

Efficient Selection of Time Samples for Higher-Order DPA with Projection Pursuits

François Durvaux¹, François-Xavier Standaert¹, Nicolas Veyrat-Charvillon²
Jean-Baptiste Mairy³, Yves Deville³.

¹ ICTEAM/ELEN/Crypto Group, Université catholique de Louvain, Belgium.

² IRISA-CAIRN, Campus ENSSAT, 22305 Lannion, France.

³ ICTEAM/INGI, Université catholique de Louvain, Belgium.

Abstract. The selection of points-of-interest in leakage traces is a frequently neglected problem in the side-channel literature. However, it can become the bottleneck of practical adversaries/evaluators as the size of the measurement traces increases, especially in the challenging context of masked implementations, where only a combination of multiple shares reveals information in higher-order statistical moments. In this paper, we describe new (black box) tools for efficiently dealing with this problem. The proposed techniques exploit projection pursuits and specialized local search algorithms, work with minimum memory requirements and practical time complexity. We validate them with two case-studies of unprotected and first-order masked implementations in an 8-bit device, the latter one being hard to analyze with previously known methods.

1 Introduction

The selection of Points-Of-Interest (POIs) in leakage traces is an important (and not very discussed) problem in the application of Side-Channel Analysis (SCA) attacks. When targeting unprotected implementations, the naive strategy that is commonly used in the literature is to test all the time samples independently. It raises two important challenges. First, how to combine these time samples efficiently, in order to maximize the amount of information extracted from each leakage trace? Second, how to extend this technique in the context of masked implementations where the sensitive data is split into d shares manipulated in different clock cycles (as it is typically the case in software), and only the combination of these shares' leakage reveals key-dependent information – which makes the complexity of an exhaustive analysis grow combinatorially with d ?

Solutions to the first problem typically include dimensionality reduction techniques such as PCA and LDA. These tools (introduced to SCA in [1, 29] and recently revisited in [3, 5]) essentially project the leakage traces into a lower-dimensional subspace that optimizes some *objective function*. Namely, PCA usually maximizes the variance between the mean leakage traces – i.e. the signal of a first-order DPA, while LDA maximizes the ratio between inter-class and intra-class variances – i.e. its Signal-to-Noise Ratio (SNR), essentially. Their main advantage is to provide a principled and intuitive solution to the problem, since the projection (i.e. eigenvectors) they produce indicate the POIs. Yet, they are

somewhat limited when moving to masked implementations for which the information lies in high-order statistical moments, since their objective function is based on a definition of signal that primarily captures first-order leakages¹.

Solutions to the second problem are even sparser. To the best of our knowledge, the usual reference for selecting POIs for masked implementations is the educated guess proposed by Oswald et al. in [18] (i.e. an exhaustive search over all d -tuples of time samples in a window selected based on engineering intuition). Next, Reparaz et al. proposed an alternative solution exploiting Mutual Information Analysis (MIA) [9], that allows gaining a constant (but practically meaningful) factor corresponding the number of key hypotheses in the attack [26]. In both cases, the proposed tools do not output a projection but a list of the most useful POIs (i.e. d -tuples) in function of the (non-profiled) attack considered.

In this paper, we investigate the use of *Projection Pursuits* (PPs), as alternative tools for the selection of POIs in leakage traces [8]. Intuitively, PPs machine-pick “interesting” low-dimensional projections of a high-dimensional data space by numerically maximizing a certain objective function. They essentially work by tracking the improvements (or lack thereof) of the projection when modifying it with small random perturbations. Their main advantage in our context is that they can deal with any objective function, which naturally fits to the problem of higher-order SCA. Their main drawback is (in general) their heuristic nature, since the convergence of the method is not guaranteed and its complexity is context-dependent. As a result, and in order to validate the interest of PPs in our SCA context, we first applied them to the simple case of an unprotected implementation of the AES. We show that different objective functions can be efficiently used for this purpose, leading to powerful subspace-based attacks, with similar informativeness as previous solutions such as LDA.

Next, we moved to the more challenging context of masking. In this case, we combined the (linear) projection with an objective function exploiting higher-order statistical moments. Initial experiments suggest that the straightforward implementation of a PP algorithm is not efficient in detecting the POIs of such protected implementations (especially as the number of useless dimensions in the traces increases). The main reason is that as long as a d -tuple of POIs is not present in the projection, the objective function essentially returns random indications. Interestingly, we then show that a specialized PP algorithm exploiting an improved *local search* could give excellent results even in this challenging context. Intuitively, it works by looking for the best size and position of d windows covering parts of the traces, again by iterating small random perturbations. Our experiments suggest that we can recover POIs with significantly less calls to the objective function than a exhaustive analysis. We further discuss the main parameters influencing the success of such a detection method, and detail the time vs. measurement complexity tradeoff resulting from these parameters.

¹ Of course, a trivial solution would be to apply PCA/LDA to “product traces” containing all the possible products of d -tuples, but this rapidly leads to unrealistic memory requirements in the masked software context that we consider next.

Cautionary note & related works. In general, a projection search algorithm can be evaluated according to two orthogonal axes, namely its time and data complexity (i.e. how many iterations and measurements do we need to obtain a projection?) or the informativeness of its outputs (which relates to the data complexity of an attack exploiting the projections obtained). Hence, it is first worth recalling that (standard) dimensionality reductions for unprotected implementations indeed optimize informativeness, whereas existing solutions to detect POIs in masked implementations focus on the complexity issue (because of the more challenging nature of the problem). In this context, we note that comparing different projection’s informativeness (e.g. PCA, LDA and the recent works in [11, 17]), in the context of unprotected implementations, is of limited interest anyway. Indeed, it has been shown in [14] that the objective functions of LDA (which improves over PCA in terms of informativeness) and [11, 17] are essentially equivalent in this case, meaning that LDA, these works and our new projections all have similar informativeness as well (up to statistical artifacts). Nevertheless, and for completeness, we show empirical evidence of the gain they provide over PCA and its impact in the DPA contest v2 [20]. As for comparisons in the case of higher-order leakages and masked implementations, the main issue is that none of the previous dimensionality reductions generalizes to such contexts². In fact, and as witnessed by the previous state-of-the-art, existing solutions for detecting POIs in masked implementations are mainly based on exhaustive approaches, which turn out to be too intensive as the size of the leakage traces increases. Hence, our main contribution is to provide a first efficient alternative for dealing with large traces containing only higher-order information leakages. For this purpose, and since our focus is on complexity issues, we will put ourselves in the most challenging scenario, i.e. a black box analysis where no information about the source code is available, and compare the gains of our optimizations over combinatorial search and other naive approaches such as signal integration [6]. Quite naturally, any engineering intuition allowing an educated guess – by focusing only on certain parts of the traces – could be exploited as well. Besides we will also discuss the differences between our work and the one of Reparaz et al. and exhibit that they are essentially complementary.

