

Higher-Order DCA against Standard Side-Channel Countermeasures

Andrey Bogdanov¹, Matthieu Rivain², Philip S. Vejrø¹, and Junwei Wang^{2,3,4}

¹ Technical University of Denmark
{anbog,psve}@dtu.dk

² CryptoExperts, France
{matthieu.rivain,junwei.wang}@cryptoexperts.com

³ University of Luxembourg

⁴ University Paris 8, France

Abstract. At CHES 2016, Bos *et al.* introduced *differential computational analysis* (DCA) as an attack on white-box software implementations of block ciphers. This attack builds on the same principles as DPA in the classical side-channel context, but uses computational traces consisting of plain values computed by the implementation during execution. It was shown to be able to recover the key of many existing AES white-box implementations.

The *DCA adversary* is *passive*, and so does not exploit the full power of the white-box setting, implying that many white-box schemes are insecure even in a weaker setting than the one they were designed for. It is therefore important to develop implementations which are resistant to this attack. We investigate the approach of applying standard side-channel countermeasures such as *masking* and *shuffling*. Under some necessary conditions on the underlying randomness generation, we show that these countermeasures provide resistance to standard (first-order) DCA. Furthermore, we introduce *higher-order DCA*, along with an enhanced *multivariate* version, and analyze the security of the countermeasures against these attacks. We derive analytic expressions for the complexity of the attacks – backed up through extensive attack experiments – enabling a designer to quantify the security level of a masked and shuffled implementation in the (higher-order) DCA setting.

Keywords: White-box cryptography · Higher-order DCA · Masking · Shuffling

1 Introduction

In the classical cryptanalytic setting, the adversary faces the challenge of breaking the security of *e.g.* an encryption algorithm while only being able to consider the algorithm as a *black box*; she can query the box with inputs and receive the corresponding outputs. While the design of the algorithm is known, the adversary cannot observe the internal state of the algorithm, or affect the execution of the algorithm.

In practice, a cryptographic algorithm has to be implemented somewhere to be useful, *i.e.* in hardware or software. Thus, the adversary has the option of physically interacting with the encryption device. In this case, the adversary has access to implementation specific *side-channel information*. If an implementation is not sufficiently protected, leaking information such as the execution time or the power consumption can be used to extract secret information, *e.g.* encryption keys. The widespread use and success of side-channel attacks show that a cryptographer has to be very careful when operating in this *gray-box* model.

1.1 Shades of Gray

For hardware implementations, the gray-box model is often the limit of what an adversary can achieve. This is not the case for software implementations that are executed in untrusted environments. If an adversary is given full access to the execution environment of the cryptographic software, she can easily observe and

manipulate the execution of the primitive, instantiated with some secret key. This setting, introduced by Chow *et al.* in [10], is called the *white-box* model.

In the white-box setting, the adversary can study the logic flow of the implementation, observe any tables the implementation uses, observe intermediate values computed during execution, alter the implementation at run-time, etc. Indeed, “*when the attacker has internal information about a cryptographic implementation, choice of implementation is the sole remaining line of defense*” [10]. Ideally, the aim of a white-box implementation would be to leave the white-box adversary with at most the same advantage as a black-box adversary, but this seems to be very difficult. Instead, the current white-box paradigm aims to provide *practical security*, in the sense that the implementation is difficult enough to attack so that an adversary is forced to attempt other attack vectors.

Several different approaches have been proposed to make white-box secure implementations for standard block ciphers, such as AES [10,9,29,17,5,6]. Sadly, all these designs have been broken by structural attacks [2,22,21,27,19,23]. Additionally, no provably secure solutions have been reported in the literature. Still, there exists an increasing industrial need for the protection of cryptographic implementations executed in untrusted environments, such as the traditional *digital rights management* use case and the mobile payment applications running on smart devices. This has driven the industry to develop home-made solutions relying on *obscurity* (i.e. secrecy of the underlying obfuscation techniques). In this paradigm, white-box cryptography acts as a moving target based on regular security updates and/or short-term key tokens and is considered a building block of wider security solutions.

1.2 Differential Computational Analysis

The mentioned attacks against *public* white-box implementations exploit flaws in the underlying white-box schemes. However, secret variants of known designs that change a few parameters or combine different techniques would likely thwart these attacks, if exact designs are kept secret. To attack such implementations, the adversary would have to perform reverse engineering, which can take considerable time and effort if various layers of obfuscation have been applied.

A more generic approach was given recently by Bos *et al.*; [4] introduced the white-box equivalent of DPA, namely *differential computational analysis* (DCA), and demonstrated how this technique is able to recover the encryption key of several existing white-box implementations of the AES. Notably, the *DCA adversary* is extremely powerful as it is *implementation agnostic* and therefore does not need to exert expensive reverse engineering efforts. Thus, DCA has been devastating for industrial solutions that leverage design secrecy to develop white-box implementations consisting of a mix of various techniques (common white-box techniques, code obfuscation, home-made encodings, etc.).

Moreover, the DCA adversary does not take advantage of the full power of the white-box model. The adversary only needs to be able to observe the addresses and values of memory being accessed during the execution of the implementation. The adversary does not need to reason about the implementation details, or modify the functionality of the code in any way, *e.g.* by disabling RNG functionalities – tasks that could require considerable effort. Thus, the DCA attack is a *passive* and *non-invasive* attack, existing in a setting which is closer to a gray-box than a white-box. Additionally, the attacks presented in [4] have very low complexities.

Following these observations, the current white-box AES implementations are not even secure in a weaker attack context than the one they were designed for, and as a consequence, designing secure white-box implementation seems out of reach. Indeed, DCA seems to be the biggest hindrance to designing practically secure white-box implementations. It is therefore of importance to first explore the design of cryptographic implementations which are secure against DCA.

1.3 Our Contributions

A natural approach when attempting to mitigate the threat of DCA attacks is to apply known countermeasures from the side-channel literature. However, it is not clear how well these countermeasures carry over to

the white-box context and what level of security can be achieved by such countermeasures against a DCA adversary. To address these issues, we achieve the following:

- **Side-channel countermeasures in the white-box setting:** In Section 3 we discuss how to apply the well known side-channel countermeasures in the white-box setting focusing on the (passive) DCA adversary. Specifically, we focus on *higher-order masking* along with *shuffling of operations* to introduce noise in the DCA traces. We show that if the source of randomness used in the implementation satisfies some specific security properties, then this approach is sufficient to achieve security against standard first-order DCA.
- **Higher-order DCA:** We develop *higher-order DCA* in Section 4 to analyze the security of the proposed protection. We show that higher-order DCA is able to break a masked implementation of any order using a couple of traces. However, by introducing noise in the form of shuffling, the security of the implementation can be dramatically increased. As a demonstration, a typical AES implementation with 2nd order masking (and shuffling degree of 16) requires 2²¹ traces to break with 3rd order DCA.
- **Multivariate higher-order DCA:** We extend the above attack by introducing a multivariate version in Section 5, which reduces the computational complexity by decreasing the number of required traces for a successful attack. Using this multivariate variant, the number of traces required to successfully attack the AES implementation mentioned above can be reduced to 2¹⁰.
- **Formal analysis and experimental verification:** We derive analytic expressions for the success probability and attack complexities of both the higher-order DCA and its multivariate variant. Using these expressions, we are able to give estimates for the security level of a masked and shuffled implementation in the DCA setting. As an example, an AES implementation with 7th order masking would have a security level of about 85 bits in this setting. Then accuracy of our expressions for the success probability of the multivariate higher-order DCA is verified in Section 5.2 through extensive experiments for a wide range of implementation and attack parameters. 2 000 attacks of up to order 4 were simulated, using as many as 30 000 traces per attack.

In summary, our result provides formal ground to the study of standard side-channel defenses in the white-box setting. We have analyzed the widely used masking and shuffling countermeasures with respect to the DCA adversary and we have quantified their security against advanced DCA attacks. From our analysis, a designer can choose an appropriate set of implementation parameters to achieve a given security level with respect to DCA, which is a first step towards building security against a stronger white-box adversary.

1.4 Related Works

Two independent and related works have been published since the first version of the present paper which both address masking in the white-box context. First, Biryukov and Udovenko [3] broadly overview how a white-box adversary could attack masked white-box implementations in several aspects, in particular, in an *active* (fault injection) attack setting. Second, Goubin *et al.* [15] proposed a method to attack an *obscure* white-box implementation (in the sense that the adversary has no/limited knowledge on the design), which was successfully applied to break the winning challenge in the recent WhibOx 2017 contest [1]. Particularly, their *linear decoding analysis* can break a noise-free masked implementation with complexity approximately cubic in the size of the computation trace.

