

Making Masking Security Proofs Concrete

Or How to Evaluate the Security of any Leaking Device

Alexandre Duc¹, Sebastian Faust^{1,2}, François-Xavier Standaert³

¹ EPFL, Lausanne, Switzerland

{alexandre.duc,sebastian.faust}@epfl.ch

² Horst Görtz Institute, Ruhr-University Bochum, Germany

³ ICTEAM/ELEN/Crypto Group, Université catholique de Louvain, Belgium
fstandae@uclouvain.be

Abstract. We investigate the relationships between theoretical studies of leaking cryptographic devices and concrete security evaluations with standard side-channel attacks. Our contributions are in four parts. First, we connect the formal analysis of the masking countermeasure proposed by Duc et al. (Eurocrypt 2014) with the Eurocrypt 2009 evaluation framework for side-channel key recovery attacks. In particular, we re-state their main proof for the masking countermeasure based on a mutual information metric, which is frequently used in concrete physical security evaluations. Second, we discuss the tightness of the Eurocrypt 2014 bounds based on experimental case studies. This allows us to conjecture a simplified link between the mutual information metric and the success rate of a side-channel adversary, ignoring technical parameters and proof artifacts. Third, we introduce heuristic (yet well-motivated) tools for the evaluation of the masking countermeasure when its independent leakage assumption is not perfectly fulfilled, as it is frequently encountered in practice. Thanks to these tools, we argue that masking with non-independent leakages may provide improved security levels in certain scenarios. Eventually, we consider the tradeoff between measurement complexity and key enumeration in divide-and-conquer side-channel attacks, and show that it can be predicted based on the mutual information metric, by solving a non-linear integer programming problem for which efficient solutions exist. The combination of these observations enables significant reductions of the evaluation costs for certification bodies.

1 Introduction

Side-channel attacks are an important concern for the security of cryptographic hardware, and masking is one of the most investigated solutions to counteract them. Its underlying principle is to randomize any sensitive data manipulated by a leaking implementation by splitting it into d shares, and to perform all the computations on these shared values only. Intuitively, such a process is expected to force the adversary to combine several leakages corresponding to the different shares in order to recover secret information. As a result, it has first been shown by Chari et al. that the measurement complexity of a specialized

attack – namely a single-bit Differential Power Analysis (DPA) [35] – against a carefully implemented masked computation (i.e. where the leakages of all the shares are independent and sufficiently noisy) increases exponentially with d [14]. Following this seminal work, a number of progresses have been made in order to state the security guarantee of masking in both general and rigorous terms. For example, Ishai, Sahai and Wagner introduced a compiler (next referred to as the ISW compiler), able to encode any circuit into an equivalent (secret-shared) one, and proved its security against so-called probing adversaries, able to read a bounded number of wires in the implementation [33]. A practical counterpart to these results was published at Asiacrypt 2010, where Standaert et al. analyzed the security of several masked implementations [61], using the information theoretic framework introduced in [60]. While this analysis was specialized to a few concrete case studies, it allowed confirming the exponential security increase provided by masking against actual leakages, typically made of a noisy but arbitrary function of the target device’s state. Following, Faust et al. attempted to analyze the ISW compiler against more realistic leakage functions, and succeeded to prove its security against computationally bounded (yet still unrealistic) ones, e.g. in the AC^0 complexity class [25]. Prouff and Rivain then made a complementary step towards bridging the gap between the theory and practice of masking schemes, by providing a formal information theoretic analysis of a wide (and realistic) class of so-called noisy leakage functions [49]. Eventually, Duc et al. turned this analysis into a simulation-based security proof, under standard conditions (i.e. chosen-message rather than random-message attacks, without leak-free components, and with reduced noise requirements) [22]. The central and fundamental ingredient of this last work was a reduction from the noisy leakage model of Prouff and Rivain to the probing model of Ishai et al.

Our contribution. In view of this state-of-the-art, one of the main remaining questions regarding the security of the masking countermeasure is whether its proofs can be helpful in the security evaluation of concrete devices. That is, can we state theorems for masking so that the hypotheses can be easily fulfilled by hardware designers, and the resulting guarantee is reflective of the actual security level of the target implementation. For this purpose, we first observe that the proofs in [22, 49] express their hypothesis for the amount of noise in the shares’ leakages based on a statistical distance. This is in contrast with the large body of published work where the mutual information metric introduced in [60] is estimated for various implementations (e.g. [4, 12, 27, 30, 32, 42, 50, 51, 57, 63, 66]). Since the latter metric generally carries more intuition (see, e.g. [3] in the context of linear cryptanalysis), and benefits from recent advances in leakage certification, allowing to make sure that its estimation is accurate and based on sound assumptions [23], we first provide a useful link between the statistical distance and mutual information, and also connect them with easy-to-interpret (but more specialized) tools such as the Signal-to-Noise Ratio (SNR). We then re-state the theorems of Duc et al. based on the mutual information metric in two relevant scenarios. Namely, we consider both the security of an idealized

implementation with a “leak-free refreshing” of the shares, and the one of a standard ISW-like encoding (i.e. capturing any type of leaking computation).

Interestingly, the implementation with leak-free refreshing corresponds to the frequently investigated (practical) context where a side-channel attack aims at key recovery, and only targets the d shares’ leakage of a so-called sensitive intermediate variable (i.e. that depends on the plaintext and key) [17]. So despite being less interesting from a theoretical point of view, this scenario allows us to compare the theorem bounds with concrete attacks. Taking advantage of this comparison, we discuss the bounds’ tightness and separate parameters that are physically motivated from more “technical ones” (that most likely result of proof artifacts). As a result, we conjecture a simplified link between the mutual information metric and the success rate of a side-channel adversary, which allows accurate approximations of the attacks’ measurement complexity at minimum (evaluation) cost. We further illustrate that the noise condition for masking has a simple and intuitive interpretation when stated in terms of SNR.

Next, we note that the published results about masking (including the previously mentioned theorems and conjecture) assume independence between the leakages corresponding to different shares in an implementation. Yet, concrete experiments have shown that small (or even large) deviations from this assumption frequently occur in practice (see, e.g. [5, 16, 41, 54]). Hence, we complete our discussion by providing sound heuristics to analyze the impact of “non-independent leakages” which allow, for the first time, to evaluate and predict the security level of a masked implementation in such imperfect conditions.

Eventually, we consider the tradeoff between measurement complexity and time complexity in the important context of divide-and-conquer attacks. Previously known approaches for this purpose were based on launching key enumeration and/or rank estimation algorithms for multiple attacks, and to average results to obtain a success rate [64, 65]. We provide an alternative solution, where success rates (possibly obtained from estimations of the mutual information metric) are estimated/bounded for all the target key bytes of the divide-and-conquer attack first, and the impact of enumeration is evaluated only once afterwards. We also connect the problem of approximating the enumeration cost for a given number of measurements with a non-linear integer programming problem, and provide simple heuristics to estimate bounds on this enumeration cost.

Summarizing, the combination of these observations highlights that the security evaluation of a masked implementation boils down to the estimation of the mutual information between its shares and their corresponding leakages. Incidentally, the tools introduced in this paper apply identically to unprotected implementations, or implementations protected with other countermeasures, as long as one can estimate the same mutual information metric for the target intermediate values. Therefore, our results clarify the long standing open question whether the (informal) link between information theoretic and security metrics in the Eurocrypt 2009 evaluation framework [60] can be proved formally. They also have important consequences for certification bodies, since they translate

the (worst-case) side-channel evaluation problem into the well-defined challenge of estimating a single metric, leading to significantly reduced evaluation costs.

Notations. We next use capital letters for random variables, small caps for their realizations and hats for estimations. Vectors will be denoted with bold notations, functions with sans serif fonts, and sets with calligraphic ones.

2 Background

2.1 Leakage traces and assumptions

Let y be a n -bit sensitive value manipulated by a leaking device. Typically, it could be the output of an S-box computation such that $y = \mathsf{S}(x \oplus k)$ with n -bit plaintext/key words x and k . Let y_1, y_2, \dots, y_d be the d shares representing y in a Boolean masking scheme (i.e. $y = y_1 \oplus y_2 \oplus \dots \oplus y_d$). In a side-channel attack, the adversary is provided with some information (aka leakage) on each share. Typically, this leakage takes the form of a random variable \mathbf{L}_{y_i} that is the output of a leakage function L with y_i and a noise variable \mathbf{R}_i as arguments:

$$\mathbf{L}_{y_i} = \mathsf{L}(y_i, \mathbf{R}_i) . \quad (1)$$

The top of Figure 1 represents a leakage trace corresponding to the manipulation of d shares. Concretely, each subtrace \mathbf{L}_{y_i} is a vector of which the elements represent time samples. Whenever accessing a single time sample t , we use the notation $\mathbf{L}_{y_i}^t = \mathsf{L}^t(y_i, \mathbf{R}_i^t)$. From this general setup, a number of assumptions are frequently used in the literature. We will consider the following three.

a. Selection of points-of-interest / dimensionality reduction. For convenience, a number of attacks start with a pre-processing in order to reduce each leakage subtrace \mathbf{L}_{y_i} to a scalar random variable L_{y_i} . Such a pre-processing is motivated both by popular side-channel distinguishers such as Correlation Power Analysis (CPA) [11], which can only deal with univariate data, and by the easier representation of small dimensional data spaces. In this respect, even distinguishers that naturally extend towards multivariate data (such as Template attacks (TA) [15], Linear Regression (LR) [58] or Mutual Information Analysis (MIA) [28]) generally benefit from some dimensionality reduction. This step can be achieved heuristically, by looking for leakage samples where one distinguisher works best, or more systematically using tools such as Principal Component Analysis (PCA) [2] or Linear Discriminant Analysis (LDA) [59]. An example of reduced leakage trace is represented at the bottom of Figure 1.

b. Additive noise. A standard assumption in the literature is to consider leakage functions made of a deterministic part $\mathsf{G}(y_i)$ and additive noise \mathbf{N}_i [40]:

$$\mathbf{L}_{y_i} = \mathsf{L}(y_i, \mathbf{R}_i) \approx \mathsf{G}(y_i) + \mathbf{N}_i . \quad (2)$$

For example, a typical setting is to assume reduced leakages to be approximately generated as the combination of a Hamming weight function (or some other simple function of the shares' bits [58]) with additive Gaussian noise.

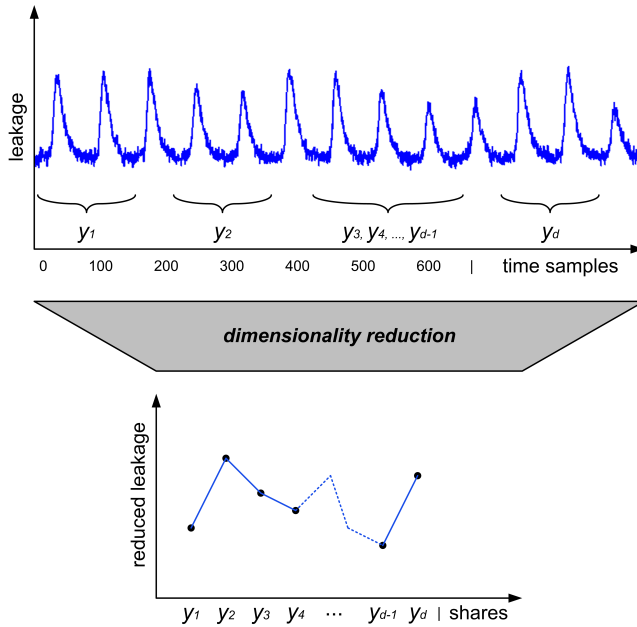


Fig. 1. Leakage trace & reduced leakage trace of a d -shared secret.

c. Independence condition. A secure implementation of the masking countermeasure requires that the leakage vectors \mathbf{L}_{y_i} are independent random variables. If respected, it implies that a d -share masking will lead to a $(d-1)^{\text{th}}$ -order secure implementation as defined in [17]. That is, it guarantees that every d -tuple of leakage vectors is independent of any sensitive variable. This means that any adversary targeting the implementation will have to “combine” the information of at least d shares, and that extracting information from these d shares will require to estimate a d^{th} -order moment of the leakage PDF (conditioned on a sensitive variable) – a task that becomes exponentially hard in d if the noise is sufficient. As witnessed by several prior works, this condition may be hard to fulfill in practice. In software implementations, it typically requires avoiding transition-based leakages (i.e. leakages that depend on the distance between shares rather than directly on the shares) [5, 16]. In hardware implementations, physical defaults such as glitches are another usual issue that can invalidate the independence assumption [41], which motivates various research efforts to mitigate this risk, both at the hardware level (e.g. [43]) and at the algorithmic level (e.g. [46]).

