

Adaptive Chosen-Message Side-Channel Attacks

Nicolas Veyrat-Charvillon* and François-Xavier Standaert**

Université catholique de Louvain, Crypto Group, Belgium
nicolas.veyrat, fstandae@uclouvain.be

Abstract. Most side-channel attacks that have been published in the open literature assume known- or chosen-message adversarial scenarios. In this paper, we analyze the increase of the attacks' efficiencies that can be obtained by adaptively selecting the messages. For this purpose, we first describe a generic strategy that allows an adversary to take advantage of this capability. We show that it can be applied to any differential power or electromagnetic analysis attack, against unprotected or protected devices and exploiting profiled or non-profiled leakage models. Then, we provide various experiments to quantify these improvements. Finally, we discuss the optimality of our strategy and its implications for the security evaluation of leakage-resilient cryptographic hardware.

1 Introduction

In classical cryptanalysis, the adaptive selection of the inputs to a cryptographic primitive is known to be a powerful ability for the adversaries. For example, blockwise-adaptive chosen-message attacks have been used to show the insecurity of different encryption modes in [9]. Similarly, Bleichenbacher has demonstrated in [3] that chosen-ciphertext attacks can be used to attack the RSA Encryption Standard PKCS #1. And in the symmetric setting, boomerang attacks are an example of how the adaptivity can be exploited to reduce the complexity of certain categories of attacks [24]. Quite surprisingly, and although it is frequently suggested as a possible improvement, very few related works have been performed in the context of side-channel attacks. In [20], Schindler presented a timing attack against RSA with the Chinese remainder theorem that requires some form of adaptivity. And more recently, Köpf and Basin provided a careful model and analysis of such an attack scenario. But the investigations in [12] are carried out in a restricted context of noiseless leakage. This typically applies to timing attacks such as [10], but is of limited interest in the case of power or electromagnetic side-channel attacks, in which noise is a typical issue adversaries have to deal with [1,11,19]. To the best of the authors's knowledge, the generalization of this previous work, from the context of deterministic leakages to the one of probabilistic (or noisy) leakages was left as an open question.

In this paper, we consequently tackle this problem and propose a careful investigation of adaptive chosen-message side-channel attacks. We describe a generic

* Work supported by the Walloon region through the project SCEPTIC.

** Associate researcher of the Belgian Fund for Scientific Research (F.R.S.-FNRS).

strategy that can be applied to improve the efficiency of any distinguisher. As an illustration, we detail its impact for correlation and template attacks [4,5], both from simulations and actual experiments. Our evaluations show significant increases of the side-channel key-recovery success rates. We additionally evaluate the application of adaptive strategies against implementations protected with masking [8] and observe very similar improvements. Eventually, we discuss the optimality of our approach and compare it with the one in [12].

These results imply direct consequences for the good security evaluations of leaking cryptographic devices. They show that when applicable, adaptive strategies require to take larger security margins than for random-message attacks. In other words, the attacks described in this paper indicate how to best exploit the physical information leakage, in the standard DPA setting formalized in [15]. Hence, they can be used to determine the worst-case number of measurements required to performed a successful key-recovery in this context. We note that considering key-recovery attacks in security evaluations can appear as too weak from a theoretical point of view. But as demonstrated in [2], there is a strong relationship between distinguishing attacks and key-recovery attacks in the context of block ciphers. Hence, this situation is not very different than the one in classical (*e.g.* linear, differential) cryptanalysis, in which one considers the best available attacks in order to approximate the security of a cipher. And the adaptive strategies presented in this paper are part of these evaluation tools.

Note also that our results have interesting connections with recent works in the area of leakage resilient cryptography. Indeed, one of the assumptions in, *e.g.* [6,16], is that the leakage function can be adaptively selected by the adversary. But as discussed in [23], this is a quite strong requirement that is rarely observed: it would require that the adversary can change his measurement setup in a constructive manner. In practice, most attacks rather rely on a fixed leakage function and measure this function for different plaintexts. It is the combination of several plaintexts that allows increasing the information leakage in such a way that the key is eventually revealed. Hence, our adaptive selection of the messages provides a more realistic counterpart of the adversarial capabilities.

2 Terminology and Notations

In a side-channel attack, an adversary tries to recover some secret information from a leaking implementation, *e.g.* a software program or an IC computing a cryptographic algorithm. In this paper, we focus on the divide-and-conquer strategies that are most frequently considered in the literature [14] and are formalized as “standard DPA attacks” in [15]. In the context of a block cipher implementation (that will be our running example), one typically targets small pieces of the master key - called subkeys in the following - one by one. The attacks then follow the different steps illustrated in Figure 1.