2 Background

Notations. We use capital letters for random variables, small caps for their realizations, sans serif fonts for functions and calligraphic letters for sets.

² More precisely, the results in [11, 17] are actually similar to ours in the first-order setting. In fact, they can be viewed as a heuristic (computationally efficient) analogue to LDA. But their application to the higher-order case would be difficult for the same reasons as mentioned at the beginning of Section 4 for our (non-specialized) projection search. In this respect, an important difference between this previous work and ours is the separation between the objective functions and optimization algorithms: we need to change both in order to deal with higher-order leakages efficiently.

2.1 Measurement setups

Our experiments are based on measurements of an AES implementation run by an 8-bit Atmel AVR (ATMega644P) microcontroller at a 20 MHz clock frequency. We monitored the voltage variations across a 22 Ω resistor introduced in the supply circuit of our target chip. Acquisitions were performed using a Lecroy HRO66ZI oscilloscope running at 200 MHz and providing 8-bit samples. For concreteness, our evaluations focused on the leakage of the first AES master key byte (but would apply identically to any other enumerable target). Leakage traces were produced according to the following procedure. Let x and s be our target input plaintext byte and subkey, and $y = x \oplus s$ denote a key addition. For each of the 256 values of y , we generated 1000 unprotected encryption traces (resp. 500 for masked traces), where the rest of the plaintext and key was random, i.e. we generated 256 000 (resp. 128 000) traces in total, with plaintexts of the shape $p = x || r_1 || \dots || r_{15}$, keys of the shape $k = s || r_{16} || \dots || r_{30}$, and the r_i 's denoting uniformly random bytes. In case of masked implementations, additional uniform randomness was used to generate the shares. In order to reduce the memory cost of our evaluations, we only stored the leakages corresponding to the 2 first AES rounds in the unprotected case (as the dependencies in our target byte $y = x \oplus s$ typically vanish after the first round, because of the strong diffusion properties of the AES). As for the protected case, we only considered a single S-box, for which the precomputation of a masked table already implies large traces with $Ns = 30,000$ time samples (vs. $Ns = 1500$ for the unprotected one). As will be clear next, these sets of measurements were large enough to emphasize the interest of our projection pursuit algorithms. In the following, we will denote the 1000 (resp. 500) encryption traces obtained from a plaintext p including the target byte x under a key k including the subkey s as: $\text{AES}_{k,s}(p_x) \rightsquigarrow l_y^i$, with $i \in [1; 1000]$ (resp. $i \in [1; 500]$). Whenever accessing the points of these traces, we will additionally use an argument t (for time), leading to $l_y^i(t)$. Our goal is to generate projections exhibiting the time samples that contain information about y . Note that since we assume the plaintext to be known by the adversary (as usual in SCAs), it directly translates into information about s – which typically occurs during the key addition $y = x \oplus s$ and S-box execution $z = S(x \oplus s)$.

2.2 Objective functions (aka evaluation metrics)

In order to “guide” the PP, we need to define criteria to determine whether some modification of the projection is positive. Any SCA evaluation metric can be used for this purpose. We list a few candidates in this section. In order to guarantee their soundness, we focused on objective functions based on profiled distinguishers (which allows mitigating biases due to incorrect a-priori choices of models – given that the profiles are well estimated and based on sound assumptions).

CPA [4]. In a profiled Correlation Power Analysis, the adversary first estimates the first-order moments corresponding to each value y from a vector of N_p profiling traces \mathbf{l}_y^p , that we denote as $\hat{\mathbf{m}}_y^1 = \hat{E}(\mathbf{l}_y^p)$, with \hat{E} the sample mean operator.

This step is performed for each time sample independently, leading to $\hat{\mathbf{m}}_y^1(t)$. Since there are 256 y values in our AES case study, it amounts to compute $256 \times N_s$ means, with N_s the number of samples per trace. Then, he computes the correlation between these mean values and the samples coming from a vector of test traces \mathbf{l}_y^t , leading to $\hat{\rho}(\hat{\mathbf{m}}_y^1(t), \mathbf{l}_y^t(t))$ with $\hat{\rho}$ denoting Pearson’s coefficient.

SNR [13]. An alternative to CPA is the SNR defined at CT-RSA 2004 as:

$$\text{SNR}(t) = \frac{\hat{\text{v}}\text{ar}_y\left(\hat{\mathbf{E}}(\mathbf{l}_y^t(t))\right)}{\hat{\mathbf{E}}_y\left(\hat{\text{v}}\text{ar}(\mathbf{l}_y^t(t))\right)},$$

with $\hat{\text{v}}\text{ar}$ the sample variance operator. Similarly to the correlation coefficient, such a criteria is discriminant for first-order information (i.e. information lying in the first-order moments of the leakage distribution). In order to deal with masked implementations, we also need objective functions that capture more general dependencies. In this context, a natural option is the information theoretic metric introduced in [31] and later refined in [25]. Its sample definition is given by:

$$\hat{\mathbf{I}}(S; X, L) = \mathbf{H}[S] - \sum_{s \in \mathcal{S}} \Pr[s] \sum_{x \in \mathcal{X}} \Pr[x] \sum_{l_y^i \in \mathcal{L}_y^i} \Pr_{\text{chip}}[l_y^i | s, x] \cdot \log_2 \hat{\Pr}_{\text{model}}[s | x, l_y^i],$$

where $\hat{\Pr}_{\text{model}}$ is a probabilistic model estimated thanks to the set of profiling traces (just as the $256 \times N_s$ mean values in the correlation case). Computing such an objective function implies (constant but significant) performance overheads, because it requires applying Bayes’ law and marginalizing over the key hypotheses. Since the objective function will typically be applied after projection in the following sections (i.e. in a univariate context), a cheaper alternative is to exploit the following “Moments-Correlating Profiled DPA” (MCP-DPA):

MCP-DPA [16]. The attack features essentially the same steps as a profiled CPA. The only difference is that the adversary will estimate d th-order moments $\hat{\mathbf{m}}_y^d(t)$ with the profiling traces. In the following, we will be particularly interested in the Moments against Moments Profiled Correlation (MMPC) criteria:

$$\text{MMPC}(t) = \hat{\rho}(\hat{\mathbf{m}}_y^d(t), \tilde{\mathbf{m}}_y^d(t)),$$

where $\tilde{\mathbf{m}}_y^d(t)$ are another vector of moments, estimated with the test traces. As detailed in [16], MCP-DPA is able to capture information in any statistical moment, while enjoying the implementation efficiency of CPA (which is highly beneficial in our context where the objective function is intensively used).

