Unifying Leakage Models on a Rényi Day

Thomas Prest[1⋆], Dahmun Goudarzi[1⋆⋆], Ange Martinelli[2], and Alain Passelègue[3,4]

[1] PQShield
{thomas.prest,dahmun.goudarzi}@pqshield.com
[2] Thales
ange.martinelli@thalesgoup.com
[3] INRIA
[4] ENS Lyon
alain.passelegue@inria.fr

Abstract In the last decade, several works have focused on finding the best way to model the leakage in order to obtain provably secure implementations. One of the most realistic models is the noisy leakage model, introduced in [PR13,DDF14] together with secure constructions. These works suffer from various limitations, in particular the use of ideal leak-free gates in [PR13] and an important loss (in the size of the field) in the reduction in [DDF14].
In this work, we provide new strategies to prove the security of masked implementations and start by unifying the different noisiness metrics used in prior works by relating all of them to a standard notion in information theory: the pointwise mutual information. Based on this new interpretation, we define two new natural metrics and analyze the security of known compilers with respect to these metrics. In particular, we prove (1) a tighter bound for reducing the noisy leakage models to the probing model using our first new metric, (2) better bounds for amplification-based security proofs using the second metric.
To support that the improvements we obtain are not only a consequence of the use of alternative metrics, we show that for concrete representation of leakage (*e.g.,* "Hamming weight + Gaussian noise"), our approach significantly improves the parameters compared to prior works. Finally, using the Rényi divergence, we quantify concretely the advantage of an adversary in attacking a block cipher depending on the number of leakage acquisitions available to it.

1 Introduction

In modern cryptography, it is common to prove the security of a construction by relying on the security of its underlying building blocks or on the hardness of standard computational problems. This approach has allowed the community

⋆ Part of this work was done when the author was an engineer at Thales.
⋆⋆ Part of this work was done when the author was a PhD student at CryptoExperts and École Normale Supérieure.

to propose a wide variety of cryptographic primitives based on only a limited number of different assumptions (e.g., factoring, learning parity with noise, existence of one-way functions, security of AES or SHA-3, etc). Unfortunately, there is still a significant gap between the *ideal security models* that are used in provable security, and the *actual environments* in which these cryptosystems are deployed. Notably, standard security models usually assume that attackers have only a black-box access to the cryptosystem: attackers do not have any information beyond the input/output behavior.

Yet, it is well known that this is generally *not true in practice*. These cryptosystems are run by physical devices, hence an adversary might be able to learn partial information such as the running-time, the power consumption, the electromagnetic emanation, or several other physical measures of the device. As revealed by Kocher et al. in [Koc96,KJJ99], these additional information, referred to as the *leakage of the computation*, are valuable and can be used to mount *side-channel attacks* against cryptographic *implementations*. Hence, a cryptosystem that is proven secure in an ideal security model can become completely vulnerable when deployed in the real-world.

Due to the fundamental importance of secure implementations of cryptographic primitives, constructing leakage-resilient cryptography has become a major area of research. Many empirical countermeasures have been proposed over the last decades and an important line of works has aimed at formalizing the notion of leakage towards obtaining provably secure implementations.

The presence of leakage in the real-world has been formalized by introducing new security models in which the attacker can obtain additional information about the computation. In a seminal work from 2003 by Ishai, Sahai, and Wagner [ISW03], the authors introduced the d-probing (or d-threshold probing) model, in which an attacker can learn a *bounded number d of intermediate results* (i.e. wire values, also called *probes*) of a computation C. A circuit is then secure in this model if any subset of at most d probes does not reveal any information about the inputs of the computation. That is, the distribution of values obtained by probing should be independent of the inputs of the computation. While this model is ideal and does not fully catch the behavior of a device in the real-world (*e.g.,* physical leakages reveal information about the whole computation), it is simple enough to get efficient compilers that transform any circuit into a secure one in the d-probing model, as shown in [ISW03]. They built secure addition and multiplication in the d-probing model based on secret-sharing techniques[1] and immunize any arithmetic circuit by replacing every gate by its secure variant. This transformation blows up the size of the circuit by a factor $O(d^2)$. A different and more realistic model was proposed by Micali and Reyzin in 2004 [MR04]. They defined a model of cryptography in presence of arbitrary forms of leakage about the whole computation. The above two works are cornerstones of leakage-resilient cryptography. In particular, the assumption that *only*