Compared to these works, the adversary we consider is *passive*, as such she would not be impacted by fault detection/correction measures. Moreover, [15] demonstrates that masking is a weak countermeasure unless it is composed with some kind of *noise*. Therefore, this work focuses on white-box implementations protected with both masking and shuffling.

2 Differential Computation Analysis

The *DCA adversary* is capable of querying a software implementation of a cryptographic primitive with arbitrary input to obtain a *computational trace* of the execution. The computational trace consists of: any

value calculated, written, or read by the implementation, and the address of any memory location read from or written to during execution. Each data point in the computational trace is further annotated with the time it occurred in the execution. When a number of traces have been collected, the adversary calculates correlations between a prediction of a key-dependent intermediate value and computed values.

2.1 DCA Setting

As for hardware side-channel attacks, DCA exploits that the (software) implementation leaks some information about intermediate variables involved in the execution of the cryptographic algorithm. Some of these intermediate variables depend on the plaintext and (part of) the secret key, and knowledge of such variables can therefore reveal the key. We denote such a *secret variable* by $s = \varphi(x, k^*)$, where φ is a deterministic function, x is a public value, *e.g.* (part of) the plaintext, and $k^* \in \mathcal{K}$ is a (secret) subkey over some subkey space \mathcal{K} . For instance, k^* could be a byte of the secret key and \mathcal{K} would then be $\{0, 1\}^8$.

The DCA attack itself consists of first obtaining a number of computational traces from the execution of the cipher implementation with secret subkey k^* for several (random) plaintexts. We denote a computational trace consisting of t time points by the ordered t -tuple $\mathbf{v} = (v_1, v_2, \dots, v_t)$, with $v_i \in \mathcal{V}$ for some set \mathcal{V} . In principle, the traces can be any value exposed to a dynamic binary analysis tool, as explained above. Usually, an attacker will obtain N computational traces, $\mathbf{v}_1, \dots, \mathbf{v}_N$, representing N executions of the implementation with different inputs, each corresponding to a value $s_i = \varphi(x_i, k^*)$ of the target secret variable. These traces could *e.g.* arise from the encryption of N different plaintexts. The attacker then performs a classic DPA, in which a *distinguisher* is used to indicate a correct guess of k^* . The distinguisher is a function D which maps the set of computational traces $(\mathbf{v}_i)_i$, and corresponding inputs $(x_i)_i$, to a *score vector*:

$$(\gamma_k)_{k \in \mathcal{K}} = D((\mathbf{v}_1, \dots, \mathbf{v}_N); (x_1, \dots, x_N)) .$$

The adversary then selects the key guess k with the highest score γ_k as candidate for the correct value of k^* . We define the success probability of the attack as

$$p_{\text{succ}} = \Pr(\operatorname{argmax}_{k \in \mathcal{K}} \gamma_k = k^*),$$

where this probability is taken over any randomness supplied to the implementation (including the randomness of the inputs).

2.2 Standard First-Order DCA

The description above does not specify the distinguisher D . Here, we briefly describe the distinguisher used in [4], which we will call the *standard first-order DCA*. Let us denote by $v_{i,j}$ the value at the j 'th time point of the i 'th trace. The standard first-order DCA attack consists of calculating a correlation coefficient between a vector of predicted values of the secret variable, (s_1^k, \dots, s_N^k) , where $s_i^k = \varphi(x_i, k)$, and the vector $(v_{1,j}, \dots, v_{N,j})$, for every time index $1 \leq j \leq t$. Then the score γ_k is defined as the maximum correlation obtained over the different time indices, *i.e.*

$$\gamma_k = \max_j C((v_{1,j}, \dots, v_{N,j}), (\psi(s_1^k), \dots, \psi(s_N^k))),$$

for some correlation measure C and some *pre-processing function* ψ . For example, C could be the Pearson correlation coefficient and ψ either the Hamming weight function or the selection of one bit of the predicted variable. If there exists a statistical correlation between the secret variable and the values of the computational trace, we would expect a large absolute value of C for some index j and the correct prediction of the secret variables, *i.e.* the vector $(s_1^{k^*}, \dots, s_N^{k^*})$. On the other hand, if $k \neq k^*$, we expect a low correlation between all s_i^k and any point in the computational trace. It was shown in [4] that this approach is very effective against a range of different AES and DES white-box implementations.

3 Side-Channel Countermeasures against DCA

The DCA adversary is highly reminiscent of the standard side-channel adversary. It is therefore natural to apply traditional side-channel countermeasures to a white-box implementation, and to evaluate their performance against a DCA adversary. In the following, we specifically study the common software countermeasures of *higher-order masking* and *operation shuffling*. We further discuss the source of randomness necessary to feed these countermeasures and state a few security properties that it should satisfy in this setting. We then show that this approach achieves security against standard first-order DCA. The rest of the paper is dedicated to the study of advanced DCA attacks against these countermeasures.

3.1 DCA is a Passive and Non-Invasive Gray-Box Attack

We start by noting that applying the mentioned countermeasures in a strict white-box context might be hazardous – indeed classical countermeasures such as masking and shuffling use fresh randomness throughout the execution of the protected implementation which is usually provided by an external random number generator (RNG). Since a white-box adversary has full control over the execution environment, such an RNG could be detected and disabled, shuffled operations could be re-synchronized (*e.g.* using memory addresses, program counter, etc.), and/or masks could be canceled (if masked variables and corresponding masks are easily identified).

In order to make such disabling difficult, the used randomness should rely on some internal PRNG (see Section 3.4) and one should further add some layers of obfuscation countermeasures on top of it. The adversary then has to invest some reverse engineering effort to bypass these countermeasures. Nevertheless, we stress that this is exactly the type of analysis an adversary performing a DCA attack is trying to avoid. While DCA attacks might not be optimal in terms of time and/or data complexity, they are very powerful due to their genericness and the fact that they can be applied in a *black-box* way, *i.e.* without requiring reverse engineering effort. This is an essential property of these attacks in the current white-box cryptography paradigm where designers aim at practical security (as provable security seems out of reach) and use the secrecy of the design as a leverage towards this goal. The main purpose of protecting against DCA is therefore to force an adversary to employ more complicated and dedicated attack techniques, which might take a long time to develop and apply, which is beneficial when combined with a moving target strategy.

3.2 Masking

A widely used countermeasure to standard DPA of hardware implementations is *masking* [7,16]. Since the DCA attack relies on the same ideas as DPA, the prospect of applying masking to secure a software implementation against DCA is promising. To mask a secret variable, it is split into several parts that are then processed independently. Specifically, each secret variable s occurring in the execution of the implementation is split into d shares s_1, \dots, s_d such that $s_1 \oplus \dots \oplus s_d = s$. The masking must be done such that any subset of less than d shares are statistically independent of s . A simple way to achieve this is by picking s_1, \dots, s_{d-1} uniformly at random (*the masks*), and setting $s_d = s \oplus s_1 \oplus \dots \oplus s_{d-1}$ (*the masked variable*). The masking is then said to be of *order* $d - 1$. One important aspect of a masked implementation is therefore that of randomness: the implementation has to use a (P)RNG to generate these $d - 1$ masks.

Knowledge of all d shares is required to recover s , but combining the d shares would reveal the secret variable to the DCA adversary. Thus, the implementation must be able to perform computations on the secret variable s without combining the shares, *i.e.* for a function f , we want to compute shares r_i such that $r_1 \oplus \dots \oplus r_d = f(s)$, from the original shares s_1, \dots, s_d . The computation is then said to be *secure at the order* τ if no τ -tuple of intermediate variables is statistically dependent on a key-dependent variable. Usually, d 'th order masking aims to provide security at the order $\tau = d - 1$. To compute any \mathbb{F}_2 -linear function on a masked variable s , we simply compute the function on each share separately. Thus, calculation of the linear components of a typical SPN can be easily implemented on the masked state. Computing the non-linear components (*i.e.* typically the S-boxes) is more involved but several masking schemes exist that achieve $(d - 1)$ 'th order security (see for instance [24,13,11]).




Fig. 1. An illustration of memory (top) and time (bottom) shuffling applied to a second-order masked implementation. The location of each share in memory and the order of iteration is randomised for each execution.