3 Projection pursuit against unprotected devices

In this section we investigate the application of PPs to the simple case of the (unprotected) AES furious implementation available as open source from [21].

In this context, our goal is to find a projection vector α that will convert the N_s samples of a leakage vector \mathbf{l}_y^t to a single (projected) sample λ_y^i , that is:

$$\lambda_y^i = \sum_{t=0}^{N_s-1} \alpha(t) \cdot l_y^i(t),$$

such that univariate attacks exploiting the λ_y^i 's will be most efficient. This essentially requires to define an objective function that measures the ‘‘informative-ness’’ of these samples. As mentioned in the previous section, this task is quite easy when first-order information is available in the leakage traces: Pearson’s correlation coefficient obtained from a CPA and Mangard’s SNR are natural candidates – we will try them both in the next subsection. Following the equivalence results in [14], they should provide similar results in this case (also similar to the ones that would be obtained with an information theoretic metric).

3.1 Projection pursuit algorithm

The pseudo-code of our projection pursuit algorithm is given in Algorithm 1.

Algorithm 1 Projection Pursuit.

```

basic_PP( $N_r, N_{it}$ )
 $\alpha = initialize()$ ;
repeat  $N_r$  times
     $r = rand\_index(N_s)$ ;
     $\alpha_{new} = max\_search(@f_{obj}, \mathcal{L}^P, \alpha, r, N_{it})$ ;
     $\alpha = \alpha_{new}$ ;
end

```

It essentially repeats (N_r times) the selection of a random index r followed by a maximization of the objective function for the corresponding time sample, based on the set of profiling traces \mathcal{L}^P (which contains traces for all the intermediate values y). For this purpose, the *max_search()* function consists in successive parabolic interpolations (illustrated in Appendix A, Figure 4), which work in two iterated steps. We first look for samples that enclose the extremum as follows. From a starting point x_1 , we add a Δ in the direction that increases f_{obj} (blue plain curve) to get x_2 . Then, we keep adding Δ 's until finding x_3 such that $y_3 < y_2$ (see Figure 4.a). As the weights assigned to each time sample are between 0 and 1, we typically take Δ 's corresponding to a couple of percents (e.g. 0.1 in our experiments) and repeat such additions at most $1/\Delta$ times. Then, based on these three points, we start interpolating (as in the dashed red curve of Figure 4.b-c). This process is iterated N_{it} times, during which we replace the ‘‘oldest’’ x -point by the x -coordinate (x_v) of the parabola vertex (y -values are re-computed accordingly). The new $\alpha(t)$ gets its value from the median x -value at the end of the last iteration. In our experiments, $N_{it} = 3$ iterations were enough to get a

good approximation of the maximum. This method has the advantage of being very fast to compute and to converge. Note finally that the number of repetitions N_r should ideally be larger than the number of samples N_s (e.g. twice, typically), because some weights benefit from being re-adjusted after the modification of other $\alpha(t)$'s. Yet, when applied in the context of an unprotected implementation, the time complexity of Algorithm 1 was never a practical limitation (it typically corresponded to a couple of minutes of computations in our experiments).

3.2 Experimental results

We implemented the PP algorithm for both the CPA and SNR objective functions, and targeted the first AES key byte for illustration. For each of the 256 values of $y = x \oplus s$, we measured $N_p = N_t = 50$ traces for the CPA objective function, and $N_t = 100$ traces for the SNR one, each of them made of $N_s = 1500$ time samples. We set N_r , N_{it} and Δ as just explained (to 3000, 3 and 0.1, respectively). The projections obtained in both cases are given in Appendix A, Figure 5, for illustration. As expected, they are very similar. We then computed success rates to compare the quality of the projections obtained with the most informative sample, by performing 2000 experimental univariate Template Attacks (TA). These results show the effectiveness of the projections as they need only 7 traces to get a 90% success rate, against 28 traces for the univariate TA. It also confirms that both objective functions are indeed equivalent in this case. It is finally interesting to compare our findings with the results in [30] that target a similar implementation (with very similar success rate for the univariate TA). In particular, we see that the univariate attack based on the single sample provided by our projections leads to approximately the same data complexities as the hexivariate template attack taking (heuristic) advantage of all the POIs in this previous work. This informally confirms the quality of our projection.




Fig. 1. Template attack success rates against unprotected device

3.3 Discussion & comparison with existing works

As mentioned in introduction, the open literature essentially provides satisfying solutions for detecting POIs and reducing the dimensionality of leakage traces obtained from unprotected devices, and this holds both from the complexity and informativeness points-of-view. Namely, from the complexity point-of-view, testing all time samples exhaustively is usually achievable in this case. And from the informativeness point-of-view, any projection using one of the criteria in [14], e.g. like LDA does with the SNR, should be optimal. In practice though, computing LDA for large traces may lead to numerical issues, as it requires to estimate a covariance matrix of the size of the leakage traces (whereas PCA enables a “small sample size variant” where the size of this covariance matrix can be limited to the number of key hypotheses [1]). In this respect, both our projection pursuit in Algorithm 1 and the proposal in [17] provide useful alternatives. Note that the approach in this previous work slightly differs from ours since its experiments are specialized to CPA with a Hamming distance leakage model (hence deviate from our optimality goal). By contrast, we rather suggest to use a profiled correlation coefficient (or the SNR) as objective function(s), since the evaluation of the objective function anyway requires some key knowledge. These minor differences highlight that projection pursuits are versatile and can be applied with a variety of objective functions. Eventually, and despite it is not the main concern in this work, we want to emphasize that deviating from an optimal objective function can lead to substantial differences in the attack success rates. For example, both PCA and our projection pursuit algorithm using the SNR as objective function have been tested in the context of the DPA contest v2 [20]. As illustrated in Appendix A, Figure 6, the data complexity of an attack using PCA is approximately doubled in this case, which can be explained by the correlation between the noise distributions taken at different time samples. Note that this projection pursuit was included amongst the best attacks of this contest, leading to a data complexity of 1173 traces (439 after key enumeration) as detailed in [7].

4 Projection pursuit against masked implementations

In contrast with the previous section, detecting POIs in leakage traces of masked implementations is a quite challenging task. From the complexity point-of-view, exhaustive approaches may grow exponentially with the number of shares (if these shares are manipulated at different time samples), making them unpractical for long traces. Furthermore, the information in the leakages of masked implementation lies in higher-order moments of their probability distribution, which are harder to estimate. As a result, the direct application of Algorithm 1 with the previous objective functions in this context does not provide successful results. In the (simple) case where the shares of a masking scheme are manipulated in parallel, adapting the objective function may be sufficient to deal with this problem. But in case of software implementations, where the shares are manipulated at different time samples, it is the algorithm itself that has to be adapted. Intuitively, this is because it works by modifying time samples one at a

time, while for such masked implementation, we require at least one meaningful d -tuple of samples to be active in the projection for an objective function to output relevant information. We now describe how to specialize PPs to take this constraint into account, and detect POIs for masked implementations.