Note that only this last (independence) assumption is strictly needed for the following proofs of Section 3 to hold. By contrast, the previous assumptions (a) and (b) will be useful to provide practical intuition in Section 4. Furthermore, it is worth noting that slight deviations from this independence assumption (i.e. slight dependencies between the shares’ leakages) may still lead to concrete se-

curity improvements, despite falling outside the proofs’ formal guarantees. Such (practically meaningful) contexts will be further analyzed in Section 4.2.

2.2 Evaluation metrics

Following [60], one generally considers two types of evaluation metrics for leaking cryptographic devices. First, information theoretic metrics aim to capture the amount of information available in a side-channel, independent of the adversary exploiting it. Second, security metrics aim to quantify how this information can be exploited by some concrete adversary. As will be clear next, the two types of metrics are related. For example, in the context of standard DPA attacks [41], they both measure the prediction of the (true) leakage function with some model, the latter usually expressed as an estimation of the leakage Probability Density Function (PDF). Yet they differ since information theoretic metrics only depend on the leakage function and model, while security metrics also depend on the adversary’s computational power. For example, the capacity to enumerate key candidates may improve security metrics, but has no impact on information theoretic ones [64, 65]. Our goal in the following is to draw a formal connection between information theoretic and security metrics, i.e. between the amount of leakage provided by an implementation and its (worst-case) security level.

In the case of masking, proofs informally state that “*given that the leakage of each share is independent of each other and sufficiently noisy, the security of the implementation increases exponentially in the number of shares*”. So we need the two types of metrics to quantify the noise condition and security level.

b. Metrics to quantify the noise condition. In general (i.e. without assumptions on the leakage distribution), the noise condition on the shares can be expressed with an information theoretic metric. The Mutual Information (MI) advocated in [60] is the most frequently used candidate for this purpose:

$$\text{MI}(Y_i; \mathbf{L}_{Y_i}) = H[Y_i] + \sum_{y_i \in \mathcal{Y}} \Pr[y_i] \cdot \sum_{\mathbf{l}_{y_i} \in \mathcal{L}} \Pr[\mathbf{l}_{y_i} | y_i] \cdot \log_2 \Pr[y_i | \mathbf{l}_{y_i}], \quad (3)$$

where we use the notation $\Pr[Y_i = y_i] =: \Pr[y_i]$ when clear from the context. Note that whenever trying to compute this quantity from an actual implementation, evaluators face the problem that the leakage PDF is unknown and can only be sampled and estimated. As a result, one then computes the Perceived Information (PI), which is the evaluator’s best estimate of the MI [54]:

$$\hat{\text{PI}}(Y_i; \mathbf{L}_{Y_i}) = H[Y_i] + \sum_{y_i \in \mathcal{Y}} \Pr[y_i] \cdot \sum_{\mathbf{l}_{y_i} \in \mathcal{L}} \Pr_{\text{chip}}[\mathbf{l}_{y_i} | y_i] \cdot \log_2 \hat{\Pr}_{\text{model}}[y_i | \mathbf{l}_{y_i}], \quad (4)$$

with \Pr_{chip} the true chip distribution that can only be sampled and $\hat{\Pr}_{\text{model}}$ the adversary’s estimated model. For simplicity, we will ignore this issue and use the MI in our discussions (conclusions would be identical with the PI).

Interestingly, when additionally considering reduced leakages with additive Gaussian noise, and restricting the evaluation to so-called “first-order information” (i.e. information lying in the first-order statistical moments of the leakage PDF, which is typically the case for the leakage of each share), simpler metrics can be considered [40]. For example, the SNR introduced by Mangard at CT-RSA 2004 in [38] is of particular interest for our following discussions:

$$\text{SNR} = \frac{\hat{\text{v}}\text{ar}_{Y_i} \left(\hat{\text{E}}_{n_i}(L_{Y_i}) \right)}{\hat{\text{E}}_{Y_i} \left(\hat{\text{v}}\text{ar}_{n_i}(L_{Y_i}) \right)}, \quad (5)$$

where $\hat{\text{E}}$ is the sample mean operator and $\hat{\text{v}}\text{ar}$ is the sample variance. Summarizing, stating the noise condition based on the MI metric is more general (as it can capture any leakage PDF). By contrast, the SNR provides a simpler and more intuitive condition in a more specific but practically relevant context.

Eventually, the previous works of Prouff–Rivain and Duc et al. [22, 49] consider the following Statistical Distance (SD) to state their noise condition:

$$\text{SD}(Y_i; Y_i | \mathbf{L}_{Y_i}) = \sum_{\mathbf{l}_{y_i} \in \mathcal{L}} \text{Pr}[\mathbf{l}_{y_i}] \cdot \mathbf{d}(Y_i; Y_i | \mathbf{l}_{y_i}), \quad (6)$$

with \mathbf{d} the Euclidean norm in [49] and $\mathbf{d}(X_1, X_2) = \frac{1}{2} \sum_{x \in \mathcal{X}} |\text{Pr}[X_1 = x] - \text{Pr}[X_2 = x]|$ in [22]. In their terminology, a leakage function \mathbf{L} is then called “ δ -noisy” if $\delta = \text{SD}(Y_i; Y_i | \mathbf{L}_{Y_i})$, which was useful to connect different leakage models.

As previously mentioned, some of these metrics can be related under certain conditions. For example, in the context of univariate Gaussian random variables, the MI can be approximated from Pearson’s correlation coefficient [40], which was also connected to the SNR by Mangard [38]. The combination of those links corresponds to the classical MI bound in Cover and Thomas [19]:

$$\text{MI}(Y_i; \mathbf{L}_{Y_i}) \approx -\frac{1}{2} \log \left(1 - \left(\frac{1}{\sqrt{1 + \frac{1}{\text{SNR}}}} \right)^2 \right) \leq \frac{1}{2} \log (1 + \text{SNR}). \quad (7)$$

In Section 3.1, we show that the MI and SD metrics can be connected as well.

c. Metrics to quantify the security result. Quantifying security requires defining the adversary’s goal. Current side-channel attacks published in the literature mostly focus on key recovery. In this context, one can easily evaluate the exploitation of the leakages with the success rate defined in [60], i.e. the probability that an adversary recovers the key given the observation of some (typically known or chosen) plaintexts, ciphertexts and leakages. We will next denote it with SR^{kr} . Key recovery is a weak security notion from a cryptographic point of view. As a result, rigorous proofs for masking such as the one of Duc et al. in [22] rather define security using the standard real/ideal world paradigm, which consider two settings: the ideal world where the adversary attacks the algorithm of a

cryptographic scheme in a black-box way, and the real world where he additionally obtains leakages. A scheme is said to be secure in the real world, if for any adversary in the real world there exists an adversary in the ideal world. In other words: any attack that can be carried out given the leakages can also be carried out in a black-box manner. A proof of security usually involves constructing an efficient simulator that is able to simulate the leakages just giving black-box access to the attacked cryptographic scheme. Whenever considering this (standard) indistinguishability-based security notion, we will denote the adversary’s success probability of distinguishing the two worlds with SR^{dist} .

3 Making proofs concrete: theory

In this section, we discuss theoretical tweaks allowing to improve the concreteness of masking proofs. For this purpose, we recall three important leakage models that are relevant for our work. First, the t -probing and ϵ -probing (aka random probing) models were introduced in [33]. In the former one, the adversary obtains t intermediate values of the computation (e.g. can probe t wires if we compute in binary fields). In the latter one, he obtains each of these intermediate values with probability ϵ , and gets \perp with probability $1 - \epsilon$ (where \perp means no knowledge). Using a Chernoff-bound it is easy to show that security in the t -probing model reduces to security in the ϵ -probing model for certain values of ϵ . Second, the noisy leakage model describes many realistic side-channel attacks and allows an adversary to obtain each intermediate value perturbed with a δ -noisy leakage function L [49]. As mentioned in the previous section, a leakage function L is called δ -noisy if for a uniformly random variable Y (over the field \mathbb{F}) we have $\text{SD}(Y; Y|L_Y) \leq \delta$. In contrast with the conceptually simpler ϵ -probing model, the adversary obtains noisy leakages on each intermediate variable. For example, in the context of masking, he obtains $L(Y_i, \mathbf{R})$ for all the shares Y_i , which is more reflective of actual implementations where the adversary can potentially observe the leakage of all these shares, since they are all present in leakage traces such as in Figure 1. Recently, Duc et al. showed that security against probing attacks implies security against noisy leakages (up to a factor $|\mathbb{F}|$, where \mathbb{F} is the underlying field in which the operations are carried out) [22]. In the rest of this section, we first connect the statistical distance SD with the mutual information metric MI , which shows that both can be used to quantify the noise condition required for masking. Next, we provide alternative forms for the theorems of Duc et al. and show (i) the security of the encoding used in (e.g. Boolean) masking and (ii) the security of a complete circuit based on the ISW compiler.

3.1 From statistical distance to MI

The results from Duc et al. require to have a bound on the SD between the shares and the shares given the leakage. For different reasons, expressing this distance based on the MI metric may be more convenient in practice (as witnessed by the numerous works where this metric has been computed, for various types

of devices, countermeasures and technologies – see the list in introduction). For example, the MI metric is useful to determine whether the leakage model used in a standard DPA is sound (see the discussion in Section 4.1) and for analyzing the impact of key enumeration in divide-and-conquer attacks (see the discussion in Section 4.3). Very concretely, Equations (3) and (4) are also expressed in a way that requires summing over the intermediate values first and on the leakages afterwards, which corresponds to the way security evaluations are performed (i.e. fix the target device’s state, and then perform measurements). Thus, we now show how to express the SD in function of the MI. We use a previous result from Dodis [21], which proofs follows [9] that we rephrase with our notations.

Lemma 1 ([21], Lemma 6). *Let Y_i and \mathbf{L}_{Y_i} be two random variables. Then:*

$$\frac{1}{2} \left(\sum_{(y \in \mathcal{Y}, \ell \in \mathcal{L})} |\Pr[Y_i = y, \mathbf{L}_{Y_i} = \ell] - \Pr[Y_i = y] \Pr[\mathbf{L}_{Y_i} = \ell]| \right)^2 \leq \text{MI}(Y_i; \mathbf{L}_{Y_i}) .$$

Using this lemma, we can now express the SD in function of the MI as follows.

