

Entropy Reduction for the Correlation-Enhanced Power Analysis Collision Attack

Andreas Wiemers, Dominik Klein

Bundesamt für Sicherheit in der Informationstechnik (BSI)
{firstname.lastname}@bsi.bund.de

Abstract. Side Channel Attacks are an important attack vector on secure AES implementations. The *Correlation-Enhanced Power Analysis Collision Attack* by Moradi et al. [MME10] is a powerful collision attack that exploits leakage caused by collisions in between S-Box computations of AES. The attack yields observations from which the AES key can be inferred. Due to noise, an insufficient number of collisions, or errors in the measurement setup, the attack does not find the correct AES key uniquely in practice, and it is unclear how to determine the key in such a scenario. Based on a theoretical analysis on how to quantify the remaining entropy, we derive a practical search algorithm. Both our theoretical analysis and practical experiments show that even in a setting with high noise or few available traces we can either successfully recover the full AES key or reduce its entropy significantly.

1 Introduction

Kocher’s [Koc96] groundbreaking paper on side channel attacks has led both science and industry to focus on attacking and hardening their implementations [AARR03]. Due to its popularity and de-facto standard w.r.t. symmetric cryptographic algorithms, AES [BCO04,GMO01,KJJ99,Nat01,QS01] is of particular interest. Despite its theoretical cryptographic strength, a secure AES implementation that does not leak information about processed data remains to be a challenge. A popular counter-measure to minimize leakage about the AES key is *masking*. Different masking schemes exist, but the general idea of masking is that whenever secret data is about to enter critical stages of operation, some reversible operation that makes the data appear to be random is applied. Any cryptanalysis of intermediate data of the processing step is thus worthless. After leaving the critical stage of operation, the operation is reversed and the processed result can be used. Appropriate masking schemes can successfully prevent several attacks.

One particular class of attacks against AES are collision attacks. In collision attacks, one exploits the fact that sometimes leakage of the device can indicate that the same intermediate value has been processed during some critical stage of operation. By using this observation, one can gather information about secret data and cryptanalyze the device. In particular attacks that detect internal

collisions are of interest. This kind of attack method was originally applied to DES [LMV04,SWP03], but later applied to AES [Bog08,SLFP04] as well.

A very powerful kind of collision attack against AES was applied in [MME10], and later improved in [CFG⁺11], and reconsidered in [MS16]. The attack works by feeding data into a device in order to create collisions. A major important observation by [MME10] is that since the S-Box in AES is mathematically the same for every key byte (as opposed to i.e. DES), in most implementations the S-Box is also the same for every key byte. For example in a software implementation there is likely only one S-Box procedure that is called for every key byte, and in the case of a hardware implementation there is a single S-Box circuit that is used to process every key byte. This implies that for two same processed values, the resulting power consumption should be the same as well. Their idea is to create collisions such that this leakage *between* S-Box computations of different key byte positions is exploitable. This makes the attack very powerful — it is shown in in [MME10] that a device with S-Boxes that are masked using the Canright S-Box implementation [Can05,CB08] can be broken with a reasonable amount of available traces. It is important to note here that in general the leakage of the device attacked in [MME10] was minimal, and in particular a typical state-of-the-art template attack [CRR03] was close to impossible to execute. In particular the amount of trace data needed to mount a successful attack was magnitudes lower for the correlation attack than for a template attack.

The attack gives some information about the correct AES key. However, the attack might not find the correct AES key uniquely in practice. There are several reasons for this: Noise, an insufficient number of collisions, errors in the measurement setup, or simply the device itself, i.e. the design and implementation of the cryptographic co-processor for a hardware implementation, or the processor design and execution flow in a software implementation. Moreover, it is not clear a priori how to find a set of key candidates that fit to the observations of the attack. A naive approach, i.e. enumerating all possible key candidates is computationally infeasible due to the large search space.

It is also unclear how to assess the leakage of the device; in particular it leaves open the question how many measurements (traces) are required to successfully mount the attack. Obviously, if the key uniquely identified for a certain amount of traces, this gives an upper bound. However what if less measurements are available?

In this paper we provide an algorithm to recover the AES key in the above scenario. The theoretical motivation of the algorithm is the basis for our analysis on how to quantify the remaining entropy, which can be used to assess the leakage of a device. Both our theoretical analysis and practical experiments show that even in a fuzzy setting with high noise or few available traces, we can either successfully recover the full AES key or reduce its entropy significantly.

This paper is structured as follows. In Section 2, we first briefly recall the attack by Moradi et al. as formulated in [MME10] and then introduce our algorithm in Section 3. For the algorithm we give a thorough theoretical justification in Section 4. Then in Section 5 we analyze the success rate of the algorithm,

i.e. its impact on the entropy of a vulnerable system w.r.t. its leakage, and give upper and lower bounds of the remaining entropy. Our theoretical findings are verified by providing experimental data in Section 6. The algorithm presented here can be seen as a particular form of a key-search algorithm. In Section 7 we show how our algorithm relates to known existing key search algorithms, and to the generalization of [MME10] described in [MS16]. Finally, we conclude our presentation in Section 8.

