

Technical analysis of the masked AES implementation

LSC, ANSSI

January 6, 2021

Contents

1	Introduction	3
2	Secured AES: Affine masking	4
3	Characterization phase	6
3.1	Key bytes	6
3.2	Plaintext bytes	7
3.3	Ciphertext bytes	7
3.4	Mask bytes	8
3.4.1	State masks	8
3.4.2	Key masks	13
3.5	Sbox output bytes	16
3.5.1	Raw bytes	16
3.5.2	Multiplied by mask	16
3.5.3	Affine masked	17
3.6	Unshuffled output bytes	17
3.6.1	Raw bytes	17
3.6.2	Multiplicately masked	18
3.6.3	Affinely masked	19
4	First-order correlation analysis	19
4.1	Privileged knowledge	19
4.1.1	Affinely masked Sbox output	20
4.1.2	Multiplicately masked Sbox output	20
4.1.3	Raw Sbox output	21
4.2	Unknown permutation, processed traces	22
4.2.1	Affinely masked Sbox output	22
4.2.2	Multiplicately masked Sbox output	22
4.2.3	Raw Sbox output	23
4.3	Unknown permutation, non-processed traces	24
4.4	Summary	24

5	Second-order correlation analysis	24
5.1	Privileged knowledge	25
5.1.1	Multiplicatively masked Sbox output	25
5.1.2	Raw Sbox output	25
5.2	Unknown permutation, processed traces	26
5.2.1	Multiplicatively masked Sbox output	26
5.2.2	Raw Sbox output	27
5.3	Unknown permutation, non-processed traces	28
5.4	Summary	28
6	Third-order correlation analysis	28
6.1	Privileged knowledge	29
6.1.1	Raw Sbox output	29
6.2	Unknown permutation, processed traces	30
6.2.1	Raw Sbox output	30
6.3	Unknown permutation, non-processed traces	31
6.4	Summary	31

1 Introduction

This document describes the side-channel characterization of a secure AES encryption implementation. We perform this study on the ChipWhisperer target based on a STM32F303RCT7 chip (using a Cortex-M4 core). Using the ChipWhisperer as a board for side-channel characterization is explained by an easy access to power consumption and clean signal acquisition chain. Furthermore, this board natively supports the *underclocking* of the Cortex-M4 target using an external oscillator: slowing down the core frequency to 4 MHz allows to use reasonable sampling rates during the acquisition.

This evaluation was performed on power consumption traces captured through an oscilloscope, sampling 100.000.000 samples by second. The obtained traces consist in 2.000.000 samples, encompassing the whole AES implementation. Figure 1 shows two main steps: the first one is the pre-processing, and the second one is the execution of the raw AES rounds. More details are given in Section 2.

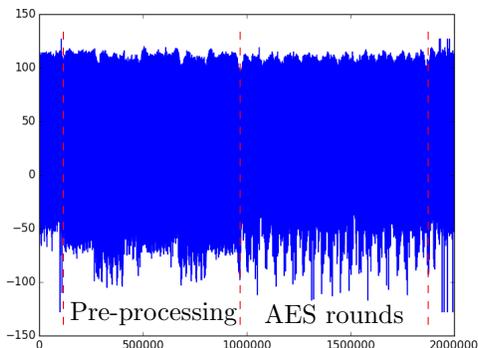


Figure 1: Power consumption trace of the AES encryption.

No resynchronization step has been performed on these traces. Nonetheless, visual inspection indicates that no significant desynchronization occurs. To illustrate this, Figure 2 displays a temporal zoom on three different traces.

A first acquisition campaign of 50.000 traces of the rolled version (*ie.*, the default version) is used to perform the characterisation step as well as the first-order resilience assessment. For the second and third order assessments, we used a second campaign of 100.000 traces of the unrolled version. The following table summarizes the results. The analyses are further detailed in the rest of this document.

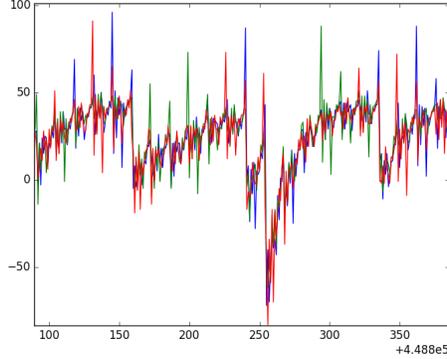


Figure 2: Three traces on a short temporal window.

	Known permutation			Unknown permutation		
Known	r_m, r_{out}	r_m	None	r_m, r_{out}	r_m	None
Order 1	≈ 1.000	> 50.000	> 50.000	≈ 20.000	> 50.000	> 50.000
Order 2	N/A	> 8.000	> 100.000	N/A	> 100.000	> 100.000
Order 3	N/A	N/A	> 100.000	N/A	N/A	> 100.000

Figure 3: Number of traces needed for a successful first-order, second-order, and thirs-order attack, depending on the masks known to the attacker.

2 Secured AES: Affine masking

The version of the AES is an affine implementation according to the paper *Affine masking against high-order side channel analysis* [?]. Each byte state of the AES state, denoted $state[i]$ for the i -th byte, is manipulated under the form $r_m \times state[i] \oplus state_M[i]$: r_m is a non zero random byte, $state_M[i]$ is initialized with 16-random bytes denoted $M[0] \dots M[15]$, and \times denotes the multiplication over the AES finite field [?]. For each round, the AES operations are applied to both $state$ and $state_M$.

Similarly, the key schedule operation is also protected by manipulating each byte of the key state, denoted $state_K[i]$, under an affine form $r_m' \times state_K[i] \oplus state_M'[i]$ where $state_M'$ is initialize with 16-random bytes denoted $M'[0] \dots M'[15]$ independent from the aes random bytes $M[0] \dots M[15]$. As the key schedule is performed independently from encryption and decryption operations, before using each masked key byte, it has to be changed into $r_m \times state_K[i] \oplus state_M'[i]$ (same multiplicative byte r_m) to have a consistent computation.

To sum up, the AES encryption/decryption and the key schedule each use 19 random bytes as input:

- 16 bytes for Boolean masking $M[0] \dots M[15]/M'[0] \dots M'[15]$ that initialized the masking state $state_M/state_K$;

- 1 multiplicative byte r_m/r_m' ;
- 2 bytes $r_{in}, r_{out} / r_{in}', r_{out}'$ (two boolean masking bytes used for SubByte operation).

To make this report more readable, we provide the correspondance Table 1 between the notations in this report (“Here”) and in the “assembly code”.

Assembly code	Here	size
<code>key_rin[0]</code>	r_{in}'	1
<code>key_rout[0]</code>	r_{out}'	1
<code>key_rmult[0]</code>	r_m'	1
<code>key_maskedState[i]</code>	$r_m \times \text{state}_K[i] \oplus \text{state}'_M[i]$	1
<code>key_masksState[i]</code>	$\text{state}'_M[i]$	1
<code>key_gtab</code>	GTab	16×16
<code>key_permIndicesMC</code>	permIndicesMC	4
<code>aes_state[i]</code>	$r_m \times \text{state}[i] \oplus \text{state}_M[i]$	1
<code>aes_state2[i]</code>	$\text{state}_M[i]$	1
<code>aes_rin[0]</code>	r_{in}	1
<code>aes_rout[0]</code>	r_{out}	1
<code>aes_rmult[0]</code>	r_m	1
<code>aes_gtab</code>	GTab	16×16
<code>aes_permIndices</code>	permIndices	16
<code>aes_permIndicesBis</code>	permIndicesBis	16
<code>aes_permIndicesMC</code>	permIndicesMC	4
<code>aes_permIndicesBisMC</code>	permIndicesMCbis	4
<code>aes_sboxMasked</code>	Sbox _m	16×16

Table 1: Notations correspondance between this report and the assembly code. Sizes are in number of bytes.