Theorem 1. *Let Y_i and \mathbf{L}_{Y_i} be two random variables. Then:*

$$2 \cdot \text{SD}(Y_i; Y_i \mid \mathbf{L}_{Y_i})^2 \leq \text{MI}(Y_i; \mathbf{L}_{Y_i}) .$$

Proof. The proof follows the proof of [8], Lemma 4.4. We have:

$$\begin{aligned} & \sum_{(y \in \mathcal{Y}, \ell \in \mathcal{L})} |\Pr[Y_i = y, \mathbf{L}_{Y_i} = \ell] - \Pr[Y_i = y] \Pr[\mathbf{L}_{Y_i} = \ell]|, \\ &= \sum_{\ell \in \mathcal{L}} \Pr[\mathbf{L}_{Y_i} = \ell] \sum_{y \in \mathcal{Y}} |\Pr[Y_i = y \mid \mathbf{L}_{Y_i} = \ell] - \Pr[Y_i = y]|, \\ &= 2 \cdot \text{SD}(Y_i; Y_i \mid \mathbf{L}_{Y_i}) . \end{aligned}$$

The final result directly derives from Lemma 1. □

3.2 Security of the encoding

In this section, we analyze the security of an encoding when m measurements are performed and the encoding is refreshed between each measurements using a leak-free gate. More precisely, we assume that a secret y is secret-shared into d shares y_1, \dots, y_d , using an additive masking scheme over a finite field \mathbb{F} . Between each measurement, we assume that we take fresh y_1, \dots, y_d values such that $y = y_1 + \dots + y_d$ (e.g. it could be the Boolean encoding of Section 2.1). We also assume that this refreshing process does not leak and first recall a previous result from [22] that relates the random probing model to the noisy model. For conciseness, we call an adversary in the random-probing model a “random-probing adversary”, an adversary in the δ -noisy model a “ δ -noisy adversary”, and an adversary having access to leakages such that $\text{MI}(Y; Y \mid \mathbf{L}_Y) \leq \delta$ a “ δ -MI-adversary”. However, note that the physical noise (and its quantification with the MI) is a property of the implementation rather than of the adversary.

Lemma 2 ([22], Lemma 3). *Let \mathcal{A} be a δ -noisy adversary on \mathbb{F}^d . Then, there exists a $\delta \cdot |\mathbb{F}|$ -random-probing adversary \mathcal{S} on \mathbb{F}^d such that for every (y_1, \dots, y_d) , \mathcal{A} and \mathcal{S} produce the same view when applied on (y_1, \dots, y_d) .*

This result enables us to work directly in the random-probing model instead of the noisy leakage model. Next, we study the security of the encoding. As mentioned in introduction, the adversary’s goal in this case is to recover the encoded value, which is equivalent to key recovery if this value is a key. In order to make it completely comparable with actual attacks, we also add the number of measurements m used by the adversary as a parameter in our bounds.

Theorem 2. *Let d be the number of shares used for a key encoding, m be the number of measurements, and $\text{MI}(Y_i, \mathbf{L}_{Y_i}) \leq t$ for some $t \leq 2/|\mathbb{F}|^2$. Then, if we refresh the encoding in a leak-free manner between each measurement, the probability of success of a key recovery adversary under independent leakage is:*

$$\text{SR}^{\text{kr}} \leq 1 - \left(1 - \left(|\mathbb{F}| \sqrt{t/2} \right)^d \right)^m . \quad (8)$$

Proof. In the random probing model with parameter ϵ , an adversary learns nothing about the secret if there is at least one share that did not leak. Since all the measurements are independent and we use leak-free refreshing gates, we have:

$$\text{SR}^{\text{kr}} \leq 1 - (1 - \epsilon^d)^m . \quad (9)$$

Let \mathbf{A} be a t -MI-adversary on \mathbb{F}^d . From Theorem 1, we know that \mathbf{A} implies a $\sqrt{t/2}$ -noisy-adversary on \mathbb{F}^d and, by Lemma 2, we obtain a $|\mathbb{F}| \sqrt{t/2}$ -random-probing adversary on \mathbb{F}^d . Letting $\epsilon := |\mathbb{F}| \sqrt{t/2}$ in (9) gives us the result. \square

Note that Equation (9) focuses on the impact of the adversary’s measurement complexity m on the success rate, which is usually the dominating factor in concrete side-channel analyses. Yet, the impact of time complexity when considering key enumeration will be discussed in Section 4.3. Besides and for readability, this equation only includes the terms corresponding to attacks taking advantage of the leakages. We ignore the additional terms corresponding to mathematical cryptanalysis (e.g. exhaustive search) that should be added for completeness. In order to allow us comparing this result with the case where we study the security of a complete circuit encoded with the ISW compiler, we also write our result according to the following corollary (which is less general than Theorem 2).

Corollary 1. *Let d be the number of shares used for a key encoding and m the number of measurements. Then, if we refresh the encoding in leak-free manner between each measurement and for any $\alpha > 0$, the probability of success of a key recovery adversary under independent leakage is:*

$$\text{SR}^{\text{kr}} \leq m \cdot \exp(-\alpha d) , \quad (10)$$

if we have:

$$\text{MI}(Y_i; \mathbf{L}_{Y_i}) \leq 2 \left(\frac{1}{e^\alpha |\mathbb{F}|} \right)^2 . \quad (11)$$

Proof. We have:

$$1 - (1 - \epsilon^d)^m \leq m\epsilon^{\log(\epsilon)d}.$$

We want $\log(\epsilon) = -\alpha$. Hence, from Theorem 2, we get our result. \square

3.3 Security of the whole circuit

In this section, we restate the theorems from Duc et al. when securing a whole circuit with the seminal ISW compiler. The main theorem from [22] bounds the probability of success of a distinguishing adversary in the noisy leakage model. We provide an alternative version of their theorem and, as in the previous section, we relate it to the mutual information instead of the statistical distance.

Theorem 3. *Suppose that we have a circuit of size $|\Gamma|$ protected with the ISW compiler with d shares. Then, the probability of success of a distinguishing adversary under independent leakage is:*

$$\text{SR}^{\text{dist}} \leq |\Gamma| \cdot \exp\left(-\frac{d}{12}\right) = |\Gamma| \cdot 2^{\left(-\frac{d \cdot \log_2(\epsilon)}{12}\right)} \leq |\Gamma| \cdot 2^{-d/9}, \quad (12)$$

if we have:

$$\text{MI}(Y_i; \mathbf{L}_{Y_i}) \leq 2 \cdot \left(\frac{1}{|\Gamma| \cdot (28d + 16)}\right)^2. \quad (13)$$

Similarly to what we did in the previous section, we also write this corollary.

Corollary 2. *Suppose that we have a circuit of size $|\Gamma|$ protected with the ISW compiler with d shares. Then, if $\text{MI}(Y_i, \mathbf{L}_{Y_i}) \leq t$, a distinguisher adversary under independent leakage needs:*

$$d \geq \frac{1 - 16|\Gamma|\sqrt{\frac{1}{2}t}}{28|\Gamma|\sqrt{\frac{1}{2}t}} \quad (14)$$

shares in order to obtain:

$$\text{SR}^{\text{dist}} \leq |\Gamma| \cdot \exp\left(-\frac{d}{12}\right) \leq |\Gamma| \cdot \exp\left(-\frac{1 - 16|\Gamma|\sqrt{\frac{1}{2}t}}{336|\Gamma|\sqrt{\frac{1}{2}t}}\right). \quad (15)$$

Note that the ISW compiler can actually be used to efficiently compute any circuit. For example, the work of Rivain and Prouff at CHES 2010 showed how to adapt the compiler to $|F| = 256$ which leads to efficient masked implementations of the AES [56] (see also various following works such as [13, 18, 31, 57]).

4 Making proofs concrete: practice

In this section, we complement the previous theoretical results with an experimental analysis. Our contributions are threefold. First, we provide an empirical evaluation of the encoding scheme in Section 3.2, which allows us to discuss the noise condition and tightness of the bounds in our proofs. We use this discussion to conjecture a simple connection between the mutual information metric and the success rate of a (worst-case) side-channel adversary, and argue that it can lead to quite accurate approximations of the attacks’ measurement complexity. Next, we discuss possible deviations from the independent leakage assumption and provide tools allowing one to approximate the security level of concrete devices in such cases. Eventually, we consider the tradeoff between measurement complexity and time complexity in the context of divide-and-conquer side-channel attacks. We show how one can build a side-channel security graph (i.e. a plot of the adversary’s success probability bounds in function of both parameters [65]), based only on the estimation of the MI metric for each share of a masking scheme. Along these lines, we eventually provide a formal justification for the physical security evaluation framework proposed at Eurocrypt 2009 [60].

4.1 Experimental validation

In order to discuss the relevance of the proofs in the previous section, we take the (usual) context of standard DPA attacks defined in [40]. More precisely, we consider the simple case where an adversary targets a single S-box from a block cipher (e.g. the AES) as specified in Section 2.1, and obtains leakage variables $\mathbf{L}_{y_i} = \mathbf{L}(y_i, \mathbf{R}_i)$ for $1 \leq i \leq d$ (the case of multiple S-boxes will be studied in Section 4.3). For convenience, we mainly consider the context of mathematically-generated Gaussian Hamming weight leakages, where $\mathbf{L}_{y_i} = \text{HW}(y_i) + N_i$, with HW the Hamming weight function and N_i a Gaussian-distributed noise, with variance σ^2 . In this respect, we note that we did not mount concrete attacks since we would have had to measure hundreds of different implementations to observe useful trends in practice. Our experiments indeed correspond to hundreds of different noise levels. Yet, we note that devices that exhibit close to Hamming weight leakages are frequently encountered in practice [39]. Furthermore, such a simulated setting is a well established tool to analyze masking schemes (see, e.g. [18] for polynomial masking, [4] for inner product masking and [12] for leakage squeezing). Besides, we also consider random Gaussian leakage functions, of which the deterministic part corresponds to random functions over \mathcal{Y} , to confirm that all the trends we put forward are also observed with leakage functions that radically differ from the usual Hamming weight one.

a. Computing the MI metric. In this DPA setting, we aim to compute the MI between the key and the plaintext and leakages. For conciseness, we use the notations $\bar{Y} = [Y_1, \dots, Y_d]$ and $\bar{\mathbf{L}} = [\mathbf{L}_{Y_1}, \dots, \mathbf{L}_{Y_d}]$ for vectors containing the d shares and their corresponding leakages. Then we compute:

$$\begin{aligned} \text{MI}(K; X, \bar{\mathcal{L}}_Y) &= \text{H}[K] + \sum_{k \in \mathcal{K}} \text{Pr}[k] \cdot \\ &\sum_{x \in \mathcal{X}, \bar{y} \in \mathcal{Y}^d} \text{Pr}[x, \bar{y}] \cdot \sum_{\bar{l}_y \in \mathcal{L}^d} \text{Pr}[\bar{l}_y | k, x, \bar{y}] \cdot \log_2 \text{Pr}[k | x, \bar{l}_y]. \end{aligned} \quad (16)$$

While this expression may look quite involved, we note that it is actually simple to estimate in practice, by sampling the target implementation. Evaluators just have to set keys k in their device and generate leakage traces corresponding to (known) plaintexts x and (unknown) shares \bar{y} . Say there are $|\mathcal{K}| = n_k$ key candidates and we generate n_t leakage traces \bar{l}_i , then, one just assigns probabilities \hat{p}_i^j to each key candidate k_j^* , for each measured trace, as in Table 1. This is typically done using TA or LR. Following, if the correct key candidate is k , the second line of (16) can be computed as $\bar{E}_i \log_2(\hat{p}_i^k)$. Note that whenever considering the standard DPA setting where the target operations follow a key addition, it is not even necessary to sum over the keys since $\text{MI}(K = k; X, \bar{\mathcal{L}}_Y)$ is identical for all k 's, thanks to the key equivalence property put forward in [40].

Table 1. Computing key candidate probabilities for MI metric estimation.