2 Correlation-Enhanced Power Analysis Collision Attack

Let K_1, \dots, K_{16} be the correct key which is used in the first round of an AES encryption. We denote by small letters k_1, \dots, k_{16} candidates for the key. We briefly recall the Correlation-Enhanced Power Analysis Collision Attack as described in [MME10].

During the measurement phase we record N power consumption traces of the first round of an AES-128 encryption¹. These traces consists of 16 single S-Box computations. The measurement of each single S-Box computation is given as a vector of T numbers. We denote by $b_{i,w,t}$ this power consumption trace of a single S-Box computation i of a known plaintext $p_{w,i}$, $1 \leq i \leq 16$, $1 \leq w \leq N$, $1 \leq t \leq T$. As a first step we compute the average value $M_{i,\beta,t}$ over all w with $\beta = p_{w,i}$. Secondly, for any i, j and t we derive the empirical correlation coefficient $C_{i,j,\alpha,t}$ between $M_{i,\beta,t}$ and $M_{i,\beta \oplus \alpha,t}$ for any byte value α , where we treat β as a random variable uniformly distributed on all 256 byte values. At last, we set $c_{i,j}(\alpha)$ for the maximum of all $C_{i,j,\alpha,t}$, where t runs over all time points.

The idea of this approach is as follows: If the measurement $b_{i,w,t}$ is slightly dependent on the input byte $p_{w,i} \oplus K_i$ of the S-Box computation i , the average $M_{i,\beta,t}$ depends on $\beta \oplus K_i$ even more significantly. Now the input bytes $\beta \oplus K_i$ of S-Box i and $\beta \oplus K_j \oplus \alpha$ of S-Box j are the same for the choice $\alpha = K_i \oplus K_j$. Therefore, we can hope that the correlation $C_{i,j,\alpha,t}$ has — at least for some t — a significantly higher value for the correct choice $K_i \oplus K_j$ of α .

3 Recovering the AES Key

In this section we formulate our algorithm for computing candidates for the full AES key. We assume that we have given $120 \cdot 256$ values in the form

$$c_{i,j}(\alpha)$$

for $1 \leq i < j \leq 16$, where α runs over all byte values. If for each i, j the value $c_{i,j}(K_i \oplus K_j)$ is always the highest among all $c_{i,j}(\alpha)$, then it is easy to derive the full key. Here we are interested in the situation, where for each i, j , the value

¹ The attack can be extended to other key sizes in a straight-forward manner by targeting a second AES round [MME10].

$c_{i,j}(K_i \oplus K_j)$ has only a tendency of being large compared to other $c_{i,j}(\alpha)$ with $\alpha \neq K_i \oplus K_j$. The idea of our approach is to consider the ad-hoc evaluation function

$$B = \sum_{i < j} c_{i,j}(k_i \oplus k_j)$$

for any key candidate (k_1, \dots, k_{16}) and choose the key candidate with the highest value in B . Since this is not feasible in a straightforward manner, we instead try to compute B via partial sums. To this end, we fix an integer W , resp. integers g_2, \dots, g_{16} .

Algorithm 1 Recovering the AES Key

1: Set $k_1 = 0$, $S_1 = \{k_1\}$ and $B_1 = 0$.

2: **for** $s = 1, \dots, 15$ **do**

3: **for** each key candidate k_1, \dots, k_s in S_s **do**

4: **for** each value of the next key bytes k_{s+1} **do**

5: compute the evaluation function

$$B_{s+1} = B_s((k_1, \dots, k_s)) + \sum_{1 \leq i \leq s} c_{i,s+1}(k_i \oplus k_{s+1})$$

6: **end for**

7: **end for**

8: select subset of candidates k_1, \dots, k_s, k_{s+1} w.r.t. some criteria and store in S_{s+1} :

9: **Variante I:** Select W candidates k_1, \dots, k_s, k_{s+1} with largest B_{s+1}

10: **Variante II:** Select all k_1, \dots, k_s, k_{s+1} with $B_{s+1} \geq g_{s+1}$

11: **end for**

12: **return**

Remarks and Observations We note some properties of Algorithm 1: Since B_s and B only depend on \oplus -sums of key bytes, we can choose one key byte as a fixed value. Here, we set $k_1 = 0$. The success probability of both variants of our algorithm for finding the correct key depends on the input parameters W , resp. g_2, \dots, g_{16} . Obviously, if we choose $g_s = B_s((K_i \oplus K_j))$, Variante II of Algorithm 1 is guaranteed to output the correct key. However, in this case S_s might become too large to store in practice. The parameter W can be treated as a measure of the workload (i.e. the number of computational steps) of Variante I of Algorithm 1. Both variants of the algorithm assume a fixed order of key byte positions. The result of the algorithm depends on that assumed order of the key bytes. One can repeat the algorithm with different orders, though. As s grows, the order becomes less important. We investigate the effect of the order on the success of the algorithm in Section 6. The choice of the order could take into account the actual distribution of the values $c_{i,j}(\alpha)$. Those i, j with significantly high values in $c_{i,j}(\alpha)$ could be considered first. In a practical setting, visual inspection of $C_{i,j,\alpha,t}$ could give a hint, cf. for example Figures 3a and 3b. In

general however, we are more interested in the situation where $c_{i,j}(K_i \oplus K_j)$ is not automatically the highest value among the $c_{i,j}(\alpha)$, but is only larger on average over all i, j .