We can distinguish three main steps:

- 1 Pre-processing: before the AES rounds, the following operations in this order are performed : loading of inputs, computation of the table GTab (multiplication by r_m), computation of the affine sbox denoted Sbox_m, key schedule with affine masking.
- 2 AES rounds: 10 AES rounds with affine masking of the state and subkeys.
- 3 Post-processing: after AES rounds, the state `state` is unmasked by removing the Boolean mask `stateM` and by removing the multiplicative factor r_m (multiplying by r_m^{-1}).

Shuffling: Many sub-operations on the AES state are performed in a random order: each byte is processed in a random order with pre-computed permutation based on the input random values. For AES rounds: `permIndicesMC` is used to shuffle MixColumns of state and `permIndicesMCbis` to shuffle MixColumns of

state_M . Permutation `perIndicesMC` (resp. `perIndicesMCbis`) is computed from the initial value $\text{state}_M[0] = M[0]$ (resp. $\text{state}_M[1] = M[1]$). The permutation `perIndices` is used to shuffle `SubBytes` and `ShiftRows` of `state` and the permutation `perIndicesBis` is used to shuffle `SubBytes` and `ShiftRows` of state_M . These permutations are computed from $M[0], M[1], M[2], M[3]$.

For key schedule: the four permutations are similar. The permutation `perIndicesMC` is used to shuffle the manipulation of the words' coordinates of the masked key state (the term word referring here to 4-bytes vector) and `perIndicesMCbis` is used to shuffle the manipulation of the words' coordinates of state'_M . The former permutation is computed from $M'[0]$, while the latter one is computed from $M'[1]$. The permutation `perIndices` is used to shuffle `SubBytes` and `ShiftRows` of the masked key state and `perIndicesBis` is used to shuffle `SubBytes` and `ShiftRows` of the key mask state state'_M . These permutations are computed from $M'[0], M'[1], M'[2], M'[3]$.

3 Characterization phase

In this section, we perform several Signal-to-Noise Ratio (SNR) computations in order to identify univariate leakage samples related to the manipulation of different sensitive values.

3.1 Key bytes

We perform a SNR targeting each key byte. We superpose the leakage and SNR curves for the first key byte in Figure 4 to locate where the SNR peaks are located during AES computation. For each key byte, the results evidence one leakage point, which highlights the key manipulation during its masking (`Load_masterKey` function).

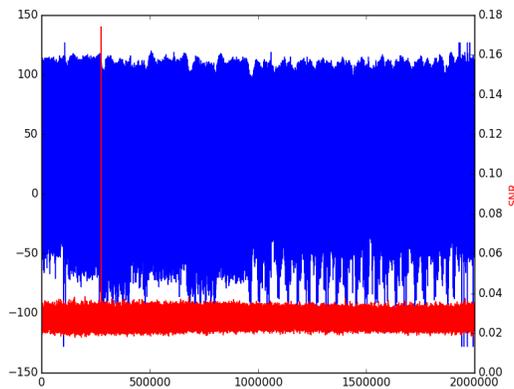


Figure 4: SNR targeting key byte 0 (in red), superposed to one leakage trace. SNR computed using 10.000 traces.

3.2 Plaintext bytes

Similarly, we perform a SNR targeting each plaintext byte as shown in Figure 5. The results evidence three leakage points. The first leakage is explained by the loading of the plaintext in the target. The second leakage is explained by the loading of the plaintext in the context (`Load_data` function). The third leakage is explained by the masking of the plaintext (`Map_in_G` function).

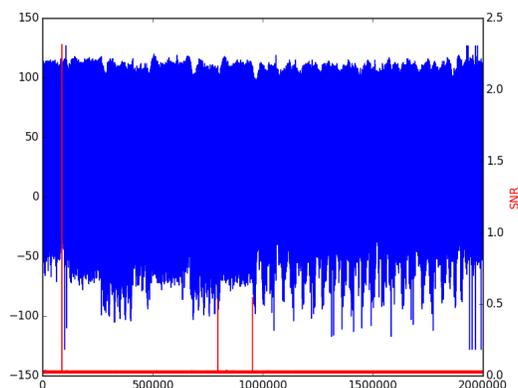


Figure 5: SNR targeting plaintext byte 0 (in red), superposed to one leakage trace. SNR computed using 10.000 traces.

3.3 Ciphertext bytes

Similarly, we perform a SNR targeting ciphertext bytes as shown in Figure 6. Leakages appear during the unmasking operation (`Multiplicative_unmasking` function) and the serial writing of the bytes: when the ciphertext is loaded in the output variable.

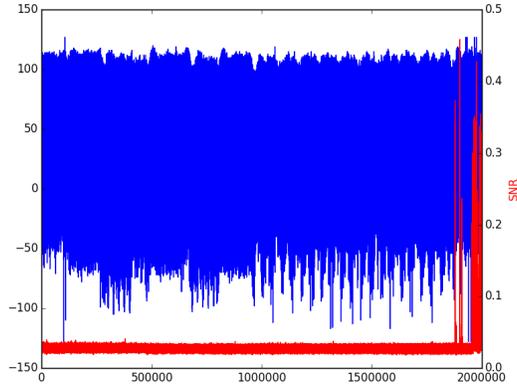


Figure 6: SNR targeting ciphertext byte 0 (in red), superposed to one leakage trace. SNR computed using 10.000 traces.

3.4 Mask bytes

Similarly, we perform a SNR targeting each mask byte. We detail the results hereafter, depending on the usages of each mask in the implementation.

3.4.1 State masks

Bytes $M[0]$ to $M[15]$, r_{in} , r_{out} and r_m are used for the masking of the state of the encryption. We characterize hereafter their leakages.

Figure 7 and Figure 8 illustrate the results when targeting bytes $M[0]$ and $M[1]$. These bytes are used for computing both the shuffling permutation used in the substitution layer of the AES, and the permutations used in the linear layer. They are also used as linear masks. Leakages appear during the loading of the bytes (`Load_random` function). Then, leakages appear during the computation of the shuffling permutation of the substitution layer, `Compute_permIndices_Tables`, and the computation of the permutations in the linear layer (`Compute_permIndices_Tables_MC` functions). Another leakage peak appears when the masks are loaded just before the encryption (`Load_data` function), and a smaller peak appears when the state is masked (`Xor_states` function). Finally, we observe nine peaks corresponding to the 9 `MixColumns` operations.

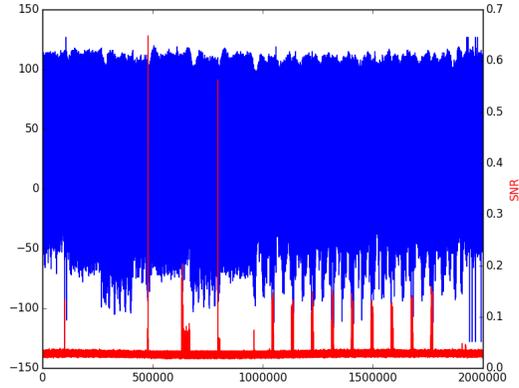


Figure 7: SNR targeting mask byte $M[0]$ (in red), superposed to one leakage trace. SNR computed using 10.000 traces.