State & leakage	Key candidates			
	k_1^*	k_2^*	...	$k_{N_k}^*$
$(k, x_1) \rightsquigarrow \bar{l}_1$	\hat{p}_1^1	\hat{p}_1^2	...	$\hat{p}_1^{n_k}$
$(k, x_2) \rightsquigarrow \bar{l}_2$	\hat{p}_2^1	\hat{p}_2^2	...	$\hat{p}_2^{n_k}$
...
$(k, x_{n_t}) \rightsquigarrow \bar{l}_{n_t}$	$\hat{p}_{n_t}^1$	$\hat{p}_{n_t}^2$...	$\hat{p}_{n_t}^{n_k}$

Intuitively, $\text{MI}(K; X, \bar{\mathcal{L}}_Y)$ measures the amount of information leaked on the key variable K . The framework in [60] additionally defines a Mutual Information Matrix (MIM) that captures the correlation between any key k and key candidates k^* . Using our sampling notations, it can be simply defined as $\text{MIM}_{k, k^*} = \text{H}[K] + \sum_i \log_2(\hat{p}_i^{k^*})$, which directly leads to $\text{MI}(K; X, \bar{\mathcal{L}}_Y) = \text{E}_k(\text{MIM}_{k, k})$.

b. Intuition behind the noise condition. Theorems 2 and 3 both require that the MI between the shares and their corresponding leakage is sufficiently small. In other words, they require the noise to be sufficiently large. In this section, we compute the MI metric for both an unprotected implementation (i.e. $d = 1$) and a masked one (i.e. $d = 2$) in function of different parameters.¹ In order to illustrate the computation of this metric, we provide a simple open source code that evaluates the MI between a sensitive variable Y and its Hamming weights, for different noise levels, both via numerical integration (that is only possible for mathematically-generated leakages) and sampling (that is more reflective of the evaluation of an actual device) [1]. In the latter case, an evaluator additionally

¹ For the masked case, we consider univariate leakages corresponding to the parallel setting in [7], for which computing the MI is slightly faster than in the serial one.

has to make sure that his estimations are accurate enough. Tools for ensuring this condition are discussed in [23]. In the following, this sufficient sampling is informally confirmed by the smooth shape of our experimental curves.

We start with the simplest possible plot, where the MI metric is computed in function of the noise variance σ^2 . Figure 2 shows these quantities, both for Hamming weight leakage functions and for random ones with output range N_l (in the latter context, the functions for different N_l 's were randomly picked up prior to the experiments, and stable across experiments). We also considered different bit sizes ($n = 2, 4, 6, 8$). Positively, we see that in all cases, the curves reach a linear behavior, where the slope corresponds to the number of shares d . Since the independent leakage condition is fulfilled in these experiments, this d corresponds to the smallest key-dependent moment in the leakage distribution. And since the measurement (aka sampling) cost for estimating such moments is proportional to $(\sigma^2)^d$, we observe that the MI decreases exponentially in d for large enough noises. Note that this behavior is plotted for $d = 1, 2$, but was experimented for d 's up to 4 in [61], and in fact holds for any d , since it exactly corresponds to Theorem 2 in a context where its assumptions are fulfilled.

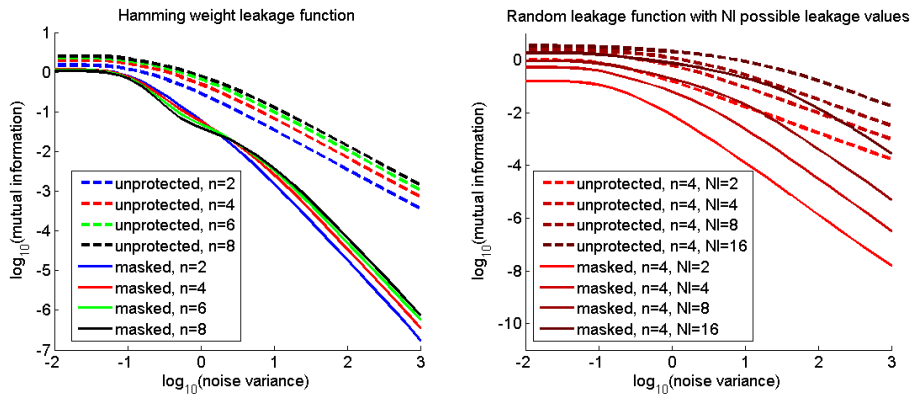


Fig. 2. MI metric in function of σ^2 . HW (left) and random (right) leakages.

Negatively, we also see that the noise level that can be considered as high enough depends on the leakage functions. For example, the random leakage functions in the right part of the figure have signals that vary from approximately $\frac{2}{4}$ for $N_l = 2$ to $\frac{16}{4}$ for $N_l = 16$. It implies that the linearly decreasing part of the curves is reached for larger noises in the latter case. Yet, this observation in fact nicely captures the intuition behind the noise condition. That is, the noise should be high enough for hiding the signal. Therefore, a very convenient way to express it is to plot the MI metric in function of shares' SNR, as in Figure 3. Here, we clearly see that as soon as the SNR is below a certain constant (10^{-1} , typically), the shape of the MI curves gets close to linear. This corroborates the condition in Theorem 2 that masking requires $\text{MI}(K_i; X, \mathbf{L}_{Y_i})$ to be smaller than

a given constant. Our experiments with different bit sizes also suggest that the $|\mathbb{F}|$ factor in this noise condition is a proof artifact. This is now formally proven by Dziembowski, Faust and Skorski in [24]. Of course, and as discussed in Section 2.2, the SNR metric is only applicable under certain conditions (univariate Gaussian leakages). So concretely, an evaluator may choose between computing it after dimensionality reduction (leading to a heuristic but intuitive condition), or to directly state the condition in function of the MI. For completeness, we also plot the MI metric for an unprotected and masked implementation in function of the share’s MI in Appendix, Figure 10. It clearly exhibits that as the share’s MI decreases, this reduction is amplified by masking (exponentially in d).

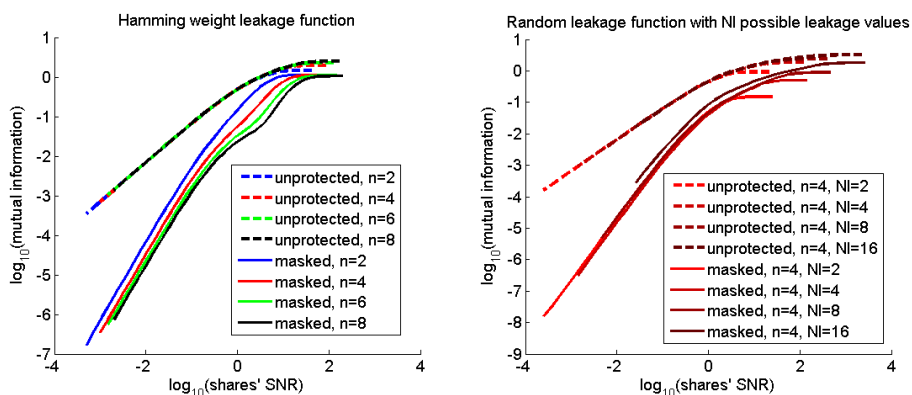


Fig. 3. MI metric in fct. of the shares’ SNR. HW (left) and random (right) leakages.

c. Tightness of the bounds. Given that the noise is high enough (as just discussed), Theorems 2 and 3 guarantee that the success rate of a side-channel adversary can be bounded based on the value of the share’s leakage, measured with $\text{MI}(K_i; X, \mathbf{L}_{Y_i})$. This directly leads to useful bounds on the measurement complexity to reach a given success rate, e.g. from (8) we can compute:

$$m \geq \frac{\log(1 - \text{SR}^{\text{kr}})}{\log\left(1 - \left(|\mathbb{F}| \sqrt{\frac{\text{MI}(K_i; X, \mathbf{L}_{Y_i})}{2}}\right)^d\right)}. \quad (17)$$

We now want to investigate how tight this bound is. For this purpose, we compared it with the measurement complexity of concrete key recovery TA (using a perfect leakage model).² As previously mentioned, the $|\mathbb{F}|$ factor in this equation can be seen as a proof artifact related to the reduction in our theorems – so we tested a bound excluding this factor. For similar reasons, we also tested

² Our attacks exploit the leakages of an S-box output, as specified in Section 2.1. We took the PRESENT S-box for $n = 4$, the AES one for $n = 8$, and picked up two random S-boxes for $n = 2, 6$, as we did for the random leakage functions.

a bound additionally excluding the square root loss in the reductions (coming from Theorem 1). As illustrated in Figure 4, the measurement complexity of the attacks is indeed bounded by Equation (17), and removing the square root loss allows the experimental and theoretical curves to have similar slopes. The latter observation fits with the upper bound $\text{MI}(Y_i; \mathbf{L}_{Y_i}) \leq \frac{|\mathbb{F}|}{\ln(2)} \cdot \text{SD}(Y_i; Y_i | \mathbf{L}_{Y_i})$ given in [49] that becomes tight as the noise increases.³ As expected, the bounds become meaningless for too low noise levels (or too large SNRs, see Appendix, Figure 11). Intuitively, this is because we reach success rates that are stuck to one when we deviate from this condition. For completeness, we added approximations obtained by normalizing the shares' MI by $H[K]$ to the figure, which provide hints about the behavior of a leaking device when the noise is too low.

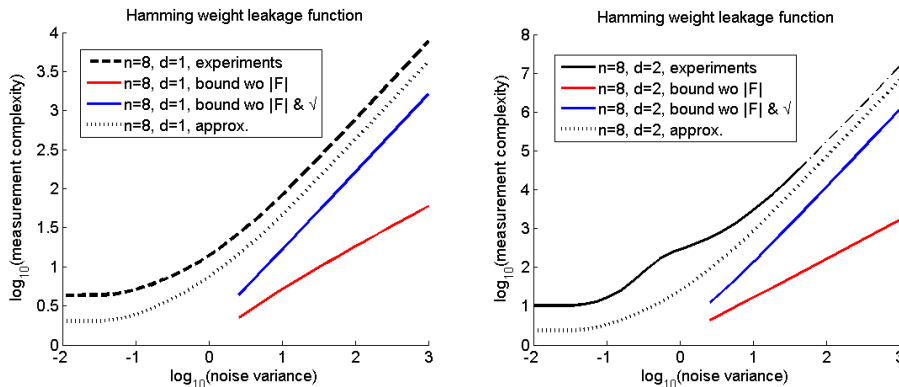


Fig. 4. Measurement complexity and bounds/approximations for concrete TA.

Interestingly, these results also allow us to reach a comprehensive view of the parameters in Theorem 3, where the security of a complete circuit encoded according to the ISW compiler is proven. That is, in this case as well we expect the $|\mathbb{F}|$ and $1/9$ factors in Equation (12) to be due to proof technicalities. By contrast, the $|T|$ factor is physically motivated, since it corresponds to the size of the circuit and fits the intuition that more computations inevitably means more exploitable leakage. The d factor appearing in the noise condition of Equation (13) can also be explained, since it directly relates to the fact that in the ISW compiler, any multiplication will require to manipulate each share d times. It typically reflects the distance between standard (divide-and-conquer) side-channel attacks (such as analyzed in this section) and more powerful (multivariate) adversaries trying to exploit the leakage of all the intermediate computations in a block cipher, e.g. based on algebraic cryptanalysis (see [52, 53] and follow up works). Taking all these observations into account, we summarize the concrete security of any masking scheme with the following informal conjecture.

³ Since their inequality comes from a $\log(1+x) < \log(x)$ inequality that gets close to an equality when x gets close to 0, which happens for large noise levels.