4.1 Specialized projection pursuit algorithm

The main tool used in our following optimization is local search, which is a collection of iterative methods that are efficient for quickly finding good solutions to optimization problems (note that the previous PP algorithm can be viewed as a simple local search). Despite heuristic, it generally works more efficiently than exhaustive analyses. Furthermore, local search has very limited storage requirements. For example, in our context, it exploits the leakage traces directly – which is a significant advantage compared to heuristics exploiting “product traces” as mentioned in footnote 1. A good reference to these methods is [12]. Their working principle is simple: they always keep a solution (called the current solution) as well as the best solution found since the beginning of the search. At each iteration of the algorithm, the current solution is perturbed, giving a set of new solutions, called its neighborhood. One of the neighboring solutions is then selected and replaces the current solution. The algorithm terminates when its convergence criterion is met (e.g. number of iterations without improvement, time limit, etc.). Intuitively, such an approach to optimization exploits *diversification* and *intensification*. The former aims at exploring a large and diverse search space, while the latter intends to improve the current solution. Their combination is expected to avoid being trapped into local optima

When applied to masking, one key element has to be taken into account by optimizations. Namely, the sensitive variables are split into d shares and the objective function should not be informative as long as a meaningful d -tuple of shares is not present in the projection. Besides, in practice it frequently happens that dimensions near a POI also contain valuable information. These two facts motivate the way we designed our improved search algorithm as follows. First, we consider a projection vector containing d *windows* of non-zero weights (all the others being zero) and denote a group of successive dimensions as a window. The weights inside these windows are uniform. In this context, and since local search only considers local modifications of the current solution, the information given by the objective function will return essentially random indications (so no reliable information) if this current solution does not cover the d shares. On the contrary, when the windows spans a d -tuple of shares, the objective function can be used to refine the current solution. For this reason, our specialized PP algorithm will be split into two parts next denoted as `find_sol` and `improve_sol`. The `find_sol` phase probes the search space with large windows and a lot of randomness until it has good indication that the windows span the d -tuples of shares. In order to detect that the windows span these d -tuples, we use two sets of profiling traces ($\mathcal{L}_{\text{tr}}^p$ and $\mathcal{L}_{\text{va}}^p$, where `tr` stands for training and `va` for validation). Then, the `improve_sol` phase refines those windows. The `find_sol` phase thus puts more emphasis on diversification and the `improve_sol`, on intensification.

Algorithm 2 Specialized projection pursuit algorithm using local search.

```
specialized_PP_Local_Search( $d, W_{len}, T_{det}, TP := TP' \cup TP''$ )
   $\alpha = \text{find\_sol\_phase}(d, W_{len}, T_{det}, TP')$ ;
  if ( $\alpha \neq \text{null}$ )
    return  $\text{improve\_sol\_phase}(\alpha, TP'')$ ;
end
```

The pseudocodes of the specialized PP algorithm using local search are given in Algorithms 2, 3 and 4. These algorithms depend on various parameters: some of them will be explicitly discussed as they hold important intuitions, the remaining ones – next denoted as technical parameters (TP) – will be fixed according to state-of-the-art strategies. Our main tool is the `specialized_PP_Local_Search` function (Algorithm 2). As just explained, it organizes the search in two main steps. The first one is the `find_sol` phase which returns a first candidate projection α (after N_r^f repetitions). If this first step is successful, the `improve_sol` phase is repeated N_r^i times to refine the solution. The `find_sol` phase is described in Algorithm 3. At each iteration, it randomly selects d windows of

Algorithm 3 Find solution phase.

```
 $\text{find\_sol\_phase}(d, W_{len}, T_{det}, TP')$ 
 $TP' := \{N_r^f, \text{num\_hops}\}$ 
 $i = 0$ ;
repeat  $N_r^f$  times
   $\alpha = \text{random\_window}(d, W_{len})$ ;
   $\text{neighborhood} = \text{get\_neighbors\_FS}(\alpha, \text{num\_hops})$ ;
   $\text{best\_neighbor} = \text{max}(@f_{obj}, \text{neighborhood}, \mathcal{L}_{tr}^p)$ ;
  if  $f_{obj}(\text{best\_neighbor}, \mathcal{L}_{tr}^p) > T_{det}$  &  $f_{obj}(\text{best\_neighbor}, \mathcal{L}_{va}^p) > T_{det}$ 
    return  $(i + 1, \text{best\_neighbor})$ ;
  end
   $i++$ ;
end
end
```

length W_{len} with non-zero weights (function `random_window`). All the neighbors of the solution are then computed with the function `get_neighbors_FS`. Each neighbor is constructed by moving one of the windows left or right (if we see the projection vector as a row vector). The lengths of the moves considered are small multiples of the window length (as set by the `num_hops` parameter). During the computation of the neighbors, the collisions between windows are avoided in order to keep d distinct windows. Next, the best neighbor is selected as the neighbor having the maximal evaluation of f_{obj} on the set \mathcal{L}_{tr}^p . This best neighbor is finally tested to detect if a d -tuple of shares is spanned by the windows. The detection is based on a threshold T_{det} on the objective function that will be

carefully discussed in the next section. In order to dodge the randomness of the objective function when the d shares are not spanned, this threshold has to be exceeded on both the training and validation sets of traces $\mathcal{L}_{tr}^p, \mathcal{L}_{va}^p$. If those two conditions are met, the projection vector is returned by the algorithm.

If the `find_sol` phase was able to find a solution spanning the d shares, the objective function is informative enough to allow a second (intensification) step, and the `improve_sol` phase (in Algorithm 4) is run for N_r^i iterations. At each iteration, the entire neighborhood is constructed with the function

Algorithm 4 Improve solution phase.

```

improve_sol_phase( $\alpha, TP''$ )
 $TP'' := \{N_r^i, move\_steps, resize\_steps, min\_WS, max\_WS, N_n, max\_stagn\}$ 
 $\alpha_{best} = \alpha;$ 
Repeat  $N_r^i$  times
   $neighborhood = get\_neighbors\_IS(\alpha, move\_steps, resize\_steps, min\_WS, max\_WS);$ 
   $\alpha = select\_neighbor(@f_{obj}, \mathcal{L}_{tr}^p, N_n);$ 
  if  $f_{obj}(\alpha, \mathcal{L}_{tr}^p) > f_{obj}(\alpha_{best}, \mathcal{L}_{tr}^p)$ 
     $\alpha_{best} = \alpha;$ 
     $num\_stagn = 0;$ 
  else
     $num\_stagn ++;$ 
  end

  if  $num\_stagn > max\_stagn$ 
    return  $\alpha_{best};$ 
  end
end
return  $\alpha_{best};$ 
end

