

A Nearly Tight Proof of Duc et al.’s Conjectured Security Bound for Masked Implementations

Loïc Masure¹[0000–0003–2978–4067], Olivier Rioul²[0000–0002–8681–8916], and François-Xavier Standaert¹[0000–0001–7444–0285]

¹ ICTEAM Institute, Université catholique de Louvain, Louvain-la-Neuve, Belgium
² LTCI, Télécom Paris, Institut Polytechnique de Paris, Palaiseau, France

Abstract. We prove a bound that approaches Duc et al.’s conjecture from EUROCRYPT 2015 for the side-channel security of masked implementations. Let Y be a sensitive intermediate variable of a cryptographic primitive taking its values in a set \mathcal{Y} . If Y is protected by masking (a.k.a. secret sharing) at order d (i.e., with $d + 1$ shares), then the complexity of any non-adaptive side-channel analysis — measured by the number of queries to the target implementation required to guess the secret key with sufficient confidence — is lower bounded by a quantity inversely proportional to the product of mutual informations between each share of Y and their respective leakage. Our new bound is nearly tight in the sense that each factor in the product has an exponent of -1 as conjectured, and its multiplicative constant is $\mathcal{O}(\log |\mathcal{Y}| \cdot |\mathcal{Y}|^{-1} \cdot C^{-d})$, where $C \leq 2 \log(2) \approx 1.38$. It drastically improves upon previous proven bounds, where the exponent was $-1/2$, and the multiplicative constant was $\mathcal{O}(|\mathcal{Y}|^{-d})$. As a consequence for side-channel security evaluators, it is possible to provably and efficiently infer the security level of a masked implementation by simply analyzing each individual share, under the necessary condition that the leakage of these shares are independent.

Keywords: Mutual Information · Masking · Convolution · Security Bound

1 Introduction

Evaluating the side-channel security of a cryptographic implementation is a sensitive task, in part due to the challenge of defining the adversary’s capabilities [ABB⁺20]. One approach to deal with this problem is to consider the worst-case security of the implementation, which is characterized by the mutual information between its sensitive intermediate variables and the leakage [SMY09]. Worst-case analysis can be viewed as a natural extension of Kerckhoffs’ laws to side-channel security, where all the implementation details are given to the evaluator who can even profile (i.e., estimate the statistical distribution of the leakage) in an offline phase where he controls the implementation (including its keys and the random coins used in countermeasures). This approach has the advantage of leading to a simple definition of security matching the standard practice of modern cryptography. It has been known for a while that the link

between the mutual information metric and the security of an unprotected implementation is nearly tight [MOS11]. It is also known that if an evaluator can estimate this metric for an implementation with countermeasures, the link with its security level is nearly tight as well [dCGRP19]. This state of the art essentially leaves evaluators with the problem of estimating the mutual information between sensitive intermediate variables and the leakage of an implementation protected with countermeasures, which can be much more challenging since countermeasures typically make leakage distributions more complex.

In this paper, we are concerned with the important case of the masking countermeasures [CJRR99]. Its high-level idea is to split any sensitive variable of an implementation into $d + 1$ shares and to compute on those shares only. As a result, evaluating the worst-case security of a masked implementation requires the characterization of high-order and multivariate distributions, which rapidly turns out to be expensive as the number of shares increases [SVO⁺10]. In order to mitigate this difficulty, a sequence of works has focused on the formal understanding of the masking countermeasure [PR13,DDF14], and its link with concrete evaluation practice [DFS15a]. In this last reference, a lower bound on the minimum number of queries N_a^* required to recover a target secret with a success rate at least β thanks to a side-channel attack was established as:

$$N_a^* \geq \frac{\log(1 - \beta)}{\log\left(1 - \left(\frac{|\mathcal{Y}|}{\sqrt{2}}\right)^{d+1} \prod_{i=0}^d \text{MI}(Y_i; \mathbf{L}_i)^{1/2}\right)} \approx \frac{\log(1 - \beta)}{\left(\frac{|\mathcal{Y}|}{\sqrt{2}}\right)^{d+1}} \prod_{i=0}^d \text{MI}(Y_i; \mathbf{L}_i)^{-1/2} ,$$

