

Template Attacks in Principal Subspaces

C. Archambeau, E. Peeters, F.-X. Standaert*, and J.-J. Quisquater

UCL Crypto Group - Université catholique de Louvain
Place du Levant 3, B-1348 Louvain-la-Neuve, Belgium
{fstandae, jjq@uclouvain.be}

Abstract. Side-channel attacks are a serious threat to implementations of cryptographic algorithms. Secret information is recovered based on power consumption, electromagnetic emanations or any other form of physical information leakage. Template attacks are probabilistic side-channel attacks, which assume a Gaussian noise model. Using the maximum likelihood principle enables us to reveal (part of) a device's secret data for each set of recordings (i.e. leakage trace). In practice, however, the major concerns are (i) how to select the points of interest of the traces, (ii) how to choose the minimal distance between these points, and (iii) how many points of interest are needed for attacking. So far, only heuristics were provided. In this work, we propose to perform template attacks in the principal subspace of the traces. This new type of attack addresses all practical issues in principled way and automatically. The approach is validated by attacking stream ciphers such as RC4. We also report analysis results of template style attacks against an FPGA implementation of the AES Rijndael. Roughly, the template attack we carried out requires five times less encrypted messages than the best reported correlation attack against similar block cipher implementations.

1 Introduction

Since their first public appearance in 1996 [6], side-channel attacks have been intensively studied by the cryptographic community. The basic principle is to monitor one (or more) unintentional channels that leak from a device such as a smart card and to match these observations with a key-dependent leakage prediction. This channel is usually monitored thanks to an oscilloscope that samples a continuous analog signal and turns it into a discrete digitalized sequence. This sequence is often referred to as a trace. Recently, a probabilistic side-channel attack, called the Template Attack (TA), was introduced [2]. This attack was originally mounted to target stream cipher implementations. In this context, the attacker can only observe a single use of the key, usually during the initialization step of the cipher. As it is not possible to generate different leakages from the same secret key (e.g. corresponding to different plaintexts), TAs were purposed for a more efficient way of retrieving information from side-channel traces.

* Postdoctoral researcher of the Belgian Fund for Scientific Research (FNRS).

There are three main reasons that make TAs more efficient than previous approaches to exploit side-channel leakages. First, TAs usually require a profiling step, in order to build a (probabilistic) noise model of the side-channel that can be used to capture the secret information leaked by a running device. Second, TAs usually exploit multivariate statistics to characterize the dependencies between the different time instant in the traces. Finally, TAs use maximum likelihood as similarity measure, that can capture any type of dependency (if the probabilistic model is found to be adequate), whereas, for example correlation analysis only captures linear dependencies [1]. In general, the cost of these improvements is a reduction of the adversarial flexibility. For example, Hamming weight leakage models can generally be used for any CMOS devices while template attacks profile the leakage function for one particular device.

Actual TAs generally rely on the hypothesis that the leakage information is located in the variability of the leakage traces. In order to recover the secret data hidden in a cryptographic device, one has thus to focus at the time instants where the variability is maximal. However, in practice it is not clear how many and which moments exactly are important. The attacks are therefore based on heuristics, which specify these quantities according to some prior belief. For example, it is common to force the successive, relevant time instants to be one clock cycle distant. The main contribution of this work is that we take TAs a step further. Instead of applying TA directly, we first transform the leakage traces such that we are able to select the relevant features (i.e. transformed time instants) and their number automatically. Meanwhile, we do not need to determine a specific feature interdistance and when performing TAs after transformation, we still take the correlations between the features into account. In order to find a suitable transformation, we assume that the secret information leakage is mainly hidden in the local variability of the mean traces. Then, we take the optimal linear combination of the relevant time samples and perform TAs in the principal subspace of the mean traces. We call this approach principal subspace-based TA (PSTA). A principal subspace can be viewed as a lower dimensional subspace embedded in the data space¹ where each coordinate axis successively indicates the direction in which the data have maximal variability (or variance).