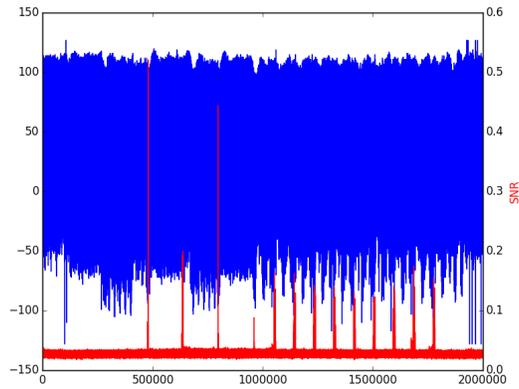


Figure 8: SNR targeting mask byte $M[1]$ (in red), superposed to one leakage trace. SNR computed using 10.000 traces.

Figures 9 and 10 illustrate the results when targeting bytes $M[2]$ and $M[3]$. These bytes are used for computing the shuffling permutation used in the substitution layer, and as linear masks. The leakages are similar to those of bytes $M[0]$ and $M[1]$, except that no leakage is observed during the `MixColumns` operations, nor during the computation of the permutations used in the linear layer.

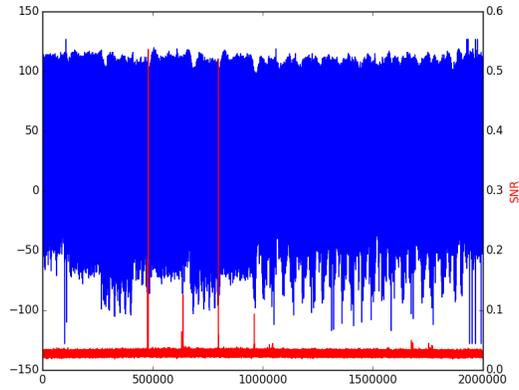


Figure 9: SNR targeting mask byte $M[2]$ (in red), superposed to one leakage trace. SNR computed using 10.000 traces.

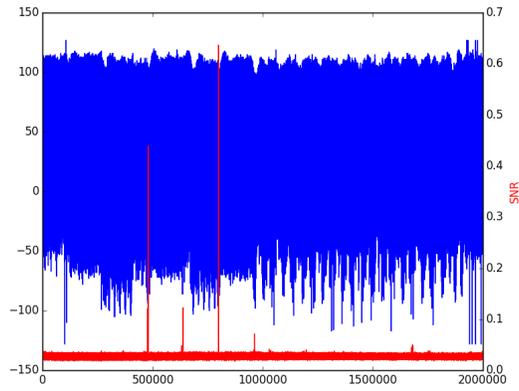


Figure 10: SNR targeting mask byte $M[3]$ (in red), superposed to one leakage trace. SNR computed using 10.000 traces.

Figures 11 and 12 illustrate the results when targeting bytes $M[4]$ and $M[15]$. These bytes are used as linear masks. Similar results are obtained for bytes $M[5]$ to $M[14]$. The leakages are similar to the one obtained for previous bytes, but are now only observed for the loadings and masking of the state.

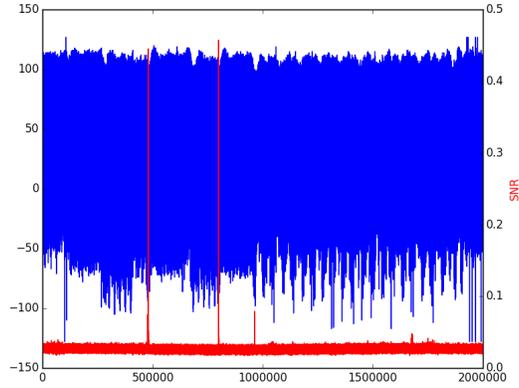


Figure 11: SNR targeting mask byte $M[4]$ (in red), superposed to one leakage trace. SNR computed using 10.000 traces.

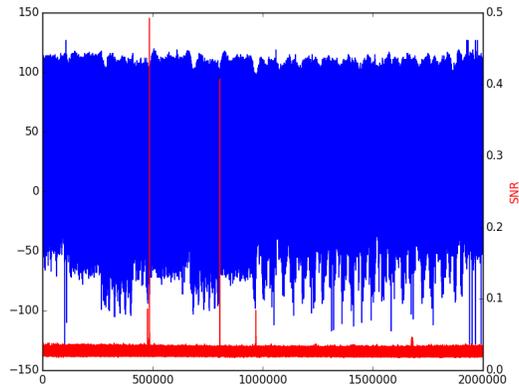


Figure 12: SNR targeting mask byte $M[15]$ (in red), superposed to one leakage trace. SNR computed using 10.000 traces.

Figure 13 illustrates the results when targeting r_{in} . This byte is used as the input mask of the substitution layer. Leakage first appears during the loading of the randoms (`Load_data`). Strong leakage can then be observed during the re-computation of the substitution table (`Compute_Affine_sboxMasked` function). Finally 10 peaks appear during the substitution steps of each round of the AES.

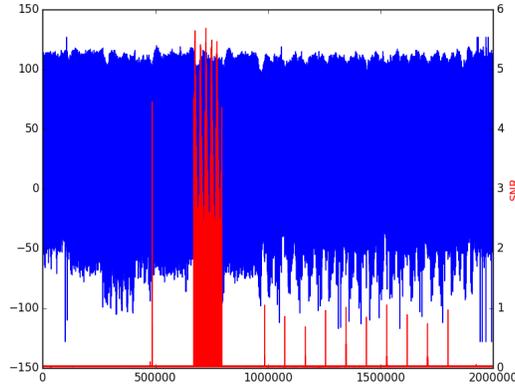


Figure 13: SNR targeting mask byte r_{in} (in red), superposed to one leakage trace. SNR computed using 10.000 traces.

Figure 14 illustrates the results when targeting r_{out} . This byte is used as the output mask of the substitution layer. Leakages appear at similar time samples as for r_{in} .

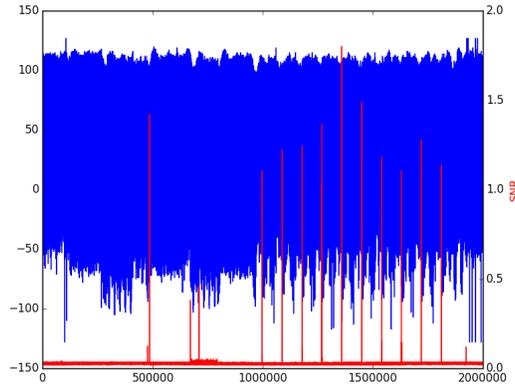


Figure 14: SNR targeting mask byte r_{out} (in red), superposed to one leakage trace. SNR computed using 10.000 traces.

Figure 15 illustrates the results when targeting r_m . This byte is used as the multiplicative mask of the substitution layer. Leakage once again appear when the randoms are loaded (`Load_data` function). Leakage also appear during the computation of the multiplicative table (`Compute_GTab` function), and during the computation of the substitution box (`Compute_Affine_sboxMasked` function).

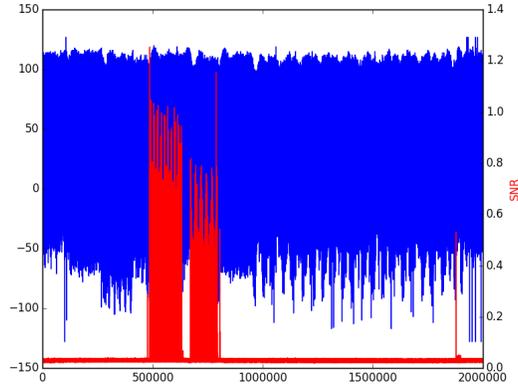


Figure 15: SNR targeting mask byte r_m (in red), superposed to one leakage trace. SNR computed using 10.000 traces.

3.4.2 Key masks

Bytes $M'[0]$ to $M'[15]$, r_{in}' , r_{out}' and r_m' are used for the masking of the key. We characterize hereafter their leakages. Figure 16 illustrates the results when targeting byte $M'[0]$. These bytes are used for computing the shuffling permutation used in the key expansion. They are also used as linear masks.