where $|\mathcal{Y}|$ stands for the size of the group over which masking is applied. Such bounds are interesting since they reduce the assessment of the success rate of an attack to the evaluation of the $d + 1$ mutual information values between the shares and their corresponding leakage $\text{MI}(Y_i; \mathbf{L}_i)$. Each of these mutual information values is substantially simpler to estimate than $\text{MI}(Y; \mathbf{L})$ since the distribution of the leakage random variable \mathbf{L}_i is a first-order one. Unfortunately, it was also shown in the same paper that this proven bound is no tight. More precisely, empirical attacks suggest that the $-1/2$ exponent for each $\text{MI}(Y_i; \mathbf{L}_i)$ factor might be decreased to -1 and that the $|\mathcal{Y}|^d$ factor would actually be a proof artifact. Halving the exponent has a strong practical impact since it implies that the required number of shares needed to provably reach a given security level might be doubled compared to what is strictly necessary, and the implementation overheads caused by masking scale quadratically in the number of shares. As a result, Duc et al. conjectured that, provided that the shares' leakage is sufficiently small, the lower bound might be tightened as:

$$N_a^* \geq f(\beta) \prod_{i=0}^d \text{MI}(Y_i; \mathbf{L}_i)^{-1} , \tag{1}$$

where $f(\beta) = H(Y) - (1 - \beta) \cdot \log_2(2^n - 1) - H_2(\beta)$ is a function of the attack success rate β , given by Fano's inequality [CT06], as shown in [dCGRP19].³

³ H_2 stands for the *binary* entropy function [dCGRP19].

In this note, we prove a lower bound on N_a^* that fulfills almost all the conditions of Duc et al.’s conjecture. More precisely, we establish a lower bound like the one in Equation 1 with the function $f(\beta)$ divided by a factor $|\mathcal{Y}| \cdot C^d$, with $C = 2 \log(2) \approx 1.38$, regardless of the nature of the group \mathcal{Y} .

The proof is simple to establish. It mixes Chérisey et al.’s inequality and Dziembowski et al.’s *XOR lemma* [dCGRP19,DFS16], and holds for any group-based masking, such as Boolean or arithmetical masking. The former is expressed with the mutual information metric while the latter is expressed with the statistical distance. We bridge the gap between both by converting Dziembowski et al.’s XOR lemma into a variant that is based on the mutual information.

Related works. Prest et al. used the Rényi divergence in order to improve the tightness of masking security proofs but do not get rid of the square root loss (i.e., the $-1/2$ exponent) on which we focus [PGMP19]. Nevertheless, their bound has a logarithmic dependency on the field size $|\mathcal{Y}|$. Liu et al. used the α -information in order to improve Chérisey’s bound [LCGR21]. It could be used to improve our results (at the cost of a slightly less readable bound). In a paper published on the IACR ePrint at the same time as ours (now accepted at CCS 2022), Ito et al. independently obtained a result very similar to ours [IUH22]. Although it is only valid for binary fields (whereas ours is valid for any finite field), their bound is slightly tighter, as the $|\mathcal{Y}|$ is replaced by a $|\mathcal{Y}| - 1$ factor. Interestingly, they also conjecture through the derivation of another bound that is leakage-dependent, and through experimental verifications similar to ours, that the obtained MI-dependent bound is not completely tight in practical cases.

2 Background

2.1 Problem Statement

Let \mathcal{Y} be a finite set. Let $Y \in \mathcal{Y}$ be a random variable denoting the sensitive intermediate variable targeted by a side-channel adversary. In the “standard DPA setting” we consider [MOS11], Y depends on both a uniformly distributed public plaintext and a secret chunk. We assume an implementation that is protected by a d -th order masking. This means that Y is encoded into $d + 1$ shares Y_0, \dots, Y_d such that Y_1, \dots, Y_d are drawn uniformly at random from the group (\mathcal{Y}, \star) and $Y_0 = Y \star (Y_1 \star \dots \star Y_d)^{-1}$, with \star the operation over which the masking is applied (e.g., \oplus for Boolean masking, modular addition $+$ for arithmetic masking, \dots). As required by masking security proofs, we further assume that the shares’ leakage vectors \mathbf{L}_i are the output of a memoryless side-channel and depend only on the realization Y_i , so that the random vectors $\mathbf{L}_0, \dots, \mathbf{L}_d$ are mutually independent.⁴ Intuitively, the goal of a worst-case side-channel security evaluation is to quantify the distance of the random variable Y

⁴ This typically captures a software implementation manipulating the shares sequentially, but as discussed in [BDF⁺17], Lemma 1, it generalizes to parallel (e.g., hardware) implementations as long as the leakage due to the manipulation of shares in parallel can be written as a linear combination of the \mathbf{L}_i vectors.

to the uniform distribution over \mathcal{Y} given the observation of $\mathbf{L} = (\mathbf{L}_0, \dots, \mathbf{L}_d)$. To simplify our computations, we use (in the proof of [Theorem 3](#)) Dziembowski et al.’s reduction to random walks [[DFS16](#), Proof of Lemma 3]. Namely, it is equivalent to consider Y_0 to be uniformly distributed over \mathcal{Y} and to quantify the distance of the variable $Y = Y_0 \star \dots \star Y_d$ to the uniform distribution given the observations of \mathbf{L} .