```

`get_neighbors_IS`. Each neighbor results from the shift (left or right) of one window or the resizing of all the windows (we keep the same size for all windows). The move steps considered are given in `move_steps`, and the resize steps in `resize_steps`. The size of the windows is constrained to remain between `min_WS` and `max_WS`. The selection of the neighbor is then performed by `select_neighbor`, as a random neighbor amongst the N_n best neighbors. Using this selection strategy allows the search to avoid being trapped into local optima, ensuring a sufficient diversification. The search also memorizes the best projection obtained since the beginning of the phase in α_{best} . This is mandatory as it is allowed to select projection vectors that decrease the objective function. Eventually, the variable `num_stagn` records the number of iterations without improvement of the best solution. Once `num_stagn` is larger than `max_stagn` or when the number of iterations reaches $N_r = N_r^f + N_r^i$, the search returns the best solution α_{best} .

As far as the technical parameters are concerned, we first set the number of hops (num_hops) in the `find_sol` phase to allow the windows covering all the dimensions of the traces. It enables an iteration to find a covering set of windows when one window is incorrectly placed. Next, in the `improve_sol` phase, the more move steps ($move_steps$) and resize steps ($resize_steps$), the quicker the algorithm converges towards the optimal windows, but the longer each iteration is. We found that a good tradeoff in our context was to use $move_steps$ of 1, 3 or 5 dimensions and $resize_steps$ of 1 dimension. Those settings allow the iterations to be fast while still covering a large part of the search space around the solution found by the `find_sol` phase. The min_WS parameter typically depends on the sampling rate of the oscilloscope used in the attack: we set it to 5 which corresponds to half a cycle in our experiments, based on the intuition that dimensions next to a POI may also contain information. max_WS was then chosen as $2 * W_{len}$, reflecting that this information can be spread on multiple clock cycles. Finally, a max_stagn value of 50 allows the local search to stop when it is unlikely to further improve the quality of the windows. And given the low span of the moves and the resizes, an exploration parameter N_n of 3 is enough to escape local optima and still converge towards the optimal solution.

4.2 Simulated experiments

We now discuss the setting of the more intuitive parameters W_{len} and T_{det} together with the performance gains obtained thanks to our specialized PP algorithm. In view of their heuristic nature, these questions are best investigated with simulated examples, where we can play with some important parameters of leaking implementations. For this purpose, we will consider a first-order masked S-box where the adversary receives N_i pairs of leakage variables of the form:

$$\begin{aligned} L_i^1 &= \text{HW}(\mathbf{S}(x \oplus s) \oplus m) + R_i^1, \\ L_i^2 &= \text{HW}(m) + R_i^2, \end{aligned} \tag{1}$$

where HW is the Hamming weight function, \mathbf{S} the AES S-box, x a plaintext byte, s a key byte, m a secret random mask, and R_i^1, R_i^2 are normally distributed noise variables with variance σ_n^2 ($1 < i \leq N_i$). For simplicity, we make sure that the N_i samples corresponding to the two shares are not overlapping. Next to these $2 \times N_i$ informative samples, we finally add $N_s - 2 \times N_i$ random samples N_j , so that N_s is the total number of samples in our simulated traces.

Setting the detection threshold An important parameter in Algorithm 3 is the threshold value used to decide whether an improvement of the objective function is significant. In this context, a particularly convenient feature of the MMPC criteria (defined in Section 2.2) is that it gradually tends to one as the number of measurements used in the detection increases. That is, given that the order of the statistical moment (e.g. $d = 2$ in our current simulations) and number of measurements used in the detection is sufficient, this criteria always reaches high values. Intuitively, it is because the MMPC relates to the statistical confidence

we have in our estimated moments rather than their informativeness (see [16] for a discussion). As a result, and using such an objective function, we are able to set the detection threshold T_{det} in a completely black box manner (i.e. independent of the implementation details). Indeed, the only thing we have to guarantee is that the MMPC as computed by the objective function is significant in front of the one that would be obtained by chance, for non-informative samples. But this essentially depends on the size of the target operations. For example, the correlation between random 256-element vectors is (roughly) Gaussian-distributed³ with mean zero. And the probability that $MMPC > 0.2$ by chance in this case is already below the one corresponding to three σ 's (i.e. below 0.1%). Of course, one can expect slight deviations from such an ideal behavior (e.g. so-called ghost peaks leading to non-zero mean MMPC for non-informative samples), but our next experiments will confirm that setting T_{det} to 0.2 is generally good.

Impact of W_{len} , σ_n^2 and N_i on the detection success. Given a detection threshold set as just explained, we can now evaluate the impact of different parameters on the success of our *find_sol* phase. In particular, the noise variance σ_n^2 , number of informative pairs of samples in the traces N_i and window length W_{len} are important in this respect. As just explained, we know that given a large enough number of measurements, the MMPC criteria should become larger than 0.2 for the informative samples. But it also means that if this number of measurements is not sufficient, the moments used in MCP-DPA will not be sufficiently well estimated and the detection may fail. As usual, the main parameter influencing the estimation complexity is the noise variance σ_n^2 . Yet, since we apply the objective function after projection in our PP algorithm, the size of the window W_{len} also matters here. Indeed, adding W_{len} samples with noise variance σ_n^2 implies a larger noise variance $W_{len} \times \sigma_n^2$ after projection. This is typically illustrated in the left part of Figure 2, where we see the impact of increasing W_{len} for two noise levels ($\sigma_n^2 = 0.1$ in the top figure, $\sigma_n^2 = 2$ in the bottom one). That is, for too large noise variances or window lengths, the estimation of the MMPC criteria is not good enough to take good decisions (i.e. is below T_{det}). In other words, more measurements are needed in this case for the PP algorithm to output meaningful results. Interestingly, we also see in the right part of the figure that adding meaningful samples in the traces (i.e. increasing N_i) quite significantly mitigates the impact of large window lengths. So intuitively, traces with multiples POIs available will better benefit from our proposed method.

Time complexity. The previous results suggest that the complexity of PP algorithms is essentially a tradeoff between time and measurement complexities. That is, increasing the windows length should decrease their time complexity⁴, but increases the noise after projection, and so the number of measurements needed to estimate the MMPC criteria with sufficient confidence. This is typically illustrated in the left part of Table 1, where we also see the benefit of having

³ More precise estimates can be obtained with Fisher's Z transform.

⁴ At most linearly since the benefit of increasing the window length W_{len} saturates whenever it is not negligible in front of the number of samples in the traces N_s .