Observed peaks correspond to the loading of the random in the target, loading of the masks (`Load_random_key` function), masking of the key (`Load_masterKey`), and the key expansion (`KeyExpansion_masked` function), due to the involvement of the mask byte in the computation of the shuffling permutation (`Compute_permIndices_Tables_MC` function).

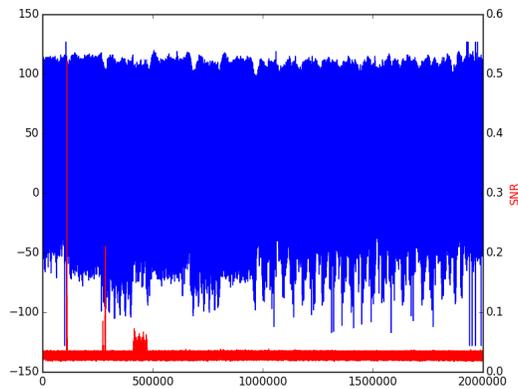


Figure 16: SNR targeting mask byte $M'[0]$ (in red), superposed to one leakage trace. SNR computed using 10.000 traces.

Figures 17 and 18 illustrate the results when targeting bytes $M'[1]$ and $M'[2]$.

Similar results are observed when targeting bytes $M'[3]$ to $M'[15]$. These bytes are not used for the computation of the shuffling permutation. However, they serve as linear masks during the key expansion, and, as such, a small peak appears at the beginning of the `KeyExpansion_masked` function. The other peaks are similar to the ones appearing for byte $M'[0]$.

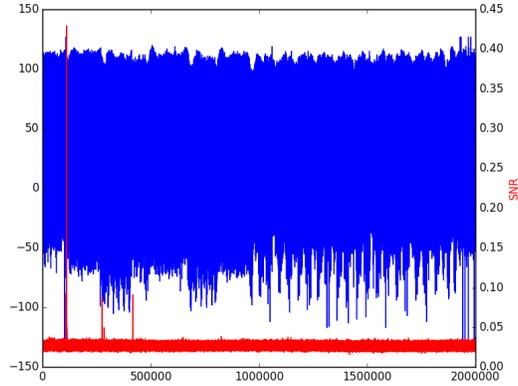


Figure 17: SNR targeting mask byte $M'[1]$ (in red), superposed to one leakage trace. SNR computed using 10.000 traces.

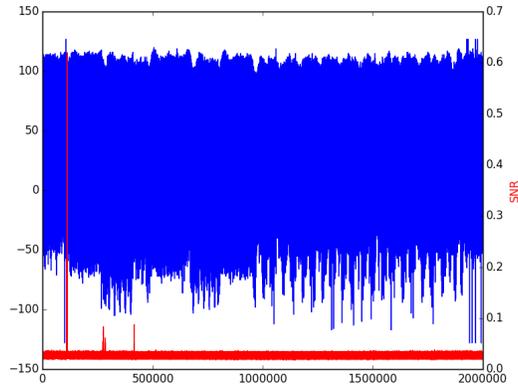


Figure 18: SNR targeting mask byte $M'[2]$ (in red), superposed to one leakage trace. SNR computed using 10.000 traces.

Figure 19 illustrates the results when targeting r_{in}' . This byte is used as the input mask of the substitution layer. Leakages appear at the same time samples as for bytes $M'[1]$ to $M'[15]$, and also during the recomputation of the Sbox and the key expansion.

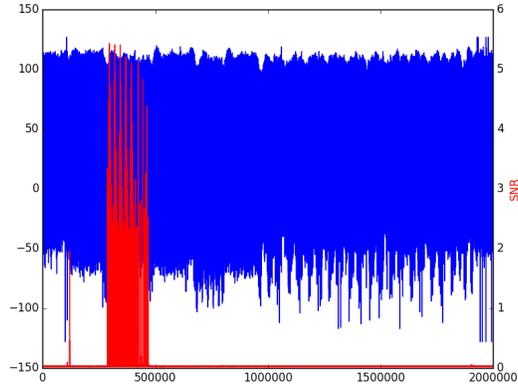


Figure 19: SNR targeting mask byte r_{in}' (in red), superposed to one leakage trace. SNR computed using 10.000 traces.

Figure 20 illustrates the results when targeting r_{out}' . This byte is used as the output mask of the substitution layer. Leakages appear at the same time samples as for bytes $M'[1]$ to $M'[15]$, and also during the recomputation of the Sbox and the key expansion.

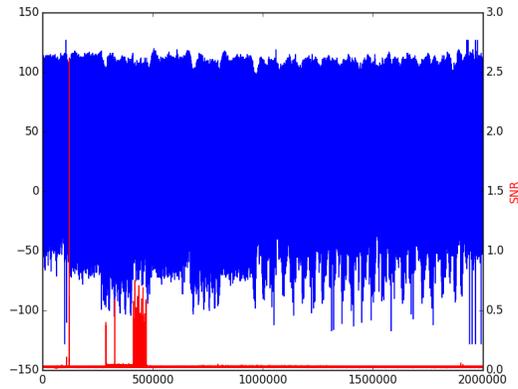


Figure 20: SNR targeting mask byte r_{out}' (in red), superposed to one leakage trace. SNR computed using 10.000 traces.

Figure 21 illustrates the results when targeting byte r_m' . This byte is used as the multiplicative mask of the substitution layer. Leakages appear at the same time samples as the previous bytes, and also during the computation of the multiplication table G_{tab} . Finally, a peak appears just before the encryption, when the value is loaded.

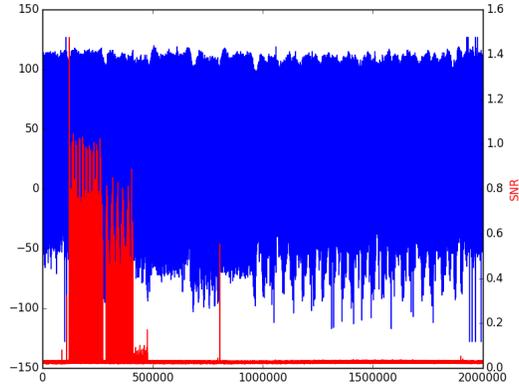


Figure 21: SNR targeting mask byte r_m' (in red), superposed to one leakage trace. SNR computed using 10.000 traces.

3.5 Sbox output bytes

3.5.1 Raw bytes

Figure 22 illustrates the results when targeting a Sbox output byte without taking into account the output mask (ie. for each byte i , the value $S[P[i] \oplus K[i]]$ is targeted) and there is no observable leakage.

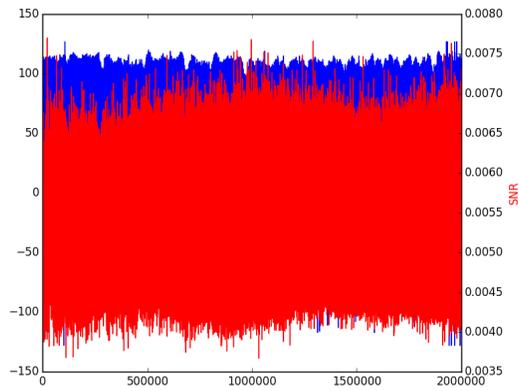


Figure 22: SNR targeting Sbox output byte 0 (in red), superposed to one leakage trace. SNR computed using 50.000 traces.

3.5.2 Multiplied by mask

Figure 23 illustrates the results when targeting a byte of the Sbox output, multiplied by the multiplicative mask, ie., for each byte i , the value $r_m \times S[P[i] \oplus K[i]]$ is targeted. No significant leakage is observed.