2.2 Quantifying the Distance to Uniform

To quantify the notion of distance to the uniform distribution over \mathcal{Y} , we will use two different metrics. The first one is widely known in information theory.

Definition 1 (Mutual Information). *Let \mathbf{p}, \mathbf{m} be two Probability Mass Functions (PMFs) over the finite set \mathcal{Y} .⁵ We denote by $D_{\text{KL}}(\mathbf{p} \parallel \mathbf{m})$ the Kullback-Leibler (KL) divergence between \mathbf{p} and \mathbf{m} :*

$$D_{\text{KL}}(\mathbf{p} \parallel \mathbf{m}) = \sum_{y \in \mathcal{Y}} \mathbf{p}(y) \log_2 \left(\frac{\mathbf{p}(y)}{\mathbf{m}(y)} \right) . \quad (2)$$

Then, we define the Mutual Information (MI) between a discrete random variable Y and a continuous random vector \mathbf{L} as follows:

$$\text{MI}(Y; \mathbf{L}) = \mathbb{E}_{\mathbf{L}} \left[D_{\text{KL}} \left(\mathbf{p}_{Y | \mathbf{L}} \parallel \mathbf{p}_Y \right) \right] , \quad (3)$$

where \mathbf{p}_Y and $\mathbf{p}_{Y | \mathbf{L}}$ respectively denote the PMF of Y and the PMF of Y given a realization \mathbf{l} of the random vector \mathbf{L} , with the expectation taken over \mathbf{L} .

The second metric is well-known in the cryptographic community.

Definition 2 (Statistical Distance). *Let \mathbf{p}, \mathbf{m} be two PMFs over the finite set \mathcal{Y} . We denote by $\text{TV}(\mathbf{p}; \mathbf{m})$ the Total Variation (TV) between \mathbf{p} and \mathbf{m} :*

$$\text{TV}(\mathbf{p}; \mathbf{m}) = \frac{1}{2} \sum_{y \in \mathcal{Y}} |\mathbf{p}(y) - \mathbf{m}(y)| . \quad (4)$$

Then, we define the Statistical Distance (SD) as follows:

$$\text{SD}(Y; \mathbf{L}) = \mathbb{E}_{\mathbf{L}} \left[\text{TV} \left(\mathbf{p}_{Y | \mathbf{L}}; \mathbf{p}_Y \right) \right] . \quad (5)$$

Note that both the MI and the SD are so-called *global* metrics and are similarly constructed as the expectation over the marginal leakage distribution of so-called *local* quantities, namely the KL divergence and the TV.

Remark 1. [Equation 5](#) is not a distance between the distributions of \mathbf{L} and Y *per se*, as both random variables are not even defined on the same support. Actually, the TV (local) metric is denoted as SD in the work of Dziembowski et al. [[DFS16](#)], whereas our definition of SD coincides with their definition of bias. Nevertheless, we keep this notation for the SD metric in order to keep consistency with the notations for the MI, as previously done by Prest et al. [[PGMP19](#)].

⁵ We assume without loss of generality that both \mathbf{p} and \mathbf{m} have full support over \mathcal{Y} .

In previous works such as the ones of Prouff et al. [PR13], Duc et al. [DFS15a] or Prest et al. [PGMP19], all the inequalities used are stated in terms of global metrics. The idea of our proof is to rely on some similar inequalities between the local quantities, since they are arguably stronger. We introduce such inequalities hereafter. The first one is the well-known Pinsker’s inequality.

Proposition 1 (Pinsker’s inequality [CT06, Lemma 11.6.1]). *Let p, m be two distributions over the — not necessarily finite — set \mathcal{Y} . Then:*

$$\text{TV}(p; m)^2 \cdot 2 \log_2(e) \leq D_{\text{KL}}(p \parallel m) . \quad (6)$$

Pinsker’s local inequality implies the global inequality

$$2 \text{SD}(Y; \mathbf{L})^2 \leq \text{MI}(Y; \mathbf{L})$$

used by Duc et al. [DFS15a, Thm. 1], and Prest et al. [PGMP19, Prop. 2].

Remark 2. It is possible to find tighter distribution-dependent constants for Equation 6 [OW05]. Nevertheless, the universal constant denoted in the inequality remains the tightest possible if m is the uniform distribution — which is our case of interest.

The second inequality we need is a reversed version of Pinsker’s inequality.