A standard statistical tool for finding the principal subspace of a data set is principal component analysis (PCA) [5]. PCA performs an eigendecomposition of the empirical data covariance matrix in order to identify both the principal directions (eigenvectors) and the variance (eigenvalues) associated to each one of them. However, practical issues may arise in the context of PSTAs, as the dimension of the traces is much larger (typically $\mathcal{O}(10^5)$) than the number of traces (typically $\mathcal{O}(10^3)$). Therefore, we propose to use a variant of PCA that is more suitable in this situation (see Section 3.1 for further details). An attractive feature of PSTAs is that the projected traces are aligned with the directions of maximal variance. These directions are nothing else than a weighted sum of all the time instants, the weights being determined such that the data variability is

¹ Here, the data space is the space in which the leakage traces live.

preserved after projection. So, in contrast to TAs, which selects a relevant subset of time instants according to a heuristic, PSTAs determines first the optimal (in terms of maximal variance) linear combination of these time instants. In other words, there is no need to determine an interdistance between the time samples anymore as the irrelevant ones will be assigned a small weight. Furthermore, based on the value of the eigenvalues, one can determine which (the largest) and how many directions are relevant. In order to validate our approach, we finally apply the described techniques to two implementation cases. First we target an implementation of RC4, similar to the one in [3] as a typical context where template attacks are necessary. Then, we target an FPGA implementation of the AES Rijndael. For this purpose, we suggest an adaptation of template attacks that allow characterizing the leakage traces of block ciphers. We finally compare the obtained results with previously reported ones and observe a significant improvement of the attacks efficiency (again to be traded with less flexibility).

2 Template Attacks

In this section, the underlying principle of Template Attacks (TAs) is first presented. Next, we introduce principal subspace TAs (PSTAs). In this approach, (linear) dimensionality reduction techniques are used to select automatically the most relevant features and their number. In this context, features can be understood as weighted sums of the most relevant trace samples. In addition, both the computational requirements as well as the prohibitive memory usage of standard TAs are reduced in a principled way. In our descriptions, we consider a divide-and-conquer adversarial strategy in which different small parts of a cryptographic key are recovered independently. We denote these small parts of secret information as key classes s (they typically correspond to one or two bytes of information). Additionally, we define a leakage vector $\mathbf{l}_q = [l_1, l_2, \dots, l_q]$ containing the leakage traces corresponding to q queries to a target device. In template attacks, an adversary typically gains access to a number of traces in order to build a tailored model for the leakages (i.e. templates) and then uses this model to perform an online attack against a given target device.

2.1 Templates

Suppose that an adversary is provided with N_t traces for a given operation, *e.g.* the encryption of q messages with a key class s . In template attacks, a Gaussian noise model is considered, meaning that these traces $\{\mathbf{l}_q^{s,i}\}_{i=1}^{N_t}$ are assumed to be drawn from the multivariate Gaussian distribution $\mathcal{N}(\mathbf{l}_q^{s,i} | \boldsymbol{\mu}_s, \boldsymbol{\Sigma}_s)$:

$$\mathcal{N}(\mathbf{l}_q^{s,i} | \boldsymbol{\mu}_s, \boldsymbol{\Sigma}_s) = \frac{1}{(2\pi)^{\frac{N}{2}} |\boldsymbol{\Sigma}_s|^{\frac{1}{2}}} \exp \left\{ -\frac{1}{2} (\mathbf{l}_q^{s,i} - \boldsymbol{\mu}_s)^\top \boldsymbol{\Sigma}_s^{-1} (\mathbf{l}_q^{s,i} - \boldsymbol{\mu}_s) \right\} \quad (1)$$

Note that the mean $\boldsymbol{\mu}_s$ and the covariance matrix $\boldsymbol{\Sigma}_s$ specify completely the noise distribution associated to each key class s . Constructing the templates consists then in estimating the sets of parameters $\{\boldsymbol{\mu}_s\}_{s=1}^{|S|}$ and $\{\boldsymbol{\Sigma}_s\}_{s=1}^{|S|}$.