3.3 Shuffling

We will show in Section 4 that if masking is the only countermeasure, the DCA adversary can easily recover the key. Indeed, the strength of a masked implementation is directly related to how noisy the adversary’s observation of the shares is. Several approaches for introducing and increasing noise in masked implementations have been proposed and analyzed, *e.g.* in [28,26,12,25]. One such approach is shuffling: instead of processing the calculations of the cipher in some fixed order, the order of execution is randomly chosen for each run of the implementation based on the value of the input (*e.g.* the plaintext). The situation is slightly more complicated in the DCA setting. Here, the adversary can make observations in two dimensions, namely *time* and *memory*. Even if the order of execution is shuffled in time, an adversary can choose to order the traces by the memory addresses accessed. Thus, we need to shuffle in both the time and memory dimension.

Memory Shuffle. In a masked implementation, we will typically have some state in which each element is shared as described in Section 3.2. The idea of the memory shuffle is to randomly rearrange the shares of the state in memory. Consider a state consisting of S elements. We assume that the shares $s_{i,j}$, $1 \leq i < S$, $1 \leq j \leq d$, are stored in an array, initially in order. That is, the implementation uses the array $(s_{1,1}, s_{2,1}, \dots, s_{S-1,d}, s_{S,d})$. Then, we randomly pick a permutation $P : [1, S] \times [1, d] \rightarrow [1, S] \times [1, d]$, based on the value of the input. Note that this can be done efficiently using the Fisher-Yates shuffle [14]. Now, instead of using the in-order array, we rearrange the array such that the implementation uses the array $(s_{P(1,1)}, s_{P(2,1)}, \dots, s_{P(S-1,d)}, s_{P(S,d)})$. Whenever the implementation needs to access share $s_{i,j}$, it simply looks up the element in position $P^{-1}(i, j)$ of the array. A similar randomisation is performed for any key shares.

Time Shuffle. In a typical SPN, there will be several operations that operate on each element of the state in each round. The order of these operations is typically suggested by the cipher designers. As an example, consider the case where we want to apply a linear operation A to each element of the state separately. Since the operation is linear, we can apply it to each share of the masked elements individually. This will normally be done in some “natural” order, *e.g.* $A(s_{1,1}), A(s_{1,2}), A(s_{1,3}), \dots, A(s_{S,d})$. However, the exact order of execution does not matter. Thus, we can shuffle in the time dimension by randomly ordering these $S \cdot d$ operations. In general, if a set of λ independent operations exists, we can freely shuffle the order in which we process the $\lambda \cdot d$ shares. Formally, we randomly pick a permutation $Q : [1, \lambda] \times [1, d] \rightarrow [1, \lambda] \times [1, d]$. Then, when we normally would have processed share $s_{i,j}$, we instead process share $s_{Q(i,j)}$. Thus, the probability that a specific share is processed in a given step is $1/(\lambda \cdot d)$. We will denote the size of the smallest maximal set of independent operations the *shuffling degree*.

3.4 On the Source of Randomness

A potential issue while applying side-channel countermeasures to the white-box context is randomness generation. Since a white-box adversary can easily get rid of an external RNG, the randomness used by a white-box implementation must be pseudo-randomly generated from the single available source of variation: the input plaintext. In other words, the white-box implementation should embed some kind of pseudo random number generator (PRNG) seeded by the input plaintext. We now (informally) state a few security properties that should be fulfilled by such a PRNG in the white-box setting:

1. *Pseudorandomness*: The output stream of the PRNG should be hard to distinguish from true randomness.
2. *Obscurity*: The design of the PRNG should be kept secret.
3. *Obfuscation*: The PRNG should be mixed with the white-box implementation so that its output stream is hard to distinguish from other intermediate variables.

The pseudorandomness property is required to ensure that the PRNG does not introduce a statistical flaw in the implemented countermeasures. It is well known that a flawed RNG can be a disaster for the security of masking (see for instance [20]). The pseudorandomness property further implies that the generated randomness is *unpredictable* provided that the obscurity property also holds. The unpredictability of the generated randomness is necessary to get DCA resistance, since without it all the intermediate variables can be expressed as (known) deterministic functions of the plaintext and the secret key, which enables the application of standard first-order DCA.⁵

Indeed, if the PRNG design was known to the adversary, then she could predict all the generated randomness from the plaintext. Therefore in order to provide unpredictability, some part of the design must be secret, even if the obscurity concept clearly clashes with the adage of Kerckhoffs’s Principle [18]. Nevertheless, it seems almost impossible to provide any security if the full design is known, and we stress that this does not imply that one should forego all good cryptographic engineering practices. One could use a keyed PRNG (or PRNG with secret initial state), but even then if the design was known to the adversary she could mount a DCA attack to recover the PRNG key and we would then face a chicken and egg problem. Alternatively, an implementation could use a known strong PRNG with some sound changes to design parameters, in order to have some confidence in its security. Another approach, which aligns with the moving target strategy, would be to have a set of different PRNG designs that are often changed.

Finally the obfuscation property is required to prevent easy detection of the PRNG output which could facilitate a DCA attack. It is for instance described in [15] how the generated randomness can be easily detected by switching the values of intermediate variables and checking whether this affects the final result. Such a detection is an *active* attack that tampers with the execution (in the same way as *fault attacks*) and is hence out of scope of the DCA adversary. However it should be made difficult (in the same way as fault attacks should be made difficult) to achieve some level of resistance in practice.

In the following, we shall consider that the above security properties are satisfied by the used PRNG so that the DCA adversary cannot easily remove or predict the generated randomness. We will then analyze which level of security is achievable by using masking and shuffling in this context.

3.5 Resistance to First-Order DCA

In Section 2 we described the capabilities of the DCA adversary and the standard first-order DCA. For the masked implementation described above with $d > 1$, the d ’th order security of the underlying masking scheme implies that any d -tuple of variables from the computation trace $\mathbf{v} = (v_1, v_2, \dots, v_t)$ is statistically independent of any key-dependent variable. Assuming that the PRNG embedded in the target implementation outputs strong and unpredictable pseudorandomness (as required in Section 3.4) the distribution of any d -tuple of variables from \mathbf{v} is indistinguishable from the same distribution with perfect randomness, which makes it (computationally) independent of any key-dependent variable. Therefore, standard first-order DCA as described in Section 2.2 is doomed to fail.

⁵ Following the DCA setting described in Section 2.1, the only impact of the countermeasures in presence of a known PRNG is to change the deterministic function φ in the expression of the secret variable s .

4 Introducing Higher-Order DCA

While masking has been proven to be an effective defense against standard DPA, and we have argued for its effectiveness against standard first-order DCA, there are ways to attack such masked implementations. For hardware implementations, it is well known that a $(d - 1)$ 'th-order masked implementation, such as the one described above, can be defeated by d 'th-order DPA, if no other protection is employed. We will therefore develop a higher-order version of DCA.

A d 'th-order DCA consists of a pre-processing step followed by a first-order DCA. The adversary first pre-processes each computational trace \mathbf{v} to obtain a d 'th-order computational trace \mathbf{w} by applying a so-called (d 'th-order) *combination function* ψ . Specifically, the d 'th-order computational trace \mathbf{w} consists of $q = \binom{t}{d}$ points (w_1, \dots, w_q) given by

$$w_j = \psi(v_{j_1}, v_{j_2}, \dots, v_{j_d}), \quad \{j_1, \dots, j_d\} = \phi(j),$$

where $\phi(j)$ is the j 'th subset of $\{1, \dots, t\}$ of size d (for some ordering). After computing the set of d 'th order traces $\mathbf{w}_1, \dots, \mathbf{w}_N$, the adversary proceeds as for first-order DCA, using the \mathbf{w}_i 's as input to the distinguisher D . Specifically, the adversary computes the score vector $(\gamma_k)_{k \in K} = D((\mathbf{w}_i)_i; (x_i)_i)$ in order to determine a candidate for k^* .

