noisy measurements, this can only be guaranteed by considering less informative classes of leakages or similar heuristics. For example, previous works in this direction typically considered Hamming weights h 's between $h - d$ and $h + d$ for increasing distances d 's, which rapidly makes the attack computationally hard. So just as the gain of TASCA over ASCA increases with the measurement noise, the gain of SASCA over ASCA does too.

4 SASCA against a concrete AES implementation

In this section, we complete the previous simulated experiments and explore whether SASCA can be transposed in the more realistic context of measured leakages. To the best of our knowledge, we describe the first uses of SASCA against a concrete AES implementation, and take advantage of this case-study to answer several questions such as *(i)* how to perform the profiling of the many target intermediate values in SASCA?, *(ii)* what happens when the implementation details (such as the source code) are unknown?, and *(iii)* are there significant differences (or even gaps) between concrete and simulated experiments?

4.1 Profiling step

We first describe how to exploit the tools from Section 2.2 in order to detect POIs for our 1230 target intermediate values (which correspond to 1262 variable nodes minus 32 corresponding to the 16 bytes of plaintext and ciphertext). In this context, directly computing the SNRs or CPAs in parallel for all our samples turns out to be difficult. Indeed, the memory requirements to compute the mean trace of an intermediate value with simple precision requires $94,000$ (samples) \times 256 (values) \times 4 (bytes) $\approx 91\text{MB}$, which means approximately 100 GB for the 1,230 values. For similar reasons, computing all these SNRs or CPAs sequentially is not possible (i.e. would require too much time). So the natural option is to trade time and memory by cutting the traces in a number of pieces that hold in RAM. This is easily done if we can assume some knowledge about the implementation (which we did), resulting in a relatively easy profiling step carried out in a dozen of hours on a single desktop computer. A similar profiling could be performed without implementation knowledge, by iteratively testing the intermediate values that should appear sequentially in an AES implementation.

A typical outcome of this profiling is given in Figure 3, where we show the SNR we observed for the intermediate value t_1 from the factor graph in Figure 1 (i.e. the value of the bitwise XOR of the first subkey and the first subplaintext). As intuitively expected, we can identify significant leakages at three different times. The first one, at $t = 20,779$, corresponds to the computation of the value t_1 , i.e. the XOR between p_1 and k_1 . The second one, at $t = 22,077$, corresponds to the computation of the value s_1 , i.e. a memory access to the look-up table of the S-box. The third one, at $t = 24,004$, corresponds to memory copies of




Fig. 3. SNR-based profiling of a single intermediate value.

s_1 during the computation of `MixColumns`. Indeed, the SNR cannot tell apart intermediate values that are bijectively related. So we used the CPA distinguisher to get rid of this limitation (taking advantage that a simple Hamming weight leakage model was applicable against our target implementation).

A summary of the results obtained after our profiling step is given in Table 2, where the most interesting observation is that the informativeness of the leakage samples strongly depends on the target intermediate values. In particular, we see that memory accesses allow SNRs over 2, while XOR operations lead to SNRs below 0.4 (and this SNR is further reduced in case of consecutive XOR operations). This is in strong contrast, with the simulated cases (in the previous section and in [32]), where all the variables were assumed to leak with the same SNR. Note that the table mentions both SNR and CPA values, though our selection of POIs was based on the (more generic) first criteria, and CPA was only used to separate the POIs of bijectively related intermediate values.

4.2 Experimental results

Taking advantage of the previous POI detection, we now want to discuss the consequences of different assumptions about the implementation knowledge. These investigations are motivated by the usual gap between Kerckhoff’s laws [13], which advises to keep the key as only secret in cryptography, and the practice in embedded security, that usually takes advantage of some obscurity regarding the implementations. For this purpose, we considered three adversaries:

1. *Informed*. The adversary has access to the implementation details (i.e. source code), and can exploit the leakages of all the target intermediate values.
2. *Informed, but excluding the key scheduling*. This is the same case as the previous one, but we exclude the key scheduling leakages as in the simulations of the previous section (e.g. because round keys are precomputed).