[1] Basically, their secure variants take as input additive shares of the input and produce additive shares of the output. Their secure multiplication that operates on additive shares is often referred to as the ISW-multiplication.

computation leaks information (thus a program can still hide some secrets) originated in these works. Following the path of [MR04], Dziembowski and Pietrzak proposed in 2008 a simplified model for *leakage-resilient cryptography* in [DP08]. In this model, any elementary operation on some input x leaks a partial information about x, modeled as the evaluation $f(x)$ of a *leakage function* whose range is bounded, so an adversary is given access to the leakage $f(x)$ for every intermediate result x of the evaluation of C. Unfortunately, this model has a drawback: the range of the leakage function is bounded and fairly small (e.g., 128-bit strings) compared to the actual amount of information that can be obtained from a device (e.g., a power trace on an AES computation can contain several megabytes of information).

To circumvent this limitation, Prouff and Rivain proposed in [PR13] a more realistic leakage model, called the *noisy leakage model*. The authors modified the above definition of leakage by making an *additional but realistic assumption*: the information $f(x)$ leaked by an elementary operation on some input x is *noisy*. Specifically, the authors assumed that f is a randomized function such that the leakage $f(x)$ only implies a bounded bias in the distribution of x, which is formally defined as distributions X and $X|f(X)$ being close (up to some fixed bound δ), where X denotes the distribution of x. The authors measured *closeness* with the Euclidean Norm (denoted EN) between the distributions (over finite sets) and propose solutions to immunize symmetric primitives in this model. Their model is inspired by the seminal work of Chari et al. [CJRR99] that considered the leakage as inherently noisy and proved that using additive secret-sharing (or masking) on a variable X decreases the information revealed by the leakage by an exponential factor in the number of shares (or masking order). This kind of proof is referred to as amplification-based, and Prouff and Rivain extended it to a whole block cipher evaluation.

A drawback of this model is the difficulty to design proofs. In addition, the constructions in [PR13] rely on a fairly strong assumption: the existence of *leak-free refresh gates* (i.e. gates that do not leak any information and refresh additive shares of x)[2]. Both limitations were solved by Duc, Dziembowski, and Faust in [DDF14]. In the latter work, using the statistical distance (denoted SD) instead of the Euclidean norm as measure of closeness, the authors showed that constructions proven secure in the ideal d-probing model of Ishai et al. are also secure in the δ-noisy leakage model, provided d is large enough (function of δ). As a consequence, the simple compilers for building d-probing secure circuits can serve for achieving security in the noisy leakage model, proving a conjecture broadly admitted for several years based on empirical observations.

The present work extends the two above results and proposes general solutions to immunize cryptographic primitives in the noisy leakage model. We start by giving a more formal overview of these two works.

[2] In practice, refresh gates are often implemented via an ISW-multiplication with additive shares of 1 (*e.g.,* shares (1,0,...,0)).

1.1 Previous Works

As explained above, two distinct approaches for immunizing cryptosystems in the noisy leakage model have been considered: (1) a direct approach, used in [PR13], that proves the security of a construction directly in the model via noise amplification, and (2) an indirect approach, used in [DDF14], that consists in reducing security in the noisy leakage model to security in ideal models (*e.g.,* the probing model) and then applying compilers for the latter models.