Informal conjecture. *Suppose that we have a circuit of size $|\Gamma|$ masked with d shares such that the information leakage on each of these shares (using all available time samples) is bounded by $\text{MI}(Y_i; \mathbf{L}_{Y_i})$. Then, the probability of success of a distinguishing adversary using m measurements and targeting a single element (e.g. gate) of the circuit under independent and sufficiently noisy leakage is:*

$$\text{SR}_1^{\text{dist}} \leq 1 - (1 - \text{MI}(Y_i; \mathbf{L}_{Y_i})^d)^m, \quad (18)$$

and the probability of success targeting all $|\Gamma|$ elements independently equals:

$$\text{SR}_{|\Gamma|}^{\text{dist}} \leq 1 - (1 - \text{SR}_1^{\text{dist}})^{|\Gamma|}. \quad (19)$$

Interestingly, Equation (19) (like Theorem 3) assumes that the leakages of the $|\Gamma|$ gates (or target intermediate values) are exploited independently. This perfectly corresponds to the probing model in which the adversary gains either full knowledge or no knowledge of such computing elements. Thanks to [22], it also implies a similar result against noisy leakages if the noise condition is fulfilled. However, as the noise level decreases, some advanced (e.g. algebraic) side-channel attacks can sometimes take advantage of different computations jointly in a more efficient manner. Note that this informal conjecture is backed up by the results in [3] (Theorem 6) where a similar bound is given in the context of statistical cryptanalysis. By using the approximation $\log(1 - x) \approx -x$ that holds for x 's close to 0, Equation (18) directly leads to the following simple approximation of a standard DPA's measurement complexity for large noise levels:

$$m \geq \frac{\log(1 - \text{SR}_1^{\text{dist}})}{\log(1 - \text{MI}(Y_i; \mathbf{L}_{Y_i})^d)} \approx \frac{c}{\text{MI}(Y_i; \mathbf{L}_{Y_i})^d}, \quad (20)$$

where c is a small constant that depends on the target success rate. A similar approximation can be obtained from Equation (19) for multi-target attacks.

d. Relation with the Eurocrypt 2009 evaluation framework. The evaluation of leaking cryptographic implementations with a combination of information and security metrics was put forward by Standaert et al. at Eurocrypt 2009. In this reference, the authors showed a qualitative connection between both metrics. Namely, they proved that the model (i.e. the approximation of the leakage PDF) used by a side-channel adversary is sound (i.e. allows key recoveries) if and only if the mutual information matrix (defined in paragraph (a) of this section) is such that its diagonal values are maximum for each line. By contrast, they left the quantitative connection between these metrics as an open problem (i.e. does more MI imply less security?). Our results provide a formal foundation for this quantitative connection. They prove that for any implementation, decreasing the MI of the target intermediate values is beneficial to security. This can be achieved by ad hoc countermeasures, in which case it is the goal of an evaluation laboratory to quantify the MI metric, or by masking, in which case we can bound security based only on the value of this metric for each share taken separately.

4.2 Beyond independent leakage

The previous section evaluated an experimental setting where the leakage of each share is independent of each other, i.e. $\mathbf{L}_{y_i} = \mathbf{G}(y_i) + N_i$. But as discussed in introduction, this condition frequently turns out to be hard to fulfill and so far, there are only limited (in)formal tools allowing to analyze the deviations from independent leakages that may be observed in practice. In order to contribute to this topic, we first launched another set of experiments (for 2-share masking), where the leakage of each share can be written as:

$$\begin{aligned} \mathbf{L}_{y_1} &= \mathbf{G}_1(y_1) + f \cdot \mathbf{G}_{1,2}(y_1, y_2) + N_1, \\ \mathbf{L}_{y_2} &= \mathbf{G}_2(y_2) + f \cdot \mathbf{G}_{2,1}(y_1, y_2) + N_2. \end{aligned}$$

Here the G_i functions manipulate the shares independently, while the $G_{i,j}$ functions depend on both shares. We additionally used the f (for flaw) parameter in order to specify how strongly we deviate from the independent leakage assumption. As in the previous section, we considered Hamming weight and random functions for all G 's (and we used $G_{i,j}(y_i, y_j) = G(y_i \oplus y_j)$ for illustration). Exemplary results of an information theoretic analysis in this context are given in Figure 5 for the $n = 4$ -, and 8-bit cases (and in Appendix, Figure 12 for the $n = 2$ - and 6-bit S-box cases). We mainly observe that as the noise increases, even small flaws are exploitable by an adversary. Indeed, breaking the independence condition makes smaller-order moments of the leakage distribution key-dependent. Consequently, for large enough noise, it is always this smaller-order moment that will be the most informative. This is empirically confirmed by the slopes of the IT curves in the figures, that gradually reach one rather than two.

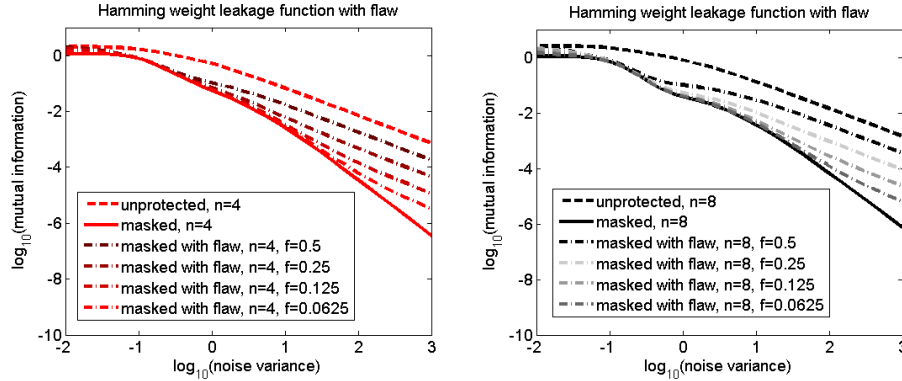


Fig. 5. MI metric for masked implementation with flaw ($n = 4, 8$).

Following these experiments, let us consider a chip that concretely exhibits such a flaw for a given noise level σ_{exp}^2 (corresponding to its actual measurements). Despite falling outside the masking proofs' guarantees, an important

question is whether we can still (approximatively) predict its security level based on sound statistical tools. In this respect, a useful observation is that the MI metric cannot directly answer the question since it captures the information lying in all the statistical moments of the leakage PDF. So we need another ingredient in order to reveal the informativeness of each moment of the leakage PDF, separately. The Moments-Correlating DPA (MC-DPA) recently introduced in [44] is a natural candidate for this purpose. We now describe how it can be used to (informally) analyze the security of a flawed masked implementation.

In this context, we first need to launch MC-DPA for different statistical moments, e.g. the first- and second-order ones in our 2-share example. They are illustrated by the circle and square markers in the left part of Figure 6. For concreteness, we take the (most revealing) case where the second-order moment is more informative than the first-order one. Assuming that the noise condition in our theorems is fulfilled, the impact of increasing the noise on the value of the MC-DPA distinguisher can be predicted as indicated by the curves of the figure. That is, with a slope of 1/2 for the first-order moment and a slope of 1 for the second-order one.⁴ Hence, we can directly predict the noise level $\sigma_{\text{exp}}^2 + \Delta$ such that the first-order moment becomes more informative. Eventually, we just observe that concrete side-channel attacks always exploit the smallest key-dependent moment in priority (which motivates the definition of the security-order for masking schemes [17]). So starting from the value of the MI at σ_{exp}^2 (represented by a circle in the right part of the figure), we can extrapolate that this MI will decrease following a curve with slope 2 until $\sigma_{\text{exp}}^2 + \Delta$ and a curve with slope 1 afterwards. Taking advantage of the theorems in the previous sections, this directly leads to approximations of the best attacks' measurement complexity. Furthermore, extending this reasoning to more shares and higher-order statistical moments is straightforward: it just requires to add MC-DPA curves in the left part of Figure 6, and to always consider the one leading to the

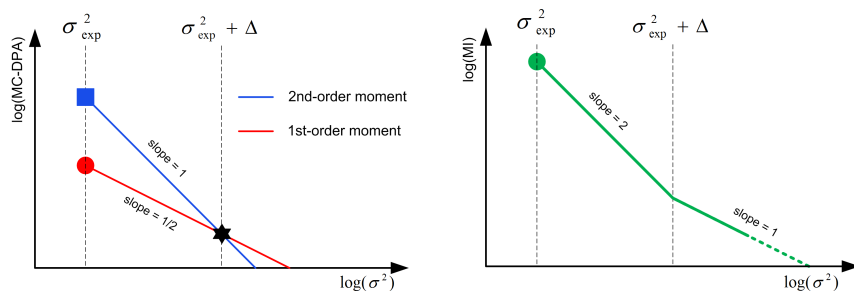


Fig. 6. Evaluating non-independent leakages with MC-DPA (left) and MI (right).

⁴ Slopes are divided by 2 when considering Pearson's correlation rather than the MI since this correlation is essentially proportional to the square root of the SNR. This is also reflected by the measurement complexity of CPA, that is proportional to the inverse of the squared correlation vs. the inverse of the MI for TA [62].

highest MC-DPA value to set the slope of the MI curves, in the right part of the figure. To the best of our knowledge, such figures (despite informal) provide the first concrete tools to approximate the security level in such contexts.

Note finally that the shape of the non-independent leakages (i.e. the $G_{i,j}$ functions) observed in practice highly depends on the implementations. For example in hardware, multiple shares can leak jointly in a hardly predictable manner [41, 54]. By contrast in software, the most usual issue (due to transition-based leakages) is easier to analyse [5]. It typically divides the order of the smallest key-dependent moment in the leakage distribution by two, which corresponds to the additional square root loss in the security bounds of Duc et al. when considering leakages that depend on two wires simultaneously (see [22], Section 5.5).

4.3 Exploiting computational power

In this section, we finally tackle the problem of divide-and-conquer DPA attacks, where the adversary aims to combine side-channel information gathered from a number of measurements, and computational power. That is, how to deal with the practically critical situation where the number of measurements available is not sufficient to exactly recover the key? As discussed in [64, 65], optimal enumeration and key ranking algorithms provide a concrete answer to this question. They allow building security graphs, where the success rate is plotted in function of a number of measurements and computing power, by repeating attacks multiple times. We next discuss more efficient and analytical strategies.

a. Why MI is not enough? Whenever trying to exploit both side-channel leakage and brute-force computation (e.g. key enumeration) the most challenging aspect of the problem is to capture how measurements and computation actually combine. This is easily illustrated with the following example. Imagine two hypothetical side-channel attacks that both succeed with probability $1/100$. In the first case, the adversary gains nothing with probability $99/100$ and the full key with probability $1/100$. In the second case, he always gains a set of 100 equally likely keys. Clearly, enumeration will be pretty useless in the first case, while extremely powerful in the second one. More generally, such examples essentially suggest that the computational cost of an enumeration does not only depend on the informativeness of the leakage function (e.g. measured with the MI) but also on its shape. For illustration, a line of the mutual information matrix computed from Hamming weight leakages for two noise levels is given in Figure 7, where we can clearly identify the patterns due to this leakage model. While $MIM_{k,k}$ only corresponds to a single value of the matrix line (here $k = 111$), which bounds the measurement complexity to recover this key without additional computation (as previously discussed), how helpful is enumeration will additionally depend on the relative distance between the $MIM_{k,k}$ and MIM_{k,k^*} values [68]. Incidentally, this example also puts forward some limitations of the probing leakage model when measuring computational cost, since it describes an all-or-nothing strategy – as already mentioned in Section 4.1, paragraph (c) – which is not the case for