Assembly code	Graph description	SNR	$\rho(W_H)$
Add Round Key			
ld H1, Y+	*	*	*
eor ST11, H1	_Xor t1 p1 k1	0.1493	0.5186
Sbox			
ldi ZH, high(sbox<<1)	*	*	*
mov ZL, ST11	*	*	*
lpm ST11, Z	_Sbox s1 t1	1.6301	0.4766
MixColumns			
ldi ZH, high(xtime<<1)	*	*	*
mov H1, ST11	*	*	*
eor H1, ST21	_Xor h1 s1 s2	0.1261	0.6158
eor H1, ST31	_Xor h2 h1 s3	0.0391	0.1449
eor H1, ST41	_Xor h3 h2 s4	0.3293	0.5261
mov H2, ST11	*	*	*
mov H3, ST11	*	*	*
eor H3, ST21	_Xor mc1 s1 s2	0.2802	0.6163
mov ZL, H3	*	*	*
lpm H3, Z	_Xtime xt1 mc1	2.8650	0.6199
eor ST11, H3	_Xor cm1 xt1 s1	0.0723	0.2508
eor ST11, H1	_Xor p17 cm1 h3	0.1064	0.3492
Key Schedule			
ldi H1, 1	*	*	*
ldi ZH, high(sbox<<1)	*	*	*
mov ZL, ST24	*	*	*
lpm H3, Z	_Sbox sk14 k14	2.2216	0.5553
eor ST11, H3	_Xor ak1 sk14 k1	0.1158	0.5291
eor ST11, H1	_XorCste k17 ak1 1	0.3435	0.5140

Table 2. Summary of profiling step results.

3. *Uninformed.* Here the adversary only knows the AES is running, and will only exploit the generic operations as described in the algorithm specifications (i.e. the inputs and outputs of `AddRoundKey`, `ByteSub`, `ShiftRows` and `MixColumns`, together with the key rounds' inputs and outputs).

In order to have fair comparisons, we used the same profiling for all three cases (i.e. we just excluded some POIs for cases 2 and 3), and we used 100 sets of 30 traces with different keys and plaintexts to calculate the success rate of SASCA in these different conditions. The results of our experiments are in Figure 4. Our first and main observation is that SASCA are applicable to actual implementations, for which the leakages observed provide more or less information (and SNR) depending on the intermediate values. As expected, the informed adversary is the most powerful. But we also see that excluding the key scheduling leakages, or considering an uninformed adversary, only marginally reduces the attack success rates. Interestingly, there is a strong correlation between this success rate and the number of leakage samples exploited, since excluding the key scheduling implies the removal of 226 leakage function nodes, and the un-

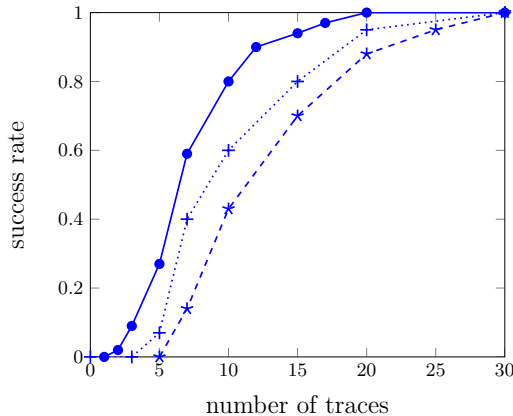


Fig. 4. Success rate in function of the # of traces for different adversaries: informed one (—●—), informed one without key scheduling leakages (⋯+⋯) and uninformed one (-*-).

informed adversary has 540 leakage function nodes less than the informed one (mostly corresponding to the `MixColumns` operation). So we can conclude that SASCA are not only a threat for highly informed adversaries, and in fact quite generically apply to software implementations with many leaking points.