4 Theoretical Justification

In this section, we give a justification of the evaluation function B . To this end, we treat $c_{i,j}(\alpha)$ for any i, j, α as a realization of a normally distributed random variable. Note that technically, this assumption is not correct. First, correlation values are bound in $[-1, 1]$ and thus cannot be normally distributed. Second, we do not consider correlation values directly, but instead $c_{i,j}(\alpha)$ is a maximum of $C_{i,j,\alpha,t}$ over several time points t . We can easily force $c_{i,j}(\alpha)$ to be normally distributed by fixing one single time-point, and apply Fisher's z transformation on the correlation values. In practice however, we can see that the correlations $C_{i,j,\alpha,t}$ for wrong α (cf. Figures 3a, 3b) lie within $[-0.25, 0.25]$. Fisher's z transformation is almost identical for values in that interval, so that we do not consider this transformation in our case. In addition we observe (cf. Figures 5a, 5b) that the distribution of $c_{i,j}(\alpha)$ is not perfectly symmetrical, but minimally skewed. This can theoretically be justified by the fact that $c_{i,j}(\alpha)$ is a maximum of $C_{i,j,\alpha,t}$ over several time points t . All in all, the deviation of the distribution $c_{i,j}(\alpha)$ from a normally distributed random variable is very small and is thus neglected in the following.

We assume the easiest scenario: For any i, j, α with $\alpha \neq K_i \oplus K_j$ the means and the standard deviations are equal and are denoted by a , resp. σ . Furthermore, for the correct $\alpha = K_i \oplus K_j$ the means are equal and are denoted by b and in addition, the standard deviations are equal to σ . For any key candidate $k = (k_1, \dots, k_{16})$ we can check whether for all key candidates \tilde{k}

$$\begin{aligned} c_{i,j}(\tilde{k}_i \oplus \tilde{k}_j) &\approx a, \text{ if } \tilde{k}_i \oplus \tilde{k}_j = k_i \oplus k_j \\ c_{i,j}(\tilde{k}_i \oplus \tilde{k}_j) &\approx b, \text{ if } \tilde{k}_i \oplus \tilde{k}_j \neq k_i \oplus k_j \end{aligned}$$

As a likelihood measure for any key candidate we seek a function in the single probability density functions as

$$\begin{aligned} \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(c_{i,j}(\tilde{k}_i \oplus \tilde{k}_j) - a)^2}{2\sigma^2}\right), \text{ resp.} \\ \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(c_{i,j}(\tilde{k}_i \oplus \tilde{k}_j) - b)^2}{2\sigma^2}\right) \end{aligned}$$

The cumulative probability density function of two *independent* random variables is just the product of the single probability density functions. Therefore, we are led to use as an evaluation function the product over all single probability

density functions. Taking logarithms we get

$$\sum_{\tilde{k}} \left[\sum_{\substack{i < j, \\ \tilde{k}_i \oplus \tilde{k}_j = k_i \oplus k_j}} (c_{i,j}(\tilde{k}_i \oplus \tilde{k}_j) - a)^2 + \sum_{\substack{i < j, \\ \tilde{k}_i \oplus \tilde{k}_j \neq k_i \oplus k_j}} (c_{i,j}(\tilde{k}_i \oplus \tilde{k}_j) - b)^2 \right]$$

An equivalent evaluation function is therefore

$$\sum_{i < j} c_{i,j}(k_i \oplus k_j).$$

5 Success Rate of the Algorithm (Variant II)

In this section we want to give theoretical estimates of the success rate and workload for the second variant of the algorithm. The purpose of this section is to find relations between those theoretical estimates and basic properties of the distributions of $c_{i,j}(\alpha)$. To make the derivation as simple as possible, we restrict ourselves to the scenario in the last section, i.e. $c_{i,j}(\alpha)$ for any i, j, α is treated as a realization of a normally distributed random variable with mean a , resp. b , and standard deviation σ . Furthermore, for any key candidate the evaluation function

$$B_s = \sum_{i < j \leq s} c_{i,j}(k_i \oplus k_j)$$

is considered as a sum of *independent* random variables. Therefore, B_s is a normally distributed random variable. For a randomly chosen key candidate we have the expectation value

$$\mathbb{E}(c_{i,j}(k_i \oplus k_j)) = \frac{255}{256}a + \frac{1}{256}b \approx a$$

since b is assumed to be only slightly larger than a . Therefore, the mean and standard deviation of B_s are $\binom{s}{2}a$, resp. $\sqrt{\binom{s}{2}}\sigma$. For the correct key, $B_s((K_1, \dots, K_s))$ is normally distributed with mean $\binom{s}{2}b$.

For having B_s near to its mean value, we want to avoid small values in $\binom{s}{2}$. In practice, we set in Variant II of the algorithm $g_s = -\infty$ for $s \leq 4$. For the ease of presentation we want to assume

$$B_s((K_1, \dots, K_s)) \geq \binom{s}{2}b \text{ for } s \geq 5$$

Therefore, we set here $g_s = \binom{s}{2}b$.