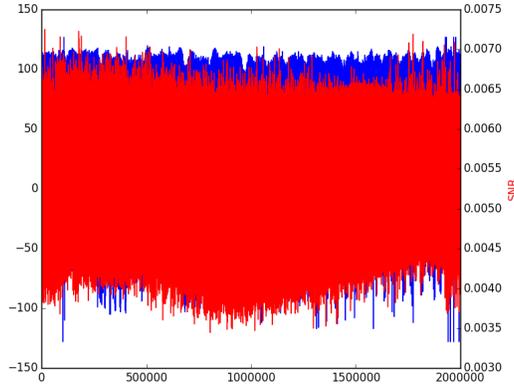


Figure 23: SNR targeting the multiplicatively masked Sbox output byte 0 (in red), superposed to one leakage trace. SNR computed using 50.000 traces.

3.5.3 Affine masked

Figure 24 illustrates the results when targeting a byte of the Sbox output, affinely masked, ie., for each byte i , the value $r_m \times S[P[i] \oplus K[i]] \oplus r_{out}$ is targeted. We remark that 50.000 traces are not enough to highlight a leakage of this value. This result is expected because of the shuffling countermeasure.

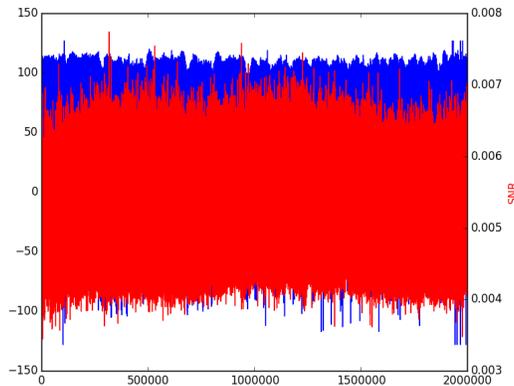


Figure 24: SNR targeting the affinely masked Sbox output byte 0 (in red), superposed to one leakage trace. SNR computed using 50.000 traces.

3.6 Unshuffled output bytes

3.6.1 Raw bytes

We use the knowledge of the mask bytes to inverse the shuffling Sh performed during the substitution phase. For each index i , we target the value $S[P[Sh[i]] \oplus$

$K[Sh[i]]$. No significant leakage can be observed. Figure 25 illustrates the results.

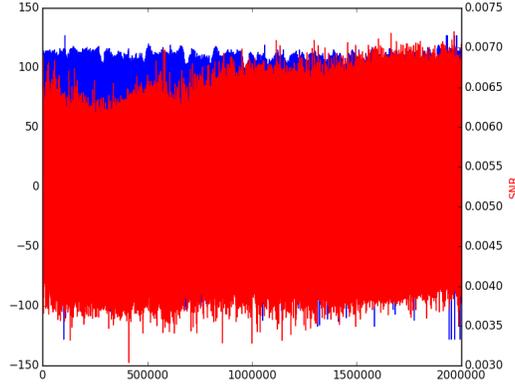


Figure 25: SNR targeting the unshuffled Sbox output byte 0 (in red), superposed to one leakage trace. SNR computed using 50.000 traces.

3.6.2 Multiplicately masked

We use the knowledge of the mask bytes to inverse the shuffling Sh performed during the substitution phase. For each index i , we target the value $r_m \times S[P[Sh(i)] \oplus K[Sh(i)]]$. A small amount of leakage is present, at the time of the manipulation of r_m . This is easily explained since the random variables $r_m \times S[P[Sh(i)] \oplus K[Sh(i)]]$ and r_m are not independent. Figure 26 illustrates the results.

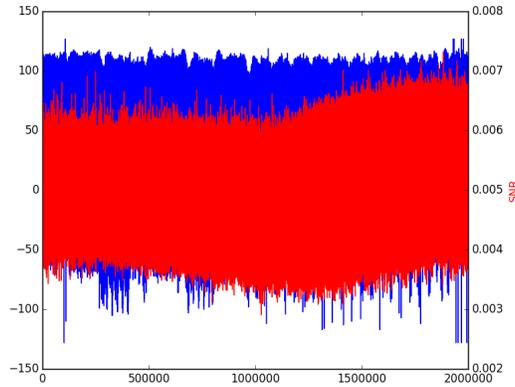


Figure 26: SNR targeting the unshuffled multiplicately masked Sbox output byte 0 (in red), superposed to one leakage trace. SNR computed using 50.000 traces.

3.6.3 Affinely masked

We use the knowledge of the mask bytes to inverse the shuffling Sh performed during the substitution phase. For each index i , we target the value $r_m \times S[P[Sh(i)] \oplus K[Sh(i)]] \oplus r_{out}$. A significant leakage is observed after the first substitution layer. Figure 27 illustrates the results.

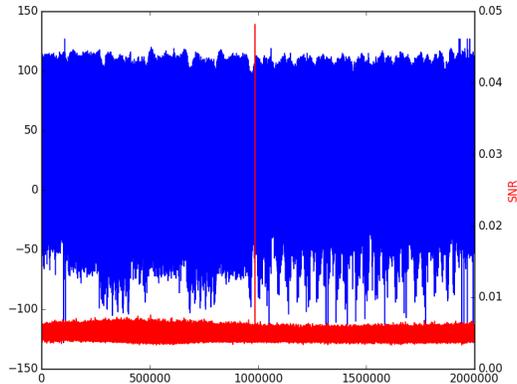


Figure 27: SNR targeting the unshuffled affinely masked Sbox output byte 0 (in red), superposed to one leakage trace. SNR computed using 50.000 traces.

4 First-order correlation analysis

In this section, we analyse the effectiveness of the implemented countermeasures against first-order side-channel analysis based on the linear correlation coefficient (Correlation Power Analysis CPA). For this study, we will target the Sbox output value.

We separate this study in two parts. In the first part, we will characterize the effectiveness of a CPA against this implementation, by considering an attacker with privileged knowledge. In the second part, we will evaluate the effectiveness of a CPA against this implementation without such knowledge.

4.1 Privileged knowledge

We first study the setting where the attacker knows the random masks used to compute the permutation.

In order to characterize the effectiveness of a CPA against this implementation, we use the random masks to recompute the permutation $Sh(i)$ for each byte index i . We then perform an attack targeting the value $Z_{\hat{K}}[i] = S[P[Sh(i)] \oplus K[Sh(i)] \oplus \hat{K}]$, where \hat{K} is an hypothesis on a value of one byte. The attack is successful if the best hypothesis returned by the attack is 0.

4.1.1 Affinely masked Sbox output

We suppose the knowledge of both the multiplicative mask and the output mask. This knowledge allows for the prediction of the value $r_m \times Z_{\hat{K}}[i] \oplus r_{out}$ for any hypothesis K . Depending on the index byte, the attack succeeds with around 1.000 traces. Figure 28 illustrates the result of the attack targeting byte 0, using 1.000 traces.

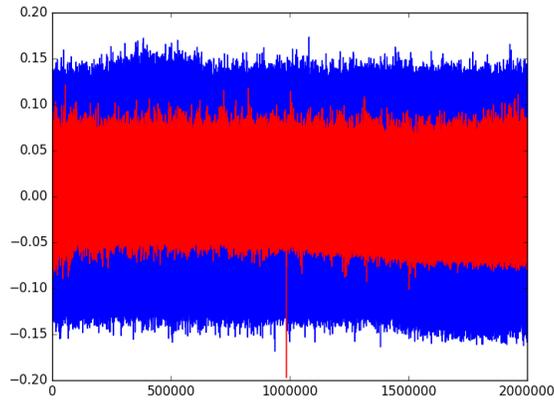


Figure 28: Correlation coefficients obtained when targeting $r_m \times Z_{\hat{K}}[i] \oplus r_{out}$, for every value of \hat{K} . Correct hypothesis is plotted in red.