Simulation vs. measurement. In view of the previous results, with information leakages depending on the target intermediate values, a natural question is whether security against SASCA was reasonably predicted with a simulated analysis. Of course, we know that in general, analytical attacks are much harder to predict than DPA [25], and do not enjoy simple formulas for the prediction of their success rates [17]. Yet, we would like to study informally the possible connection between simple simulated analyses and concrete ones. For this purpose, we compare the results obtained in these two cases in Figure 5. For readability, we only report results for the informed and uninformed cases, and consider different SNRs for the simulated attacks. In this context, we first recall Table 2 where the SNRs observed for our AES implementation vary between 2^1 and 2^{-2} . Interestingly, we see from Figure 5 that the experimental success rate is indeed bounded by these extremes. (Tighter and more rigorous bounds are probably hard to obtain for such heuristic attacks). Besides, we also observe that the success rates of the measurements and simulations are closer in the case of the uninformed adversary, which can be explained by the fact that we essentially ignore `MixColumns` leakages in this case, for which the SNRs are lower.

5 Comparison with DPA and enumeration

In this section, we start from the observation that elegant approaches to side-channel analysis generally require more computational power than standard




Fig. 5. Experimental results for SASCA for an informed adversary (a) and uninformed adversary (b). Red curves are for simulated cases (\times , $+$, $*$, \circ) for SNR ($2^1, 2^{-1}, 2^{-2}, 2^{-3}$). Blue curves (\bullet) are for experiments on real traces.

DPA. Thus, a fair comparison between both approaches should not only look at the success rate in function of the number of traces, but also take into account the resolution time as a parameter. As a result, and in order to compare SASCA and the pragmatic DPA on a sound basis, this section investigates the result of DC attacks combined with computational power for key enumeration.

5.1 Evaluation of profiled template attacks

In order to be as comparable as possible with the previous SASCA, our comparison will be based on the profiled TA described in Section 2.3.⁴ More precisely, we considered a quite pragmatic DC attack exploiting the bivariate leakages corresponding to the `AddRoundKey` and `ByteSub` operations (i.e. $\{s_i\}_{i=1}^{16}$ and $\{t_i\}_{i=1}^{16}$ in Figure 1). We can take advantage of the same detection of POIs as described in the previous section for this purpose. As previously mentioned, this choice allows us to keep the computational complexity of the TA itself very minimal (since relying only on 8-bit hypotheses). Following, we built the security graph of our bivariate TA, as represented in Figure 6, where the white (resp. black) curve corresponds to the maximum (resp. minimum) rank observed, and the red curve is for the average rank. It indicates that approximately 60 plaintexts are required to recover the key without any enumeration. But more interestingly, the graph also highlights that allowing enumeration up to ranks (e.g.) 2^{30} allows to reduce the required number of traces down to approximately 10.

⁴ We considered TA for our DPA comparison because they share the same profiled setting as ASCA. Comparisons with a non-profiled CPA can only be beneficial to SASCA. More precisely, we expect a typical loss factor of 2 to 5 between (W_H -based) CPA and TA, according to the results in [29] obtained on the same device.

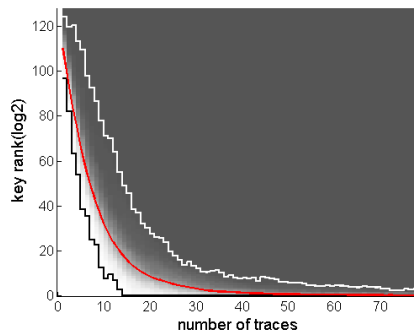


Fig. 6. Security graph of a bivariate TA.

5.2 Comparing SASCA and DPA with enumeration

In our prototype implementation running on a desktop computer, SASCA requires roughly one second per plaintext, and reaches a success rate of one after 20 plaintexts (for the informed adversary). In order to allow reasonably fair comparisons, we first measured that the same desktop computer can perform a bit more than 2^{20} AES encryptions in 20 seconds. So this is typically the amount of enumeration that we should grant the bivariate TA for comparisons with SASCA.⁵ For completeness, we also considered the success rates of bivariate TA without enumeration and with 2^{30} enumeration power. The results of these last experiments are in Figure 7. Overall, they bring an interesting counterpart to our previous investigations. On the one hand, we see that SASCA remains the most powerful attack when the adversary has enough knowledge of the implementation. By contrast in the uninformed case, the gain over the pragmatic TA with enumeration is lower. So as expected, it is really the amount and type of leakage samples exploitable by the adversary that make SASCA more or less powerful, and determine their interest (or lack thereof) compared to DC attacks. In this respect, a meaningful observation is that the gap between SASCA and DPA without enumeration (here approximately 5) is lower than the approximate factor 10 that was observed in the previous simulations of [32]. This difference is mainly due to the lower SNRs observed in the MixColumns transform.