For side-channel analysis of hardware implementations, it has been shown that a good combination function for higher-order DPA is the centered product $\psi : (v_1, \dots, v_d) \mapsto \prod_j (v_j - \mu_j)$, where μ_j is the average of the leakage point v_j over several encryptions. Nevertheless, since the measurements in this setting are inherently noisy, a larger masking degree will require a larger number of traces to obtain a good success probability. Note that this is not the case in the DCA context, if no noise is introduced in the implementation, *e.g.* by using shuffling as described in Section 3.3. In this case, the exact value of each variable that appears in an execution of the implementation appears at the same position of every computational trace \mathbf{v}_i . Then there exists a fixed j^* , such that for $\phi(j^*) = (j_1^*, \dots, j_d^*)$, the elements $v_{j_1^*}, \dots, v_{j_d^*}$ of the trace are the shares of the target secret variable s . In that case, an optimal choice for the combination function is the XOR sum of the trace values, that is

$$\psi(v_{j_1}, v_{j_2}, \dots, v_{j_d}) = v_{j_1} \oplus v_{j_2} \oplus \dots \oplus v_{j_d} .$$

For this combination function, we have that $\psi(v_{j_1^*}, v_{j_2^*}, \dots, v_{j_d^*}) = s$ for all the d 'th-order traces. By counting the number of times this equality holds, we can easily determine the correct key. That is, we set

$$\gamma_k = \max_j (C_k(\mathbf{v}_{\phi(j)}, (x_i)_i)), \quad \text{with } C_k(\mathbf{v}_{\phi(j)}, (x_i)_i) = \left| \left\{ i ; \bigoplus_{l \in \phi(j)} v_{i,l} = \varphi(x_i, k) \right\} \right|.$$

For the correct key k^* , we deterministically have that $\gamma_{k^*} = N$. Thus, if no noise is present, the higher-order DCA is successful when $\gamma_{k^\times} < N$ for all $k^\times \neq k^*$. The probability of this happening is quite close to 1, even for small N . Thus, the introduction of some noise in the traces is required to secure a masked white-box implementation against DCA.

4.1 Higher-Order DCA against Masking and Shuffling

We now consider how well the masked and shuffled implementation resists the higher-order DCA attack described above. Due to the shuffling, the adversary is no longer guaranteed that her prediction for the correct key guess will correspond to a single time point for all traces. Thus, she must compensate by increasing the number of traces. The higher the degree of shuffling, the more traces needs to be collected.

Attack Analysis In the following, we assume that the adversary knows exactly where in the computational trace to attack. That is, for a masking order $d - 1$ and a shuffling degree λ , she knows the range of the $t = \lambda \cdot d$ time points that contain the shares of the target secret variable. In other words, the length of each

Table 1. The number of traces N and the time needed to successfully attack an implementation with $(d - 1)$ -order masking and shuffling of degree λ with d 'th-order DCA. Here, $|\mathcal{K}| = 256$, and we fix the success probability at 90%. The parameters chosen would be typical for a protected AES implementation.

d	λ	$\log_2 N$	$\log_2 \text{time}$	d	λ	$\log_2 N$	$\log_2 \text{time}$	d	λ	$\log_2 N$	$\log_2 \text{time}$
2	8	8.6	23.5	3	8	15.7	34.7	4	8	23.6	46.7
2	16	11.0	28.0	3	16	21.6	43.7	4	16	31.7	59.0

computational trace \mathbf{v} is t . This, intuitively, represents the optimal situation for the adversary.⁶ We seek an expression for the success probability of the attack, *i.e.* the probability that the correct key has a higher score than all other key candidates.

The adversary proceeds as above and computes the d 'th-order computational trace. However, there will no longer be a single value j such that $C_{k^*}(\mathbf{v}_{\phi(j)}, (x_i)_i) = N$ deterministically for the correct key k^* . Thus, we need to know the distribution of γ_k , both for a wrong and a right guess of the key.

Theorem 1. Consider a masked white-box implementation of order $d - 1$ with shuffling degree λ . Let $p = \binom{t}{d}^{-1}$ where $t = \lambda \cdot d$, and let $F(x; n, q)$ be the CDF (cumulative distribution function) of the binomial distribution with parameters n and q . Let $|\mathcal{K}|$ be the number of possible key values and define

$$F_{\max}^{\times}(x) = F(x; N, (1 - p) \frac{1}{|\mathcal{K}|}) \binom{t}{d},$$

$$F_{\max}^*(x) = F(x; N, p + (1 - p) \frac{1}{|\mathcal{K}|}) \binom{t}{d}.$$

Then the probability of recovering a key using d 'th-order DCA with N traces is

$$p_{\text{succ}} = \left(\sum_{i=0}^N (F_{\max}^*(i) - F_{\max}^*(i - 1)) \cdot F_{\max}^{\times}(i - 1) \right)^{|\mathcal{K}| - 1}.$$

We prove Theorem 1 in Appendix A. We can use this formula to calculate the required N to obtain a desired probability of success. The number of traces required to obtain 90% success probability for a range of parameters is shown in Table 1. Here, $|\mathcal{K}| = 256$, and the parameters would be typical choices for *e.g.* a protected AES implementation.

Attack Complexity We consider the time complexity of recovering the secret key k^* using the higher-order DCA attack. For a fixed probability of success p_{succ} , let N_d be the number of computational traces required to obtain this probability for a d 'th-order implementation. We again assume that $t = \lambda \cdot d$. The cost of computing the higher-order trace is $N_d \cdot \binom{t}{d}$. Then, for each key guess k and each time point in the higher-order trace, the adversary computes C_k . The complexity of this is $|\mathcal{K}| \cdot N_d \cdot \binom{t}{d}$. Thus, the time complexity is $\mathcal{O}(|\mathcal{K}| \cdot N_d \cdot \binom{t}{d})$. Table 1 shows the time complexity of the attack for a range of parameters.

5 Multivariate Higher-Order DCA

In the higher-order DCA, presented in Section 4, the adversary tries to correlate each sample of the higher-order trace with the predicted variable independently, finally taking the maximum over the obtained correlation scores. Such an approach is not optimal, as successive samples may carry joint information on the secret. As in the side-channel context, one can take advantage of this joint information by performing a *multivariate attack*, namely an attack in which the distinguisher exploits the multivariate distribution of different samples in the higher-order trace. Emblematic multivariate attacks in the classical side-channel context are the so-called *template attacks* [8]. In the following section, we describe a similar attack in the setting of the DCA adversary.

⁶ In practice, the adversary could exhaustively search the correct location of the $(\lambda \cdot d)$ -length subtrace in the full computation trace of length t_{full} , which increases the complexity at most t_{full} times.

5.1 Multivariate Higher-Order DCA against Masking and Shuffling

Our proposed multivariate higher-order DCA attack is based on the principle of maximum likelihood. Similar techniques have been adopted in side-channel template attacks. Let K , $(X_i)_i$, and $(\mathbf{V}_i)_i$ be random variables representing the subkey k , the public inputs $(x_i)_i$, and the computational traces $(\mathbf{v}_i)_i$. The likelihood distinguisher is then defined as

$$\begin{aligned} \mathbf{L} : ((\mathbf{v}_i)_i, (x_i)_i) &\mapsto (\ell_k)_{k \in \mathcal{K}}, \\ \ell_k &\propto \Pr(K = k \mid (\mathbf{V}_i)_i = (\mathbf{v}_i)_i \wedge (X_i)_i = (x_i)_i), \end{aligned} \quad (1)$$

where \propto means equal up to some factor constant w.r.t. k . To evaluate this likelihood function, we need a model for the distribution of the traces (also called a *template* in the side-channel context). It is well known that if Equation 1 is evaluated from the true distributions of $(X_i)_i$ and $(\mathbf{V}_i)_i$, then the above distinguisher is optimal. This is sound, as in this case, the score is the exact probability that the target subkey equals a key guess k , for all $k \in \mathcal{K}$.

In the following, we will assume that \mathbf{V}_i is composed of t uniformly distributed random variables $V_{i,1}, V_{i,2}, \dots, V_{i,t}$, with the constraint that for a uniformly chosen j , we have $\bigoplus_{l \in \phi(j)} V_{i,l} = \varphi(X_i, K)$. This assumption matches the setting of a masked and shuffled implementation. The public inputs X_i and the subkey K are also assumed to be uniformly distributed and mutually independent. Under this model, we have the following result (see proof in Appendix B):

Proposition 1. *The likelihood distinguisher, Equation 1, satisfies:*

$$\ell_k \propto \prod_{i=1}^N C_k(\mathbf{v}_i, x_i),$$

where $C_k(\mathbf{v}, x)$ is the number of d -tuples in a trace \mathbf{v} with bitwise sum equals to $\varphi(x, k)$, that is $C_k(\mathbf{v}, x) = |\{(v_{j_1}, \dots, v_{j_d}) ; v_{j_1} \oplus \dots \oplus v_{j_d} = \varphi(x, k)\}|$.

Remark 1. For practical reasons, it is more convenient to evaluate the log-likelihood, that is $\log \ell_k = \sum_{i=1}^N \log C_k(\mathbf{v}_i, x_i)$. Note that this does not affect the ranking of the key guesses (as the logarithm is a monotonically increasing function) and therefore has no impact on the success probability of the attack.