Namely, we consider a device performing several cryptographic computations $E_k(x_i)$ on different plaintexts x_i drawn from the text space \mathcal{X} , using some fixed key k drawn from the key space \mathcal{K} . While computing $E_k(x_i)$, the device handles

some intermediate values that depend on the known input x_i and the unknown key k (defined as sensitive variables in [18]). In practice, the interesting sensitive variables in a DPA attack are the ones that depend on an enumerable subkey s : we denote them as v_i^s , for a plaintext x_i . Any time such a sensitive intermediate value is computed, the device generates some physical leakage, denoted as l_i^k (where the k superscript indicates that the leakage potentially depend on all the key k , including the subkey s). Hence, in order to perform a key-recovery, an adversary first has to select a sensitive value. Given that this variable only depends on a subkey s , he can then predict its result for the plaintexts x_i that generated l_i^k and enumerate every possible subkey candidate $s^* \in \mathcal{S}$. This leads to different hypothetical intermediate variables $v_i^{s^*}$. Afterwards, the adversary exploits a leakage model to map these values from their original space \mathcal{V} towards a modeled leakage space \mathcal{M} . As a result, he obtains $|\mathcal{S}|$ different models, denoted

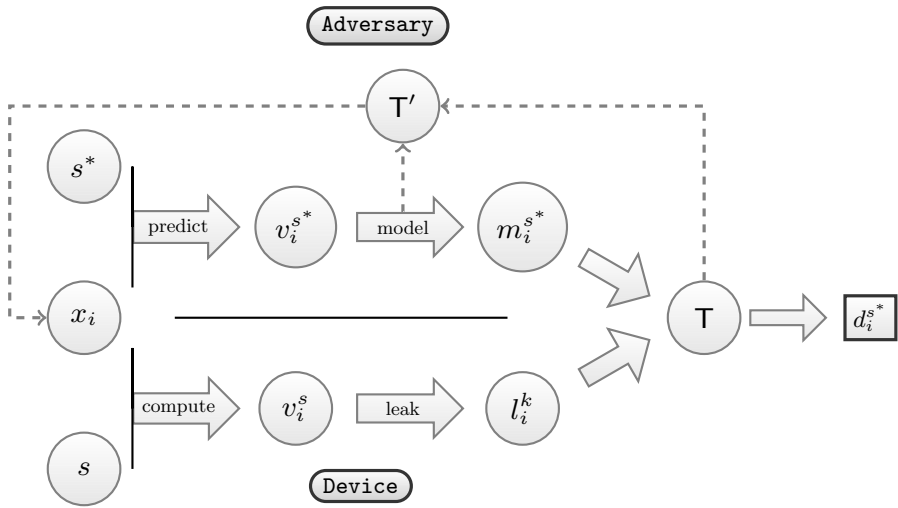


Fig. 1. Schematic description of a side-channel key-recovery attack

as $m_i^{s^*}$, again corresponding to the different subkey candidates. Eventually, he uses a statistical test T to compare the different models $m_i^{s^*}$ with the actual leakages l_i^k . If the attack is successful, the highest value for this test should occur for the correct subkey candidate $s^* = s$. This procedure can be repeated for different subkeys in order to recover the complete key k .

In view of this description, there are several important parameters that determine the efficiency of a DPA. First, the choice of an intermediate computation and leakage model have a significant impact. For example, it is well known that predicting the first round S-boxes' outputs in a block cipher leads to a better discrimination of the subkeys than predicting their inputs [17]. As for the leakage models, it mainly relates to the a-priori knowledge of the adversary about the

device he targets. One generally distinguishes profiled and non-profiled attacks. In the first ones (*e.g.* template attacks [5]), the adversary can characterize the leakage probability density functions (pdf for short) prior to the online attack. In the second ones, he exploits simpler models (*e.g.* predicting only certain moments of the leakage pdf, as in correlation attacks [4]) or performs the profiling “on-the-fly” [7]. Second, and closely related, the choice of a statistical test is usually determined by the type of models available to the adversary.

Another parameter that is less frequently considered (and evaluated) in the literature is the selection of the plaintexts. That is, in most experimental settings, one generally considers attacks with random input messages. But as illustrated in Figure 1, a more powerful scenario is to adaptively select the plaintexts, in function of the prior knowledge about the secret subkey and an hypothetical leakage model. In this paper, we consequently investigate the statistical tests T' that can be used in order to best exploit the available leakage.

3 Adaptive Template Attacks

In this section, we present the principles of our adaptive chosen-message strategy. We first describe it in the (profiled) context of template attacks. Then, we discuss how to generalize our solution to non-profiled distinguishers.