Theorem 1 (Reversed Pinsker’s Inequality [SV15, Thm. 1]). *Let p be a PMF over the finite set \mathcal{Y} , and let u denote the uniform PMF over \mathcal{Y} . Then:*

$$D_{\text{KL}}(p \parallel u) \leq \log_2 \left(1 + 2 |\mathcal{Y}| \text{TV}(p; u)^2 \right) \leq 2 \log_2(e) |\mathcal{Y}| \cdot \text{TV}(p; u)^2 . \quad (7)$$

Again, the reversed Pinsker’s inequality is stronger than some previous results from Prest et al. as it implies the following global inequality established by the authors of [PGMP19]: $\text{MI}(Y; \mathbf{L}) \leq 2 \text{RE}(Y; \mathbf{L}) \cdot \text{SD}(Y; \mathbf{L})$, where $\text{RE}(Y; \mathbf{L})$ stands for the Relative Error (RE) between Y and \mathbf{L} , another global distance metric introduced in this reference. The reversed Pinsker’s (local) inequality is even strictly stronger than Prest et al.’s global inequality, as the former one enables to show that the KL divergence and the squared TV are equivalent metrics (up to a multiplicative constant) whereas, to the best of our knowledge, it is not possible to state that the MI and the squared SD are equivalent global metrics, as the latter one is always bounded by 1, whereas the former one is only bounded by $\log_2 |\mathcal{Y}|$, which can be arbitrarily high. The third inequality we need is the so-called XOR lemma stated by Dziembowski et al. at TCC 2016 [DFS16, Thm. 2]. We next provide a slightly looser version of this result with a simpler proof.

Theorem 2 (XOR Lemma). *Let Y_1, Y_2 be independent random variables on a group \mathcal{Y} , and let u denote the uniform PMF over \mathcal{Y} . Then:*

$$\text{TV}(p_{Y_1 \star Y_2}; u) \leq 2 \cdot \text{TV}(p_{Y_1}; u) \cdot \text{TV}(p_{Y_2}; u) . \quad (8)$$

Proof. Equation 8 is actually a corollary of Young’s convolution inequality [Gra14, Thm. 1.2.10]. Denoting by $*$ the convolution product between two PMFs over \mathcal{Y} , we have:

$$\text{TV}(\mathbf{p}_{Y_1 \star Y_2}; \mathbf{u}) = \frac{1}{2} \|\mathbf{p}_{Y_1 \star Y_2} - \mathbf{u}\|_1 = \frac{1}{2} \|\mathbf{p}_{Y_1} * \mathbf{p}_{Y_2} - \mathbf{u}\|_1 \quad (9)$$

$$= \frac{1}{2} \|(\mathbf{p}_{Y_1} - \mathbf{u}) * (\mathbf{p}_{Y_2} - \mathbf{u})\|_1 \quad (10)$$

$$\leq \frac{1}{2} \|(\mathbf{p}_{Y_1} - \mathbf{u})\|_1 \cdot \|(\mathbf{p}_{Y_2} - \mathbf{u})\|_1 \quad (11)$$

$$= 2 \text{TV}(\mathbf{p}_{Y_1}; \mathbf{u}) \cdot \text{TV}(\mathbf{p}_{Y_2}; \mathbf{u}) \quad (12)$$

where Equation 10 comes from the fact that the uniform PMF is absorbing for the convolution, and Equation 11 comes from Young’s convolution inequality. \square

Theorem 2 is the core of the noise amplification result of Dziembowski et al. and will also be used to argue about it in our proofs. More precisely, we will use the following corollary that is straightforwardly implied by Theorem 2.

Corollary 1. *Let Y_0, \dots, Y_d be independent random variables on a group \mathcal{Y} . Denote $Y_0 \star \dots \star Y_d$ by Y . Then, we have:*

$$\text{TV}(\mathbf{p}_Y; \mathbf{u}) \leq 2^d \cdot \prod_{i=0}^d \text{TV}(\mathbf{p}_{Y_i}; \mathbf{u}) \quad (13)$$

3 Nearly Tight Bounds

We now provide new provable bounds for the worst-case side-channel security of masked cryptographic implementations. We start with an upper bound on the mutual information and follow with a lower bound on the security level.

3.1 Upper Bounding the Mutual Information

We first establish noise amplification in terms of KL divergence.

Proposition 2. *Let Y_0, \dots, Y_d be independent but not necessarily identically distributed random variables over \mathcal{Y} , with PMFs respectively worth \mathbf{p}_{Y_i} . Let $C = 2 \log(2) \approx 1.3862$. Denote the PMF of $Y_0 \star \dots \star Y_d$ as \mathbf{p}_Y . Then:*

$$D_{\text{KL}}(\mathbf{p}_Y \parallel \mathbf{u}) \leq \log_2 \left(1 + |\mathcal{Y}| \cdot \prod_{i=0}^d (C \cdot D_{\text{KL}}(\mathbf{p}_{Y_i} \parallel \mathbf{u})) \right) \quad (14)$$