5.1 An upper Bound of the Remaining Entropy

Variant II of the algorithm only finds key candidates for which $B_s \geq g_s$ for *all* s , $5 \leq s \leq 16$. In every step, the set S_s is a subset of all key candidates for

which the condition $B_s \geq g_s$ is fulfilled. The size A_s of this larger set can be approximated by

$$\#S_s \leq A_s = 2^{(s-1)8} \mathbf{P} \left(B_s \geq \binom{s}{2} b \right)$$

and $\log_2(A_{16})$ is an upper bound for the remaining entropy.² $A_{16} \approx 1$ means that the correct key has been found more or less uniquely, and $\max_s A_s$ is an upper bound for the workload of variant II of Algorithm 1. The inequality of integrals

$$\int_x^\infty e^{-t^2/2} dt \leq \int_x^\infty \frac{t}{x} e^{-t^2/2} dt = \frac{1}{x} e^{-x^2/2}$$

can be used to give an upper bound for the standardized normal distribution $\mathcal{N}_{0,1}$:

$$\mathcal{N}_{0,1}(x, \infty) \leq \frac{1}{x\sqrt{2\pi}} e^{-x^2/2}$$

We set

$$\tau = \frac{b-a}{\sigma}$$

and derive

$$\begin{aligned} A_s &= 2^{(s-1)8} \mathbf{P} \left(\frac{B_s - \binom{s}{2} a}{\sigma \sqrt{\binom{s}{2}}} \geq \tau \sqrt{\binom{s}{2}} \right) \\ &\leq 2^{(s-1)8} \frac{1}{\tau \sqrt{2\pi \binom{s}{2}}} e^{-\frac{1}{2} \binom{s}{2} \tau^2} = \frac{1}{\tau \sqrt{2\pi \binom{s}{2}}} 2^{(s-1)8 - \frac{1}{2\ln(2)} \binom{s}{2} \tau^2} \end{aligned}$$

This approximation of $\log_2(A_s)$ has roughly the form of a parabola in s . The condition $A_{16} \approx 1$ corresponds to the equation

$$\tau = \frac{b-a}{\sigma} \approx \sqrt{2\ln(2)} \approx 1.2$$

Some results are provided in Table 1. This can be interpreted in that we can expect that for $\frac{b-a}{\sigma} \geq 1$, Variant II of Algorithm 1 is successful with workload $\leq 2^{36}$ and remaining entropy ≤ 29 .

5.2 A lower Bound of the Remaining Entropy

We want to analyze variant II of Algorithm 1 step by step. We expect that the size of $\#S_{s+1}$ can be approximated by a *conditional* probability of the form

$$\begin{aligned} \#S_{s+1} &\approx \#S_s 2^8 \mathbf{P} \left(B_{s+1} \geq \binom{s+1}{2} b \mid B_s \geq \binom{s}{2} b, \right. \\ &\quad \left. B_{s-1} \geq \binom{s-1}{2} b, \dots, B_5 \geq \binom{5}{2} b \right) \end{aligned}$$

² Since one key byte cannot be determined by the algorithm, the accurate remaining entropy is more properly $\log_2(A_{16}) + 8$.

Table 1. Upper Bounds for the Remaining Entropy

$\frac{b-a}{\sigma}$	$\log_2(A_{16})$	$\log_2(A_4)$	$\max_{s \geq 5} \log_2(A_s)$
1.4	0	24	15
1.2	0	24	23
1.1	10	24	29
1.0	29	24	36
0.9	45	24	46

Table 2. Bounds for the Remaining Entropy

$\tau = \frac{b-a}{\sigma}$	$\log_2(A_{16})$	$\log_2(A_4)$	$\max_{s \geq 5} \log_2(A_s)$	Lower bound of $\log_2(\#S_{16})$	Lower bound of $\log_2(\#S_s)$
1.4	0	24	15	0	14
1.2	0	24	23	0	21
1.1	10	24	29	2	26
1.0	29	24	36	21	32
0.9	45	24	46	37	41

Note that

$$X_s = \frac{B_s - \binom{s}{2}a}{\sigma}$$

is the sum of $\binom{s}{2}$ $\mathcal{N}_{0,1}$ -distributed independent random variables. Therefore, X_s represents a Gaussian random walk. We write the conditional probability in the form

$$\begin{aligned} & \mathbf{P} \left(B_{s+1} \geq \binom{s+1}{2}b \mid B_s \geq \binom{s}{2}b, B_{s-1} \geq \binom{s-1}{2}b, \dots, B_5 \geq \binom{5}{2}b \right) \\ &= \mathbf{P} \left(X_{s+1} \geq \binom{s+1}{2}\tau \mid X_s \geq \binom{s}{2}\tau, X_{s-1} \geq \binom{s-1}{2}\tau, \dots, X_5 \geq \binom{5}{2}\tau \right) \end{aligned}$$

The probability

$$\mathbf{P} \left(X_{s+1} \geq \binom{s+1}{2}\tau, X_s \geq \binom{s}{2}\tau, X_{s-1} \geq \binom{s-1}{2}\tau, \dots, X_5 \geq \binom{5}{2}\tau \right)$$

can be interpreted as the probability of a Gaussian random walk with at least linear growth at special steps. We get a lower bound of the conditional probability if we omit the conditions on all $s' < s$.