A standard approach is to use the maximum likelihood principle. In this approach, we seek for the parameters that maximize the likelihood of the observations (traces) under the chosen noise model. Maximizing the likelihood is equivalent to maximizing the log-likelihood, which is given by:

$$\log \mathcal{L}_s \equiv \log \prod_{i=1}^{N_t} \Pr[\mathbf{l}_q^{s,i} | \mathbf{E}_s] = \sum_{i=1}^{N_t} \log \mathcal{N}(\mathbf{l}_q^{s,i} | \boldsymbol{\mu}_s, \boldsymbol{\Sigma}_s), \quad (2)$$

where $\Pr[\mathbf{l}_q^{s,i} | \mathbf{E}_s]$ is the probability density of observing a trace $\mathbf{l}_q^{s,i}$ if we assume that an encryption operation \mathbf{E}_s using the class s was carried out on the device. Direct maximization of (2) is straightforward and leads to the following estimates (corresponding to the empirical mean and covariance matrix associated to the observations $\{\mathbf{l}_q^{s,i}\}_{i=1}^{N_t}$):

$$\hat{\boldsymbol{\mu}}_s = \frac{1}{N_t} \sum_{i=1}^{N_t} \mathbf{l}_q^{s,i} \quad (3)$$

$$\hat{\boldsymbol{\Sigma}}_s = \frac{1}{N_t} \sum_{i=1}^{N_t} (\mathbf{l}_q^{s,i} - \hat{\boldsymbol{\mu}}_s)(\mathbf{l}_q^{s,i} - \hat{\boldsymbol{\mu}}_s)^\top \quad (4)$$

2.2 Attack

Assume that there are $|\mathcal{S}|$ possible secret key classes. In order to determine by which secret signal a new trace \mathbf{l}_{new} was generated, we apply Bayes' rule. This leads to the following classification rule:

$$\begin{aligned} \tilde{s} &= \operatorname{argmax}_{s^*} \hat{\Pr}[s^* | \mathbf{l}_{\text{new}}] \\ &= \operatorname{argmax}_{s^*} \hat{\Pr}[\mathbf{l}_{\text{new}} | s^*] \Pr[s^*], \end{aligned} \quad (5)$$

where $\hat{\Pr}[\mathbf{l}_{\text{new}} | s^*] = \mathcal{N}(\mathbf{l}_{\text{new}} | \hat{\boldsymbol{\mu}}_{s^*}, \hat{\boldsymbol{\Sigma}}_{s^*})$ and $\Pr[s^*]$ is the prior probability of the key class candidate s^* . The classification rule assigns \mathbf{l}_{new} to the candidate s^* with the highest posterior probability. When the secret key classes are equiprobable (which is generally assumed), we have that $\Pr[s^*] = \frac{1}{|\mathcal{S}|}$.

3 Principal Component Analysis

In practice, the number of samples N per trace is very large, typically $\mathcal{O}(10^5)$ as it depends on the sampling rate of the recording device. A high sampling rate is usually mandatory in order to retain the frequency content of the side-channel. This leads to excessive computational loads and a prohibitively large memory usage. Furthermore, it is expected that only a limited number of time samples are relevant for template attacks. Therefore, several attempts were made in the past to address these practical issues. Chari, *et al.* [2] select time samples

showing the largest difference between the mean traces $\{\hat{\boldsymbol{\mu}}_s\}_{s=1}^{|\mathcal{S}|}$. Rechberger and Oswald [8] used a similar method. Their selection rule is based on the cumulative difference between the mean traces. In addition, the traces are pre-processed by a Fast Fourier Transform (FFT) in order to remove high frequency noise. Another, simple rule is to select the points (after pre-processing) where the the largest variance of the mean traces occur. All these approaches assume that the relevant samples are the ones with the highest variability. However, they only provide heuristics and are therefore not optimal. Furthermore, they require to chose an arbitrary minimum distance between successive points (for example the clock cycle) in order to avoid redundancy and there is no satisfactory rule to determine how many such samples are needed to attack optimally.

Another, more systematic approach is to select the relevant points based on Principal Component Analysis (PCA) (see for example [4, 5]). PCA is a standard statistical tool for dimensionality reduction. It looks for a linear transformation that projects high-dimensional data into a low-dimensional subspace while preserving the data variance (*i.e.* it minimizes the mean squared reconstruction error). In order to minimize the loss of relevant information, PCA works in two steps. First, it looks for a rotation of the original axes such that the new coordinate system indicates the successive directions in which the data have maximal variance. Second, it only retains the N_c most important directions in order to reduce the dimensionality. It therefore assumes that the variability in the discarded directions corresponds to noise.

An example is shown in Appendix A, Figure 6.