Fig. 2. Incidence of the window length W_{len} on the information detection.

more informative samples in the traces (i.e. increasing N_i). Furthermore, the right part of the table highlights the impact of increasing the size of the traces N_s . As in a combinatorial search, the time complexity of the PP algorithm should increase quadratically with it (more generally, it depend on N_s^d with d the number of shares in the masking scheme). Yet, increasing W_{len} or N_i can make this increase quasi-linear for some (not too large) values of N_s . Besides, note that Table 1 includes all the constant factors related to the technical parameters in the previous section, which sometimes amortizes these asymptotic predictions. Note also that this table counts the calls to the objective function for readability, but this count is not fully reflective of the PP's time complexity when changing the size of the profiling sets \mathcal{L}_{tr}^p and \mathcal{L}_{va}^p , since larger sets also increase the complexity of each evaluation of the objective function. Yet, thanks to the parallelism of MCP-DPA attacks, the impact of these increases was limited in our experiments, leaving us with strong concrete results, as the next section will show.

4.3 Measured experiments

The previous simulated experiments suggest that a specialized PP algorithm can be an efficient way to find POIs in the leakage traces of masked implementations. We now would like to confirm this hope in front of a real case-study. For this purpose, we will consider the actual measurements of a first-order masked AES S-box based on table lookups [23, 28]. For every pair of input/output masks

Table 1. Impact of W_{len} , N_i and N_s on the average number of f_{obj} calls.

$N_s = 1000$		N_i	
		5	10
W_{len}	10	7306	4681
	20	3920	3008
	30	3266	2782
	50	-	2138
	100	-	1020
	150	-	-

$W_{len} = 50, N_i = 10$	N_s		
	500	1000	2000
	905	2138	4673

(m, q) , it pre-computes an S-box S^* such that $S^*(x \oplus s \oplus m) = S(x \oplus s) \oplus q$. Since this pre-computation is part of the adversary’s measurements, it leads to quite memory-consuming traces of $N_s = 30,000$ samples (which would be a challenging target for a combinatorial search). Furthermore, we verified empirically that our implementation does not lead to any (easy-to-detect) first-order information leakage, by running template attacks for all the time samples, and making sure that the success rate remained negligible (which should be guaranteed by the use of independent masks m and q , in order to prevent leakages based on the transitions between the the S-box input and output). Our motivation for using this setup was twofold. First, we selected a masking countermeasure based on pre-computed tables in view of the difficulty to obtain a first-order secure implementation based on other standard masking schemes such as [27] – see [2] for a recent discussion of this problem. Second, we purposely put ourselves in a challenging scenario with large traces, without trying to compress them (e.g. by reducing the sampling frequency or through educated guess). While we agree that concrete adversaries would try to exploit these possibilities, we assume that they would not always be able to compress traces up to feasible combinatorial search, and the experiments in this section aim to reflect this possibility.

We then analyzed our set of profiling and test traces, in order to evaluate the success and efficiency of our POI detection tool. We used the same MMPC criteria and detection threshold of 0.2 as previously discussed, and selected a window length W_{len} of 25, corresponding to approximately two clock cycles in our measurements: this is the only physical intuition used in our experiments. With these parameters, it turned out that the estimation of the objective function was sufficiently accurate (for our detection threshold to make sense) with 50 profiling traces per template (i.e. 50×256 among the 500×256 measured). Based on our 1500 test traces, we then evaluated that the local search algorithm was able to return a solution within an average of 12 000 calls to f_{obj} (roughly corresponding to 7 minutes of execution time on our desktop computer). We then repeated this search multiple times in order to find several pairs of informative windows. We finally used these windows to launch multivariate (Gaussian) template attacks using 2, 4 and 8 dimensions. For this purpose, we selected the smallest windows (which turned out to contain 5 samples) and built templates for their mean values (so that each pair of window provided us with 2 dimensions). The results of

these attacks are illustrated in Figure 3 and confirm that our tool successfully detected POIs in this challenging case⁵. Interestingly, we see that the gain due to increased dimensionalities vanishes when moving from 4-dimension templates to 8-dimension ones. We conjecture that this mainly relates to template estimation issues. Note anyway that, as mentioned in introduction, these attacks are not aimed to be optimal from the data complexity point-of-view (since we have no guarantee to find the most informative samples). Our main goal was to provide a time-efficient POI detection tool, in a black box setting. To the best of our knowledge, previous methods for this purpose would not have been able to deal with 30,000-sample traces without an educated guess (For illustration, the product traces mentioned in footnote 1 would correspond to $900 \cdot 10^6$ samples).




Fig. 3. 4th-order success rates of multivariate template attacks.

4.4 Discussion & comparison with existing works

The previous (simulated and actual) experiments underline that the detection of POIs in the leakage traces of masked implementations is a tradeoff between the time complexity of the PPs and the amount of traces available for this purpose. In this context, it is first important to remark that the selection of a good objective function is essential. In theory, any objective function that captures higher-order statistical moments of the leakage distribution can lead to successful detections. For example, the information theoretic metric of Section 2.2 is a possibility. Yet, concretely, it would be more expensive to compute and estimate, so the MMPC criteria used in this section appeared more convenient for the purposes of this paper (where we want to detect POIs, not to optimize their informativeness). Next, this tradeoff also draws pretty intuitive connections between our work and naive approaches to the detection of POIs. On the one hand, (randomized) combinatorial search corresponds to one extreme scenario, where we use windows

⁵ For convenience, and in order to limit our measurement needs, we estimated a 4th-order success rate which corresponds to an adversary able to enumerate 2^{32} keys.

of length one, implying minimum trace requirements (so the simplest possible estimation) at the cost of maximum time complexity. On the other hand, signal integration corresponds to the other extreme scenario, where we use a single window of maximum length, implying very large trace requirements (since it implies estimating very noisy projected samples) but minimum time complexity (since the window covers the full trace). Hence, PPs can be viewed as a natural and practically useful way to explore the tradeoffs between these extremes.