$$\#S_{s+1} \geq V_s 2^8 \mathbf{P} \left(B_s + \sum_{1 \leq i \leq s} c_{i,s+1}(k_i \oplus k_{s+1}) \geq \binom{s+1}{2}b \mid B_s \geq \binom{s}{2}b \right)$$

for $s \geq 5$ and $\#S_5 = A_5$. Since $\binom{5}{2} = 10$, we use for $\#S_5 = A_5$ the formula

$$\#S_5 = 2^{32} \mathcal{N}_{0,1} \left(\sqrt{10} \frac{b-a}{\sigma}, \infty \right)$$

The probability $\mathbf{P} \left(B_s + \sum_{1 \leq i \leq s} c_{i,s+1}(k_i \oplus k_{s+1}) \geq \binom{s+1}{2} b \mid B_s \geq \binom{s}{2} b \right)$ only depends on s and $\tau = \frac{b-a}{\sigma}$. This probability and therefore all lower bounds of $\#S_{s+1}$ can be calculated numerically. Table 2 extends Table 1 above. We can expect that the remaining entropy and the workload of variant II of Algorithm 1 are within the limits of this table.

5.3 Probability of the Event $B_s(K_1, \dots, K_s) \geq \binom{s}{2} b$

We consider the event

$$B_s((K_1, \dots, K_s)) \geq \binom{s}{2} b \text{ for all } s = 16, 15, \dots, 5$$

Note, that the probability of this event does not depend on b and σ , it is just a real number. On first sight, one could believe that the probability of this

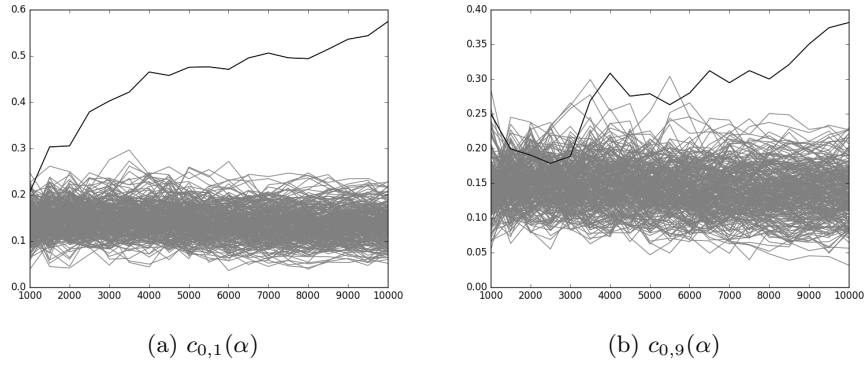


Fig. 1. Correlation of $c_{i,j}(\alpha)$ vs # of traces. The correct value of α is shown in black.

event is 2^{-12} . But since $B_{s-1}((K_1, \dots, K_{s-1}))$ is a subsum of $B_s((K_1, \dots, K_s))$, the probability of $B_s(K_1, \dots, K_s) \geq \binom{s}{2} b$ is larger than $\frac{1}{2}$ if we already know that $B_{s-1}((K_1, \dots, K_{s-1})) \geq \binom{s-1}{2} b$. We compute an approximation of this probability by a simulation of normally distributed random variables. We get

$$\mathbf{P} \left(B_s((K_1, \dots, K_s)) \geq \binom{s}{2} b \text{ for all } s = 16, 15, \dots, 5 \right) \approx 0.15.$$

To this end, the assumption $B_s(K_1, \dots, K_s) \geq \binom{s}{2} b$ for all s is not too restrictive.

6 Experiments

Experimental Setup Our setup consists of an AES-128 based software implementation running on an Atmel ATMEGA328P-PU. The S-Boxes are realized as lookup tables and stored in the program memory of the ATMEGA. The S-Boxes are masked using the method presented in [AG01]. This masking is known to be not leakage-free, and as shown in [CFG⁺11,MME10], also Canright S-Boxes [Can05] are susceptible to correlation attacks. Hence we anticipate that our results are representative. Moreover, the masking used [AG01] is straightforward to implement. The ATMEGA was setup on a custom prototype board, and powered by a lab-grade power supply at 3.3 Volt. We used a LeCroy HDO6104 to record the power consumption of the ATMEGA with a resistor against ground. We recorded N traces of AES-128 encryption. The plaintext was chosen at random, and we attacked the `masked_subbytes` procedure of the first AES round.




Fig. 2. Number of Traces vs. Ranking Positions of $c_{i,j}(\alpha)$ for correct α .

Practical Results We recorded 10000 traces with randomly generated plaintext values. Our implementation follows closely [MME10] and we compute the correlation w.r.t. each possible value of one $C_{i,j,\alpha,t}$. Figure 3a shows the resulting correlation of $C_{0,1,\alpha,t}$ for one AES round, i.e. in our setup $t = 0, \dots, 25809$. Correlation peaks at the end of the S-Box for the correct value are clearly visible. However a correlation peak is not apparent for every pair of key byte positions i, j ; for example considering the correlation values $C_{9,11,\alpha,t}$ as depicted in Figure 3b, no clear peak is observable for the correct value α . Nevertheless when ranking all possible values $c_{9,11}(\alpha)$, the correct value is still at position 18. Fewer traces result in less collisions and more noise, and the rankings become more fuzzy. To give two examples, the rankings for the correct value of α for 2500

and 10000 traces are as shown in Figure 4a and 4b. If more traces are available, the correlation values $c_{i,j}(\alpha)$ for the correct value α become very distinct from those for incorrect values of α , as illustrated in Figure 1a and Figure 1b. For $N = 10000$, the correct value of α often shows on rank 1, but there are some outliers. For $N = 2500$, there are few rankings with position 1.