Proof. Using the inequalities introduced in [subsection 2.2](#), we get:

$$D_{\text{KL}}(\mathbf{p}_Y \parallel \mathbf{u}) \stackrel{(7)}{\leq} \log_2 \left(1 + 2 |\mathcal{Y}| \text{TV}(\mathbf{p}_Y; \mathbf{u})^2 \right) \quad (15)$$

$$\stackrel{(13)}{\leq} \log_2 \left(1 + |\mathcal{Y}| \cdot \prod_{i=0}^d \left(2 \text{TV}(\mathbf{p}_{Y_i}; \mathbf{u}) \right)^2 \right) \quad (16)$$

$$\stackrel{(6)}{\leq} \log_2 \left(1 + |\mathcal{Y}| \cdot \prod_{i=0}^d \left(\frac{2}{\log_2(e)} D_{\text{KL}}(\mathbf{p}_{Y_i} \parallel \mathbf{u}) \right) \right) . \quad (17)$$

□

Having established an amplification result at a local scale, we can now extend it towards the (global) MI metric by taking the expectation over the marginal distribution of the leakage, as stated by the following theorem.

Theorem 3 (MI upper bound (main result)). *Let Y_0, \dots, Y_d be $d + 1$ Independent and Identically Distributed (IID) shares uniformly distributed over \mathcal{Y} . Let $\mathbf{L}_0, \dots, \mathbf{L}_d$ be the leakages occurred by each share. Denote $Y = Y_0 \star \dots \star Y_d$, and $\mathbf{L} = (\mathbf{L}_0, \dots, \mathbf{L}_d)$. Then:*

$$\text{MI}(Y; \mathbf{L}) \leq \log_2 \left(1 + |\mathcal{Y}| \cdot \prod_{i=0}^d C \cdot \text{MI}(Y_i; \mathbf{L}_i) \right) . \quad (18)$$

Proof. We apply [Proposition 2](#) to the random variables

$$Y'_0 = (Y_0 \mid \mathbf{L}_0), \dots, Y'_d = (Y_d \mid \mathbf{L}_d).$$

Therefore, we introduce $Y' = Y'_0 \star \dots \star Y'_d$. As a consequence, each term $D_{\text{KL}}(\mathbf{p}_{Y'_i} \parallel \mathbf{u})$ becomes a random variable depending only on the realization of \mathbf{L}_i . Furthermore, as stated in [subsection 2.1](#), thanks to Dziembowski et al.'s reduction to random walks, the random variables $\mathbf{L}_0, \dots, \mathbf{L}_d$ are mutually independent. As a consequence:

$$\text{MI}(Y; \mathbf{L}) = \mathbb{E}_{\mathbf{L}} \left[D_{\text{KL}}(\mathbf{p}_{Y \mid \mathbf{L}} \parallel \mathbf{p}_Y) \right] = \mathbb{E}_{\mathbf{L}} [D_{\text{KL}}(\mathbf{p}_{Y'} \parallel \mathbf{u})] \quad (19)$$

$$\stackrel{(14)}{\leq} \mathbb{E}_{\mathbf{L}} \left[\log_2 \left(1 + |\mathcal{Y}| \cdot \prod_{i=0}^d \left(C \cdot D_{\text{KL}}(\mathbf{p}_{Y'_i} \parallel \mathbf{u}) \right) \right) \right] \quad (20)$$

$$\leq \log_2 \left(1 + |\mathcal{Y}| \cdot \mathbb{E}_{\mathbf{L}} \left[\prod_{i=0}^d \left(C \cdot D_{\text{KL}}(\mathbf{p}_{Y'_i} \parallel \mathbf{u}) \right) \right] \right) \quad (21)$$

$$\leq \log_2 \left(1 + |\mathcal{Y}| \cdot \prod_{i=0}^d \left(C \cdot \mathbb{E}_{\mathbf{L}_i} [D_{\text{KL}}(\mathbf{p}_{Y'_i} \parallel \mathbf{u})] \right) \right) \quad (22)$$

$$\leq \log_2 \left(1 + |\mathcal{Y}| \cdot \prod_{i=0}^d C \cdot \text{MI}(Y_i; \mathbf{L}_i) \right) , \quad (23)$$

where Equation 21 comes from Jensen’s inequality applied to the logarithm, as it is concave, Equation 22 comes from the independence of the leakages, and Equation 19, Equation 23 come from the definition of MI in Equation 3. \square

Corollary 2. *Using the same notations as in Theorem 3, we have:*

$$\text{MI}(Y; \mathbf{L}) \leq 2 \cdot |\mathcal{Y}| C^d \prod_{i=0}^d \text{MI}(Y_i; \mathbf{L}_i) . \quad (24)$$