Eventually, we note that in view of these results, another natural approach would be to use enumeration for SASCA. Unfortunately, our experiments have shown that enumeration is much less effective in the context of analytical attacks. This is essentially caused by the fact that DC attacks consider key bytes independently, whereas SASCA decode the full key at once, which implies that the

⁵ We omit to take the (time and memory) resources required for the generation of the list of the most probable keys to enumerate into account in our comparisons, since these resources remain small in the total enumeration cost. Using the state-of-the-art enumeration algorithm [30], we required 2.7MB + 0.55 seconds to generate a list of 2^{20} keys, and 1.8GB + 3130 seconds to generate a list of 2^{32} keys.

subkey probabilities are not independent in this case. Possible tracks to improve this issue include the use of list decoding algorithms for LDPC codes (as already mentioned in [9]), or enumeration algorithms that can take subkey dependencies into account (as suggested in [14] for elliptic curve implementations).

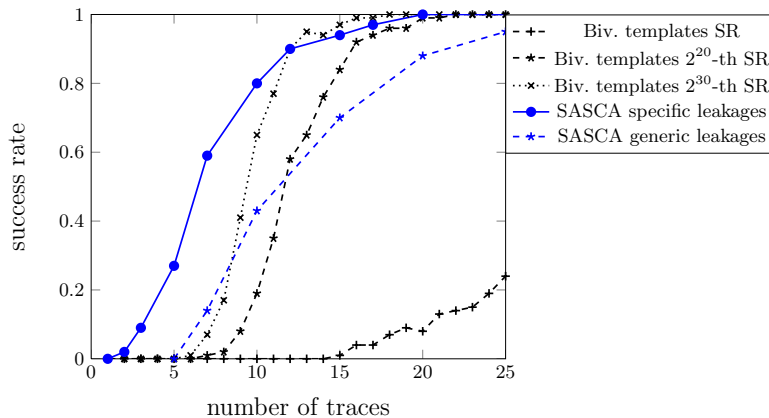


Fig. 7. Comparison between elegant and pragmatic approaches.

6 Conclusion and open problems

This paper puts forward that the technicalities involved in elaborate analytical side-channel attacks, such as the recent SASCA, are possible to solve in practice. In particular, our results show that the intensive profiling of many target intermediate values within an implementation is achievable with the same (SNR & CPA) tools as any profiled attack (such as the bivariate TA we considered). This profiling only requires a dozen of hours to complete, and then enables very efficient SASCA that recover the key of our AES implementation in a couple of seconds and traces, using a single desktop computer. Furthermore, these successful attacks are even possible in a context where limited knowledge about the target implementation is available, hence mitigating previous intuitions regarding analytical attacks being “only theoretical”. Besides this positive conclusion, a fair comparison with DC attacks also highlights that the gap between a bivariate TA and a SASCA can be quite reduced in case enumeration power is granted to the DC adversary, and several known plaintexts are available. Intuitively, the important observation in this respect is that the advantage of SASCA really depends on the amount and type of intermediate values leaking information, which highly depends on the algorithms and implementations analyzed.