Table 3. Full Key Recovery for $W = 1500$ using Algorithm 1 (Variant I).

	1500	...	3500	4000	4500	5000	5500	6000	6500
1 \rightarrow 16	-		✓	✓	✓	✓	✓	✓	✓
16 \rightarrow 1	-		-	-	-	-	-	-	-
random	-		1/5	3/5	3/5	3/5	3/5	3/5	3/5
	7000	7500	8000	8500	9000	9500	10000		
1 \rightarrow 16	-	✓	✓	✓	✓	✓	✓		
16 \rightarrow 1	-	-	-	-	-	-	-		
random	3/5	4/5	3/5	4/5	5/5	5/5	5/5		

For Variant II of Algorithm 1 one needs to choose appropriate input values g_2, \dots, g_{16} . This requires prior knowledge about the quality of the rankings $c_{i,j}(\alpha)$. If such knowledge is not available, appropriate values could also be estimated by manual analysis of the rankings $c_{i,j}(\alpha)$ and/or by visual inspection of $C_{i,j,\alpha,t}$. For example the comparison of Figure 3a and Figure 3b indicates that for $c_{0,1}(\alpha)$ the ranking for the correct value of α is very likely at one of the top positions, whereas for $c_{9,11}(\alpha)$ this is likely not the case. On the other hand, Variant I of Algorithm 1 requires no prior knowledge at all. Also, if Variant I succeeds, we can always choose the bound g_{s+1} in step s for Variant II in such a way that the choice of Variant I is simulated. This is why we have here chosen to implement Variant I of the algorithm.

As mentioned in Section 3, the success of the algorithm depends on the order in which the next key byte position is chosen, the parameter W of candidates that are kept in each iteration, and of course the number of available traces N . As mentioned in Section 2, in particular the order of choosing the next key byte position is important, as the next example illustrates:

Example 1. For simplicity, suppose we have only a key consisting of three bytes, and suppose one key byte can take only value 0 or 1. Let $c_{i,j}(\alpha)$ be as follows:

$$\begin{array}{lll}
 c_{0,1}(0) = 0.4 & c_{0,1}(1) = 0.1 & c_{1,2}(0) = 0.4 \\
 c_{1,2}(1) = 0.1 & c_{0,2}(0) = 0.1 & c_{0,2}(1) = 0.8
 \end{array}$$

Suppose that we set $W = 2$. Assume the order $0 < 1 < 2$. We start with $S_1 = \{0, 1\}$. Since $c_{0,1}(0) > c_{0,1}(1)$ and $0 \oplus 0 = 0$ as well as $1 \oplus 1 = 0$, we yield the set $S_2 = \{00, 11\}$. Algorithm 1 terminates with $S_3 = \{001, 110\}$. On the other hand, it is not difficult to verify that for the key byte order $2 < 1 < 0$,




Fig. 3. Correlation $C_{i,j,\alpha,t}$ for each timepoint t within one trace. The correct value of α is denoted in black.




Fig. 4. Ranking positions for correct α .




Fig. 5. Distribution of $c_{i,j}(\alpha)$ for incorrect values of α .

Algorithm 1 terminates with the set $S_3 = \{100, 011\}$. Note also that Algorithm 1 is nondeterministic in general: If we set $W = 1$ in the second step during the run with order $0 < 1 < 2$, we have $B((00)) = B((11)) = 0.4$, and it is open which partial key to keep.

We executed Algorithm 1 for $N = 1500$ up to $N = 10000$ in steps of 500 traces, and for both $W = 1500$ and $W = 10000$. As for the dependency on the order, we considered the natural order of starting at key byte position 1 and moving upward to position 16 (denoted by $1 \rightarrow 16$ in the following), the reverse order of starting at position 16 and moving down to 1 (denoted by $16 \rightarrow 1$), and on five randomly chosen orders for each value of N . Table 3 and Table 4³ show results

Table 4. Full Key Recovery for $W = 10000$ using Algorithm 1 (Variant I).

	1500	...	3500	4000	4500	5000	5500	6000	6500
$1 \rightarrow 16$	–		✓	✓	✓	✓	✓	✓	✓
$16 \rightarrow 1$	–		–	–	–	✓	✓	✓	✓
random	–		1/5	3/5	3/5	3/5	3/5	3/5	3/5
		7000	7500	8000	8500	9000	9500	10000	
$1 \rightarrow 16$	–	✓	✓	✓	✓	✓	✓	✓	✓
$16 \rightarrow 1$	✓	✓	–	✓	✓	✓	✓	✓	✓
random	3/5	4/5	3/5	4/5	5/5	5/5	5/5	3/5	

for the full recovery of the key for $N = 1500 \dots 10000$ and the various orders. As one can see, the full key is in the computed set with good probability if at least 4000 traces are available. The probability can be increased, if a larger parameter $W = 10000$ is chosen. For less than 4000 traces, Table 5 shows the maximum number of correctly recovered key bytes in the computed set. For example, when choosing the order $1 \rightarrow 16$, then in the case of 2500 traces, the computed set contains a key where 12 key bytes are correctly identified. In other words, the entropy is significantly reduced, and it is not difficult to devise an algorithm that exploits the fact that one can assume that a certain amount of key byte position are correctly identified.