3.1 Traces Principal Subspaces

In the following, we aim to apply PCA in order to find a single linear transform that maximizes the *inter-class distance* (or inter-class variance) between the different empirical mean traces $\{\hat{\boldsymbol{\mu}}_s\}_{s=1}^{|\mathcal{S}|}$ associated to the set of keys $[1, 2, \dots, s]$. PCA looks for the first principal directions $\{\mathbf{w}_m\}_{m=1}^{N_c}$ such that $N \geq N_c$ and which form an orthonormal basis of the N_c -dimensional subspace capturing maximal variance of $\{\hat{\boldsymbol{\mu}}_s\}_{s=1}^{|\mathcal{S}|}$. It can be shown [5] that the principal directions are the eigenvectors of the empirical covariance matrix, given by:

$$\bar{\mathbf{S}} = \frac{1}{|\mathcal{S}|} \sum_{s=1}^{|\mathcal{S}|} (\hat{\boldsymbol{\mu}}_s - \bar{\boldsymbol{\mu}})(\hat{\boldsymbol{\mu}}_s - \bar{\boldsymbol{\mu}})^\top. \quad (6)$$

The quantity $\bar{\boldsymbol{\mu}} = \frac{1}{|\mathcal{S}|} \sum_{s=1}^{|\mathcal{S}|} \hat{\boldsymbol{\mu}}_s$ is the average of the mean traces. In template attacks, N is typically very large, meaning that $\bar{\mathbf{S}} \in \mathbb{R}^{N \times N}$ is beyond computation capabilities. Furthermore, the total number of mean traces $|\mathcal{S}|$ is much smaller than N . The matrix $\bar{\mathbf{S}}$ is of rank $|\mathcal{S}| - 1$ (or less) and has therefore only $|\mathcal{S}| - 1$ eigenvectors. Fortunately, one can compute the first $|\mathcal{S}| - 1$ eigenvectors without having to compute the complete covariance matrix $\bar{\mathbf{S}}$ [4].

Let $\mathbf{T} = (\hat{\boldsymbol{\mu}}_1 - \bar{\boldsymbol{\mu}}, \dots, \hat{\boldsymbol{\mu}}_{|\mathcal{S}|} - \bar{\boldsymbol{\mu}}) \in \mathbb{R}^{N \times |\mathcal{S}|}$ be the matrix of the centered empirical mean traces. By definition the empirical covariance matrix is given by $\frac{1}{|\mathcal{S}|} \mathbf{T} \mathbf{T}^\top$. Let us denote the matrix of eigenvectors and eigenvalues of $\frac{1}{|\mathcal{S}|} \mathbf{T}^\top \mathbf{T}$ by respectively \mathbf{U} and $\boldsymbol{\Delta}$, the latter being diagonal. We have $(\frac{1}{|\mathcal{S}|} \mathbf{T}^\top \mathbf{T}) \mathbf{U} = \mathbf{U} \boldsymbol{\Delta}$. Left multiplying both sides by \mathbf{T} and rearranging directly leads to:

$$\bar{\mathbf{S}}(\mathbf{T}\mathbf{U}) = (\mathbf{T}\mathbf{U})\boldsymbol{\Delta}. \quad (7)$$

We see that $\mathbf{T}\mathbf{U}$ is the matrix of the $|\mathcal{S}|$ eigenvectors of $\bar{\mathbf{S}}$. In order to form an orthonormal basis, they need to be normalized. The normalized principal directions are given by:

$$\mathbf{V} = \frac{1}{\sqrt{|\mathcal{S}|}} (\mathbf{T}\mathbf{U}) \boldsymbol{\Delta}^{-\frac{1}{2}}. \quad (8)$$

The principal directions $\{\mathbf{w}_m\}_{m=1}^{N_c}$ are the columns of \mathbf{V} corresponding to the N_c largest eigenvalues of $\boldsymbol{\Delta}$. We denote these eigenvalues by the diagonal matrix $\boldsymbol{\Lambda} \in \mathbb{R}^{N_c \times N_c}$ and the corresponding matrix of principal directions by $\mathbf{W} \in \mathbb{R}^{N \times N_c}$. As discussed above, PCA can be performed when the number of data vectors is (much) lower than their dimension. Still, one may question the pertinence of the solution, as a subspace of dimensionality $|\mathcal{S}| - 1$ goes exactly through $|\mathcal{S}|$ points. The solution found by PCA only makes sense if the intrinsic dimension of the data manifold is much lower than number of observations. In other words, the solution is valid if most of the relevant information can be summarized in very few principal directions. Fortunately, this is generally the case in the context of practical template attacks. Note that the same problematic arises in computer vision in the context of automatic face recognition. Here, the very high dimensional vectors are the face images. The principal characteristics are then found by following a similar approach, known as *eigenfaces* [12].