4.1.2 Multiplicatively masked Sbox output

We suppose the knowledge of the multiplicative mask. This knowledge allows for the prediction of the value $r_m \times Z_{\hat{K}}[i]$ for any hypothesis \hat{K} . The attack is unsuccessful using 50.000 traces. Figure 29 illustrates the result of the attack targeting byte 0, using 50.000 traces.

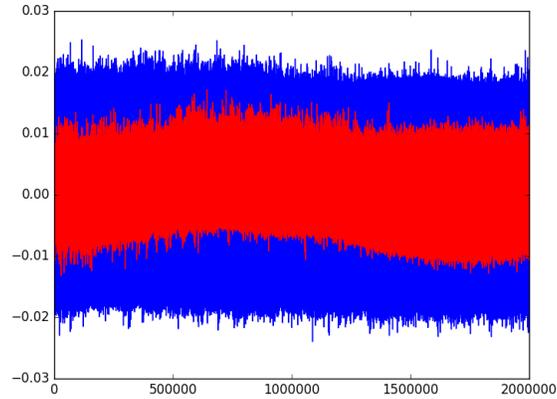


Figure 29: Correlation coefficients obtained when targeting $r_m \times Z_{\hat{K}}[i]$, for every value of \hat{K} . Correct hypothesis is plotted in red.

4.1.3 Raw Sbox output

We suppose no prior knowledge (except the permutation), and predict the value $Z_{\hat{K}}[i]$ for any hypothesis \hat{K} . The attack is unsuccessful using 50.000 traces. Figure 30 illustrates the result of the attack targeting byte 0, using 50.000 traces.

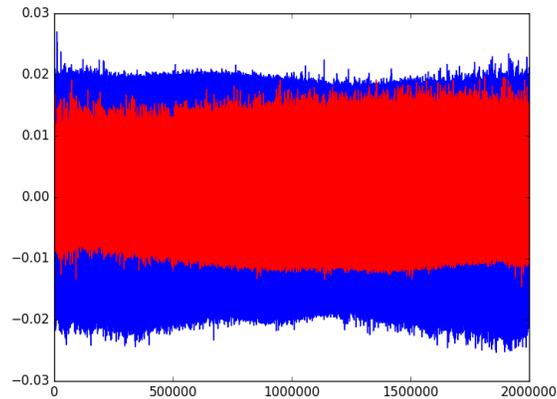


Figure 30: Correlation coefficients obtained when targeting $r_m \times Z_{\hat{k}}[i]$, for every value of \hat{k} . Correct hypothesis is plotted in red.

4.2 Unknown permutation, processed traces

We now study the setting where the attacker does not know the random masks used to compute the permutation.

We preprocess the traces in order to average the leakage over the different byte indices manipulation:

- for each index i in $[0, 15]$, we define a small window w_i of ℓ points around the SNR peak corresponding to the manipulation of $r_m \times S[P[Sh(i)] \oplus K[Sh(i)]] \oplus r_{out}$ in the characterization phase. In our experiments, the size of the window was arbitrarily fixed to $\ell = 11$.
- for each trace in our acquisition campaign, we compute the average window m such that for each time sample j , $m[j] = \frac{1}{16} \sum_{i=0}^{15} w_i[j]$. We then consider m as our reduced averaged trace of size ℓ .

4.2.1 Affinely masked Sbox output

We suppose the knowledge of both the multiplicative mask and the output mask. This knowledge allows for the prediction of the value $r_m \times S[P[i] \oplus \hat{K}] \oplus r_{out}$ for any hypothesis \hat{K} on one byte of the secret key. Depending on the byte index, the CPA is successful using around 20.000 to 30.000 traces.

Figure 31 illustrates the results using 50.000 traces.

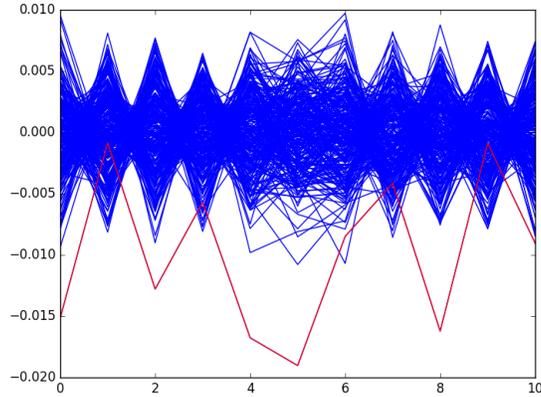


Figure 31: Correlation coefficients obtained when targeting $r_m \times S[P[i] \oplus \hat{K}] \oplus r_{out}$, on averaged traces, for every value of \hat{K} . Correct hypothesis is plotted in red.

4.2.2 Multiplicatively masked Sbox output

With 50.000 traces, the attack does not succeed when targeting $r_m \times S[P[i] \oplus \hat{K}]$. Figure 32 illustrates the results using 50.000 traces.

Figure 32 illustrates the results using 50.000 traces.

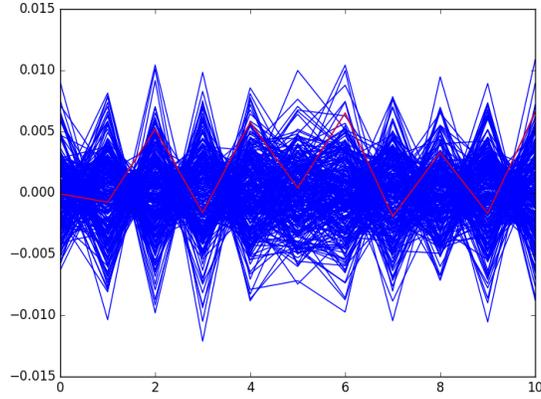


Figure 32: Correlation coefficients obtained when targeting $r_m \times S[P[i] \oplus \hat{K}] \oplus r_{out}$, on averaged traces, for every value of \hat{K} . Correct hypothesis is plotted in red.

4.2.3 Raw Sbox output

With 50.000 traces, the attack does not succeed when targeting $S[P[i] \oplus \hat{K}]$. Figure 33 illustrates the results using 50.000 traces. Figure 33 illustrates the results using 50.000 traces.

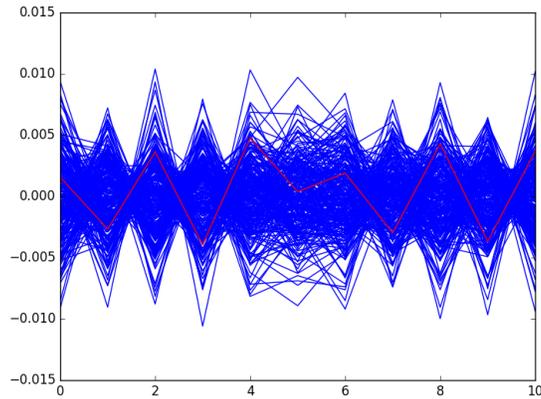


Figure 33: Correlation coefficients obtained when targeting $r_m \times S[P[i] \oplus \hat{K}] \oplus r_{out}$, on averaged traces, for every value of \hat{K} . Correct hypothesis is plotted in red.

4.3 Unknown permutation, non-processed traces

For the sake of completeness, we also perform these experiments on raw unprocessed traces. No success is obtained when targeting $Z_{\hat{K}}[i]$ or $r_m \times Z_{\hat{K}}[i]$ using 50.000 traces. Similarly, by targeting $r_m \times S[P[i] \oplus \hat{K}] \oplus r_{out}$, the attack does not succeed.