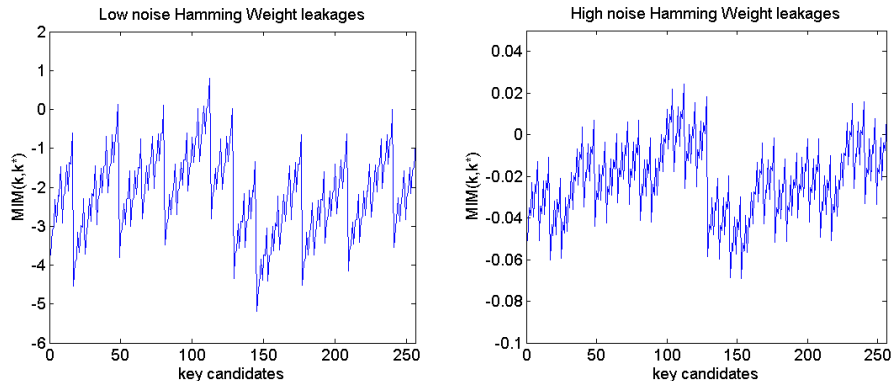


Fig. 7. Exemplary line of the mutual information matrix (for $k = 111$).

the noisy leakage setting. Hence, whereas the probing model is easier to manipulate in proofs, and therefore useful to obtain asymptotic results, noisy leakages are a more accurate tool to quantify concrete security levels as in this section.

b. Measurement and computational bounds per S-box. Interestingly, one can easily derive bounds for attacks combining side-channel measurements and enumeration power against a single S-box, by re-using exactly the same material as we anyway need to estimate $\text{MI}(K; X, \mathbf{L}_{Y_i})$ for a single secret share. In the following, we will assume that the key equivalence property mentioned in Section 4.1, paragraph (a) holds, and focus on a single line of the mutual information matrix (if it does not, evaluators simply have to compute all its lines), next denoted as $\text{MIM}_{k,-}$. In order to characterize the distance between a key k and its close candidates k^* , we first sort this line and produce $s = \text{sort}(\text{MIM}_{k,-})$. As a result, the key candidate $k_{s(1)}^*$ is the best rated one (i.e. the correct k if the leakage model is sound), $k_{s(2)}^*$ is the second best, \dots . From there, we compute a “computational version” of the mutual information matrix as:

$$\text{MIM}_{k,k}^c = H[K_i] + E_j \left(\log \left(\sum_{l=1}^c \hat{p}_j^{s(l)} \right) \right). \quad (21)$$

It essentially corresponds to the amount of information an adversary obtains about a random variable that aggregates the c most likely key candidates. Assuming that these c key candidates are equally likely (which can only be pessimistic), it directly provides simple bounds on the success rate of an attack combining m measurements with the enumeration of c keys:

$$\text{SR}^{\text{kr}}(m, c) \leq 1 - \left(1 - (\text{MIM}_{k,k}^c)^d \right)^m, \quad (22)$$

For illustration, a couple of such bounds are given in Figure 8, where we see the impact of increasing the number of shares d and number of measurements m . Note that despite requiring similar characterization efforts, these bounds are

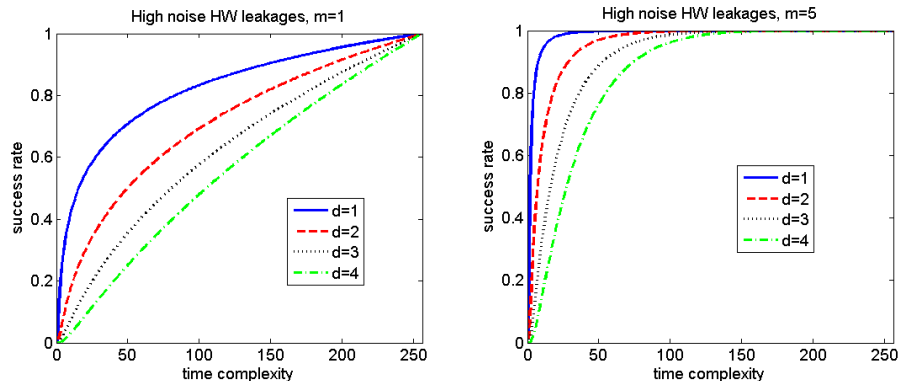


Fig. 8. Key recovery success rate against a single S-box, in function of the time complexity for the leakage function of Figure 7 (right), after m measurements.

conceptually different from the previous approaches to approximate the success rate of side-channel attacks. In particular, works like [20, 26, 37, 55] are specific to popular distinguishers (and usually require specialized assumptions about the distribution of these distinguishers), while our results directly connect to security proofs that are independent of the adversarial strategy and hold for any leakage distribution. Nevertheless, the only requirement to analyze the combination of multiple S-boxes in the next paragraph (c) is to have success rates curves for each S-box. So while this paragraph (b) describes an efficient way to build such curves, the following contribution is in fact general, and could be used as a complement to any security evaluation obtained for separate S-boxes.

c. Combining multiple S-boxes. We finally generalize our analysis of the previous paragraph to the case where we target n_s S-boxes (e.g. $n_s = 16$ for the AES), gained information about their respective input key bytes, and want to recover the full master key. We assume that we perform the same amount of measurements m on each S-box. This can be easily justified in practice, since a leakage trace usually contains samples corresponding to all S-boxes. By contrast, we make no assumption about how informative the leakages of each S-box are. For example, it could completely happen that one S-box is very leaky, and another one perfectly protected (so that enumeration is the only option to recover its corresponding key byte). As just explained, we then characterize the measurement vs. complexity tradeoff with n_s success rate curves $\text{SR}_i^{\text{kr}}(m, c_i)$ with $1 \leq i \leq n_s$. Typically, we will then set a limit β to the adversary's computational power and try to solve the following optimization problem:

$$\begin{aligned}
 & \max_{c_1, \dots, c_{n_s}} \prod_{i=1}^{n_s} \text{SR}_i^{\text{kr}}(m, c_i), \\
 & \text{subject to } \prod_{i=1}^{n_s} c_i \leq \beta.
 \end{aligned} \tag{23}$$

Taking the logarithm of both products, we get:

$$\begin{aligned} & \max_{c_1, \dots, c_{n_s}} \sum_{i=1}^{n_s} \log \left(\text{SR}_i^{\text{kr}}(m, c_i) \right), \\ & \text{subject to } \sum_{i=1}^{n_s} \log(c_i) \leq \log(\beta). \end{aligned} \tag{24}$$

For general functions SR_i^{kr} , this problem is known as a “separable, non-linear integer programming problem”. Surveys about non-linear integer programming problems are various (e.g. [10, 36]). There exist many well-studied heuristics to solve them, including branch-and-bounds and convex envelop techniques. Note that the problem generally becomes easier when dealing with convex functions.

We conclude this section with a simple and cheap heuristic algorithm which approximates well the optimal solution for the problem sizes and leakage functions we considered. The approach we propose is inspired by [29], and based on a tradeoff between the computational cost and accuracy of the solutions found, that is controlled by downsampling the success rate curves and keeping track of quantization errors. Intuitively, enumerating the combination of the possible success rates for two n -bit S-boxes requires the computation of 2^{2n} product complexities $c_i \cdot c_j$. Since combining more S-boxes exhaustively will increase the complexity exponentially (i.e. $2^{n_s \cdot n}$ for n_s n -bit S-boxes), the idea of our heuristic is simply to ignore some samples. Namely, we will fix a bound N_{max} which will designate the maximum number of samples we save per success rate curve (or combination of them). Such a well-known downsampling process is informally illustrated in the left part of Figure 9, where we can see that the original curve can be easily upperbounded and lowerbounded, since it is increasing.

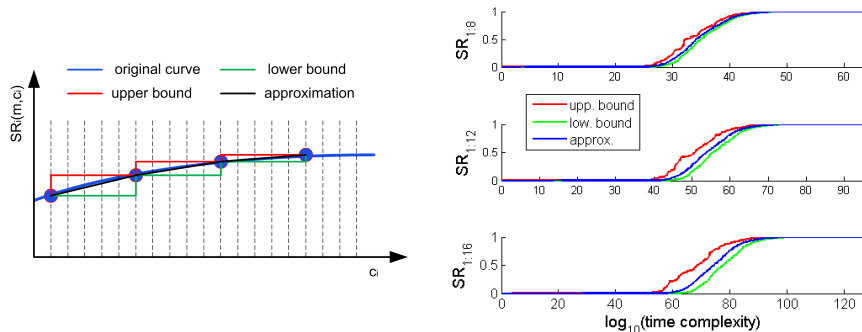


Fig. 9. Downsampling SR curves (left) and enumeration complexity bounds (right).

This solution is described more formally in Algorithm 1 and works as follows. First, we downsample each success rate curve $\text{SR}_i(m, c)$ to N_{max} linearly spaced points that we can write as N_{max} pairs $(s_{i,1}, c_{i,1}), \dots, (s_{i,N_{\text{max}}}, c_{i,N_{\text{max}}})$. Next,

we take the first S-Box and combine it with the second one, obtaining N_{\max}^2 values. These values are then downsampled again to N_{\max} linearly spaced points, so that we can iteratively combine them with the next S-boxes. We denote the aggregation of the i first success rate curves with $SR_{1:i}$. We also add an additional output to our algorithm, namely a list of complexities ℓ_i , describing how the effort is distributed among the S-boxes. Indeed, suppose for example that we combine the success rate pair $(0.1, 2^4)$ of S-box 1 with the success rate pair $(0.2, 2^5)$ of S-box 2. We obtain a success rate of 0.02 for a complexity of 2^9 , but nothing tells us how the effort is distributed between S-box 1 and S-box 2. Hence, we write the result as $(0.02, 2^9, \{2^4, 2^5\})$ which shows how the complexities are shared.

Algorithm 1 Heuristic to combine the SR curves of n_s S-boxes.

Require: Pairs $[(s_{i,1}, c_{i,1}), \dots, (s_{i,N_{\max}}, c_{i,N_{\max}})] =: SR_i$ for each S-box i .
Require: A bound N_{\max} on the number of samples and a bound β on the complexity.
Ensure: Triplets $(s_1, c_1, \ell_1), \dots, (s_{N_{\max}}, c_{N_{\max}}, \ell_{N_{\max}})$ approximating the success rate curve of the combination of the n_s S-Boxes, where the ℓ_i -s are ordered lists of complexities showing how they should be distributed among the S-boxes.

- 1: $SR_{1:1} \leftarrow [(s_{1,1}, c_{1,1}, \{c_{1,1}\}), \dots, (s_{1,N_{\max}}, c_{1,N_{\max}}, \{c_{1,N_{\max}}\})]$;
- 2: **for** $i = 2$ to n_s **do**
- 3: $SR_{1:i} \leftarrow \emptyset$;
- 4: \triangleright *Combination of aggregated S-Boxes 1 : i-1 with S-Box i.*
- 5: **for** $(s_j, c_j, \ell_j) \in SR_{1:i-1}$ **do** \triangleright *For all N_{\max} values in aggregated curve.*
- 6: **for** $(s_{i,k}, c_{i,k}) \in SR_i$ **do** \triangleright *For all N_{\max} values of the new S-Box.*
- 7: **if** $c_j \cdot c_{i,k} < \beta$ **then** \triangleright *If we did not reach the bound β .*
- 8: $SR_{1:i} \leftarrow SR_{1:i} \cup (s_j \cdot s_{i,k}, c_j \cdot c_{i,k}, \ell_j \cup c_{i,k})$; \triangleright *Merge.*
- 9: **end if**
- 10: **end for**
- 11: **end for**
- 12: \triangleright *Downsampling.*
- 13: Sort $SR_{1:i}$ and keep only N_{\max} linearly spaced pairs.
- 14: **end for**
- 15: **return** $SR_{1:n_s}$

For illustration, the right part of Figure 9 provides such bounds for the combination of 8, 12 and 16 AES S-boxes, for a noise level and number of measurement such that the rank estimation problem is challenging (i.e. with full key rank for the 16-byte master key around 2^{80}). The complexity of this heuristic is proportional to $n_s \cdot (N_{\max}^2 + \log(N_{\max}))$ and the results in the figure were obtained within seconds of computation on a desktop computer, using a simple Matlab prototype code. We leave the investigation of better solutions to obtain accurate time complexity bounds with minimum efforts as a scope for further research.