3.1 Template Attacks

Template attacks, first published in [5], are usually considered as the most powerful type of side-channel attacks, in an information theoretic sense. They work in two main steps. In a first profiling phase, the adversary builds key-dependent templates, *i.e.* he estimates the leakage pdf for different internal configurations of his target device. Then, in a second (online attack) phase, he uses these templates to perform a maximum-likelihood key-recovery. In this paper, we focus on the (most frequently considered) case of Gaussian templates.

Templates construction. Gaussian template attacks assign a Gaussian distribution to a number of different configurations of the target device. In their most generic form, they perform this assignment exhaustively. For example, if an adversary targets the 8 first bits of an AES master key, he will use one Gaussian for any pair (x_i, s^*) , out of the 2^{16} possible ones. In practice, different tricks can be used to reduce this number of templates, in order to increase the efficiency of the profiling, *e.g.* by taking advantage of symmetry properties and stochastic models [21]. In this section, we describe the generic approach for simplicity. Suppose that the adversary is provided with N_p traces to estimate the pdf corresponding to a state (x_i, s) . He will then assume that the leakage traces $\{l_i^{k,j}\}_{j=1}^{N_p}$ are drawn from the multivariate normal distribution:

$$\mathcal{N}(l_i^{k,j} | \mu_{x_i}^s, \Sigma_{x_i}^s) = \frac{1}{(2\pi)^{\frac{N}{2}} |\Sigma_{x_i}^s|^{\frac{1}{2}}} \exp \left\{ -\frac{1}{2} (l_i^{k,j} - \mu_{x_i}^s)^\top (\Sigma_{x_i}^s)^{-1} (l_i^{k,j} - \mu_{x_i}^s) \right\},$$

where the mean $\boldsymbol{\mu}_{x_i}^s$ and the covariance matrix $\boldsymbol{\Sigma}_{x_i}^s$ specify completely the noise distribution associated to the leakage trace of each pair (x_i, s) . The templates are built by estimating the sets of parameters $\boldsymbol{\mu}_{x_i}^s$ and $\boldsymbol{\Sigma}_{x_i}^s$ for $x_i \in \mathcal{X}$ and $s \in \mathcal{S}$. Maximum likelihood estimators can be used for this purpose: $\hat{\boldsymbol{\mu}}_{x_i}^s = \frac{1}{N_p} \sum_{j=1}^{N_p} l_i^{k,j}$, and $\hat{\boldsymbol{\Sigma}}_{x_i}^s = \frac{1}{N_p} \sum_{j=1}^{N_p} (l_i^{k,j} - \hat{\boldsymbol{\mu}}_{x_i,s}) (l_i^{k,j} - \hat{\boldsymbol{\mu}}_{x_i,s})^\top$.

Online attack. Assume now that there are $|\mathcal{S}|$ possible subkeys. To determine which one is the most likely to have generated a new trace l_i^k , we compute:

$$\tilde{s} = \underset{s^*}{\operatorname{argmax}} \hat{\Pr}[s^* | l_i^k] = \underset{s^*}{\operatorname{argmax}} \hat{\Pr}[l_i^k | s^*, x_i] \cdot \hat{\Pr}^{(0)}[s^*],$$

where $\hat{\Pr}[l_i^k | s^*, x_i] = \mathcal{N}(l_i^k | \hat{\boldsymbol{\mu}}_{x_i}^{s^*}, \hat{\boldsymbol{\Sigma}}_{x_i}^{s^*})$ and $\hat{\Pr}^{(0)}[s^*]$ is the a priori probability of the subkey candidate s^* , that we assume to be uniform in the following (*i.e.* equal to $1/|\mathcal{S}|, \forall s^*$). In other words, the classification rule assigns l_i^k to the candidate s^* with the highest a posteriori probability. Since in practice, a single trace is usually not enough to recover the subkey with high confidence, the adversary finally combines several plaintexts and computes $\tilde{s} = \underset{k^*}{\operatorname{argmax}} \hat{\Pr}^{(q)}[s^*]$, with:

$$\hat{\Pr}^{(q)}[s^*] = \frac{\prod_{i=1}^q \Pr[s^* | l_i^k]}{\sum_{s' \in \mathcal{S}} \prod_{i=1}^q \Pr[s' | l_i^k]},$$

and q the number of traces used in the online attack. Note that in the following sections, we will denote as univariate (*resp.* multivariate) the attacks in which the the traces l_i^k contain one (*resp.* several) samples.

3.2 Adaptive Selection of the Plaintexts

Let us now assume that a template attack has been performed with i traces, corresponding to different plaintexts x_1 to x_i , and giving rise to a certain knowledge about the subkey candidates summarized as $\hat{\Pr}^{(i)}[s^*]$. The objective of this paper, illustrated in Figure 1, is to select the next plaintext x_{i+1} in such a way that it will best discriminate the correct subkey. Ideally, this plaintext could be obtained by computing the success rate of the adversary in step $i + 1$ or, similarly, by computing the residual entropy of this correct subkey s (*i.e.* one of the metrics in [22]). But while running an attack, the adversary obviously does not know the value of this correct subkey yet. As a consequence, the only applicable strategy is to exploit a criteria that can be estimated “on-the-fly”.