A direct approach. In [PR13], the authors propose a way to immunize block ciphers of a particular form (succession of linear functions and substitution boxes, a.k.a. s-boxes, *e.g.,* AES). Their approach consists in replacing elementary operations of such block ciphers by subprotocols that operate on masked inputs and produce a masked output. They bound the leakage on each subprotocol and as a consequence are able to bound the leakage of a single evaluation of the *masked block cipher* (i.e. the block cipher obtained by replacing every elementary operation by the corresponding subprotocol and applying leak-free refresh gates between each subprotocol). They conclude by proving an upper bound on the information (in an information-theoretic sense) revealed by the leakage about the input (plaintext/key) from evaluations of the masked block cipher, in particular proving that it decreases exponentially in the masking order.

While this paper makes great progress towards constructing provably-secure leakage-resilient block ciphers, it suffers from a few limitations. First, as already mentioned, the security proof relies on leak-free refresh gates. Second, the fact that the final analysis relies on the mutual information implies a rather paradoxal situation: from an information theory perspective, a single pair of plaintext-ciphertext can reveal the key. To get around this problem, the authors assume that both the plaintext and the ciphertext are secret, which is fairly unrealistic compared to standard security models for block ciphers. Finally, to offer strong security guarantees, the mutual information should be upper bounded by $2^{-O(\lambda)}$, with λ being the security parameter. Hence, the masking order for reaching this bound only depends on λ, which is independent of the number of queries (and therefore the amount of leakage) the adversary makes.

An indirect approach. In [DDF14], the authors propose an elegant approach that applies to any form of computation. Their main result proves that any information obtained in the δ-noisy leakage model (so information of the form $f(x)$ for any intermediate result x of the computation) can be simulated from a sufficiently large number d of probes. As such a set of probes does not carry any information about the inputs if the circuit is secure in the $(d+1)$-probing model, this guarantees that the information obtained in the δ-noisy leakage model does not carry any useful information either. Hence, using standard compilers to secure a cryptosystem in the $(d+1)$-probing model makes it secure when deployed in the real-world, assuming the leakage is δ-noisy. Unfortunately, this reduction incurs an important blow-up in the parameters ($\delta \rightarrow d$). Notably d has to be at

least N times larger than δ to guarantee security, where N is the size of the field on which the circuit operates. This loss appears in an intermediate step of their reduction when first reducing the noisy leakage model to the random probing model[3]. Typically, for AES, we have $N = 256$, so the required order d of security is very large (and so is the size of the masked circuit since applying the ISW compiler increases the size by a factor d^2).

This loss is seemingly an artifact of the reduction and has not been observed in empirical measures [DFS15a]. A first attempt to circumvent this issue was made in [DFS15b] by introducing a new model, called the average random probing model, which is a tweak of the random probing model. The authors prove a tight equivalence between the noisy leakage and the average random probing models and show that the ISW compiler is secure in their model.

Yet, there are two caveats. First, their proof of security of the ISW compiler introduce leak-free gates, whereas [DDF14] does not. Second, [DFS15b] does not establish a reduction from the average random probing to the threshold probing model, hence leaving open the question of improving the reductions provided in [DDF14]. In this paper, we overcome these two issues and provide a tight reduction from a[4] noisy leakage model to the threshold probing model without leak-free gates nor a loss in the size of the field.

1.2 Our Contributions

We extend the previous studies of leakage-resilient cryptography in several directions. Our approach starts by relating the noisiness of a leakage to a standard notion in information theory: the pointwise mutual information (PMI).

From pointwise mutual information to noisiness metrics. Our first observation is that the two metrics used in prior works to measure the distance between X and $X|f(X)$, namely the Euclidean norm (EN) and the statistical distance (SD), can be easily expressed as different averages of the pointwise mutual information of the same distributions. Given this interpretation, it is easy to see that these two measures are *average-case* metrics of noisiness.

We investigate the benefits of considering the problem of building leakage-resilient cryptography based on two other *worst-case* metrics that naturally follow from the pointwise mutual information: the Average Relative Error (ARE) and the Relative Error (RE). Using these two metrics, we propose tighter proofs for immunizing cryptosystems in the noisy leakage model. We emphasize that

[3] In the ϵ-random probing model, the adversary learns each exact wire value with probability ϵ (and nothing about it with probability $1 - \epsilon$).