Besides, we would like to end this paper by discussing its differences (and complementarity) with the recent work of Reparaz et al. In summary, the results in [26] start from a different context than ours, as they primarily focus on a non-profiled attack setting. In this context, they made the useful observation that, while naive approaches to POI detection usually launch the (non-profiled) attacks for all key candidates independently, it is in fact possible to (approximately) detect these points independent of the key material first, and to launch the attacks for all the key candidates on a subset of relevant points afterwards (leading to a reduction of the time complexity by a factor roughly corresponding the number of key hypotheses in the attack). We take the example of a masked S-box to illustrate this idea, where we manipulate a masked value $V = S(x \oplus s) \oplus M$ (with the plaintext x fixed/chosen) and its corresponding mask M . The proposal of Reparaz et al. intuitively works because the informative pairs of samples for which $\hat{I}(S; X, L(V), L(M))$ is non negligible also correspond to pairs of samples for which $\hat{I}(L(V), L(M))$ is significantly larger than 0 (for a chosen x – the CHES 2012 paper also describes known plaintext variants). Interestingly, while the first quantity can only be computed using key knowledge, the estimation of the second one does not require s to be known. So overall, this previous work does not reduce the complexity of the POI detection (which is still performed based on a combination of educated guess and combinatorial search over 800 samples in [26]). It rather makes it possible to perform it only once for all key candidates in a non-profiled setting. As a result, PP algorithms could be directly combined with this proposal, by simply using Reparaz et al.’s criteria (i.e. $\hat{I}(L(V), L(M))$) as objective function. This would extend the applicability of our work to the non-profiled setting, at the cost of a less favorable “number of measurements vs. time complexity tradeoff”. Indeed, estimating such a (non-profiled) objective function would be more expensive than our profiled MMPC criteria, for the same reasons as previously mentioned for the information theoretic metric.

This discussion eventually allows us to clarify what we mean by “black box evaluation”. In our experiments, we did not require implementation details nor mask knowledge, but assumed key knowledge to speed up estimations. We believe this realistically captures the constraints of an evaluation lab. But extending our work to a non-profiled setting may be achieved by changing the objective function, as just discussed. In this context, it is eventually important to remark that the POI detection not only depends on the tradeoff between number of measurements and time complexity, but also on the “engineering intuition available” and “type of information detected”. Starting with the engineering intuition, we can observe that directly using $\hat{I}(L(V), L(M))$ as objective function would make the

setting of the detection threshold more challenging (i.e. require some engineering intuition). Indeed, the asymptotic value of this objective function depends on the leakage function, in contrast with the MMPC criteria that asymptotically tends to one independent of the leakage function (which allows a well motivated and black box choice of the detection threshold, as discussed in Section 4.2). Note that using the confidence level of a hypothesis test would be an interesting alternative for this purpose. Next, the type of information detected is another parameter that may allow speeding up POI detection. Namely, one of the reasons making an objective function costly to estimate is that its argument is an intermediate computation of which the output typically takes 256 possible values. This means that we want to learn information about a random variable with (relatively) large input range. In this respect, a natural solution is to consider a target with smaller range, and at the extreme, a binary random variable. This is in fact what is advertised in the leakage detection test introduced in [10] and further discussed in [15]. Here, the authors try to detect leakage points where there is a significant difference between two classes: the first corresponding to fixed plaintexts, the second corresponding to random plaintexts. On the positive side, this allows detecting leaking points without any profiling, with a simple (binary) hypothesis test, which indeed provides easy-to-compute confidence levels. On the negative side, there is no guarantee that the detected samples all correspond to useful (key-dependent) information, nor that all POIs can be detected in this way. In practice though, such tests have been shown quite powerful for detecting POIs in concrete implementations of, e.g. the AES. Note that leakage detection tests have mainly been applied to unprotected implementations so far, but their extension to masking is conceptually simple (see, e.g. [2]).

Summarizing, this work puts forward the interest of PP algorithms in the context of profiled black box evaluations, where they allow detecting leakage points which are certainly “of interest” (since they are key-dependent), with reasonable measurement vs. estimation cost and no implementation details. Many variations are possible in order to move to a non-profiled context, with a price to pay in estimation cost and/or need of implementation details and/or type of information detected. But all these changes of conditions will only be reflected by a change of objective function, hence preserve the relevance of our main contribution (i.e. the specialized PP algorithm). In this respect, it is important to note that independent of the evaluation context and objective function, trading time complexity for more measurements anyway becomes necessary at some point, when traces become too large for exhaustive combinatorial search. Last but not least, we recall that despite black box in the sense described in this section, our techniques still rely on some assumptions (e.g. whether the leakages are value-based or transition-based [2]). Using non-specific tests such as [10, 15] can mitigate this requirement (although it does not totally suppress it).

Before concluding, let us recall that in all these cases, the POI detection is heuristic (which is especially clear in our algorithms taking advantage of local search). Therefore, they do not guarantee to find the best (i.e. most informative) POIs. An exhaustive analysis remains the only option for this purpose.

5 Conclusions

In this work we proposed an efficient method for finding POIs in the leakage traces of cryptographic implementations. We exploit a combination of PP and local search for this purpose, and discussed the how to adapt it to the side-channel cryptanalysis problem. One of the main advantages of the method is its genericity, as it can be applied to any implementation, by simply adapting its objective function. Besides, it has very low memory requirements compared to state-of-the-art solutions and (although heuristic) works in practical time complexity. We applied our basic and specialized PP algorithms to two case studies of unprotected and 2-share masked implementations to validate our claims. Extending the specialized version to more shares would be straightforward, since this number of shares (i.e. d) is a parameter in our search algorithms.

Among the interesting open problems, we believe investigating the informativeness of the projected samples obtained with PPs in the context of protected implementations is promising – it was essentially left out of our analysis so far. Different approaches could be considered for this purpose. One would be to refine the projection vectors, possibly based on an information theoretic objective function that would better reflect the resulting attacks’ data complexity. An alternative one would be to exploit non-linear projections, e.g. inspired by the “product combining” frequently used in second-order DPA [24, 32]. Yet, preliminary results suggest that non-linear projections may be hard(er) to exploit because the addition of non-informative samples when computing the objective function has higher impact on the (non-Gaussian) noise in this case. Besides, testing new objective functions that are cheap to compute and estimate, in the profiled and non-profiled settings, is another interesting research direction.

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. Cédric Archambeau, Eric Peeters, François-Xavier Standaert, and Jean-Jacques Quisquater. Template attacks in principal subspaces. In Louis Goubin and Mitsuru Matsui, editors, *CHES*, volume 4249 of *LNCS*, pages 1–14. Springer, 2006.
2. Josep Balasch, Benedikt Gierlichs, Vincent Grosso, Oscar Reparaz, and François-Xavier Standaert. On the cost of lazy engineering for masked software implementations. *IACR Cryptology ePrint Archive*, 2014:413, 2014.
3. Lejla Batina, Jip Hogenboom, and Jasper G. J. van Woudenberg. Getting more from PCA: First results of using principal component analysis for extensive power analysis. In Orr Dunkelman, editor, *CT-RSA*, volume 7178 of *LNCS*, pages 383–397. Springer, 2012.
4. Eric Brier, Christophe Clavier, and Francis Olivier. Correlation power analysis with a leakage model. In Marc Joye and Jean-Jacques Quisquater, editors, *CHES*, volume 3156 of *LNCS*, pages 16–29. Springer, 2004.