5.2 Analysis of the Likelihood Distinguisher

In this section we analyze the success probability of the likelihood distinguisher. For the sake of simplicity, we only consider two key guesses, namely the right key guess k^* and a wrong key guess k^\times . We then consider their likelihood scores ℓ_{k^*} and ℓ_{k^\times} random variables, since

$$\ell_k = \prod_{i=1}^N C_k(\mathbf{V}_i, X_i),$$

for $k \in \{k^*, k^\times\}$, where $(\mathbf{V}_i)_i$ and $(X_i)_i$ are the random variables defined above for the computational traces and the corresponding public inputs. We then consider the probability $p_{\text{succ}} = \Pr(\ell_{k^*} > \ell_{k^\times})$ in Theorem 2.

Theorem 2. *For a multivariate d 'th-order DCA attack using the likelihood distinguisher on N traces of length t , the probability that a correct key guess is ranked higher than an incorrect key guess is approximately given by*

$$p_{\text{succ}} \approx p_{\mathcal{U}} + (1 - p_{\mathcal{U}}) \left(\frac{1}{2} + \frac{1}{2} \operatorname{erf} \left(\frac{\sqrt{N|\mathcal{V}|}}{2\sqrt{q}} \right) \right)$$

where $q = \binom{t}{d}$ and $p_{\mathcal{U}} = 1 - (1 - (1 - |\mathcal{V}|^{-1})^q)^N$.

The total success probability of the attack $p_{\text{full-succ}}$, *i.e.* the probability that the correct key guess has the largest likelihood, is then heuristically $p_{\text{full-succ}} \approx p_{\text{succ}}^{|\mathcal{K}|-1}$. Moreover, it can be checked that $p_{\mathcal{U}} \approx N \cdot (1 - |\mathcal{V}|^{-1})^q$ becomes negligible as q grows. Theorem 2 then implies

$$p_{\text{succ}} = \Theta \left(\text{erf} \left(\frac{\sqrt{N|\mathcal{V}|}}{2\sqrt{q}} \right) \right),$$

from which we deduce that the data complexity of the attack is $N = \Theta(q)$. Namely, the number of required traces N to achieve certain p_{succ} is linear in the number of combinations $q = \binom{t}{d}$. We also have $N = \Theta(q/|\mathcal{V}|)$ to make appear the impact of the definition set \mathcal{V} .

In order to prove Theorem 2, we introduce the concept of the *zero-counter event*. Denoted by \mathcal{U}_k , this is the event that $C_k(\mathbf{v}_i, x_i) = 0$ for at least one $i \in [1, N]$ for a key guess k . Note that this event can never happen for $k = k^*$, since for all i , there exists a j such that $\bigoplus_{l \in \phi(j)} \mathbf{v}_{i,l} = \varphi(x_i, k^*)$. Thus, $\Pr(\ell_{k^*} > \ell_{k^\times} \mid \mathcal{U}_{k^\times}) = 1$, since in this case the likelihood ℓ_{k^\times} equals zero (or equivalently, the log-likelihood equals $-\infty$). This is intuitively sound, as the right key guess could not give rise to a zero counter for any of the N computational traces. Then, by the law of total probability, we can write

$$p_{\text{succ}} = \Pr(\mathcal{U}_{k^\times}) + \Pr(\neg \mathcal{U}_{k^\times}) \cdot \Pr(\ell_{k^*} > \ell_{k^\times} \mid \neg \mathcal{U}_{k^\times}). \quad (2)$$

We are therefore interested in the probabilities $\Pr(\mathcal{U}_{k^\times})$ and $\Pr(\ell_{k^*} > \ell_{k^\times} \mid \neg \mathcal{U}_{k^\times})$. These are given in the following lemmas.

Lemma 1. *Given N traces of length t , the probability of the zero-counter event for a wrong key guess k^\times in a d 'th-order attack is approximately given by*

$$\Pr(\mathcal{U}_{k^\times}) \approx 1 - \left(1 - (1 - |\mathcal{V}|^{-1})^q \right)^N,$$

where $q = \binom{t}{d}$.

Lemma 2. *Given N traces of length t , let $q = \binom{t}{d}$, and assume that the zero-counter event does not occur. The probability that a correct key guess has a higher likelihood score than a wrong key guess in a d 'th-order attack is approximately*

$$\Pr(\ell_{k^*} > \ell_{k^\times} \mid \neg \mathcal{U}_{k^\times}) \approx \frac{1}{2} + \frac{1}{2} \text{erf} \left(\frac{\sqrt{N|\mathcal{V}|}}{2\sqrt{q}} \right).$$

We prove Lemma 1 in Appendix C and Lemma 2 in Appendix D. Theorem 2 then follows directly from Equation 2 and these two results.

6 Experimental Verification and Security Evaluation

The proof of Theorem 2 relies on a number of approximations. We therefore verified the accuracy of the estimate by simulating the multivariate higher-order DCA attack for various choices of the parameters d and t . We chose to simulate traces of a masked and shuffled AES implementation, that is, the target secret variable was taken to be $\varphi(x, k^*) = \text{Sbox}_{\text{AES}}(x \oplus k^*)$. The computational traces were generated according to the model described at the beginning of Section 5.1, namely by sampling random values v_j over $\mathcal{V} = \mathbb{F}_{2^8}$ with the constraint that one randomly chosen d -tuple of each trace has XOR-sum $\varphi(x, k^*)$.

We generated traces for $d \in \{2, 3, 4\}$ and $t \in \{8, 16, 24, 32, 40, 48, 56, 64, 72\}$, and calculated the log-likelihood scores for the correct key and a randomly chosen wrong key. This was repeated 2000 times, and the probability $\Pr(\ell_{k^*} > \ell_{k^\times})$ was calculated for varying values of N . The results are shown in Figure 2. The figure shows that the estimate of Theorem 2 is quite accurate in most cases, only deviating from the

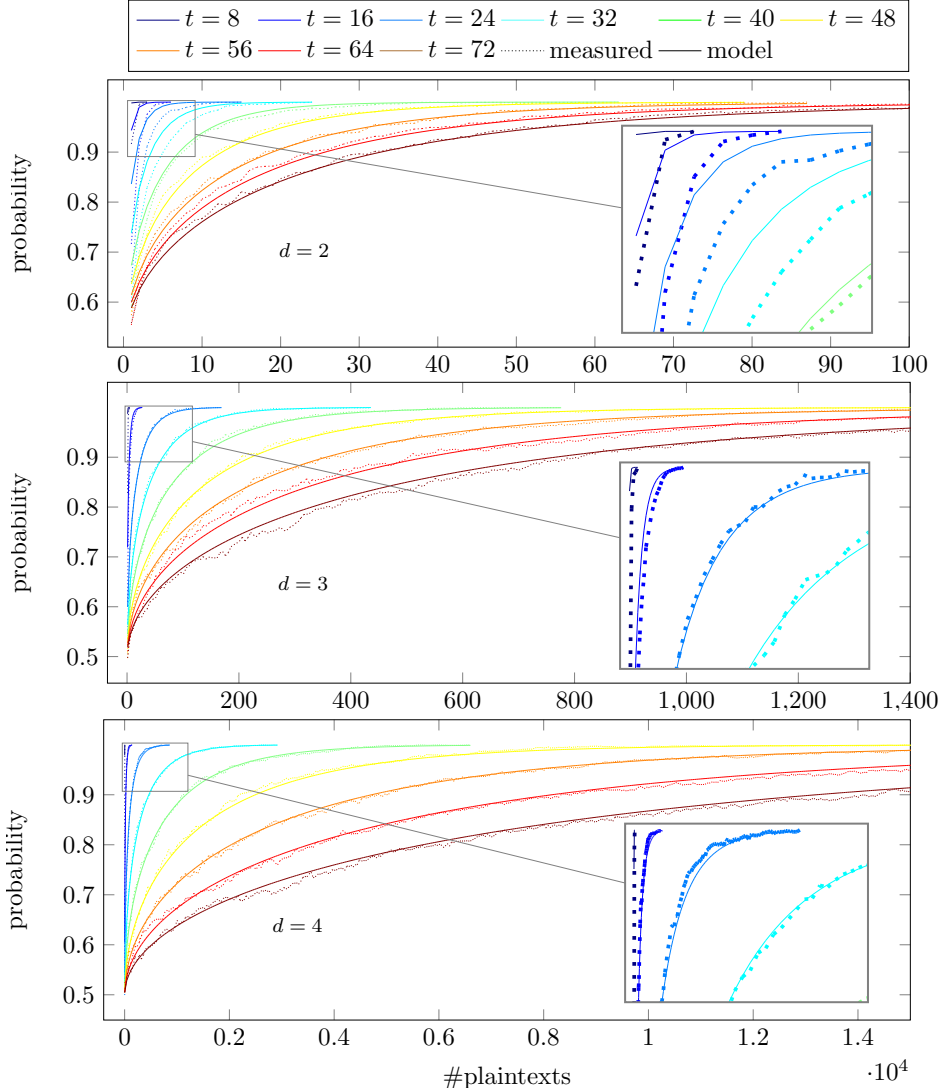


Fig. 2. The measured probability of ranking a correct key higher than an incorrect key in the multivariate higher-order DCA attack, compared to Theorem 2. The measurement is based on 2000 simulations of the attack. Here, d is the attack order and t is the length of the obtained traces.

experimental measurements for very small values of $q = \binom{t}{d}$ ($q < 300$). Note that in practice, this would rarely be a problem. For example, if all shares of the full AES state were shuffled in a first order masked implementation, as described in Section 3.3, the smallest trace that would always contain the correct shares would have $q = \binom{2^{16}}{2} = 496$.