Proof. Direct by applying the inequality $\log_2(1 + x) \leq x \log_2(e)$ to Equation 18. \square

Verification on Simulated Measurements. To verify the soundness of the previous bound, we consider a standard simulated setting where the leakage of each share corresponds to its Hamming weight with additive Gaussian noise with variance σ^2 . We first estimate the exact $\text{MI}(Y; \mathbf{L})$ with Monte-Carlo simulations for one, two, four and eight shares [MDP20]. Then, we use the estimated $\text{MI}(Y_i; \mathbf{L}_i)$ of one share (assuming it is equal for all the shares) to derive an upper bound for two, four and eight shares. Figure 1 shows the resulting information theoretic curves in function of the variance of the additive Gaussian noise. It confirms that the bound is nearly tight for binary targets. By contrast, as the the size of the masking field increases, the factor $|\mathcal{Y}|$ of Equation 24 makes it less tight. Based on the results of [DFS15b], we expect it to be a proof artifact, of which the removal in our bound is an interesting open problem.




Fig. 1: MI (plain) and new MI upper bound (dashed) for different field sizes.

3.2 From a MI Upper Bound to a Security Lower Bound

Combining Chérisey’s bound in [dCGRP19] with Equation 18 leads to following corollary that bounds the number of measurement queries needed to guess a target secret with sufficient confidence thanks to side-channel leakage.

Corollary 3. Let \mathcal{A} be any random-plaintext side-channel adversary against an implementation masked at the order d , with each share leaking respectively an amount of information $\text{MI}(Y_i; \mathbf{L}_i)$. Let $\frac{1}{|\mathcal{Y}|} \leq \beta \leq 1$. Then, for \mathcal{A} to succeed in guessing the secret Y with probability higher than β , at least

$$N_a^* \geq \frac{f(\beta)}{\log_2\left(1 + |\mathcal{Y}| \cdot \prod_{i=0}^d C \cdot \text{MI}(Y_i; \mathbf{L}_i)\right)} \geq \frac{f(\beta)}{2|\mathcal{Y}|C^d} \prod_{i=0}^d \text{MI}(Y_i; \mathbf{L}_i)^{-1} \quad (25)$$

measurement queries to the target leaking implementation are needed.

Verification on Simulated Measurements. To verify the soundness of the proposed bound, we simulate a bit recovery using the same leakage model as in our previous simulations, for one, two and three shares, and different levels of noise — captured by the Gaussian noise variance σ^2 .⁶ The results are depicted in plain curves on Figure 2. Based on the Monte-Carlo simulation of the MI for one share — still assuming that the shares verify the same leakage model — we also compute the right hand-side of Equation 25, for $\beta \in [0.5, 1]$. This gives the dotted upper bounds of β depicted on Figure 2. Figure 2a shows that the bound




Fig. 2: Success rate of concrete bit recoveries and MI-based upper bounds.

derived in Equation 25 may not be tight when the MI is high. Nevertheless, we can see from Figure 2b, Figure 2c and Figure 2d that the higher the noise variance (and the lower the MI) the tighter the expected upper bounds.

4 Conclusions

We prove a new bound approaching by Duc et al.’s conjecture for the security of masked implementations. Our result is tight in \mathbb{F}_2 , which makes it practically-relevant since bitslice masking is currently the most efficient way to implement

⁶ The success rate is estimated with bootstrapping, which gives good estimations with a negligible bias provided that the number of simulated traces is far higher than the value of N_a^* such that $\beta = 1$. Due to memory constraints, not enough samples could be drawn to get a consistent simulation in the case where $\sigma^2 = 100, d = 2$.

masking for binary ciphers (especially lightweight ones) [GR17]. For larger field sizes, a factor corresponding to the field size $|\mathcal{Y}|$ makes it less tight. Getting rid of this last source of non-tightness therefore remains as an interesting direction for further improvements. We finally note that we can improve another bound from TCC 2016 — which is stated in terms of statistical distance — as a side-effect of our investigations. We detail this last result in [section A](#).

Acknowledgments. François-Xavier Standaert is a Senior Research Associate of the Belgian Fund for Scientific Research (FNRS-F.R.S.). This work has been funded in part by the ERC project number 724725 (acronym SWORD).