Following the previous section, it appears that a natural criteria is to look at the entropy of the subkey candidates rather than the one of the correct subkey. Indeed, in a successful attack, the entropy of these subkey candidates should eventually be null (*i.e.* we should determine only the correct subkey with probability one). For example, at step i , this entropy can be estimated as:

$$\hat{H}^{(x_i)}[S^*] = - \sum_{s^*} \hat{\Pr}^{(i)}[s^*] \cdot \log_2 \hat{\Pr}^{(i)}[s^*],$$

where $\mathbf{x}_i = [x_1, x_2, \dots, x_i]$ is the vector of plaintexts used in the attack. Using this entropy as a criteria for our adaptive chosen-message attacks implies selecting the plaintext x_{i+1} as the one minimizing $\hat{H}^{(\mathbf{x}_{i+1})}[S^*]$. This can be done as follows. First, let us observe that for every plaintext candidate x_{i+1}^* and subkey candidate s^* , one can define a random variable $\hat{L}_{x_{i+1}^*}^{s^*}$, corresponding to the simulated leakage trace that perfectly follows the leakage model obtained from the templates construction phase (*i.e.* a normal curve with mean vector $\hat{\boldsymbol{\mu}}_{x_{i+1}^*}^{s^*}$ and covariance matrix $\hat{\boldsymbol{\Sigma}}_{x_{i+1}^*}^{s^*}$). We can then construct a random variable $\hat{L}_{x_{i+1}^*}^{S^*}$, as a mixture of $\hat{L}_{x_{i+1}^*}^{s^*}$'s, for different plaintext candidates x_{i+1}^* , with probability:

$$\Pr[\hat{L}_{x_{i+1}^*}^{S^*}] = \sum_{s^*} \hat{\Pr}^{(i)}[s^*] \cdot \Pr[\hat{L}_{x_{i+1}^*}^{s^*}]. \quad (1)$$

That is, we have one $\hat{L}_{x_{i+1}^*}^{S^*}$ per plaintext candidate x_{i+1}^* . Exemplary mixtures are represented in Figure 2, for two different plaintexts and in a simple context with only four possible subkeys. The definition of this variable is motivated by the fact that at step i in an attack, the only available knowledge about the subkeys is stored in $\hat{\Pr}^{(i)}[s^*]$. Hence, Equation (1) is the best available estimation of the leakage pdf at this step. For a given mixture and a fixed (simulated) leakage value $\hat{l}_{x_{i+1}^*}^{S^*}$, it is possible to compute the conditional entropy $\hat{H}^{(\mathbf{x}_{i+1}^*)}[S^* | \hat{l}_{x_{i+1}^*}^{S^*}]$, as illustrated in Figure 2 for three exemplary leakage values l_0, l_1 and l_2 .




Fig. 2. Adaptive selection of the plaintexts in a simplified context with $|S| = 4$

Integrating this entropy over the leakages yields the estimations:

$$\hat{H}^{(\mathbf{x}_{i+1}^*)}[S^*] = \int \Pr[\hat{l}_{x_{i+1}^*}^{S^*}] \cdot \hat{H}^{(\mathbf{x}_{i+1}^*)}[S^* | \hat{l}_{x_{i+1}^*}^{S^*}] dl_{x_{i+1}^*}^{S^*}.$$

And since we have one such entropy value for every possible choice of x_{i+1}^* in $|\mathcal{X}|$, we finally obtain the following rule to select the plaintexts:

$$\tilde{x}_{i+1} = \underset{x_{i+1}^*}{\operatorname{argmin}} \hat{H}^{\mathbf{x}_{i+1}^*}[S^*]$$

Summarizing, we use the available a-priori subkey information at step i and the leakage model (*i.e.* the templates) to predict how the entropy of the subkey candidates would evolve at step $i + 1$, for different plaintext candidates x_{i+1}^* .

3.3 Generalization to Non-profiled Attacks

As detailed in the previous section, an important requirement when applying an adaptive strategy (*e.g.* in the case of template attacks) is the availability of a good leakage model. Therefore, an interesting question is to know if such strategies can still help in the context of non-profiled side-channel attacks, where the model is usually less precise. As an illustration, we discuss this problem for the frequently considered correlation power analysis using Pearson's coefficient.