The attack complexity of the multivariate higher-order DCA is the same as that of the higher-order DCA, namely $\mathcal{O}(|\mathcal{K}| \cdot N_d \cdot \binom{t}{d})$. Using this and Theorem 2 we can provide an estimate of the security level obtained by a masked and shuffled implementation against the DCA adversary. Table 2 shows the complexities of attacking *e.g.* a protected AES implementation where the operations are shuffled among all 16 state bytes (the shuffling degree is $\lambda = 16$ implying $t = 16 \cdot d$). When fixing $p_{\text{full-succ}}$ at 90%, we see that an implementation which uses 7th order masking will obtain an estimated security level of 85 bits.

Table 2. The number of traces N and the time needed to successfully attack an implementation with $(d - 1)$ -order masking and shuffling of degree $\lambda = 16$ ($t = 16 \cdot d$) using multivariate d 'th-order DCA. Here, $|\mathcal{K}| = 256$, and we fix the success probability at 90%.

d	$\log_2 N$	$\log_2 \text{time}$	d	$\log_2 N$	$\log_2 \text{time}$	d	$\log_2 N$	$\log_2 \text{time}$
3	10.6	32.7	5	21.0	53.5	7	31.6	74.6
4	15.8	43.1	6	26.3	64.1	8	36.9	85.3

Acknowledgment

The fourth author was supported by European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No. 643161.

References

- CHES 2017 Capture the Flag Challenge - The WhibOx Contest, An ECRYPT White-Box Cryptography Competition. <https://whibox.cr.yt.to/>, accessed: October 2017
- Billet, O., Gilbert, H., Ech-Chatbi, C.: Cryptanalysis of a White Box AES Implementation. In: Selected Areas in Cryptography, 11th International Workshop, SAC 2004. pp. 227–240 (2004). https://doi.org/10.1007/978-3-540-30564-4_16
- Biryukov, A., Udovenko, A.: Attacks and countermeasures for white-box designs. In: Peyrin, T., Galbraith, S. (eds.) ASIACRYPT 2018, Part II. LNCS, vol. 11273, pp. 373–402. Springer, Heidelberg (Dec 2018). https://doi.org/10.1007/978-3-030-03329-3_13
- Bos, J.W., Hubain, C., Michiels, W., Teuwen, P.: Differential Computation Analysis: Hiding Your White-Box Designs is Not Enough. In: Cryptographic Hardware and Embedded Systems - CHES 2016. pp. 215–236 (2016). https://doi.org/10.1007/978-3-662-53140-2_11
- Bringer, J., Chabanne, H., Dottax, E.: Perturbing and Protecting a Traceable Block Cipher. In: Communications and Multimedia Security, 10th IFIP TC-6 TC-11 International Conference, CMS 2006. pp. 109–119 (2006). https://doi.org/10.1007/11909033_10
- Bringer, J., Chabanne, H., Dottax, E.: White Box Cryptography: Another Attempt. IACR Cryptology ePrint Archive **2006**, 468 (2006)
- Chari, S., Jutla, C.S., Rao, J.R., Rohatgi, P.: Towards Sound Approaches to Counteract Power-Analysis Attacks. In: Advances in Cryptology - CRYPTO '99. pp. 398–412 (1999). https://doi.org/10.1007/3-540-48405-1_26
- Chari, S., Rao, J.R., Rohatgi, P.: Template Attacks. In: Cryptographic Hardware and Embedded Systems - CHES 2002. pp. 13–28 (2002). https://doi.org/10.1007/3-540-36400-5_3
- Chow, S., Eisen, P.A., Johnson, H., van Oorschot, P.C.: A White-Box DES Implementation for DRM Applications. In: Security and Privacy in Digital Rights Management, ACM CCS-9 Workshop, DRM 2002. pp. 1–15 (2002). https://doi.org/10.1007/978-3-540-44993-5_1
- Chow, S., Eisen, P.A., Johnson, H., van Oorschot, P.C.: White-Box Cryptography and an AES Implementation. In: Selected Areas in Cryptography, 9th Annual International Workshop, SAC 2002. pp. 250–270 (2002). https://doi.org/10.1007/3-540-36492-7_17
- Coron, J.: Higher Order Masking of Look-Up Tables. In: Advances in Cryptology - EUROCRYPT 2014. pp. 441–458 (2014). https://doi.org/10.1007/978-3-642-55220-5_25
- Coron, J., Kizhvatov, I.: Analysis and Improvement of the Random Delay Countermeasure of CHES 2009. In: Cryptographic Hardware and Embedded Systems, CHES 2010. pp. 95–109 (2010). https://doi.org/10.1007/978-3-642-15031-9_7
- Coron, J., Prouff, E., Rivain, M., Roche, T.: Higher-Order Side Channel Security and Mask Refreshing. In: Fast Software Encryption - 20th International Workshop, FSE 2013. pp. 410–424 (2013). https://doi.org/10.1007/978-3-662-43933-3_21
- Fisher, R.A., Yates, F., et al.: Statistical tables for biological, agricultural and medical research. Statistical tables for biological, agricultural and medical research. (1938)
- Goubin, L., Paillier, P., Rivain, M., Wang, J.: How to reveal the secrets of an obscure white-box implementation. Cryptology ePrint Archive, Report 2018/098 (2018), <https://eprint.iacr.org/2018/098>
- Goubin, L., Patarin, J.: DES and Differential Power Analysis (The "Duplication" Method). In: Cryptographic Hardware and Embedded Systems, First International Workshop, CHES'99. pp. 158–172 (1999). https://doi.org/10.1007/3-540-48059-5_15

17. Karroumi, M.: Protecting White-Box AES with Dual Ciphers. In: Information Security and Cryptology - ICISC 2010. pp. 278–291 (2010). https://doi.org/10.1007/978-3-642-24209-0_19
18. Kerckhoffs, A.: La Cryptographic Militaire. Journal des Sciences Militaires pp. 5–38 (1883)
19. Lepoint, T., Rivain, M., Mulder, Y.D., Roelse, P., Preneel, B.: Two Attacks on a White-Box AES Implementation. In: Selected Areas in Cryptography - SAC 2013. pp. 265–285 (2013). https://doi.org/10.1007/978-3-662-43414-7_14
20. Mangard, S., Oswald, E., Popp, T.: Power analysis attacks - revealing the secrets of smart cards. Springer (2007)
21. Michiels, W., Gorissen, P., Hollmann, H.D.L.: Cryptanalysis of a Generic Class of White-Box Implementations. In: Selected Areas in Cryptography, 15th International Workshop, SAC 2008. pp. 414–428 (2008). https://doi.org/10.1007/978-3-642-04159-4_27
22. Mulder, Y.D., Roelse, P., Preneel, B.: Cryptanalysis of the Xiao - Lai White-Box AES Implementation. In: Selected Areas in Cryptography, 19th International Conference, SAC 2012. pp. 34–49 (2012). https://doi.org/10.1007/978-3-642-35999-6_3
23. Mulder, Y.D., Wyseur, B., Preneel, B.: Cryptanalysis of a Perturbed White-Box AES Implementation. In: Progress in Cryptology - INDOCRYPT 2010. pp. 292–310 (2010). https://doi.org/10.1007/978-3-642-17401-8_21
24. Rivain, M., Prouff, E.: Provably Secure Higher-Order Masking of AES. In: Cryptographic Hardware and Embedded Systems, CHES 2010. pp. 413–427 (2010). https://doi.org/10.1007/978-3-642-15031-9_28
25. Rivain, M., Prouff, E., Doget, J.: Higher-Order Masking and Shuffling for Software Implementations of Block Ciphers. In: Cryptographic Hardware and Embedded Systems - CHES 2009. pp. 171–188 (2009). https://doi.org/10.1007/978-3-642-04138-9_13
26. Strobel, D., Paar, C.: An Efficient Method for Eliminating Random Delays in Power Traces of Embedded Software. In: Information Security and Cryptology - ICISC 2011. pp. 48–60 (2011). https://doi.org/10.1007/978-3-642-31912-9_4
27. Tolhuizen, L.: Improved cryptanalysis of an aes implementation. In: Proceedings of the 33rd WIC Symposium on Information Theory, 2012. WIC (Werkgemeenschap voor Inform.-en Communicatietheorie) (2012)
28. Veyrat-Charvillon, N., Medwed, M., Kerckhof, S., Standaert, F.: Shuffling against Side-Channel Attacks: A Comprehensive Study with Cautionary Note. In: Advances in Cryptology - ASIACRYPT 2012. pp. 740–757 (2012). https://doi.org/10.1007/978-3-642-34961-4_44
29. Xiao, Y., Lai, X.: A secure implementation of white-box aes. In: Computer Science and its Applications, 2009. CSA'09. pp. 1–6. IEEE (2009)