The latter observation suggests two interesting directions for further research. On the one hand, the AES Rijndael is probably among the most challenging targets for SASCA. Indeed, it includes a strong linear diffusion layer, with many

XOR operations through which the information propagation is rapidly amortized. Besides, it also relies on a non-trivial key scheduling, which prevents the direct combination of information leaked from multiple rounds. So it is not impossible that the gap between SASCA and standard DPA could be larger for other ciphers (e.g. with permutation based diffusion layers [3], and very minimum key scheduling algorithms [4]). On the other hand, since the the propagation of the leakage information through the MixColumns operation is hard(er), one natural solution to protect the AES against such attacks would be to enforce good countermeasures for this part of the cipher, which would guarantee that SASCA do not exploit more information than the one of a single round. Ideally, and if one can prevent any information propagation beyond the cipher rounds, we would then have a formal guarantee that SASCA is equivalent to DPA.

Acknowledgements. F.-X. Standaert is a research associate of the Belgian Fund for Scientific Research (FNRS-F.R.S.). This work has been funded in parts by the European Commission through the ERC project 280141 (CRASH).

References

1. Valentina Banciu and Elisabeth Oswald. Pragmatism vs. elegance: Comparing two approaches to simple power attacks on AES. In Emmanuel Prouff, editor, *Constructive Side-Channel Analysis and Secure Design - 5th International Workshop, COSADE 2014, Paris, France, April 13-15, 2014. Revised Selected Papers*, volume 8622 of *Lecture Notes in Computer Science*, pages 29–40. Springer, 2014.
2. Sonia Belaïd, Vincent Grosso, and François-Xavier Standaert. Masking and leakage-resilient primitives: One, the other(s) or both? *Cryptography and Communications*, 7(1):163–184, 2015.
3. Andrey Bogdanov, Lars R. Knudsen, Gregor Leander, Christof Paar, Axel Poschmann, Matthew J. B. Robshaw, Yannick Seurin, and C. Vikkelsoe. PRESENT: an ultra-lightweight block cipher. In Pascal Paillier and Ingrid Verbauwhede, editors, *Cryptographic Hardware and Embedded Systems - CHES 2007, 9th International Workshop, Vienna, Austria, September 10-13, 2007, Proceedings*, volume 4727 of *Lecture Notes in Computer Science*, pages 450–466. Springer, 2007.
4. Andrey Bogdanov, Lars R. Knudsen, Gregor Leander, François-Xavier Standaert, John P. Steinberger, and Elmar Tischhauser. Key-alternating ciphers in a provable setting: Encryption using a small number of public permutations - (extended abstract). In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, volume 7237 of *Lecture Notes in Computer Science*, pages 45–62. Springer, 2012.
5. Eric Brier, Christophe Clavier, and Francis Olivier. Correlation power analysis with a leakage model. In Marc Joye and Jean-Jacques Quisquater, editors, *Cryptographic Hardware and Embedded Systems - CHES 2004: 6th International Workshop Cambridge, MA, USA, August 11-13, 2004. Proceedings*, volume 3156 of *Lecture Notes in Computer Science*, pages 16–29. Springer, 2004.
6. Claude Carlet, Jean-Charles Faugère, Christopher Goyet, and Guénaél Renault. Analysis of the algebraic side channel attack. *J. Cryptographic Engineering*, 2(1):45–62, 2012.