[4] Noisy-leakage models are inherently associated to the metric used to measure noisiness. [PR13] is based on the Euclidean norm, [DDF14] on the statistical distance. We introduce new metrics, therefore new models. Yet, the overall result remain comparable as only the noisiness of the leakage is impacted by the metric, but not the leakage itself, so the metric is just a tool to argue the security (security against the leakage being independent of the metric).

even though we introduce new metrics (and therefore new noisy leakage models), our goal remains to prove that we can simulate *perfectly* the leakage (which depends on the intermediate values but *does not depend* on the metric) from a certain amount of probes. The metric only plays a role in determining the amount of probes that is needed for simulating the leakage, i.e. the sufficient masking order to immunize the computation, but does not play any role in measuring the quality of the simulation (which remains perfect). We are able (in general) to prove better bounds for the amount of probes needed for the simulation. In particular, combining our results with known compilers is particularly interesting for typical forms of concrete leakage such as the "Hamming weight + Gaussian noise" model.

A tighter reduction from noisy leakage to random probing. We propose a reduction from the noisy leakage model to the random probing model, when the noise is measured with the ARE metric. Our reduction is analogous to the reduction proposed from [DDF14]. Once reduced to the random probing model, it is easy to go to the threshold probing model by a simple probabilistic argument (observed in [DDF14]). Using the ARE metric, we are able to reduce the δ-noisy leakage model (where the noise is measured with the ARE metric) to the δ-random probing model (instead of the $\delta \cdot N$-random probing model for prior work using the SD metric). Again, we emphasize that, despite using different metrics, these reductions allows to simulate the exact distribution of the leakage, which is completely independent of the underlying metric.

This tighter reduction has immediate, tangible consequences when considering compilers which are proven secure in the threshold probing model [ISW03] or in the random probing model [ADF16,GJR17,AIS18]: for a specific form of noisy leakage, as long as the ARE-noisiness is smaller than N times than the SD-noisiness, our reduction guarantees security using a smaller masking order than the reduction based on the SD metric. In particular, we show for the concrete "Hamming weight + Gaussian noise" model of leakage that our result reduces the required masking order by a factor $O(N/\sqrt{\log N})$ compared to [DDF14].

Actually, even though we do not start from the same metrics (and then from the exact same noisy leakage model), we prove that the ARE-noisiness of any function is upper bounded (up to a factor $2 \cdot N$) by its SD-noisiness. Then, even in the worst case, our reduction (which is tighter by a factor N) gives as good results (up to a factor 2) as the reduction in [DDF14]. Reversely the SD-noisiness is upper bounded by the ARE-noisiness (up to a factor 2), so the loss of a factor N in the reduction is not compensated, which explains the large improvement we gain from our approach in certain cases such as the aforementioned one.

As a side contribution, and perhaps surprisingly, we are also able to prove a converse reduction: we show that the random probing model reduces to the ARE-noisy leakage model (though it incurs a loss of a factor $N-1$). This follows from observing that the random probing model is a special instance of the ARE-noisy leakage model. This implies that the SD-noisy leakage, ARE-noisy leakage

and (average) random probing models are all equivalent. We believe that this result is of independent interest and could find applications in future works.

While we focus on using a compiler introduced in [ISW03], which has also been studied in [PR13,DDF14,DFS15b], other compilers also benefit from our work in obvious ways (e.g., the compilers described in [ADF16,GJR17,AIS18] are secure in the random probing model, hence benefit from our reduction to the noisy leakage model).