5. Omar Choudary and Markus G. Kuhn. Efficient template attacks. In Aurélien Francillon and Pankaj Rohatgi, editors, *Smart Card Research and Advanced Applications - 12th International Conference, CARDIS 2013, Berlin, Germany, November 27-29, 2013. Revised Selected Papers*, volume 8419 of *LNCS*, pages 253–270. Springer, 2013.
6. Christophe Clavier, Jean-Sébastien Coron, and Nora Dabbous. Differential power analysis in the presence of hardware countermeasures. In Çetin Kaya Koç and Christof Paar, editors, *Cryptographic Hardware and Embedded Systems - CHES 2000, Second International Workshop, Worcester, MA, USA, August 17-18, 2000, Proceedings*, volume 1965 of *LNCS*, pages 252–263. Springer, 2000.
7. Christophe Clavier, Jean-Luc Danger, Guillaume Duc, M.Abdelaziz Elaabid, Benot Grard, Sylvain Guilley, Annelie Heuser, Michael Kasper, Yang Li, Victor Lomn, Daisuke Nakatsu, Kazuo Ohta, Kazuo Sakiyama, Laurent Sauvage, Werner Schindler, Marc Stttinger, Nicolas Veyrat-Charvillon, Matthieu Walle, and Antoine Wurcker. Practical improvements of side-channel attacks on AES: feedback from the 2nd DPA contest. *Journal of Cryptographic Engineering*, pages 1–16, 2014.
8. J.H. Friedman and J.W. Tukey. A projection pursuit algorithm for exploratory data analysis. *IEEE Transactions on Computers*, 23(9):881–890, 1974.
9. Benedikt Gierlich, Lejla Batina, Pim Tuyls, and Bart Preneel. Mutual information analysis. In Oswald and Rohatgi [19], pages 426–442.
10. Gilbert Goodwill, Benjamin Jun, Josh Jaffe, and Pankaj Rohatgi. A testing methodology for side channel resistance validation. NIST non-invasive attack testing workshop, 2011. http://csrc.nist.gov/news/_events/non-invasive-attack-testing-workshop/papers/08_Goodwill.pdf.
11. Suvadeep Hajra and Debdeep Mukhopadhyay. On the optimal pre-processing for non-profiling differential power analysis. 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 161–178. Springer, 2014.
12. Holger H Hoos and Thomas Stützle. *Stochastic local search: Foundations & applications*. Elsevier, 2004.
13. Stefan Mangard. Hardware countermeasures against DPA ? A statistical analysis of their effectiveness. In Tatsuaki Okamoto, editor, *CT-RSA*, volume 2964 of *LNCS*, pages 222–235. Springer, 2004.
14. 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.
15. Luke Mather, Elisabeth Oswald, Joe Bandenburg, and Marcin Wójcik. Does my device leak information? an a priori statistical power analysis of leakage detection tests. In Kazue Sako and Palash Sarkar, editors, *Advances in Cryptology - ASIACRYPT 2013 - 19th International Conference on the Theory and Application of Cryptology and Information Security, Bengaluru, India, December 1-5, 2013, Proceedings, Part I*, volume 8269 of *Lecture Notes in Computer Science*, pages 486–505. Springer, 2013.
16. Amir Moradi and François-Xavier Standaert. Moments-correlating DPA. *IACR Cryptology ePrint Archive*, 2014:409, 2014.
17. David Oswald and Christof Paar. Improving side-channel analysis with optimal linear transforms. In Stefan Mangard, editor, *CARDIS*, volume 7771 of *LNCS*, pages 219–233. Springer, 2012.

18. Elisabeth Oswald, Stefan Mangard, Christoph Herbst, and Stefan Tillich. Practical second-order DPA attacks for masked smart card implementations of block ciphers. In Pointcheval [22], pages 192–207.
19. 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 *LNCS*. Springer, 2008.
20. Telecom ParisTech. DPA contest v2. <http://www.dpacontest.org/v2/>.
21. B. Poettering. Fast AES implementation for Atmel’s AVR microcontrollers. <http://point-at-infinity.org/avraes/>.
22. David Pointcheval, editor. *Topics in Cryptology - CT-RSA 2006, The Cryptographers’ Track at the RSA Conference 2006, San Jose, CA, USA, February 13-17, 2006, Proceedings*, volume 3860 of *LNCS*. Springer, 2006.
23. Emmanuel Prouff and Matthieu Rivain. A generic method for secure SBox implementation. In Seun Kim, Moti Yung, and Hyung-Woo Lee, editors, *WISA*, volume 4867 of *LNCS*, pages 227–244. Springer, 2007.
24. Emmanuel Prouff, Matthieu Rivain, and Régis Bevan. Statistical analysis of second order differential power analysis. *IEEE Trans. Computers*, 58(6):799–811, 2009.
25. Mathieu Renaud, François-Xavier Standaert, Nicolas Veyrat-Charvillon, Dina Kamel, and Denis Flandre. A formal study of power variability issues and side-channel attacks for nanoscale devices. In Kenneth G. Paterson, editor, *EURO-CRYPT*, volume 6632 of *LNCS*, pages 109–128. Springer, 2011.
26. Oscar Reparaz, Benedikt Gierlichs, and Ingrid Verbauwhede. Selecting time samples for multivariate DPA attacks. In Emmanuel Prouff and Patrick Schaumont, editors, *CHES*, volume 7428 of *LNCS*, pages 155–174. Springer, 2012.
27. Matthieu Rivain and Emmanuel Prouff. Provably secure higher-order masking of AES. In Stefan Mangard and François-Xavier Standaert, editors, *CHES*, volume 6225 of *LNCS*, pages 413–427. Springer, 2010.
28. Kai Schramm and Christof Paar. Higher order masking of the AES. In Pointcheval [22], pages 208–225.
29. François-Xavier Standaert and Cédric Archambeau. Using subspace-based template attacks to compare and combine power and electromagnetic information leakages. In Oswald and Rohatgi [19], pages 411–425.
30. 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, *ICISC*, volume 5461 of *LNCS*, pages 253–267. Springer, 2008.
31. François-Xavier Standaert, Tal Malkin, and Moti Yung. A unified framework for the analysis of side-channel key recovery attacks. In Antoine Joux, editor, *EURO-CRYPT*, volume 5479 of *LNCS*, pages 443–461. Springer, 2009.
32. François-Xavier Standaert, Nicolas Veyrat-Charvillon, Elisabeth Oswald, Benedikt Gierlichs, Marcel Medwed, Markus Kasper, and Stefan Mangard. The world is not enough: Another look on second-order DPA. In Masayuki Abe, editor, *ASIACRYPT*, volume 6477 of *LNCS*, pages 112–129. Springer, 2010.

A Additional figures




Fig. 4. Successive parabolic interpolations.




Fig. 5. Projection profiles




Fig. 6. Template attack success rates from the DPA contest traces.