3.2 Principal Subspace based Template Attacks

In the previous section, we showed how standard PCA can be modified in order to be used with very high-dimensional vectors such as side-channel leakage traces. This provides us with the projection matrix \mathbf{W} , which identifies successively the directions with maximal variance between the empirical mean traces $\{\hat{\boldsymbol{\mu}}_s\}_{s=1}^{|\mathcal{S}|}$. Now, in order to build principal subspace templates, we estimate the projected means $\{\boldsymbol{\nu}_s\}_{s=1}^{|\mathcal{S}|}$ and the covariance matrices of the projected traces along the (retained) principal directions $\{\mathbf{A}_s\}_{s=1}^{|\mathcal{S}|}$. These parameters are given by:

$$\boldsymbol{\nu}_s = \mathbf{W}^\top \hat{\boldsymbol{\mu}}_s, \quad \mathbf{A}_s = \mathbf{W}^\top \widehat{\boldsymbol{\Sigma}}_s \mathbf{W}. \quad (9)$$

As in standard template attacks, the noise model is given by a multivariate Gaussian distribution. However, the number of principal directions N_c is much smaller than N . Note that a direction can be considered as not being principal when the associated eigenvalue is small compared to the largest one. This will

be further discussed in Section 4. Finally, in order to classify a new trace \mathbf{l}_{new} , we apply Bayes' rule. This leads to the following classification rule:

$$\tilde{s} = \underset{s^*}{\operatorname{argmax}} \hat{\Pr}[\mathbf{W}^\top \mathbf{l}_{\text{new}} | s^*] \Pr[s^*], \quad (10)$$

with $\hat{\Pr}(\mathbf{W}^\top \mathbf{l}_{\text{new}} | s^*) = \mathcal{N}(\mathbf{W}^\top \mathbf{l}_{\text{new}} | \boldsymbol{\nu}_{s^*}, \mathbf{A}_{s^*})$.

4 Experimental results

In the experiments, the recorded traces are power leakages. We validate PSTAs both on stream ciphers (RC4) and block ciphers (AES Rijndael). Two examples of leakage traces for each encryption algorithm are shown in the Figures of Appendix B. From a practical point of view, considering a very small number $|\mathcal{S}|$ of different operations/keys can lead to a degenerate solution as only very few principal directions can be identified. This in turn may lead to poorly performing attacks. Therefore, it is convenient to augment the number of mean traces artificially in this case. For example, one can compute for each operation a pre-defined number of mean traces by picking several traces at random in the training set. Another approach is to use resampling techniques (e.g. see [3]).

4.1 RC4

The first experiments were carried out on a PIC 16F877 8-bit RISC-based microprocessor [7]. The microchip was clocked at a frequency around 4 MHz. This microprocessor requires four clock cycles to process an instruction. Each instruction is divided into four steps: (i) fetch (update of the address bus), (ii) decode and operands fetch (driven by the bus), (iii) execute and (iv) write back. We monitored the power consumption of a device by inserting a small resistor at its ground pin or power pin. The resistor value is chosen such that it disrupts the voltage supply by at most 5% from its reference². The 1-Ohm method³ was used to attack the device at the ground pin and a differential probe in the case of targeting the power pin. RC4 is a stream cipher working on a 256-byte state table denoted T hereafter. It generates a pseudo-random stream of bits which is mixed with the plaintext using a XOR function to yield a ciphertext. The state T is initialized with a variable key length (typically between 40 and 256 bytes) using the following key-scheduling algorithm:

² This is advised in IEC 61967-3: Integrated circuits - Measurement of electromagnetic emissions, 150kHz to 1GHz Part 3: Measurement of radiated emissions, surface scan method (10kHz to 3GHz), 47A/620/NP, New Work Item Proposal (July 2001).