4.4 Summary

The results in this section evidence the effectiveness of the implementation of the countermeasures against first-order side-channel attack. Assuming knowledge of the mask values, the shuffling countermeasure alone increases the number of necessary traces by a factor ~ 20 . Furthermore, masking countermeasures seem to completely thwart first order side-channel attacks using 50.000 traces.

	Known permutation			Unknown permutation		
Known	r_m, r_{out}	r_m	None	r_m, r_{out}	r_m	None
# Traces	≈ 1.000	> 50.000	> 50.000	≈ 20.000	> 50.000	> 50.000

Figure 34: Number of traces needed for a successful first-order attack, in the different settings.

5 Second-order correlation analysis

In this section, we analyse the effectiveness of the implemented countermeasures against second-order side-channel analysis based on the linear correlation coefficient (CPA). For this study, we will target the Sbox output value and consider a Hamming weight (HW) univariate leakage model.

In the remainder of this section, we will use:

- the centered product combination function $\mathcal{C}(x, y) = (x - \bar{x})(y - \bar{y})$ to preprocess the traces, where \bar{x} (resp. \bar{y}) stands for the mean value of x (resp. y), computed on all traces in our campaign;
- the corresponding function $f_{HW}(z) = \sum_m (HW(z \oplus m) - \overline{HW(z \oplus m)}) (\overline{HW(m)} - HW(m))$ to compute the predictions, where $\overline{HW(m)}$ (resp. $\overline{HW(z \oplus m)}$) stands for the mean value of $HW(m)$ (resp. $HW(z \oplus m)$) computed on all possible values of m ¹.

As in our first-order correlation analysis, we will split this study in two parts. The first part will consider priviledged knowledge on the mask values. The second part will not take this knowledge into account.

¹Note that all possible values for m are in $[0, 255]$.

5.1 Privileged knowledge

We first study the setting where the attacker knows the random masks used to compute the permutation. In order to characterize the effectiveness of a CPA against this implementation, we use the random masks to recompute the permutation $Sh(i)$ for each byte index i .

For each byte index i , we use the characterization phase to isolate a window $w_{z \oplus m}$ of 11 points of interest corresponding to the manipulation of the Sbox output $r_m \times S[P[Sh(i) \oplus K[Sh(i)]] \oplus r_{out}]$, and a window w_m of 11 points of interest corresponding to the manipulation of the mask value r_{out} . For each couple (i, j) of points in $w_{z \oplus m} \times w_m$, we apply the combination function $\mathcal{C}(i, j)$. We then perform an attack targeting the value $Z_{\hat{k}}[i] = S[P[Sh(i) \oplus K[Sh(i) \oplus \hat{k}]]$, where \hat{k} is an hypothesis on a value of one byte.

5.1.1 Multiplicatively masked Sbox output

We suppose the knowledge of the multiplicative mask r_m . This knowledge allows for the prediction of the value $r_m \times Z_{\hat{k}}[i]$ for any hypothesis \hat{k} . The attack succeeds with around 8.000 traces to 20.000 traces. Figure 35 illustrates the results using 20.000 traces.

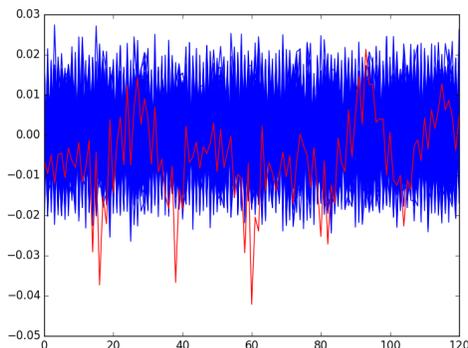


Figure 35: Correlation coefficients obtained when targeting $r_m \times S[P[i] \oplus \hat{k}]$, for every value of \hat{k} . Correct hypothesis plotted in red. X-axis represents all 121 points combinations.

5.1.2 Raw Sbox output

We target the value $Z_{\hat{k}}[i]$ for any hypothesis k . The attack does not succeed using the 100.000 traces. Figure 36 illustrates the results using 100.000 traces.

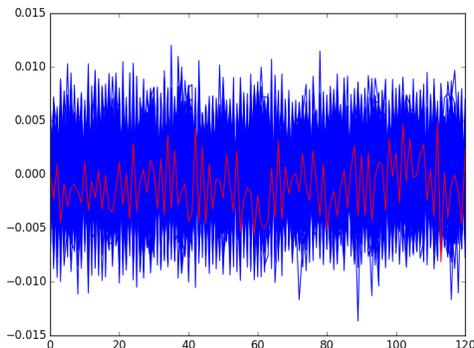


Figure 36: Correlation coefficients obtained when targeting $S[P[i] \oplus \hat{k}]$, for every value of \hat{k} . Correct hypothesis plotted in red. X-axis represents all 121 points combinations.

5.2 Unknown permutation, processed traces

We now study the setting where the attacker does not know the random masks used to compute the permutation. We preprocess the traces in order to average the leakage over the different byte indices manipulation:

- for each index i in $[0, 15]$, we define a small window w_i of ℓ points around the SNR peak corresponding to the manipulation of $r_m \times S[P[Sh(i)] \oplus K[Sh(i)]] \oplus r_{out}$ in the characterization phase. In our experiments, the size of the window was arbitrarily fixed to $\ell = 11$.
- for each trace in our acquisition campaign, we compute the average window m such that for each time sample j , $m[j] = \frac{1}{16} \sum_{i=0}^{15} w_i[j]$. We then consider m as our reduced averaged trace of size ℓ .
- we concatenate to this trace a small window of ℓ points around the SNR peak corresponding to the manipulation of r_{out} .

5.2.1 Multiplicatively masked Sbox output

We suppose the knowledge of the multiplicative mask r_m . This knowledge allows for the prediction of the value $r_m \times Z_{\hat{k}}[i]$ for any hypothesis \hat{k} . The attack is unsuccessful using 100.000 traces.

Figure 37 illustrates the results using 100.000 traces.

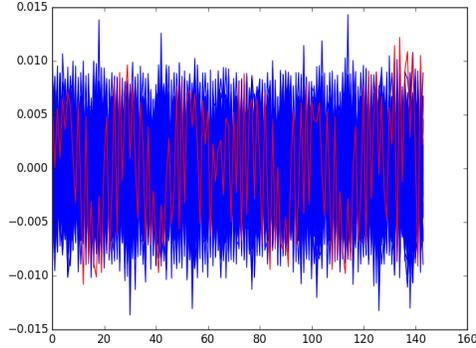


Figure 37: Correlation coefficients obtained when targeting $r_m \times S[P[i] \oplus \hat{k}]$, for every value of \hat{k} . Correct hypothesis plotted in red. X-axis represents all 121 points combinations.

However, we observe that, for certain combinations of points (around abscissa 135), the correct hypothesis is indeed the most likely. Furthermore, the rank convergence of the key, plotted on Figure 38, might indicate that the attack could be successful when using more traces.

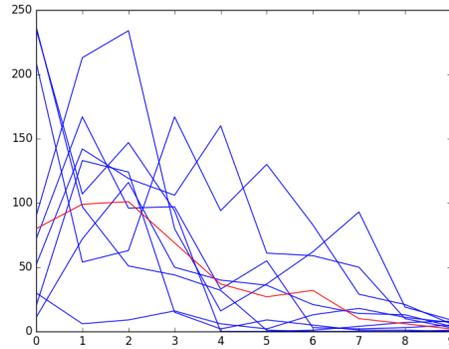


Figure 38: Rank of the ten best-ranked keys, by steps of 10000 measurements, from 10.000 to 100.000. Correct key plotted in red.