Summarizing, our results show that the (complex) task of evaluating the worst-case security level of a masked implementation against (divide-and-conquer) DPA can be simplified to the evaluation of a couple of MI values, even in contexts where the independence assumption is not fulfilled. This provides a solid

foundation for the Eurocrypt 2009 evaluation framework. It also makes it easier to implement, since success rate curves for full keys can now be derived from the MI values, rather than sampled experimentally by repeating (many) subkey recovery experiments and key rank estimations, which is an expensive task. Taking advantage of the tools in this paper therefore allow reducing both the number of measurements and the time needed to evaluate leaking devices.

Acknowledgements. Alexandre Duc is supported by the Swiss National Science Foundation, grant 200021 143899/1. Sebastian Faust received funding from the Marie Curie IEF/FP7 project GAPS (grant 626467). François-Xavier Standaert is a research associate of the Belgian Fund for Scientific Research. This work has been funded in parts by the ERC project 280141 (CRASH).

References

1. <http://perso.uclouvain.be/fstandae/PUBLIS/154.zip>.
2. Cédric Archambeau, Eric Peeters, François-Xavier Standaert, and Jean-Jacques Quisquater. Template Attacks in Principal Subspaces. In Louis Goubin and Mitsuru Matsui, editors, *CHES*, volume 4249 of *LNCS*, pages 1–14. Springer, 2006.
3. Thomas Baignères, Pascal Junod, and Serge Vaudenay. How Far Can We Go Beyond Linear Cryptanalysis? In Pil Joong Lee, editor, *ASIACRYPT*, volume 3329 of *LNCS*, pages 432–450. Springer, 2004.
4. Josep Balasch, Sebastian Faust, Benedikt Gierlichs, and Ingrid Verbauwhede. Theory and Practice of a Leakage Resilient Masking Scheme. In Wang and Sako [67], pages 758–775.
5. Josep Balasch, Benedikt Gierlichs, Vincent Grosso, Oscar Reparaz, and François-Xavier Standaert. On the Cost of Lazy Engineering for Masked Software Implementations. *IACR Cryptology ePrint Archive*, 2014:413, 2014.
6. Lejla Batina and Matthew Robshaw, editors. *Cryptographic Hardware and Embedded Systems - CHES 2014 - 16th International Workshop, Busan, South Korea, September 23-26, 2014. Proceedings*, volume 8731 of *LNCS*. Springer, 2014.
7. Sonia Belaïd, Vincent Grosso, and François-Xavier Standaert. Masking and Leakage-Resilient Primitives: One, the Other(s) or Both? *Cryptography and Communications*, 7(1):163–184, 2015.
8. Mihir Bellare, Stefano Tessaro, and Alexander Vardy. A Cryptographic Treatment of the Wiretap Channel. *IACR Cryptology ePrint Archive*, 2012:15, 2012.
9. Mihir Bellare, Stefano Tessaro, and Alexander Vardy. Semantic Security for the Wiretap Channel. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*, volume 7417 of *LNCS*, pages 294–311. Springer, 2012.
10. Dimitri P Bertsekas. *Nonlinear Programming*. Athena Scientific, 1999.
11. Eric Brier, Christophe Clavier, and Francis Olivier. Correlation Power Analysis with a Leakage Model. In Marc Joye and Jean-Jacques Quisquater, editors, *CHES*, volume 3156 of *LNCS*, pages 16–29. Springer, 2004.
12. Claude Carlet, Jean-Luc Danger, Sylvain Guilley, and Housseem Maghrebi. Leakage Squeezing: Optimal Implementation and Security Evaluation. *J. Mathematical Cryptology*, 8(3):249–295, 2014.

13. Claude Carlet, Louis Goubin, Emmanuel Prouff, Michaël Quisquater, and Matthieu Rivain. Higher-Order Masking Schemes for S-Boxes. In Anne Canteaut, editor, *Fast Software Encryption - 19th International Workshop, FSE 2012, Washington, DC, USA, March 19-21, 2012*, volume 7549 of *LNCS*, pages 366–384. Springer, 2012.
14. Suresh Chari, Charanjit S. Jutla, Josyula R. Rao, and Pankaj Rohatgi. Towards Sound Approaches to Counteract Power-Analysis Attacks. In Wiener [69], pages 398–412.
15. Suresh Chari, Josyula R. Rao, and Pankaj Rohatgi. Template Attacks. In Burton S. Kaliski Jr., Çetin Kaya Koç, and Christof Paar, editors, *CHES*, volume 2523 of *LNCS*, pages 13–28. Springer, 2002.
16. Jean-Sébastien Coron, Christophe Giraud, Emmanuel Prouff, Soline Renner, Matthieu Rivain, and Praveen Kumar Vadnala. Conversion of Security Proofs from One Leakage Model to Another: A New Issue. In Werner Schindler and Sorin A. Huss, editors, *COSADE*, volume 7275 of *LNCS*, pages 69–81. Springer, 2012.
17. Jean-Sébastien Coron, Emmanuel Prouff, and Matthieu Rivain. Side Channel Cryptanalysis of a Higher Order Masking Scheme. In Pascal Paillier and Ingrid Verbauwhede, editors, *CHES*, volume 4727 of *LNCS*, pages 28–44. Springer, 2007.
18. Jean-Sébastien Coron, Emmanuel Prouff, Matthieu Rivain, and Thomas Roche. Higher-Order Side Channel Security and Mask Refreshing. In Shiho Moriai, editor, *Fast Software Encryption - 20th International Workshop, FSE 2013, Singapore, March 11-13, 2013*, volume 8424 of *LNCS*, pages 410–424. Springer, 2013.
19. Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (2. ed.)*. Wiley, 2006.
20. A. Adam Ding, Liwei Zhang, Yunsi Fei, and Pei Luo. A Statistical Model for Higher Order DPA on Masked Devices. In Batina and Robshaw [6], pages 147–169.
21. Yevgeniy Dodis. Shannon Impossibility, Revisited. In Adam Smith, editor, *Information Theoretic Security - 6th International Conference, ICITS 2012, Montreal, QC, Canada, August 15-17, 2012. Proceedings*, volume 7412 of *LNCS*, pages 100–110. Springer, 2012.
22. Alexandre Duc, Stefan Dziembowski, and Sebastian Faust. Unifying Leakage Models: From Probing Attacks to Noisy Leakage. In Nguyen and Oswald [45], pages 423–440.
23. François Durvaux, François-Xavier Standaert, and Nicolas Veyrat-Charvillon. How to Certify the Leakage of a Chip? In Nguyen and Oswald [45], pages 459–476.
24. Stefan Dziembowski, Sebastian Faust, and Maciej Skorski. Noisy Leakage Revisited. *To appear in the proceedings of EUROCRYPT 2015*.
25. Sebastian Faust, Tal Rabin, Leonid Reyzin, Eran Tromer, and Vinod Vaikuntanathan. Protecting Circuits from Leakage: the Computationally-Bounded and Noisy Cases. In Henri Gilbert, editor, *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, French Riviera, May 30 - June 3, 2010. Proceedings*, volume 6110 of *LNCS*, pages 135–156. Springer, 2010.
26. Yunsi Fei, Qiasi Luo, and A. Adam Ding. A Statistical Model for DPA with Novel Algorithmic Confusion Analysis. In Emmanuel Prouff and Patrick Schaumont, editors, *Cryptographic Hardware and Embedded Systems - CHES 2012 - 14th International Workshop, Leuven, Belgium, September 9-12, 2012. Proceedings*, volume 7428 of *LNCS*, pages 233–250. Springer, 2012.

27. Guillaume Fumaroli, Ange Martinelli, Emmanuel Prouff, and Matthieu Rivain. Affine Masking against Higher-Order Side Channel Analysis. In Alex Biryukov, Guang Gong, and Douglas R. Stinson, editors, *Selected Areas in Cryptography - 17th International Workshop, SAC 2010, Waterloo, Ontario, Canada, August 12-13, 2010*, volume 6544 of *LNCS*, pages 262–280. Springer, 2010.
28. Benedikt Gierlichs, Lejla Batina, Pim Tuyls, and Bart Preneel. Mutual Information Analysis. In Oswald and Rohatgi [47], pages 426–442.
29. Cezary Glowacz, Vincent Grosso, Romain Poussier, Joachim Schueth, and François-Xavier Standaert. Simpler and More Efficient Rank Estimation for Side-Channel Security Assessment. *IACR Cryptology ePrint Archive*, 2014:920, 2014.
30. Louis Goubin and Ange Martinelli. Protecting AES with Shamir’s Secret Sharing Scheme. In Preneel and Takagi [48], pages 79–94.
31. Vincent Grosso, Emmanuel Prouff, and François-Xavier Standaert. Efficient Masked S-Boxes Processing - A Step Forward -. In David Pointcheval and Damien Vergnaud, editors, *Progress in Cryptology - AFRICACRYPT 2014 - 7th International Conference on Cryptology in Africa, Marrakesh, Morocco, May 28-30, 2014. Proceedings*, volume 8469 of *LNCS*, pages 251–266. Springer, 2014.
32. Vincent Grosso, François-Xavier Standaert, and Emmanuel Prouff. Low Entropy Masking Schemes, Revisited. In Aurélien Francillon and Pankaj Rohatgi, editors, *Smart Card Research and Advanced Applications - 12th International Conference, CARDIS 2013, Berlin, Germany, November 27-29, 2013*, volume 8419 of *LNCS*, pages 33–43. Springer, 2013.
33. Yuval Ishai, Amit Sahai, and David Wagner. Private Circuits: Securing Hardware against Probing Attacks. In Dan Boneh, editor, *Advances in Cryptology - CRYPTO 2003, 23rd Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 2003, Proceedings*, volume 2729 of *LNCS*, pages 463–481. Springer, 2003.
34. Thomas Johansson and Phong Q. Nguyen, editors. *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, volume 7881 of *LNCS*. Springer, 2013.
35. Paul C. Kocher, Joshua Jaffe, and Benjamin Jun. Differential Power Analysis. In Wiener [69], pages 388–397.
36. Duan Li and Xiaoling Sun. Nonlinear Knapsack Problems. In *Nonlinear Integer Programming*, volume 84 of *International Series in Operations Research & Management Science*, pages 149–207. Springer US, 2006.
37. Victor Lomné, Emmanuel Prouff, Matthieu Rivain, Thomas Roche, and Adrian Thillard. How to Estimate the Success Rate of Higher-Order Side-Channel Attacks. In Batina and Robshaw [6], pages 35–54.
38. Stefan Mangard. Hardware Countermeasures against DPA ? A Statistical Analysis of Their Effectiveness. In Tatsuaki Okamoto, editor, *CT-RSA*, volume 2964 of *LNCS*, pages 222–235. Springer, 2004.
39. Stefan Mangard, Elisabeth Oswald, and Thomas Popp. *Power Analysis Attacks - Revealing the Secrets of Smart Cards*. Springer, 2007.
40. Stefan Mangard, Elisabeth Oswald, and François-Xavier Standaert. One for All - All for One: Unifying Standard Differential Power Analysis Attacks. *IET Information Security*, 5(2):100–110, 2011.
41. Stefan Mangard, Thomas Popp, and Berndt M. Gammel. Side-Channel Leakage of Masked CMOS Gates. In Alfred Menezes, editor, *CT-RSA*, volume 3376 of *LNCS*, pages 351–365. Springer, 2005.