³ See IEC 61967-4: Integrated circuits - Measurement of electromagnetic emissions, 150 kHz to 1 GHz - Part 4: Measurement of conducted emissions $1\Omega / 150\Omega$. Direct coupling method, 47A/636/FDIS, Final Draft International Standard, 2002-01-18.

```

for i from 0 to 255
  T[i] := i
j := 0
for i from 0 to 255
  j := (j + T[i] + key[i mod keylength]) mod 256
  swap(T[i],T[j])

```

The power consumption of the first iteration was monitored; the dependence on the first byte of the key is here obvious. The 256-byte state was placed in the data memory by allocating 64 bytes per bank. Therefore, it is expected to be easier to distinguish the keys located in different banks even if they have the same Hamming weight. In the RC4 experiments, 10 keys that are believed to be “close” are considered. For each one, 500 traces are used to construct the models and 300 to validate them. In other words, 500 traces are used to estimate the parameters and 300 to assess the performance. For each trace, there are 300,000 time samples. Figure 1 shows the eigenvalues in decreasing order. Clearly, most of the variance is located in very few components. In practice, 7 components are sufficient to ensure an average rate of correct classification of 93.3% (see Figure 2), meaning that most of the test traces are correctly classified at once.

By contrast, in [2] 42 test samples were selected according to some heuristic. The noise model was chosen to be multivariate Gaussian as in Equation (1). When considering a diagonal covariance matrix (i.e. the time samples are considered independent) the classification errors reported by [2] were up to 35% for similar keys. Since the power of the attack strongly depends on the implementation and the measurement noise, we also reproduced the experiments for a fully multivariate Gaussian noise model (i.e. for full covariance matrices) for comparison purposes. The samples were selected as the ones where maximal variance occurred. The minimal distance between successive samples was chosen to be equal to the clock cycle. For 42 time samples, the average classification success was 91.8%, which is already considerable. However, note that this approach requires to choose a particular distance between the samples a priori, which affects the performances considerably. For example here, a distance of half the clock cycle leads to an average classification error of only 80.5%.

A similar loss of performance is observed when choosing too few samples to construct the multivariate noise model, but when too many samples are taken, the model reliability might be questionable. Indeed, when the dimension of the data space increases, the number of observations to reliably estimate the parameters needs to increase as well. In the case of standard TA with a 42 points of interest, estimating the mean and the covariance matrix of the multivariate Gaussian noise model requires to fit $M(M + 3)/2 = 945$ parameters. However, there is only a limited number of measurements (or traces), typically few hundreds. The number of constraints increases linearly with the dimension M . There are thus only very few measurements to estimate each model parameter.




Fig. 1. Eigenvalues in descending order for RC4.




Fig. 2. Average classification rate for RC4 as a function of the number of components.

Again, an important advantage of PSTA over TA is that the number of relevant features can be inferred from the eigenvalues. Only the significant ones need to be retained; the remaining ones are thought of as being noise. Clearly, from Figure 1, it can be observed that only the first two components are important, and indeed, the average correct classification rate for two components is already 88.7% (see Figure 2). The next few components only slightly increase the power of the attack. Furthermore, in the 7-dimensional principal subspace of the traces only 70 parameters need to be estimated (as opposed to 945), while the number of data is the same. The model parameters are thus expected to be more reliably estimated. Note also that a minimal distance between the features needs not to be chosen in the case of subspace TA. As a matter of fact, the principal components are a weighted sum of many time samples, the weights being determined as the ones minimizing the loss of variance in the data.

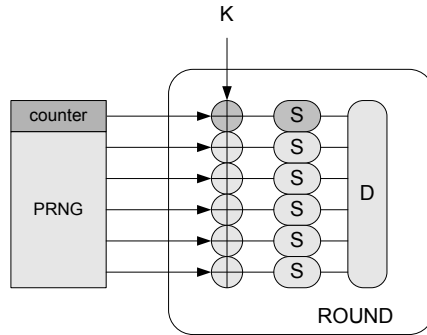


Fig. 3. Simplified view of one round in AES Rijndael. The counter feeds a particular sequence of messages to the device. The PRNG produces random message sequences. K is the encryption key, S denotes an s-box and D is the diffusion layer of the round.