In order to further interpret these results, we investigated the distribution $c_{i,j}$ in α for all (i, j) . For each (i, j) we computed the expected value as well the standard deviation, which were very similar. Figure 5a and Figure 5b show histograms of all $255 \cdot 120$ $c_{i,j}$ with incorrect value α for $N = 10000$ and $N = 2500$, respectively. Expected value and standard deviation are 0.139 ± 0.0374 for $N = 10000$ and 0.146 ± 0.0363 for $N = 2500$. As derived above, we expect a theoretical bound $\frac{b-a}{\sigma} \approx \sqrt{2 \ln(2)} \approx 1.2$. This is the smallest value b , for which we can

³ Note that the miss for $N = 8000$ and $16 \rightarrow 1$ in Table 4 is precisely due to the nondeterministic behavior of Algorithm 1, as illustrated in Example 1.

expect that the evaluation function B succeeds. For our experimental data we yield $\frac{b-a}{\sigma} = \frac{0.336-0.139}{0.0374} \approx 5.3$ for $N = 10000$ and $\frac{b-a}{\sigma} = \frac{0.195-0.146}{0.0363} \approx 1.3$ for $N = 2500$. Apparently, the parameters for $N = 2500$ are very close to the theoretical threshold. This fits with our theoretical observations in previous sections and the experimental data. Aside from the above results, we also experimentally verified our findings using a hardened security controller with a hardware based AES implementation, and the results were comparable. However, we refrain from providing more specific data here.

7 Related Work

Suppose we are given a key K_1, \dots, K_n with n independent subkeys, and a side channel attack yields for each K_i ranks for possible key values. Using these ranks, one can define a probability measurement function for each subkey. The *key enumeration problem* [VCGRS13] is then to enumerate complete keys from the most probable to the least probable one, in order to find the key of a device as quickly as possible. A direct solution is to enumerate all possible keys, compute their probability (multiplying the probabilities of the subkeys) and then sort all keys according to their probability. As this is infeasible in practice, more efficient algorithms have been proposed [VCGRS13].

In our scenario, we do not have ranks (resp. probability measurement functions) for all sixteen K_i . Rather, we obtain 120 ranks for each $K_i \oplus K_j$ for $1 \leq i, j \leq 16$. Our task is to reconstruct the key from these values. Algorithms that solve the key enumeration problem are thus not directly applicable. One naive way to still apply existing key search algorithms in our scenario would be to (1) select a number of $K_i \oplus K_j$ with their ranks which have a low remaining entropy, i.e. select $K_1 \oplus K_2, K_2 \oplus K_3, \dots, K_{15} \oplus K_{16}$. (2) recover $K_1 \oplus K_2, K_2 \oplus K_3, \dots, K_{15} \oplus K_{16}$ using an efficient key enumeration algorithm, and for each candidate (3) search among the remaining 256 possibilities for the correct key.

Note that it is not clear how to effectively select $K_i \oplus K_j$. One could for example simply take $K_1 \oplus K_2, K_2 \oplus K_3, \dots, K_{15} \oplus K_{16}$. How to assess the leakage then? The problem of *key ranking* [GGP⁺15, MOOS15, MMOS16] is concerned with quantifying the time complexity required to brute force a key given the leakages, especially if the number of keys to enumerate is beyond practical computational limits. Applying e.g. the algorithm of [MMOS16] for our test set of traces with $N = 4000$ when we interpret $K_1 \oplus K_2, K_2 \oplus K_3, \dots, K_{15} \oplus K_{16}$ as our key, the estimated rank is approximately 2^{50} . However with our algorithm we are able to recover the actual AES key within less than two minutes (cf. Table 3). Directly recovering the key (and not $K_i \oplus K_j$) by making use of all $K_i \oplus K_j$ was the motivation for our algorithm. We anticipate however, that by either employing voting techniques [Bog08] or by a more sophisticated choice of $K_i \oplus K_j$ and repeating steps (1)-(3) one could obtain similar results as with our algorithm, but we did not further pursue this direction.

Table 5. Partial Key Recovery for $W = 1500$ using Algorithm 1 (Variant I).

	1500	2000	2500	3000	3500
1 \rightarrow 16	3	5	12	13	14
16 \rightarrow 1	2	2	4	1	5
random #1	1	5	3	7	9
random #2	1	5	2	5	14
random #3	2	3	5	12	5
random #4	2	3	5	7	8
random #5	2	2	1	8	12

In [MS16] the attack of [MME10] is generalized into *Moments-Correlating Collision DPA*. There, moments are correlated with samples rather than moments with moments, as done in the original attack. Using that, they derive a measure N_{sr} (equation 1 in [MS16]) that provides a metric to quantify the number of measurements needed to perform a key recovery with a given success rate. Note however that this measure quantifies the leakage w.r.t. $K_i \oplus K_j$. As seen in the experiments, the leakage can be quite unsteady, i.e. the correct value of α is in the lower ranks for some i, j , but for other combinations the rank is much higher. Here we provide an algorithm that still works in this scenario and reveal the correct key, and our measure τ thus gives a indicates on whether the key can be recovered or not.