5.2.2 Raw Sbox output

We target the value $Z_{\hat{k}}[i]$ for any hypothesis \hat{k} . The attack does not succeed using the 100.000 traces. Figure 39 illustrates the results using 100.000 traces.

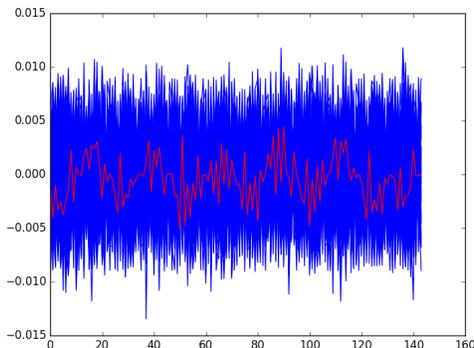


Figure 39: Correlation coefficients obtained when targeting $S[P[i] \oplus \hat{k}]$, for every value of \hat{k} . Correct hypothesis plotted in red. X-axis represents all 121 points combinations.

5.3 Unknown permutation, non-processed traces

For the sake of completeness, we also perform these experiments on raw unprocessed traces. No success is obtained when targeting $Z_{\hat{k}}[i]$ or $r_m \times Z_{\hat{k}}[i]$.

No attack is successful using 100.000 traces.

5.4 Summary

This section evidences the effectiveness of the implemented countermeasures. The shuffling countermeasures increases the number of needed measurements by a factor of at least $\frac{100.000}{8.000} = 12,5$. It is however likely that this countermeasure does not suffice when increasing slightly the number of measurements. The effectiveness of the affine countermeasure is evidenced by the fact that a 2OCPA is achievable targeting the value $r_m \times S[P[i] \oplus \hat{k}]$ but no attack is found targeting the value $S[P[i] \oplus \hat{k}]$ using 100.000 traces.

	Known permutation		Unknown permutation	
Known	r_m	None	r_m	None
# Traces	≈ 8.000	> 100.000	> 100.000	> 100.000

Figure 40: Number of traces needed for a successful second-order attack, in the different settings.

6 Third-order correlation analysis

In this section, we analyse the resistance of the implementation against third order side-channel analysis based on the linear correlation coefficient (CPA).

For this study, we will target the Sbox output value and consider a Hamming weight (HW) univariate leakage model.

In the remainder of this section, we will use:

- the centered product combination function $\mathcal{C}(x, y, z) = (x - \bar{x})(y - \bar{y})(z - \bar{z})$ to preprocess the traces, where \bar{x} (resp. \bar{y}, \bar{z}) stands for the mean value of x (resp. y, z), computed on all traces in our campaign;
- the corresponding function $f_{HW}(z) = \frac{\sum_{m, m'} (HW(m' \times z \oplus m) - \overline{HW(m' \times z \oplus m)})(HW(m) - \overline{HW(m)})(HW(m') - \overline{HW(m')})}{\overline{HW(m)} \cdot \overline{HW(m')}} - \overline{HW(m)}$ to compute the predictions, where $\overline{HW(m)}$ (resp. $\overline{HW(m' \times z \oplus m)}, \overline{HW(m')}$) stands for the mean value of $HW(m)$ (resp. $HW(m' \times z \oplus m), HW(m')$) computed on all possible values of m, m'^2 .

As in our first-order and second-order correlation analyses, we will split this study in two parts. The first part will consider privileged knowledge on the mask values. The second part will not consider this knowledge.

6.1 Privileged knowledge

We first study the setting where the attacker knows the random masks used to compute the permutation.

In order to characterize the effectiveness of a CPA against this implementation, we use the random masks to recompute the permutation $Sh(i)$ for each byte index i .

6.1.1 Raw Sbox output

We target the value $Z_k[i]$ for any hypothesis k . The attack does not succeed using the 100.000 traces. Figure 41 illustrates the results using 100000 traces.

²Note that all possible values for m are in $[0, 255]$, while all possible values for m' are in $[1, 255]$.

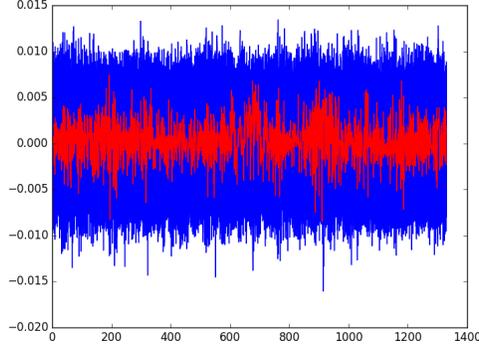


Figure 41: Correlation coefficients obtained when targeting $S[P[i] \oplus \hat{k}]$, for every value of \hat{k} . Correct hypothesis plotted in red. X-axis represents all 1331 points combinations.

6.2 Unknown permutation, processed traces

We now study the setting where the attacker does not know the random masks used to compute the permutation.

We preprocess the traces in order to average the leakage over the different byte indices manipulation:

- for each index i in $[0, 15]$, we define a small window w_i of ℓ points around the SNR peak corresponding to the manipulation of $r_m \times S[P[Sh(i)] \oplus K[Sh(i)]] \oplus r_{out}$ in the characterization phase. In our experiments, the size of the window was arbitrarily fixed to $\ell = 11$.
- for each trace in our acquisition campaign, we compute the average window m such that for each time sample j , $m[j] = \frac{1}{16} \sum_{i=0}^{15} w_i[j]$. We then consider m as our reduced averaged trace of size ℓ .
- we concatenate to this trace a small window of ℓ points around the SNR peak corresponding to the manipulation of r_{out} .

6.2.1 Raw Sbox output

We target the value $Z_{\hat{k}}[i]$ for any hypothesis k . The attack does not succeed using the 100.000 traces. The correct key rank does not seem to indicate that this number of traces is close to allow a success attack. Figure 42 illustrates the results using 100.000 traces.

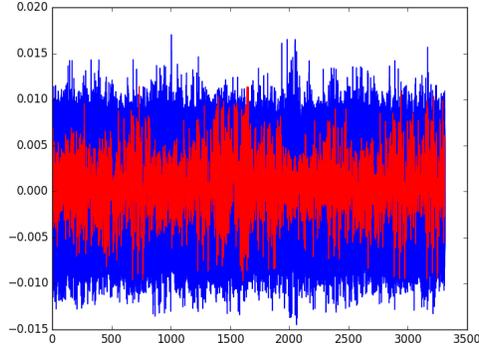


Figure 42: Correlation coefficients obtained when targeting $S[P[i] \oplus \hat{k}]$, for every value of \hat{k} . Correct hypothesis plotted in red. X-axis represents all 1331 points combinations.

6.3 Unknown permutation, non-processed traces

For the sake of completeness, we also perform these experiments on raw un-processed traces. No success is obtained when targeting $Z_{\hat{k}}[i]$. No attack is successful using 100.000 traces. Figure 43 illustrates the results using 100.000 traces.

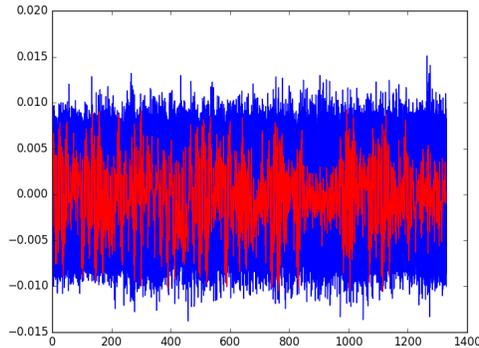


Figure 43: Correlation coefficients obtained when targeting $S[P[i] \oplus \hat{k}]$, for every value of \hat{k} . Correct hypothesis plotted in red. X-axis represents all 1331 points combinations.

6.4 Summary

No third order attack has been achieved using 100.000 traces.