4.2 AES Rijndael

Template attacks are usually applied to stream ciphers, key scheduling algorithms and pseudo-random number generators. This is motivated by the fact that such primitives are difficult to target with standard side-channel attacks like the DPA, since the attacker can only observe a single use of the key. However, in general, one could apply template attacks to any kind of cryptographic primitive in order to take advantage of a more efficient information extraction from side-channel observations. For example, in this section we show that an adaptation of subspace based TAs can be applied to FPGA implementations of block ciphers. Such a context is practically interesting since it allows to evaluate how the construction of templates may be affected by (large) amounts of algorithmic noise. It also yields particular constraints since the objective is to characterize only a part of the implemented design.

For illustration purposes, let us observe the simplified block cipher of Figure 3, where only one round is represented. In this picture, let us also assume that we want to build templates for the key bits entering the first (upper) substitution box S . Clearly, if we only want to identify the power consumption patterns of this S -box (more specifically, we want to identify the dark grey computations in the scheme, before the application of a diffusion layer), it is important to randomize all the other points in the implementation. They will then contribute to the overall leakage as random noise source. That is, all the inputs to the other S -boxes should be feed with a random number generator. Therefore, we will construct our templates according to the following procedure:

1. Select the target key bits in the implementation.
2. For each key class candidate:
 - Feed the S -box corresponding to these target key bits with the same deterministic sequence of plaintexts (e.g. a counter, typically).

- Feed the other S-boxes in the scheme with random inputs⁴.
- Build the templates from the measurement of these computations.

An important feature of this process is that each key candidate will be characterized by a number of encryptions. This is because every value in the counter will give rise to a computation that identifies these candidates. As a matter of fact, this will allow us to evaluate the efficiency of our template attack, by checking the number of encryptions required to reach a successful classification and therefore to compare our results with previous attacks against similar implementations. In practice, we targeted an FPGA implementation of the AES Rijndael [11]. Basically, we selected a loop architecture with only one round implemented in the circuit. The key scheduling was not implemented on-the-fly, but executed once, before the execution of our encryptions. However, note that the possible implementation of an on-the-fly key scheduling would not affect the construction of the templates as long as the key is fixed and therefore, once initialized, the key scheduling does not lead to any switching activity anymore.

In the experiments, 10 different keys were considered. For each one, 500 traces were used to estimate the model parameters and 500 to validate the resulting models. The number of samples per trace is equal to 500,000. Figure 4 shows the eigenvalues for AES Rijndael. Again, it can be observed that most of the variance in the data can be summarized with relatively few components. For example, with 20 components and for 128 encrypted messages the average classification success is equal to 86.7% (see Figure 5). Compared to the results with RC4, a higher number of components is necessary for a comparable classification accuracy. This result can be explained by the fact that the power traces are here much noisier (due to the parallel hardware implementation).

Although, there are relatively few significant components needed with respect to the number of encrypted messages, it is important to realize that it does not mean that the information in most of them is discarded. Indeed, in PSTAs, the PCA-step seeks of the optimal projection in the feature space. Each component corresponds thus to a weighted sum of a possibly high number of time samples. Therefore, the information leakage due to a possibly high number of encrypted messages is summarized in a single component. Figure 5 shows the average correct classification rate as a function of the number of retained components and the number of messages. As expected, when the number of encryptions decreases, the performances drops. This is due to the fact that there is less information leakage available. Similarly, when the number of components is small, there is only a poor capacity to classify correctly, as too many relevant features have been discarded. When the number of messages and the number of components increases the average correct classification rate rapidly increases.

⁴ Random inputs are used not only when constructing the templates, but also when evaluating the performance of the attack. Therefore, this set up mimics a device with unknown inputs for the other S-boxes as desired. Note that a convenient way to generate these random inputs is to use the feedback from the block cipher outputs.




Fig. 4. Eigenvalues in descending order for AES Rijndael.




Fig. 5. Average correct classification rate for AES, as a function of the number of encrypted messages and the number of retained components.

Compared to recent correlation-based power analysis attacks of AES Rijndael (also on FPGAs), the number of message required to recover the correct key bytes is much smaller. The factor of proportionality ranges from 2 to 5 depending on the fact that the attack uses trace averaging [10] or not [9]. Note also that correlation attacks require in general to carefully preprocess the traces, for example using several filters. By contrast, PSTA is much more practical as it exploits the information in the raw data directly and does not require to adjust any tuning parameters, but the number of components to retain.