References

- ABB⁺20. Melissa Azouaoui, Davide Bellizia, Ileana Buhan, Nicolas Debande, Sébastien Duval, Christophe Giraud, Éliane Jaulmes, François Koeune, Elisabeth Oswald, François-Xavier Standaert, and Carolyn Whitnall. A systematic appraisal of side channel evaluation strategies. In *SSR*, volume 12529 of *Lecture Notes in Computer Science*, pages 46–66. Springer, 2020. [1](#)
- BDF⁺17. Gilles Barthe, François Dupressoir, Sebastian Faust, Benjamin Grégoire, François-Xavier Standaert, and Pierre-Yves Strub. Parallel implementations of masking schemes and the bounded moment leakage model. In *EUROCRYPT (1)*, volume 10210 of *Lecture Notes in Computer Science*, pages 535–566, 2017. [3](#)
- CJRR99. Suresh Chari, Charanjit S. Jutla, Josyula R. Rao, and Pankaj Rohatgi. Towards sound approaches to counteract power-analysis attacks. In *CRYPTO*, volume 1666 of *Lecture Notes in Computer Science*, pages 398–412. Springer, 1999. [2](#)
- CT06. Thomas M. Cover and Joy A. Thomas. *Elements of information theory (2. ed.)*. Wiley, 2006. [2](#), [5](#)
- dCGRP19. Eloi de Chérisey, Sylvain Guilley, Olivier Rioul, and Pablo Piantanida. Best information is most successful mutual information and success rate in side-channel analysis. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2019(2):49–79, 2019. [2](#), [3](#), [8](#)
- DDF14. Alexandre Duc, Stefan Dziembowski, and Sebastian Faust. Unifying leakage models: from probing attacks to noisy leakage. *IACR Cryptol. ePrint Arch.*, page 79, 2014. [2](#)
- DFS15a. Alexandre Duc, Sebastian Faust, and François-Xavier Standaert. Making masking security proofs concrete - or how to evaluate the security of any leaking device. In *EUROCRYPT (1)*, volume 9056 of *Lecture Notes in Computer Science*, pages 401–429. Springer, 2015. [2](#), [5](#)
- DFS15b. Stefan Dziembowski, Sebastian Faust, and Maciej Skorski. Noisy leakage revisited. In *EUROCRYPT (2)*, volume 9057 of *Lecture Notes in Computer Science*, pages 159–188. Springer, 2015. [8](#)
- DFS16. Stefan Dziembowski, Sebastian Faust, and Maciej Skórski. Optimal amplification of noisy leakages. In *TCC (A2)*, volume 9563 of *Lecture Notes in Computer Science*, pages 291–318. Springer, 2016. [3](#), [4](#), [5](#), [11](#), [12](#)
- GR17. Dahmun Goudarzi and Matthieu Rivain. How fast can higher-order masking be in software? In *EUROCRYPT (1)*, volume 10210 of *Lecture Notes in Computer Science*, pages 567–597, 2017. [10](#)

- Gra14. Loukas Grafakos. Classical fourier analysis, 2014. [6](#)
- IUH22. Akira Ito, Rei Ueno, and Naofumi Homma. On the success rate of side-channel attacks on masked implementations: Information-theoretical bounds and their practical usage. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, CCS 2022, Los Angeles, CA, USA, November 7-11, 2022*, pages 1521–1535. ACM, 2022. [3](#)
- LCGR21. Yi Liu, Wei Cheng, Sylvain Guilley, and Olivier Rioul. On conditional alpha-information and its application to side-channel analysis. In *ITW*, pages 1–6. IEEE, 2021. [3](#)
- MDP20. Loïc Masure, Cécile Dumas, and Emmanuel Prouff. A comprehensive study of deep learning for side-channel analysis. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2020(1):348–375, 2020. [8](#)
- MOS11. Stefan Mangard, Elisabeth Oswald, and François-Xavier Standaert. One for all - all for one: unifying standard differential power analysis attacks. *IET Inf. Secur.*, 5(2):100–110, 2011. [2](#), [3](#)
- OW05. Erik Ordentlich and Marcelo J. Weinberger. A distribution dependent refinement of pinsker’s inequality. *IEEE Trans. Inf. Theory*, 51(5):1836–1840, 2005. [5](#)
- PGMP19. Thomas Prest, Dahmun Goudarzi, Ange Martinelli, and Alain Passelègue. Unifying leakage models on a rényi day. In *CRYPTO (1)*, volume 11692 of *Lecture Notes in Computer Science*, pages 683–712. Springer, 2019. [3](#), [4](#), [5](#)
- PR13. Emmanuel Prouff and Matthieu Rivain. Masking against side-channel attacks: A formal security proof. In *EUROCRYPT*, volume 7881 of *Lecture Notes in Computer Science*, pages 142–159. Springer, 2013. [2](#), [5](#)
- SMY09. François-Xavier Standaert, Tal Malkin, and Moti Yung. A unified framework for the analysis of side-channel key recovery attacks. In *EUROCRYPT*, volume 5479 of *Lecture Notes in Computer Science*, pages 443–461. Springer, 2009. [1](#)
- SV15. Igal Sason and Sergio Verdú. Upper bounds on the relative entropy and rényi divergence as a function of total variation distance for finite alphabets. In *ITW Fall*, pages 214–218. IEEE, 2015. [5](#)
- SVO⁺10. François-Xavier Standaert, Nicolas Veyrat-Charvillon, Elisabeth Oswald, Benedikt Gierlich, Marcel Medwed, Markus Kasper, and Stefan Mangard. The world is not enough: Another look on second-order DPA. In *ASIACRYPT*, volume 6477 of *Lecture Notes in Computer Science*, pages 112–129. Springer, 2010. [2](#)