7. Suresh Chari, Josyula R. Rao, and Pankaj Rohatgi. Template attacks. In Burton S. Kaliski Jr., Çetin Kaya Koç, and Christof Paar, editors, *Cryptographic Hardware and Embedded Systems - CHES 2002, 4th International Workshop, Redwood Shores, CA, USA, August 13-15, 2002, Revised Papers*, volume 2523 of *Lecture Notes in Computer Science*, pages 13–28. Springer, 2002.
8. Robert G. Gallager. Low-density parity-check codes. *IRE Transactions on Information Theory*, 8(1):21–28, 1962.
9. Benoît Gérard and François-Xavier Standaert. Unified and optimized linear collision attacks and their application in a non-profiled setting: extended version. *J. Cryptographic Engineering*, 3(1):45–58, 2013.
10. Benedikt Gierlichs, Lejla Batina, Pim Tuyls, and Bart Preneel. Mutual information analysis. In Elisabeth Oswald and Pankaj Rohatgi, editors, *Cryptographic Hardware and Embedded Systems - CHES 2008, 10th International Workshop, Washington, D.C., USA, August 10-13, 2008. Proceedings*, volume 5154 of *Lecture Notes in Computer Science*, pages 426–442. Springer, 2008.
11. Cezary Glowacz, Vincent Grosso, Romain Poussier, Joachim Schueth, and François-Xavier Standaert. Simpler and more efficient rank estimation for side-channel security assessment. *IACR Cryptology ePrint Archive*, 2014:920, 2014.
12. Shize Guo, Xinjie Zhao, Fan Zhang, Tao Wang, Zhijie Jerry Shi, François-Xavier Standaert, and Chujiao Ma. Exploiting the incomplete diffusion feature: A specialized analytical side-channel attack against the AES and its application to microcontroller implementations. *IEEE Transactions on Information Forensics and Security*, 9(6):999–1014, 2014.
13. Auguste Kerckhoffs. *La cryptographie militaire, ou, Des chiffres usités en temps de guerre: avec un nouveau procédé de déchiffrement applicable aux systèmes à double clef*. Librairie militaire de L. Baudoin, 1883.
14. Tanja Lange, Christine van Vredendaal, and Marnix Wakker. Kangaroos in side-channel attacks. *IACR Cryptology ePrint Archive*, 2014:565, 2014.
15. David JC MacKay. *Information theory, inference, and learning algorithms*, volume 7. Citeseer, 2003.
16. Stefan Mangard. Hardware countermeasures against DPA ? A statistical analysis of their effectiveness. In Tatsuaki Okamoto, editor, *Topics in Cryptology - CT-RSA 2004, The Cryptographers' Track at the RSA Conference 2004, San Francisco, CA, USA, February 23-27, 2004, Proceedings*, volume 2964 of *Lecture Notes in Computer Science*, pages 222–235. Springer, 2004.
17. Stefan Mangard, Elisabeth Oswald, and François-Xavier Standaert. One for all - all for one: unifying standard differential power analysis attacks. *IET Information Security*, 5(2):100–110, 2011.
18. Luke Mather, Elisabeth Oswald, and Carolyn Whitnall. Multi-target DPA attacks: Pushing DPA beyond the limits of a desktop computer. In Sarkar and Iwata [27], pages 243–261.
19. Mohamed Saied Emam Mohamed, Stanislav Bulygin, Michael Zohner, Annelie Heuser, Michael Walter, and Johannes Buchmann. Improved algebraic side-channel attack on AES. *J. Cryptographic Engineering*, 3(3):139–156, 2013.
20. Yossef Oren, Mario Kirschbaum, Thomas Popp, and Avishai Wool. Algebraic side-channel analysis in the presence of errors. In Stefan Mangard and François-Xavier Standaert, editors, *Cryptographic Hardware and Embedded Systems, CHES 2010, 12th International Workshop, Santa Barbara, CA, USA, August 17-20, 2010. Proceedings*, volume 6225 of *Lecture Notes in Computer Science*, pages 428–442. Springer, 2010.