A Success probability of Higher-Order DCA (Proof of Theorem 1)

Consider a specific value w_j of the higher-order trace \mathbf{w} . Denote by \mathcal{A} the event that w_j corresponds to the combination of the correct shares. The probability of \mathcal{A} occurring, *i.e.* of choosing the correct d shares out of the t elements of the original computational trace \mathbf{v} , is $p = \binom{t}{d}^{-1}$.

Fix some plaintext and the corresponding trace. By the law of total probability, the probability that a value w_j of the d 'th order trace is equal to a prediction $s = \varphi(x, k)$ for some key guess k is

$$\Pr(w_j = s) = \Pr(w_j = s|\mathcal{A}) \cdot \Pr(\mathcal{A}) + \Pr(w_j = s|\neg\mathcal{A}) \cdot \Pr(\neg\mathcal{A}).$$

For a wrong key guess, $k^\times \neq k^*$, $\Pr(w_j = s^\times|\mathcal{A}) = 0$, while for a right key guess $\Pr(w_j = s^*|\mathcal{A}) = 1$. In both cases, we have $\Pr(w_j = s|\neg\mathcal{A}) = 1/|\mathcal{K}|$. In total:

$$\begin{aligned} p^\times &= \Pr(w_j = s^\times) = (1 - p)/|\mathcal{K}|, \\ p^* &= \Pr(w_j = s^*) = p + (1 - p)/|\mathcal{K}|. \end{aligned}$$

Thus, for N traces,

$$C_{k^\times}(\mathbf{v}_{\phi(j)}, (x_i)_i) \sim \text{Bin}(N, p^\times)$$

for a wrong key guess, and

$$C_{k^*}(\mathbf{v}_{\phi(j)}, (x_i)_i) \sim \text{Bin}(N, p^*)$$

for a right key guess. Note that $|\mathbf{w}| = \binom{t}{d}$. Let $X_1, \dots, X_{|\mathbf{w}|}$ be distributed as $\text{Bin}(N, p^\times)$. Then $\gamma_{k^\times} \sim \max X_i$, and we denote the CDF by $F_{\max}^\times(x)$. If the X_i were independent, we would have

$$F_{\max}^\times(x) = F(x; N, p^\times)^{\binom{t}{d}}.$$

While the X_i are pairwise independent, they are not mutually independent. However, we find that in practice, the dependence is so weak that γ_{k^\times} approximately has CDF F_{\max}^\times , even for small values of $|\mathbf{w}|$ and N . We define $F_{\max}^*(x)$ similarly.

The attack is successful if $\gamma_{k^*} > \gamma_{k^\times}$ for all k^\times . As there are $|\mathcal{K}| - 1$ wrong keys, and all γ_{k^\times} are independent and identically distributed, we have $p_{\text{succ}} = \Pr(\gamma_{k^*} > \gamma_{k^\times})^{|\mathcal{K}|-1}$, where

$$\Pr(\gamma_{k^*} > \gamma_{k^\times}) = \sum_{i=0}^N (F_{\max}^*(i) - F_{\max}^*(i-1)) \cdot F_{\max}^\times(i-1).$$

which concludes the proof.

B Proof of Proposition 1

Proof. By applying the Bayes' rule, one gets (we skip random variables for the sake of clarity):

$$\Pr(k \mid (\mathbf{v}_i)_i \wedge (x_i)_i) = \frac{\Pr((\mathbf{v}_i)_i \mid k \wedge (x_i)_i) \cdot \Pr(k \wedge (x_i)_i)}{\Pr((\mathbf{v}_i)_i \wedge (x_i)_i)} \quad (3)$$

By mutual independence of the X_i 's and K , we have $\Pr(k \wedge (x_i)_i) = \frac{1}{|\mathcal{K}|} \left(\frac{1}{|\mathcal{X}|}\right)^N$ for every $k \in \mathcal{K}$. Moreover, $\Pr((\mathbf{v}_i)_i \wedge (x_i)_i)$ is constant with respect to k . We hence get

$$\Pr(k \mid (\mathbf{v}_i)_i \wedge (x_i)_i) \propto \Pr((\mathbf{v}_i)_i \mid k \wedge (x_i)_i). \quad (4)$$

By mutual independence of the \mathbf{V}_i 's and the X_i 's we further deduce

$$\Pr((\mathbf{v}_i)_i \mid k \wedge (x_i)_i) = \prod_{i=1}^N \Pr(\mathbf{v}_i \mid k \wedge x_i). \quad (5)$$

For the sake of simplicity we skip the index i in the following. By the law of total probability, we have

$$\Pr(\mathbf{v} \mid k \wedge x) = \sum_{\phi(j)} \Pr(\mathcal{S}_{\phi(j)}) \cdot \Pr(\mathbf{v} \mid k \wedge x \wedge \mathcal{S}_{\phi(j)}), \quad (6)$$

where $\mathcal{S}_{\phi(j)}$ denotes the event that the set $\phi(j)$ is selected for the sharing of $\varphi(X, K)$. By definition, we have

$$\Pr(\mathcal{S}_{\phi(j)}) = \frac{1}{\binom{t}{d}} \quad (7)$$

and

$$\Pr(\mathbf{v} \mid k \wedge x \wedge \mathcal{S}_{\phi(j)}) = \begin{cases} \left(\frac{1}{|\mathcal{V}|}\right)^{t-1} & \text{if } \bigoplus_{l \in \phi(j)} v_l = \varphi(x, k) \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

which finally gives

$$\Pr(\mathbf{v} \mid k \wedge x) \propto C_k(\mathbf{v}, x). \quad (9)$$

C Probability of the Zero-Counter Event (Proof of Lemma 1)

We first define \mathcal{Z}_k as the zero-counter event for key k for a *single* computational trace \mathbf{V} . Formally,

$$\mathcal{Z}_k = \left\{ \forall j \subseteq \left\{ 1, \dots, \binom{t}{d} \right\} : \bigoplus_{i \in \phi(j)} V_i \neq \varphi(X, k) \right\}.$$

The zero-counter event \mathcal{Z}_k occurs if and only if none of the $q = \binom{t}{d}$ combinations $\bigoplus_{i \in \phi(j)} V_i$ match the predicted value $\varphi(X, k)$. As discussed, \mathcal{Z}_{k^*} never occurs for the correct key guess k^* . For the incorrect key guess k^\times , intuitively, the zero-counter probability $\Pr(\mathcal{Z}_{k^\times})$ should quickly become negligible as the number of combinations q grows. While all q combinations are not strictly independent, we can approximate the probability of \mathcal{Z}_{k^\times} by:

$$\Pr(\mathcal{Z}_{k^\times}) \approx \left(1 - \frac{1}{|\mathcal{V}|} \right)^q. \quad (10)$$

We verified this approximation by estimating the zero-counter probability over some sampled computation traces. As illustrated in Table 3, the obtained estimations match the approximation pretty well.

Table 3. Approximation and estimation of the zero-counter probability.