A Side-Effect: Improving TCC 2016’s Bounds

In [subsection 3.1](#), we have proven that the MI between Y and the whole leakage vector is bounded (up to a multiplicative constant) by the product of the shares’ MIs. Since the squared TV and the KL divergence are consistent for finite sets like \mathcal{Y} , we may wonder whether an upper bound implying the SD as global metric can be derived from the XOR lemma as well. Dziembowski et al. actually provided such an upper bound as recalled hereafter.⁷

⁷ The original version of the theorem is stated for non-uniform secrets. In order to avoid an unfair comparison with respect to Dziembowski et al.’s work, we present an intermediate result of their proof [[DFS16](#), Sec. 3.1.5].

Proposition 3 ([DFS16, Thm. 1(i)], restated). *Let Y be a uniform random variable on a group \mathcal{Y} , encoded by the $(d+1)$ -sharing Y_0, \dots, Y_d . Suppose that all the leakages are δ -noisy for $i = 0, \dots, d$, i.e., for $0 \leq \delta < 1/2$,*

$$\text{SD}(Y_i; \mathbf{L}_i) \leq \delta .$$

Define the noise parameter $\theta = 1/2 - \delta$. Then, for all $\epsilon > 0$, in order to get $\text{SD}(Y; \mathbf{L}) \leq \epsilon$, it is sufficient that the masking order verifies:

$$d \geq 8\theta^{-2} \log\left(\frac{3}{2}\epsilon^{-1}\right) . \quad (26)$$

Informally, [Equation 26](#) gives the sufficient masking order d in order to achieve a desired security level ϵ (expressed in terms of [Statistical Distance \(SD\)](#)), depending on the noise level θ that the developer may leverage — the higher θ , the noisier the leakage model. Unfortunately, the bound [\(26\)](#) is not tight, as the authors also derive the following necessary condition [[DFS16](#), Eq. (12)]:

$$d \geq \frac{\log((2\epsilon)^{-1})}{\log((1-2\theta)^{-1})} , \quad (27)$$

which is asymptotically linear in θ^{-1} when $\theta \rightarrow 0$, whereas it is quadratic in [Equation 26](#). Actually, this is mostly due to the overhead term that occurs from the authors’ so-called “reduction to unconditional random walks”. We shall show that this reduction is not necessary, by leveraging the independence of the leakages to compute the expectation, as in our proof of [Theorem 3](#). As a result, we end up with a tight upper bound, no longer involving the overhead term.

Proposition 4 (Improved bound). *With the same notations as in [Prop. 3](#), it is sufficient that the masking order verifies [Equation 27](#):*

Proof. Starting from the definition of SD ([Definition 2](#)), we have:

$$\text{SD}(Y; \mathbf{L}) = \mathbb{E}_{\mathbf{L}} \left[\text{TV}(\mathbf{p}_Y; \mathbf{p}_{Y|\mathbf{L}}) \right] \quad (28)$$

$$\stackrel{(13)}{\leq} 2^d \mathbb{E}_{\mathbf{L}} \left[\prod_{i=0}^d \text{TV}(\mathbf{p}_{Y_i}; \mathbf{p}_{Y_i|\mathbf{L}_i}) \right] \quad (29)$$

$$= 2^d \prod_{i=0}^d \mathbb{E}_{\mathbf{L}_i} \left[\text{TV}(\mathbf{p}_{Y_i}; \mathbf{p}_{Y_i|\mathbf{L}_i}) \right] \quad (30)$$

$$= 2^d \prod_{i=0}^d \text{SD}(Y_i; \mathbf{L}_i) , \quad (31)$$

where [Equation 30](#) comes from the mutual independence of the leakages \mathbf{L}_i . Now, assuming that for all $i \in \llbracket 0, d \rrbracket$ we have $\text{SD}(Y_i; \mathbf{L}_i) \leq \delta < \frac{1}{2}$ and defining $\theta = \frac{1}{2} - \delta$, we have $\text{SD}(Y; \mathbf{L}) \leq \epsilon$ if $(1-2\theta)^d \leq 2\epsilon$. Hence, the inequality $\text{SD}(Y; \mathbf{L}) \leq \epsilon$ holds if [Equation 27](#) holds. \square