Correlation attacks, as described in [4], use the following distinguisher:

$$\hat{\rho}(\mathbf{M}_q^{s^*}, \mathbf{L}_q) = \frac{\hat{\mathbf{E}}\left((l_i - \hat{\mathbf{E}}(\mathbf{L}_q)) \cdot (m_i^{s^*} - \hat{\mathbf{E}}(\mathbf{M}_q^{s^*}))\right)}{\hat{\sigma}(\mathbf{L}_q) \cdot \hat{\sigma}(\mathbf{M}_q^{s^*})},$$

where $\hat{\mathbf{E}}$ and $\hat{\sigma}$ denote the sample means and standard deviations of a random variable, respectively. In this context, the models $m_i^{s^*}$ are not the complete leakage pdf (as in template attacks) but only their mean values (*i.e.* the first-order moments of the pdf). In general, these mean values are not estimated with profiling, but rather taken from engineering intuition. For example, a usual assumption is to use the so-called Hamming Weight or distance leakage models [14].

Adaptive correlation. When trying to apply the strategy of the previous section to correlation attacks, two main problems arise, that we now detail. First, the subkey probability estimation is not straightforward. Whereas template attacks rate these subkey candidates using their probabilities, correlation attacks return a set of scores, corresponding to the value of Pearson's coefficient. In order to mount an adaptive attack, the adversary consequently needs to use heuristics in order to estimate the subkey distribution $\hat{\Pr}^{(i)}[s^*]$, *e.g.* by:

- using the absolute value of the estimated coefficient $\hat{p}_i^{s^*} = |\hat{\rho}(\mathbf{M}_q^{s^*}, \mathbf{L}_q)|$,
- applying Fisher's transform on this correlation coefficient (in order to get a normal distribution), *i.e.* computing $\hat{p}_i^{s^*} = |\operatorname{arctanh}(\hat{\rho}(\mathbf{M}_q^{s^*}, \mathbf{L}_q))|$,

- computing the p-values associated with each correlation in a hypothesis test. For example, one could estimate the p-value obtained when stating that the subkey candidate is not correlated with the model.

In each case, we then need to normalize the $\hat{p}_i^{s^*}$'s in order to get an estimated probability distribution, as the values we obtain are not actual probabilities, and some wrong subkeys may give a non-zero score (*aka* ghost peaks [4]):

$$\hat{P}_1^{(i)}[s = s^*] = \frac{\hat{p}_i^{s^*}}{\sum_{s' \in \mathcal{S}} \hat{p}_i^{s'}}$$

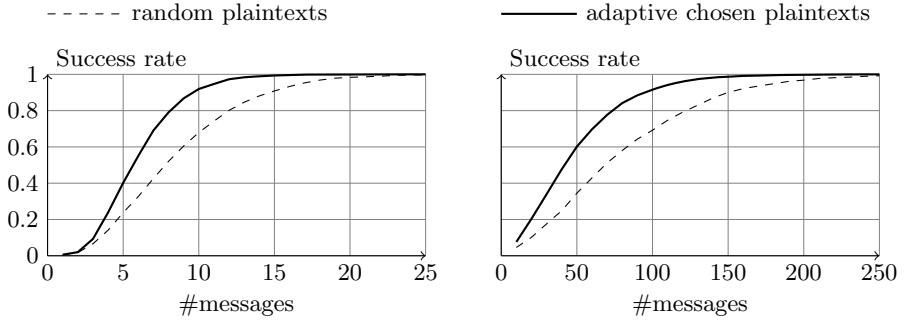
Second, and more critically, the selection procedure of Section 3.2 requires to build a random variable $\hat{L}_{x_{i+1}^*}^{S^*}$ as a mixture of $\hat{L}_{x_{i+1}^*}^{s^*}$, that estimates the leakage distribution given the subkey probabilities at step i in an attack. This requires an estimate of the leakage pdf that is given if the leakage model is probabilistic (as in template attacks), but is not directly available in a correlation attack. Again, a number of heuristics are possible. The simplest one, that we considered in this work, is to combine the (Hamming weight or distance) power models with a Gaussian assumption, *i.e.* to paste a Gaussian curve to the different Hamming weights, of which the variance is estimated “on-the-fly” during the attack.