(t, d)	(16,2)	(16,3)	(16,4)	(24,2)	(24,3)	(32,2)	(32,3)
Approximation (10)	0.625	0.112	$8 \cdot 10^{-4}$	0.340	$4 \cdot 10^{-4}$	0.144	$4 \cdot 10^{-9}$
Estimation (prec. $\sim 10^{-3}$)	0.628	0.135	$< 10^{-3}$	0.342	$< 10^{-3}$	0.145	$< 10^{-3}$

Then, by definition, the zero-counter event for N traces is the union

$$\mathcal{U}_k = \mathcal{Z}_k^{(1)} \vee \mathcal{Z}_k^{(2)} \vee \dots \vee \mathcal{Z}_k^{(N)},$$

where $\mathcal{Z}_k^{(i)}$ denotes the zero-counter event for k on trace \mathbf{V}_i . Taking the negation we obtain $\neg \mathcal{U}_{k^\times} = (\neg \mathcal{Z}_{k^\times}^{(1)}) \wedge (\neg \mathcal{Z}_{k^\times}^{(2)}) \wedge \dots \wedge (\neg \mathcal{Z}_{k^\times}^{(N)})$, and since the zero events $\mathcal{Z}_{k^\times}^{(i)}$ are mutually independent, we get

$$\Pr(\mathcal{U}_{k^\times}) = 1 - \prod_{i=1}^N \Pr(\neg \mathcal{Z}_{k^\times}^{(i)}) = 1 - (1 - \Pr(\mathcal{Z}_{k^\times}))^N.$$

This finishes the proof of Lemma 1.

D Success Probability with no Zero Counters (Proof of Lemma 2)

If the zero counter event does not occur, we can think of each trace \mathbf{V}_i as a random variable uniformly distributed over \mathcal{V}^t . Since the public input X_i is also random, the counters $C_k(\mathbf{V}, X)$ follow some probability distribution. In order to prove Lemma 2, we first prove the following result regarding these distributions.

Lemma 3. *Let k^* and k^\times be a right and wrong key guess. Let $q = \binom{t}{d}$ and $\kappa = (q-1) \frac{1}{|\mathcal{V}|}$. Then for a trace of length t and a d 'th-order attack,*

$$C_{k^*}(\mathbf{V}, X) \sim \mathcal{N}(\kappa + 1, \kappa) \quad \text{and} \quad C_{k^\times}(\mathbf{V}, X) \sim \mathcal{N}(\kappa, \kappa),$$

where $\mathcal{N}(\mu, \sigma^2)$ denotes the normal distribution with mean μ and variance σ^2 .

Proof. Let $\delta : \mathcal{V}^2 \rightarrow \{0, 1\}$ be the function defined as

$$\delta(v_1, v_2) = \begin{cases} 1 & \text{if } v_1 = v_2, \\ 0 & \text{otherwise.} \end{cases}$$

The counter $C_k(\mathbf{V}, X)$ can be rewritten as a sum $C_k(\mathbf{V}, X) = \sum_{j=1}^q \delta(W_j, \varphi(X, k))$, where the variables $(W_j)_j$ are defined as the $q = \binom{t}{d}$ combinations $\bigoplus_{i \in \phi(j)} V_i$. We recall that for one index j we have $W_j = \varphi(X, k^*)$, whereas for the other indices the W_j are randomly distributed independently of X . The counter expectation then satisfies

$$\mathbb{E}(C_k(\mathbf{V}, X)) = \sum_{j=1}^q \mathbb{E}(\delta(W_j, \varphi(X, k))) = \begin{cases} (q-1) \frac{1}{|\mathcal{V}|} & \text{if } k \neq k^*, \\ (q-1) \frac{1}{|\mathcal{V}|} + 1 & \text{if } k = k^*. \end{cases}$$

On the other hand, the counter variance can be expressed as:

$$\begin{aligned} \text{Var}(C_k(\mathbf{V}, X)) &= \sum_{j=1}^q \text{Var}(\delta(W_j, \varphi(X, k))) \\ &\quad + 2 \sum_{1 \leq j < j' \leq q} \text{Cov}(\delta(W_j, \varphi(X, k)), \delta(W_{j'}, \varphi(X, k))). \end{aligned}$$

It can be checked that the covariances will be equal to 0 most of the time. Indeed, the covariances are non-zero only when $W_j \oplus W_{j'} = \varphi(X, k^*)$, which never happens when d is odd and which happens for few pairs (j, j') when d is even. Therefore these covariance terms will only have a small impact on the overall variance. Moreover, it can be checked that this impact is negative, *i.e.* it reduces the variance.⁷ Therefore we will ignore the sum of covariances, which yields a correct result when d is odd and a slight overestimation when d is even. We then have

$$\text{Var}(\delta(W_j, \varphi(X, k))) = \begin{cases} \frac{1}{|\mathcal{V}|} \left(1 - \frac{1}{|\mathcal{V}|}\right) & \text{if } j \neq j^*, \\ 0 & \text{if } j = j^*, \end{cases}$$

where j^* denotes the index of the right combination matching $\varphi(X, k^*)$. Combining the two above equations gives:

$$\text{Var}(C_k(\mathbf{V}, X)) = (q-1) \frac{1}{|\mathcal{V}|} \left(1 - \frac{1}{|\mathcal{V}|}\right) \approx (q-1) \frac{1}{|\mathcal{V}|}.$$

Since the counter is defined as a sum of *somewhat independent* random variables, we can soundly approximate its distribution by a Gaussian, and setting $\kappa = (q-1) \frac{1}{|\mathcal{V}|}$ concludes the proof. \square

In the above proof, we use that the $\delta(W_j, \varphi(X, k))$ are somewhat independent. By *somewhat independent* we mean that these variables are pairwise independent (for most or all of them, as discussed). Note that variants of the central limit theorem exist that take some form of dependence between the summed variables into account. We have experimentally verified that the Gaussian approximation is sound for various parameters (t, d) .

Using Lemma 3, we can now prove Lemma 2. Following Remark 1, we will focus on the log-likelihood, *i.e.* we consider

$$\begin{aligned} \Pr(\ell_{k^*} > \ell_{k^\times} \mid \neg \mathcal{U}_{k^\times}) &= \Pr(\log \ell_{k^*} - \log \ell_{k^\times} > 0 \mid \neg \mathcal{U}_{k^\times}), \\ \log \ell_{k^*} - \log \ell_{k^\times} &= \sum_{i=1}^N \underbrace{\log C_{k^*}(\mathbf{V}_i, X_i) - \log C_{k^\times}(\mathbf{V}_i, X_i)}_{Y_i}. \end{aligned}$$

⁷ Most of the time we have $\varphi(X, k^*) \neq 0$ so that the pairs (j, j') with $W_j \oplus W_{j'} = \varphi(X, k^*)$ are such that $W_j \neq W_{j'}$ with high probability. In that case $\delta(W_j, \varphi(X, k)) = 1$ implies $\delta(W_{j'}, \varphi(X, k)) = 0$ and conversely which yields a negative covariance.

As introduced above, we denote by Y_i the difference between the log-counters for the trace \mathbf{V}_i . Since the Y_i are mutually independent and identically distributed, the central limit theorem implies that, for N sufficiently large,

$$\frac{1}{N}(\log \ell_{k^*} - \log \ell_{k^\times}) \sim \mathcal{N}(\mu_Y, \sigma_Y^2 N^{-1}) \quad \text{with} \quad \begin{cases} \mu_Y = \mathbb{E}(Y), \\ \sigma_Y^2 = \text{Var}(Y), \end{cases}$$

for $Y = \log C_{k^*}(\mathbf{V}, X) - \log C_{k^\times}(\mathbf{V}, X)$. Thus

$$\Pr(\ell_{k^*} > \ell_{k^\times} \mid \neg \mathcal{U}_{k^\times}) = 1 - \Phi_{\mu_Y, \sigma_Y^2/N}(0) = \frac{1}{2} + \frac{1}{2} \text{erf}\left(\frac{\sqrt{N} \mu_Y}{\sqrt{2} \sigma_Y}\right), \quad (11)$$

where $\Phi_{\mu, \sigma}$ is the CDF of $\mathcal{N}(\mu, \sigma^2)$. By the heuristic assumption that $C_{k^*}(\mathbf{V}, X)$ and $C_{k^\times}(\mathbf{V}, X)$ are mutually independent, and using the Taylor expansion of the logarithm at $\mathbb{E}(C)$, as well as [Lemma 3](#), we have

$$\mu_Y \approx \log(\kappa + 1) - \frac{\kappa}{2(\kappa + 1)^2} - \log \kappa + \frac{\kappa}{2\kappa^2} \approx \frac{1}{\kappa}, \quad \text{and} \quad \sigma_Y^2 \approx 2 \frac{\kappa}{\kappa^2} = \frac{2}{\kappa},$$

where the approximation of the mean is sound if κ is large enough (*e.g.* $\kappa > 10$). Inserting these approximations into [Equation 11](#), remembering that $\kappa = (q - 1) \frac{1}{|\mathcal{V}|} \approx \frac{q}{|\mathcal{V}|}$, finishes the proof.