8 Conclusion and Future Work

We have shown how to reduce the remaining key entropy of the attack introduced by [MME10] by providing a practical, easy-to-implement algorithm. Our theoretical analysis shows that this algorithm exploits the leakage in a natural way. Moreover, we provide a way to assess the leakage of a device w.r.t. the attack, which could be used e.g. in a Common Criteria security evaluation. Our practical evaluation supports the theoretical analysis. In particular we show that using our algorithm, a full recovery of the AES key is possible with only few available traces. That is, the key can be recovered in a setting where no visual clues w.r.t. the correct ranking are available and the attack as described by [MME10] would not have been applicable. A practical comparison with the metric derived in [MS16] remains future work.

Acknowledgements We would like to thank Sven Freud for creating the circuit board for power analysis, and Tobias Senger for his help with implementing the masking scheme. We also thank the anonymous reviewers for their helpful comments.

References

- [AARR03] D. Agrawal, B. Archambeault, J. R. Rao, and P. Rohatgi. The EM Side—Channel(s). In *Proc. 4th CHES*, pages 29–45, 2003.
- [AG01] M.-L. Akkar and C. Giraud. An Implementation of DES and AES, Secure against Some Attacks. In *Proc. 3rd CHES*, pages 309–318, 2001.
- [BCO04] E. Brier, C. Clavier, and F. Olivier. Correlation Power Analysis with a Leakage Model. In *Proc. 6th CHES*, 2004.
- [Bog08] A. Bogdanov. Multiple-Differential Side-Channel Collision Attacks on AES. In *Proc. 10th CHES*, pages 30–44, 2008.
- [Can05] D. Canright. A Very Compact S-Box for AES. In *Proc. 7th CHES.*, pages 441–455, 2005.
- [CB08] D. Canright and L. Batina. A Very Compact “Perfectly Masked” S-Box for AES. In *Proc. 6th ACNS*, pages 446–459, 2008.
- [CFG⁺11] C. Clavier, B. Feix, G. Gagnerot, M. Roussellet, and V. Verneuil. Improved Collision-Correlation Power Analysis on First Order Protected AES. In *Proc. 13th CHES*, pages 49–62, 2011.
- [CRR03] S. Chari, J. R. Rao, and P. Rohatgi. Template Attacks. In *Proc. 4th CHES*, pages 13–28, 2003.
- [GGP⁺15] C. Glowacz, V. Grosso, R. Poussier, J. Schüth, and F. Standaert. Simpler and more efficient rank estimation for side-channel security assessment. In *Proc. FSE 2015*, pages 117–129, 2015.
- [GMO01] K. Gandolfi, C. Mourtel, and F. Olivier. Electromagnetic Analysis: Concrete Results. In *Proc. 3rd CHES*, pages 251–261, 2001.
- [KJJ99] P. Kocher, J. Jaffe, and B. Jun. Differential Power Analysis. In *Proc. 19th CRYPTO*, pages 388–397, 1999.
- [Koc96] P. C. Kocher. *Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems*, pages 104–113. 1996.
- [LMV04] H. Ledig, F. Muller, and F. Valette. Enhancing Collision Attacks. In *Proc. 6th CHES*, pages 176–190, 2004.
- [MME10] A. Moradi, O. Mischke, and T. Eisenbarth. Correlation-Enhanced Power Analysis Collision Attack. In *Proc. 12th CHES*, pages 125–139, 2010.
- [MMOS16] D.P. Martin, L. Mather, E. Oswald, and M. Stam. Characterisation and estimation of the key rank distribution in the context of side channel evaluations. In *Proc. 22nd ASIACRYPT*, pages 548–572, 2016.
- [MOOS15] D. P. Martin, J. F. O’Connell, E. Oswald, and M. Stam. Counting keys in parallel after a side channel attack. In *Proc. 21st ASIACRYPT*, pages 313–337, 2015.
- [MS16] A. Moradi and F. Standaert. Moments-correlating dpa. In *Proc. 2016 TIS Sec. Workshop*, pages 5–15, 2016.
- [Nat01] National Institute of Standards and Technology. FIPS PUB 197. Advanced Encryption Standard. Technical report, 2001.
- [QS01] J.-J. Quisquater and D. Samyde. EMA: Measures and counter-measures for smart cards. In *Proc. E-smart*, pages 200–210, 2001.
- [SLFP04] K. Schramm, G. Leander, P. Felke, and C. Paar. A Collision-Attack on AES. In *Proc. 6th CHES*, pages 163–175, 2004.
- [SWP03] K. Schramm, T. Wollinger, and C. Paar. A New Class of Collision Attacks and Its Application to DES. In *Proc. 10th FSE*, pages 206–222, 2003.
- [VCGRS13] N. Veyrat-Charvillon, B. Gérard, M. Renaud, and F. Standaert. An optimal key enumeration algorithm and its application to side-channel attacks. pages 390–406, 2013.