4 Simulated Experiments


In order to validate our adaptive message selection, we first conducted software simulations. These attacks target the output of a single AES S-box in the first encryption round. Excepted if mentioned otherwise, physical leakages are simulated as the Hamming weight of the S-box outputs, to which is added a normally distributed noise with standard deviation σ_n . The efficiency of an attack is then measured with the success rate, averaged over 1000 independent key recoveries. The results of our experiments are in Figure 3 from which we observe:

1. In all cases, the adaptive strategy leads to increased success rates. It noticeable that the impact of this adaptivity becomes significant as soon as a slight a-priori knowledge is known about the target subkey. Also, and as illustrated in Figure 3.(a), this improvement holds for different noise levels.
2. The same observation also holds for different leakage functions. For example, Figure 3.(b) shows the success rates of attacks exploiting three different side-channels of the form: $L(x) = \sum_i \alpha_i x[i] + n$, where $x[i]$ is the i th bit of the target S-box output and n a Gaussian noise. Interestingly, these examples directly connect with the framework in [22]. They show that as in a non-adaptive context, a more informative leakage function (measured with the conditional entropy $H[S|\mathbf{L}_1]$) leads to more efficient attacks.
3. Although more computationally intensive (because they require to deal with mixtures of probability distributions, *e.g.* as described in [13]), attacks against masked implementations exhibit similar improvements (see Figure 3.(c)).


(a) Hamming weight leakage function, $\sigma_n = 1$ (left) and $\sigma_n = 4$ (right).



(b) Different leakage functions with the same noise level $\sigma_n = 1$ and the conditional entropy $H[S|L_1] = 7.7$ (left), $H[S|L_1] = 7.4$ (middle), $H[S|L_1] = 6.9$ (right).



(c) Masked S-box, Hamming weight leakages, $\sigma_n = 0.5$ (left) and $\sigma_n = 1$ (right).



(d) Correlation attacks, Hamming weight leakages, $\sigma_n = 0.5$ (left) and $\sigma_n = 1$ (right).




Fig. 3. Success rates of different simulated experiments

4. Eventually, the results of the heuristics proposed to exploit adaptivity in the context of correlation power analysis are given in Figure 3.(d). As expected, the imperfect approximations of the pdf imply smaller improvements.

This last point implies interesting scopes for further research. For example, it would be interesting to apply adaptive strategies to other non-profiled tools such as the MIA [7], in which an estimation of the leakage pdf is computed as part of the attack. In the same line, it could also be possible to exploit stochastic models in order to obtain a leakage model “on-the-fly”. In this respect, it is worth recalling that such distinguishers can also be used for profiling a device, without a-priori knowledge of the key (*i.e.* to obtain templates in a flexible way).

5 Experiments Using Actual Measurements

In order to confirm the previous simulations, we additionally performed actual experiments against an implementation of the AES Rijndael in an Atmel Atmega 644p chip. Such actual measurements are interesting because they allow exploiting the leakage of several time samples, contrary to the simulated case where a single point of interest was considered. In other words, actual experiments allow easily evaluating the impact of multivariate templates. In practice, we compared attacks with up to three samples, for adaptive and random message selection. The points of interest were selected as part of the profiling phase, two of them corresponding to the S-box computation, and one to the first key addition. Again, we estimated the success rates over 1000 independent key recoveries, excepted for the trivariate attack which was only launched against 50 different keys. The smaller number of attacks in this case is due to their computational cost, that grows exponentially with the number of dimensions, and makes the exhaustive analysis of Section 3.2 too intensive to be performed.

The results of these experiments are in Figure 4. They show that the adaptive strategy holds for real world implementations. That is, the leakage models built during profiling can be precise enough¹ so that the estimation of the “next-step entropy” $\hat{H}^{x_{i+1}^*}[S^*]$ leads to a meaningful selection of the next plaintext x_{i+1} . It is worth noting the large difference between univariate random-message attacks and trivariate chosen-message ones. It illustrates the variability that can be observed between different attack scenarios in physically observable cryptography.

6 Discussion and Concluding Remarks

Is our strategy optimal? Following the previous sections, a first natural question is to know if the proposed strategy is optimal. For this purpose, it is interesting to relate our work with the one of [12]. The authors estimate the number of queries required for a key-recovery, in the context of deterministic side-channel leakages. For each encryption step, the key candidates are partitioned in

¹ We used a 1000 traces to characterize each template.