5 Conclusion

In this work, we introduced principal subspace template attacks and showed that they can be successfully applied to both stream and block ciphers. Preprocessing the leakage traces beforehand by PCA allows avoiding the practical issues of ordinary template attacks. Principal subspace template attacks are motivated by the fact that template attacks consider the time instants having a great variability as being important to discriminate. If this assumption is correct, then PCA is the optimal (linear) transformation to identify the most relevant features. Besides, the eigenvalues provide a systematic rule for determining how many and which features should be selected to mount a powerful attack. Finally, it is also important to realize that the main difference between both attacks resides in the way they extract information from traces. In template attacks M of the N samples are used to mount the noise model, the selection being based on heuristics, while in principal subspace template attacks M linear combinations (preserving maximal variance) of these N samples are used. Note finally that PCA optimizes the inter-class variances without considering the intra-class variances. It is therefore not an optimal tool and still relies on heuristics that have to be experimentally confirmed (as in this work) and could be theoretically improved.

References

1. 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 *Lecture Notes in Computer Science*, pages 16–29. Springer, 2004.
2. Suresh Chari, Josyula R. Rao, and Pankaj Rohatgi. Template attacks. In Burton S. Kaliski Jr., Çetin Kaya Koç, and Christof Paar, editors, *4th International Workshop on Cryptographic Hardware and Embedded Systems (CHES)*, volume 2523 of *Lecture Notes in Computer Science*, 13–28. Springer, 2002.
3. B. Efron and R.J. Tibshirani. *An introduction to the Bootstrap*. Chapman and Hall, London, 1993.
4. K. Fukunaga. *Introduction to Statistical Pattern Recognition*. Elsevier, NY, 1990.
5. I. T. Jolliffe. *Principal Component Analysis*. Springer-Verlag, New York, 1986.
6. Paul C. Kocher. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In Neal Koblitz, editor, *16th Annual International Cryptology Conference (CRYPTO)*, volume 1109 of *Lecture Notes in Computer Science*, 104–113. Springer, 1996.

7. Microship, PIC16F877 datasheet, www.microchip.com.
8. Christian Rechberger and Elisabeth Oswald. Practical template attacks. In Chae Hoon Lim and Moti Yung, editors, *5th International Workshop on Information Security Applications (WISA)*, volume 3325 of *Lecture Notes in Computer Science*, 440–456. Springer, 2004.
9. F.-X. Standaert, S.B. Ors, and B. Preneel. Power analysis of an FPGA implementation of Rijndael: Is pipelining a DPA countermeasure? In Marc Joye and Jean-Jacques Quisquater, editors, *6th International Workshop Cryptographic Hardware and Embedded Systems (CHES)*, volume 3156 of *Lecture Notes in Computer Science*, 30–44. Springer, 2004.
10. F.-X. Standaert, E. Peeters, F. Macé, and J.-J. Quisquater. Updates on the security of FPGAs against power analysis attacks. In *proceedings of ARC 2006*, Lecture Notes in Computer Science, vol. 3985, pp. 335–346, 2006.
11. F.-X. Standaert, G. Rouvroy, J.-J. Quisquater, and J.-D. Legat. Efficient implementation of Rijndael encryption in reconfigurable hardware: Improvements and design tradeoffs. In Colin D. Walter, Çetin Kaya Koç, and Christof Paar, editors, *5th International Workshop Cryptographic Hardware and Embedded Systems*, volume 2779 of *Lecture Notes in Computer Science*, 334–350. Springer, 2003.
12. M. Turk and A. Pentland. Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, 3(1):71–86, 1991.

A Illustration of PCA

An illustration of PCA is shown Figure 6. The data is drawn from a 2-dimensional Gaussian distribution. The two principal directions \mathbf{v}_1 and \mathbf{v}_2 are shown by the solid lines. The length of the lines is proportional to the variance of the projected data onto the corresponding direction. If we remove the second dimension (after rotation) and describe the data only by the first one, then we will minimize the loss of information (i.e., loss of variance) due to this new representation.




Fig. 6. Illustration of principal component analysis (PCA).

B Exemplary leakage traces




Fig. 7. Example of a RC4 power trace.




Fig. 8. Example of an AES Rijndael power trace.