21. Yossef Oren, Mathieu Renauld, François-Xavier Standaert, and Avishai Wool. Algebraic side-channel attacks beyond the hamming weight leakage model. In Emmanuel Prouff and Patrick Schaumont, editors, *Cryptographic Hardware and Embedded Systems - CHES 2012 - 14th International Workshop, Leuven, Belgium, September 9-12, 2012. Proceedings*, volume 7428 of *Lecture Notes in Computer Science*, pages 140–154. Springer, 2012.
22. Yossef Oren, Ofir Weisse, and Avishai Wool. A new framework for constraint-based probabilistic template side channel attacks. In Lejla Batina and Matthew Robshaw, editors, *Cryptographic Hardware and Embedded Systems - CHES 2014 - 16th International Workshop, Busan, South Korea, September 23-26, 2014. Proceedings*, volume 8731 of *Lecture Notes in Computer Science*, pages 17–34. Springer, 2014.
23. Judea Pearl. Reverend bayes on inference engines: A distributed hierarchical approach. In David L. Waltz, editor, *Proceedings of the National Conference on Artificial Intelligence. Pittsburgh, PA, August 18-20, 1982.*, pages 133–136. AAAI Press, 1982.
24. Mathieu Renauld and François-Xavier Standaert. Algebraic side-channel attacks. In Feng Bao, Moti Yung, Dongdai Lin, and Jiwu Jing, editors, *Information Security and Cryptology - 5th International Conference, Inscrypt 2009, Beijing, China, December 12-15, 2009. Revised Selected Papers*, volume 6151 of *Lecture Notes in Computer Science*, pages 393–410. Springer, 2009.
25. Mathieu Renauld and Francois-Xavier Standaert. Representation-, leakage-and cipher-dependencies in algebraic side-channel attacks. *industrial track of ACNS*, 2010:1–18, 2010.
26. Mathieu Renauld, François-Xavier Standaert, and Nicolas Veyrat-Charvillon. Algebraic side-channel attacks on the AES: why time also matters in DPA. In Christophe Clavier and Kris Gaj, editors, *Cryptographic Hardware and Embedded Systems - CHES 2009, 11th International Workshop, Lausanne, Switzerland, September 6-9, 2009. Proceedings*, volume 5747 of *Lecture Notes in Computer Science*, pages 97–111. Springer, 2009.
27. Palash Sarkar and Tetsu Iwata, editors. *Advances in Cryptology - ASIACRYPT 2014 - 20th International Conference on the Theory and Application of Cryptology and Information Security, Kaoshiung, Taiwan, R.O.C., December 7-11, 2014. Proceedings, Part I*, volume 8873 of *Lecture Notes in Computer Science*. Springer, 2014.
28. Werner Schindler, Kerstin Lemke, and Christof Paar. A stochastic model for differential side channel cryptanalysis. In Josyula R. Rao and Berk Sunar, editors, *Cryptographic Hardware and Embedded Systems - CHES 2005, 7th International Workshop, Edinburgh, UK, August 29 - September 1, 2005. Proceedings*, volume 3659 of *Lecture Notes in Computer Science*, pages 30–46. Springer, 2005.
29. François-Xavier Standaert, Benedikt Gierlichs, and Ingrid Verbauwhede. Partition vs. comparison side-channel distinguishers: An empirical evaluation of statistical tests for univariate side-channel attacks against two unprotected CMOS devices. In Pil Joong Lee and Jung Hee Cheon, editors, *Information Security and Cryptology - ICISC 2008, 11th International Conference, Seoul, Korea, December 3-5, 2008, Revised Selected Papers*, volume 5461 of *Lecture Notes in Computer Science*, pages 253–267. Springer, 2008.
30. Nicolas Veyrat-Charvillon, Benoît Gérard, Mathieu Renauld, and François-Xavier Standaert. An optimal key enumeration algorithm and its application to side-channel attacks. In Lars R. Knudsen and Huapeng Wu, editors, *Selected Areas in Cryptography, 19th International Conference, SAC 2012, Windsor, ON, Canada,*

August 15-16, 2012, Revised Selected Papers, volume 7707 of *Lecture Notes in Computer Science*, pages 390–406. Springer, 2012.

31. Nicolas Veyrat-Charvillon, Benoît Gérard, and François-Xavier Standaert. Security evaluations beyond computing power. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, volume 7881 of *Lecture Notes in Computer Science*, pages 126–141. Springer, 2013.
32. Nicolas Veyrat-Charvillon, Benoît Gérard, and François-Xavier Standaert. Soft analytical side-channel attacks. In Sarkar and Iwata [27], pages 282–296.
33. Xin-jie Zhao, Fan Zhang, Shize Guo, Tao Wang, Zhijie Shi, Huiying Liu, and Keke Ji. MDASCA: an enhanced algebraic side-channel attack for error tolerance and new leakage model exploitation. In Werner Schindler and Sorin A. Huss, editors, *Constructive Side-Channel Analysis and Secure Design - Third International Workshop, COSADE 2012, Darmstadt, Germany, May 3-4, 2012. Proceedings*, volume 7275 of *Lecture Notes in Computer Science*, pages 231–248. Springer, 2012.