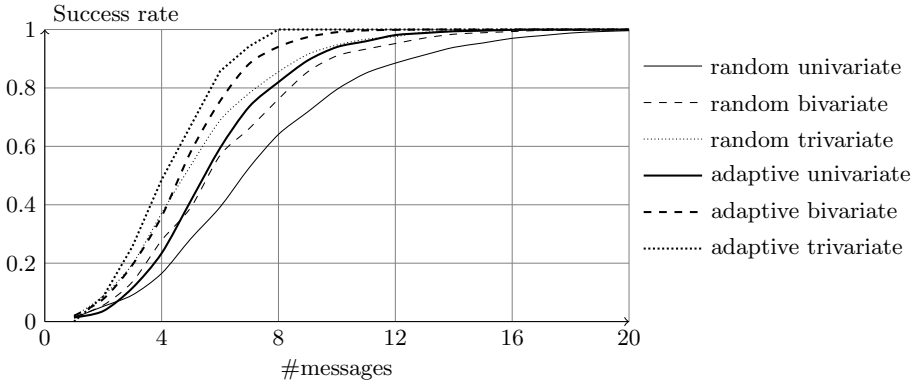


Fig. 4. Success rates of experiments carried out against an AES implementation

r sets, and the side-channel leakage allows the adversary to discriminate one set containing the correct key. The optimal strategy minimizes the number of steps required to reduce the number of key candidates to one. The main limitation is that its computational cost is doubly exponential in the number of attack steps. This is because in this optimal strategy, it is in fact several next plaintexts (x_{i+1} , x_{i+2} , ...) that have to be predicted in order to minimize the entropy of the key candidates. Hence, this strategy is hardly applicable, even for small parameters size. In order to get rid of this limitation, Köpf and Basin propose an alternative greedy heuristic, which predicts only one next plaintext at a time.

The procedure presented in this paper can be seen as the extension of such a greedy strategy, from the deterministic case towards the more general probabilistic case. The main difference is that deterministic leakages allow the adversary to effectively eliminate subkeys, whereas probabilistic leakages only help the adversary to update the subkey candidates' distribution. This extension allows an application of adaptive strategies to a broader class of attacks, including power and electromagnetic leakages, typically. But it comes at a computational cost, since we had to turn deterministic sums into integrals (that are multidimensional in the case of multivariate attacks). Summarizing, our strategy is not optimal. But as indicated in [12], greedy heuristics can provide close to (or even equal to) optimal results in practice. The exact evaluation of the greedy approach with respect to the optimal one and the investigation of alternative solutions to reduce the computational cost of adaptive attacks is a scope for further research.

Implications. Next to optimality, another important question is to determine whether the application of adaptive strategies may have practical impact in certain applications. Looking at the figures in the previous sections indicates that the improvements are not huge, but can be significant. For example, Table 1 shows that the number of measurements required to reach a certain success rate is improved, in particular when combining adaptive attacks with trivariate leakages. But in fact, the consequences of adaptivity are best observed with respect to the global success rates of the attacks. That is, because standard DPA

Table 1. Approximated data complexities for different attacks against an 8-bit subkey

Target success rate	> 20%	> 40%	> 60%	> 80%	$\approx 100\%$
random messages - 1D	5	7	8	11	20
adaptive messages - 1D	4	5	7	8	16
adaptive strategy - 3D	3	4	5	6	8

Table 2. Approximated success rates for different attacks against a 128-bit key

Number of messages	2	3	4	5	6	7	8
random messages - 1D	2^{-69}	2^{-55}	2^{-42}	2^{-29}	2^{-21}	2^{-15}	2^{-10}
adaptive messages - 1D	2^{-64}	2^{-50}	2^{-33}	2^{-20}	2^{-12}	2^{-7}	2^{-5}
adaptive strategy - 3D	2^{-58}	2^{-32}	2^{-16}	2^{-9}	2^{-4}	2^{-2}	-

attacks exploit a divide-and-conquer strategy, the overall success rate against the full AES master key can be estimated by simply raising the success rate against an 8-bit byte to the power 16. This assumes that all key bytes are equally difficult to recover, which is reasonable in most applications, in particular software ones as in Section 5. In the case of adaptive attacks, it also means that the selection of all the plaintext bytes are performed concurrently. Table 2 shows these estimated success rates in function of the number of messages in the attack. It clearly illustrates the strong impact that adaptive strategies may have. For example, one can imagine a re-keying scheme where the secret is updated every four encryptions. Our results suggest that the resulting security level would differ by a factor of 2^9 depending on the use or not of adaptive messages. This factor increases to 2^{26} if multivariate leakages are considered. And in the case of attacks against the AES-256, these factors would additionally be squared.

It is worth mentioning that targeting hardware implementations, in which all the subkeys are manipulated in parallel, would imply additional questions. For example, in a context where a single key byte has to be recovered with high efficiency, one could also take advantage of chosen plaintexts so that the remaining input bits are constant, in order to reduce the algorithmic noise. But an adaptive strategy would still apply to the target key byte. Extending the experiments of this paper towards more devices and countermeasures against side-channel attacks is anyway another interesting direction for further research.

Eventually, and as discussed in [23], the success rates of adaptive attacks can, when applicable, be used as rough (but only available ones) estimations of the bounded leakage² that is necessary to prove the security of certain leakage resilient constructions. Our results can also be directly integrated in the evaluation framework of Eurocrypt 2009 [22]: they exhibit a new type of distinguisher that can take advantage of the information leakage in a close to optimal manner. Summarizing, this paper brings an important contribution to the exploitation of side-channel leakages in both theoretical and practical settings.