42. Marcel Medwed and François-Xavier Standaert. Extractors against Side-Channel Attacks: Weak or Strong? *J. Cryptographic Engineering*, 1(3):231–241, 2011.
43. Amir Moradi and Oliver Mischke. Glitch-Free Implementation of Masking in Modern FPGAs. In *HOST*, pages 89–95. IEEE, 2012.
44. Amir Moradi and François-Xavier Standaert. Moments-Correlating DPA. *IACR Cryptology ePrint Archive*, 2014:409, 2014.
45. Phong Q. Nguyen and Elisabeth Oswald, editors. *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, volume 8441 of *LNCS*. Springer, 2014.
46. Svetla Nikova, Vincent Rijmen, and Martin Schl affer. Secure Hardware Implementation of Nonlinear Functions in the Presence of Glitches. *J. Cryptology*, 24(2):292–321, 2011.
47. Elisabeth Oswald and Pankaj Rohatgi, editors. *Cryptographic Hardware and Embedded Systems - CHES 2008, 10th International Workshop, Washington, D.C., USA, August 10-13, 2008. Proceedings*, volume 5154 of *LNCS*. Springer, 2008.
48. Bart Preneel and Tsuyoshi Takagi, editors. *Cryptographic Hardware and Embedded Systems - CHES 2011 - 13th International Workshop, Nara, Japan, September 28 - October 1, 2011. Proceedings*, volume 6917 of *LNCS*. Springer, 2011.
49. Emmanuel Prouff and Matthieu Rivain. Masking against Side-Channel Attacks: A Formal Security Proof. In Johansson and Nguyen [34], pages 142–159.
50. Emmanuel Prouff and Thomas Roche. Attack on a Higher-Order Masking of the AES Based on Homographic Functions. In Guang Gong and Kishan Chand Gupta, editors, *Progress in Cryptology - INDOCRYPT 2010 - 11th International Conference on Cryptology in India, Hyderabad, India, December 12-15, 2010. Proceedings*, volume 6498 of *LNCS*, pages 262–281. Springer, 2010.
51. Mathieu Renauld, Dina Kamel, Fran ois-Xavier Standaert, and Denis Flandre. Information Theoretic and Security Analysis of a 65-Nanometer DDSLL AES S-Box. In Preneel and Takagi [48], pages 223–239.
52. Mathieu Renauld and Fran ois-Xavier Standaert. Algebraic Side-Channel Attacks. In Feng Bao, Moti Yung, Dongdai Lin, and Jiwu Jing, editors, *Information Security and Cryptology - 5th International Conference, Inscrypt 2009, Beijing, China, December 12-15, 2009*, volume 6151 of *LNCS*, pages 393–410. Springer, 2009.
53. Mathieu Renauld, Fran ois-Xavier Standaert, and Nicolas Veyrat-Charvillon. Algebraic Side-Channel Attacks on the AES: Why Time also Matters in DPA. In Christophe Clavier and Kris Gaj, editors, *Cryptographic Hardware and Embedded Systems - CHES 2009, 11th International Workshop, Lausanne, Switzerland, September 6-9, 2009, Proceedings*, volume 5747 of *LNCS*, pages 97–111. Springer, 2009.
54. Mathieu Renauld, Fran ois-Xavier Standaert, Nicolas Veyrat-Charvillon, Dina Kamel, and Denis Flandre. A Formal Study of Power Variability Issues and Side-Channel Attacks for Nanoscale Devices. In Kenneth G. Paterson, editor, *Advances in Cryptology - EUROCRYPT 2011 - 30th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tallinn, Estonia, May 15-19, 2011. Proceedings*, volume 6632 of *LNCS*, pages 109–128. Springer, 2011.
55. Matthieu Rivain. On the Exact Success Rate of Side Channel Analysis in the Gaussian Model. In Roberto Maria Avanzi, Liam Keliher, and Francesco Sica, editors, *Selected Areas in Cryptography, 15th International Workshop, SAC 2008, Sackville, New Brunswick, Canada, August 14-15*, volume 5381 of *LNCS*, pages 165–183. Springer, 2008.

56. Matthieu Rivain and Emmanuel Prouff. Provably Secure Higher-Order Masking of AES. In Stefan Mangard and François-Xavier Standaert, editors, *Cryptographic Hardware and Embedded Systems, CHES 2010, 12th International Workshop, Santa Barbara, CA, USA, August 17-20, 2010. Proceedings*, volume 6225 of *LNCS*, pages 413–427. Springer, 2010.
57. Thomas Roche and Emmanuel Prouff. Higher-order Glitch Free Implementation of the AES using Secure Multi-Party Computation Protocols - Extended Version. *J. Cryptographic Engineering*, 2(2):111–127, 2012.
58. Werner Schindler, Kerstin Lemke, and Christof Paar. A Stochastic Model for Differential Side Channel Cryptanalysis. In Josyula R. Rao and Berk Sunar, editors, *CHES*, volume 3659 of *LNCS*, pages 30–46. Springer, 2005.
59. François-Xavier Standaert and Cédric Archambeau. Using Subspace-Based Template Attacks to Compare and Combine Power and Electromagnetic Information Leverages. In Oswald and Rohatgi [47], pages 411–425.
60. François-Xavier Standaert, Tal Malkin, and Moti Yung. A Unified Framework for the Analysis of Side-Channel Key Recovery Attacks. In Antoine Joux, editor, *EUROCRYPT*, volume 5479 of *LNCS*, pages 443–461. Springer, 2009.
61. François-Xavier Standaert, Nicolas Veyrat-Charvillon, Elisabeth Oswald, Benedikt Gierlichs, Marcel Medwed, Markus Kasper, and Stefan Mangard. The World is Not Enough: Another Look on Second-Order DPA. In Masayuki Abe, editor, *ASIACRYPT*, volume 6477 of *LNCS*, pages 112–129. Springer, 2010.
62. François-Xavier Standaert, Eric Peeters, Gaël Rouvroy, and Jean-Jacques Quisquater. An Overview of Power Analysis Attacks against Field Programmable Gate Arrays. *Proceedings of the IEEE*, 94(2):383–394, 2006.
63. François-Xavier Standaert, Christophe Petit, and Nicolas Veyrat-Charvillon. Masking with Randomized Look Up Tables - Towards Preventing Side-Channel Attacks of All Orders. In David Naccache, editor, *Cryptography and Security: From Theory to Applications*, volume 6805 of *LNCS*, pages 283–299. Springer, 2012.
64. Nicolas Veyrat-Charvillon, Benoît Gérard, Mathieu Renaud, and François-Xavier Standaert. An Optimal Key Enumeration Algorithm and Its Application to Side-Channel Attacks. In Lars R. Knudsen and Huapeng Wu, editors, *Selected Areas in Cryptography*, volume 7707 of *LNCS*, pages 390–406. Springer, 2012.
65. Nicolas Veyrat-Charvillon, Benoît Gérard, and François-Xavier Standaert. Security Evaluations beyond Computing Power. In Johansson and Nguyen [34], pages 126–141.
66. Nicolas Veyrat-Charvillon, Marcel Medwed, Stéphanie Kerckhof, and François-Xavier Standaert. Shuffling against Side-Channel Attacks: A Comprehensive Study with Cautionary Note. In Wang and Sako [67], pages 740–757.
67. Xiaoyun Wang and Kazue Sako, editors. *Advances in Cryptology - ASIACRYPT 2012 - 18th International Conference on the Theory and Application of Cryptology and Information Security, Beijing, China, December 2-6, 2012. Proceedings*, volume 7658 of *LNCS*. Springer, 2012.
68. Carolyn Whitnall and Elisabeth Oswald. A Comprehensive Evaluation of Mutual Information Analysis Using a Fair Evaluation Framework. In Phillip Rogaway, editor, *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings*, volume 6841 of *LNCS*, pages 316–334. Springer, 2011.
69. Michael J. Wiener, editor. *Advances in Cryptology - CRYPTO '99, 19th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 1999, Proceedings*, volume 1666 of *LNCS*. Springer, 1999.

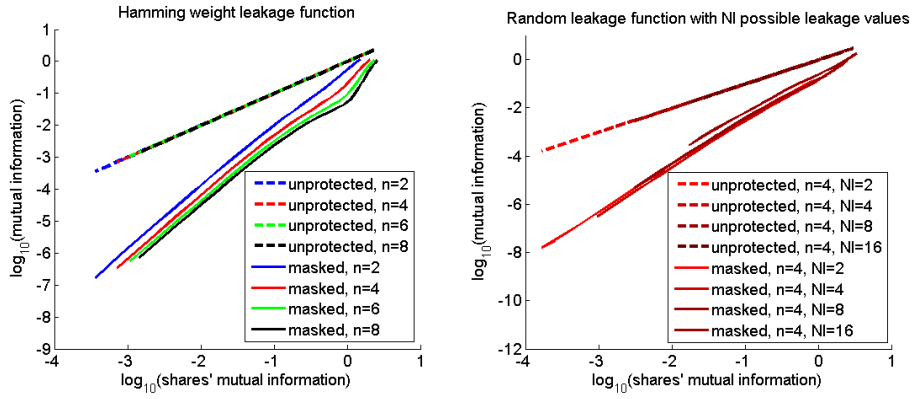


Fig. 10. MI metric in fact. of the shares' MI. HW (left) and random (right) leakages.

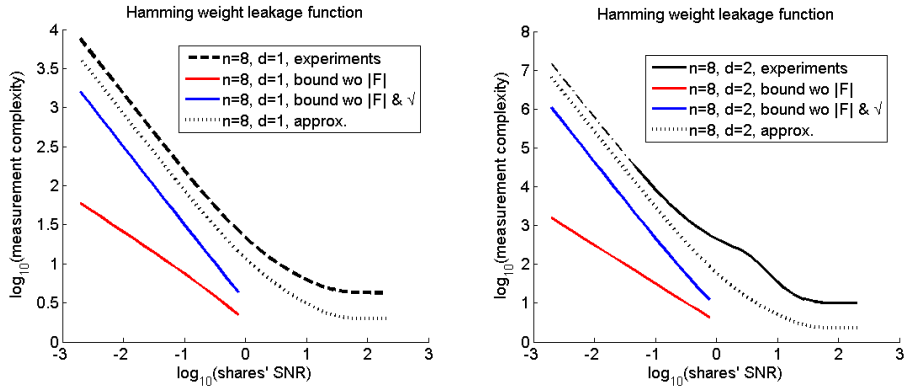


Fig. 11. Measurement complexity and bounds/approximations for concrete TA.

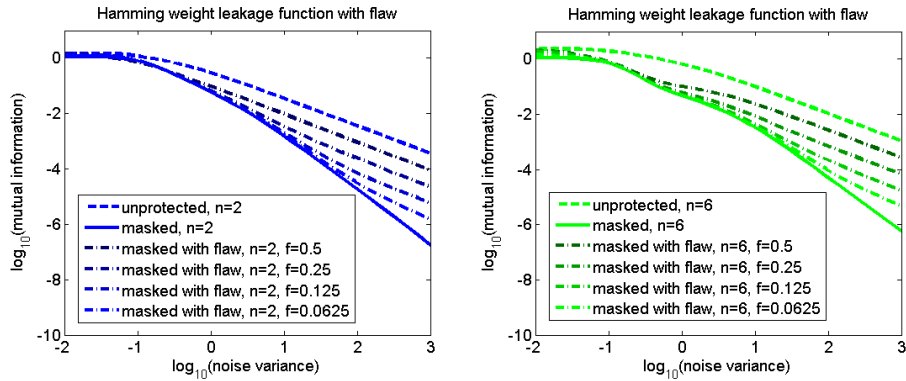


Fig. 12. MI metric for masked implementation with flaw ($n = 2, 6$).