² Given that the plaintext selection is granted to adversaries. As previously said, it is anyway a more reasonable abstraction than the adaptivity of the leakage function.

References

1. Agrawal, D., Archambeault, B., Rao, J., Rohatgi, P.: The EM Side-Channel(s). In: Kaliski Jr., B.S., Koç, Ç.K., Paar, C. (eds.) CHES 2002. LNCS, vol. 2523, pp. 29–45. Springer, Heidelberg (2003)
2. Baignères, T.: Quantitative Security of Block Ciphers: Design and Cryptanalysis Tools. PhD Thesis, EPFL, Lausanne, Switzerland (November 2008)
3. Bleichenbacher, D.: Chosen Ciphertext Attacks Against Protocols Based on the RSA Encryption Standard PKCS #1. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 1–12. Springer, Heidelberg (1998)
4. 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)
5. Chari, S., Rao, J., Rohatgi, P.: Template Attacks. In: Kaliski Jr., B.S., Koç, Ç.K., Paar, C. (eds.) CHES 2002. LNCS, vol. 2523, pp. 13–28. Springer, Heidelberg (2003)
6. Dziembowski, S., Pietrzak, K.: Leakage-Resilient Cryptography. In: The proceedings of FOCS 2008, Washington, DC, USA, October 2008, pp. 293–302 (2008)
7. Gierlichs, B., Batina, L., Tuyls, P., Preneel, B.: Mutual Information Analysis: A Generic Side-Channel Distinguisher. In: Oswald, E., Rohatgi, P. (eds.) CHES 2008. LNCS, vol. 5154, pp. 396–410. Springer, Heidelberg (2008)
8. Goubin, L., Patarin, J.: DES and Differential Power Analysis. In: Koç, Ç.K., Paar, C. (eds.) CHES 1999. LNCS, vol. 1717, pp. 158–172. Springer, Heidelberg (1999)
9. Joux, A., Martinet, G., Valette, F.: Blockwise-Adaptive Attackers: Revisiting the (In)Security of Some Provably Secure Encryption Models: CBC, GEM, IACBC. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 17–30. Springer, Heidelberg (2002)
10. Kocher, P.C.: Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In: Kobitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 104–113. Springer, Heidelberg (1996)
11. Kocher, P., Jaffe, J., Jun, B.: Differential Power Analysis. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 388–412. Springer, Heidelberg (1999)
12. Köpf, B., Basin, D.A.: An Information-Theoretic Model for Adaptive Side-Channel Attacks. In: The proceedings of the ACM Conference on Computer and Communications Security, Alexandria, Virginia, USA, October 2007, pp. 286–296 (2007)
13. Lemke-Rust, K., Paar, C.: Gaussian Mixture Models for Higher-Order Side Channel Analysis. In: Paillier, P., Verbauwhede, I. (eds.) CHES 2007. LNCS, vol. 4727, pp. 14–27. Springer, Heidelberg (2007)
14. Mangard, S., Oswald, E., Popp, T.: Power Analysis Attacks. Springer, Heidelberg (2007)
15. Mangard, S., Oswald, E., Standaert, F.-X.: One for All, All for One: Unifying Standard DPA Attacks, Cryptology ePrint Archive, Report 2009/449
16. Pietrzak, K.: A Leakage-Resilient Mode of Operation. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 462–482. Springer, Heidelberg (2010)
17. Prouff, E.: DPA Attacks and S-Boxes. In: Gilbert, H., Handschuh, H. (eds.) FSE 2005. LNCS, vol. 3557, pp. 424–441. Springer, Heidelberg (2005)
18. Rivain, M., Dottax, E., Prouff, E.: Block Ciphers Implementations Provably Secure Against Second Order Side Channel Analysis. In: Nyberg, K. (ed.) FSE 2008. LNCS, vol. 5086, pp. 127–143. Springer, Heidelberg (2008)

19. 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)
20. Schindler, W.: A Timing Attack against RSA with the Chinese Remainder Theorem. In: Paar, C., Koç, Ç.K. (eds.) CHES 2000. LNCS, vol. 1965, pp. 109–124. Springer, Heidelberg (2000)
21. 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)
22. 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 (2010); extended version available on the Cryptology ePrint Archive, Report 2006/139
23. Standaert, F.-X., Pereira, O., Yu, Y., Quisquater, J.-J., Yung, M., Oswald, E.: Leakage Resilient Cryptography in Practice. Cryptology ePrint Archive, report 2009/341
24. Wagner, D.: The Boomerang Attack. In: Knudsen, L.R. (ed.) FSE 1999. LNCS, vol. 1636, pp. 156–170. Springer, Heidelberg (1999)