Our reductions and previously known reductions are summarized in Figure 1. This diagram represents the interactions between various leakage models (from very concrete ones, like "Hamming weight with Gaussian noise", to theoretical models such as the threshold probing model) and circuit compilers. The physical noise model is displayed on the first line, noisy leakage models on the second line, probing models on the third line, and circuit compilers are displayed on the fourth line. Arrows from a model M to a compiler C means that C is proven secure in the model M. An arrow from a model M_1 to a model M_2 means that an adversary in M_1 can be simulated in M_2 with the overhead indicated next to the arrow. Our contributions (models and reductions) are displayed in bold. For the sake of clarity, constant factors are omitted. N denotes the size of the underlying finite field, and λ denotes the security parameter of the scheme to protect.

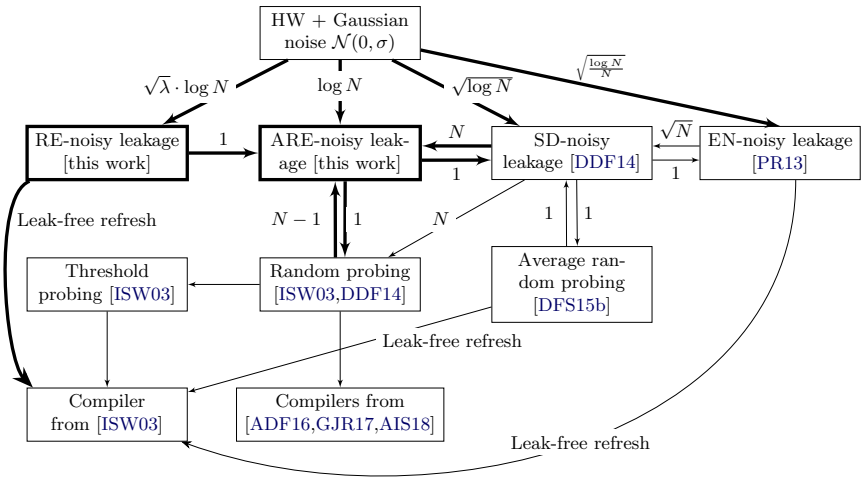

Figure 1: From concrete leakages to secure circuit compilers: an overview of reduction-based proofs and our contributions.

An amplification-based proof with the Rényi divergence. Our second main contribution is a new amplification-based proof which improves over existing ones in some aspects. Once again, we put our result in perspective with concrete noisy leakage models where the noise follows a Gaussian distribution $\mathcal{N}(0, \sigma)$ with standard deviation σ, *e.g.,* the "Hamming weight + Gaussian noise" model. In the context of leakage-resilient cryptography, known amplification-based proofs show that if σ is large enough, then the leakage of a masked circuit decreases exponentially in the masking order; equivalently (and we will use this perspective for convenience), it shows that the required amount of Gaussian noise decreases when the masking order increases.

The most notable amplification-based proofs of masked circuits are due to [PR13], which uses the EN-noisy leakage model, and [DFS15b], which uses the average random probing model (or equivalently, the SD-noisy leakage model). Both works yield a condition on σ, precisely they impose $\sigma = \Omega(d \times f \times g^{1/(d+1)})$, where the functions f and g are constant in the masking order d. Here, f acts like a factor of σ which is *fixed* (it does not depend of d), whereas g acts like a *compressible* part whose impact on σ can be decreased by increasing the masking order. Both terms are important, because f cannot be compressed, but g can be very large in practice. Our new amplification-based proof relies on the RE-noisiness, and can be seen as revisiting the proof of [PR13]. Compared to the previous works, it provides several qualitative and quantitative gains:

- Whereas in the previous works, σ was exponential in the security level λ (more precisely, larger than $2^{\lambda/(d+1)}$), in our case it is only proportional to $\sqrt{\lambda}$; This is thanks to our use of the Rényi divergence, which allows to replace $2^{\lambda/(d+1)}$ by $q^{1/(d+1)}$, where q denotes the number of traces (i.e. the number of evaluations with known leakage) obtained by the attacker. This is a far lighter constraint, since in cryptography it is typical to take $\lambda = 256$, whereas it is extremely rare to have more than 2^{32} traces available.
- Our Rényi divergence-based proof shows that the view of a black-box adversary is not significantly different from the view of an adversary which has access to leakage, and we relate the distance between these two views to the masking order and the number of traces available to the adversary (in particular upper bounded by the number of queries).
- Compared to [DFS15b], our fixed part f is larger, but our compressible part g is much smaller: for the above values of q and λ, g will be 2^{32} in our case, whereas it would be larger than 2^{1024} in the case of [DFS15b]. In addition, [DFS15b, Lemma 14 and Theorem 1] implicitly impose d to be linear in $\lambda + \log N$, which gives an extremely high masking order. Our proof imposes no such bound.

In Figure 1, amplification proofs correspond to *Leak-Free Refresh* arrows.

Finally, in Table 1, we compare our results with the state-of-the-art approaches in the case of Hamming weight + Gaussian noise for both reduction-based proofs and amplification-based proofs. Our bound for the noisiness are taken from Proposition 3. The conditions on the Gaussian noise level σ are given, as well as additional conditions when they exist. *LFR* indicates whether

Table 1: Comparison with prior works (combined with Proposition 3).

Work	Condition on σ	Other condition	LFR	Model	Tool
[DDF14, Thm 1]	$\Omega\left(dN\sqrt{\ln N}\right)$	$d = \Omega(\lambda + \ln \lvert\Gamma\rvert)$	No	CPA	Δ_{SD}
This work (sec. 5)	$\Omega\left(d \ln N\right)$	$d = \Omega(\lambda + \ln \lvert\Gamma\rvert)$	No	CPA	Δ_{SD}
[PR13, Cor 2, Thm 4]	$\Omega\left(dN\sqrt{\ln N} \times (N^3 2^\lambda \lvert\Gamma\rvert)^{1/(d+1)}\right)$	$d = \Omega(dN^{3/2}\sqrt{\ln N})$	Yes	RPA	MI
[DFS15b, Cor 4]	$\Omega\left(d\sqrt{\ln N} \times \left[(Nd2^\lambda)^4 \lvert\Gamma\rvert\right]^{1/d}\right)$	$d = \Omega(\lambda + \ln(N\lvert\Gamma\rvert))$	Yes	CPA	Δ_{SD}
This work (sec. 6)	$\Omega\left(d\sqrt{\lambda} \ln N \times (q\lvert\Gamma\rvert)^{1/(d+1)}\right)$	-	Yes	CPA	R_∞

leak-free refresh gates are required in the security proof. *Model* states the model of attacker (random-plaintext or chosen-plaintext). The model of attack is actually not considered in [DFS15b], but [DFS16, Lemma 2] shows that in the case of [DFS15b], random plaintext attacks reduce to chosen plaintext attacks and that it is therefore sufficient to consider only the former. *Tool* indicates the main notion the security proof relies on (statistical distance, mutual information or Rényi divergence of order infinity). λ denotes the security parameter of the scheme, d the masking order, N the size of the underlying field, and q the number of traces available to an attacker.

Organization of the paper. The remainder of the paper is organized as follows. Section 2 presents some theoretical background and notation. Section 3 provides a unifying background for the metrics used in prior works as well as those we introduce. Section 4 builds the bridge from a standard, concrete model of leakage (Hamming weight with Gaussian noice) to noisy leakage models. In Section 5, we detail our tight reduction from the noisy leakage model to the probing model. Our amplification-based proofs are described Section 6.

2 Preliminaries

In this section we recall basic notation and notions used throughout the paper.

2.1 Notation

For any $\ell \geq 1$, we denote by $[\ell]$ the set $\{1, \ldots, \ell\}$. We denote by $\mathcal{X}$ a finite set, by x an element of $\mathcal{X}$, by X a random variable over $\mathcal{X}$, and by $\mathcal{P}_X$ the corresponding probability mass function (i.e. the function $\mathcal{P}_X : x \mapsto \mathbb{P}[X = x]$). We often abuse notation and denote by P the distribution defined by a probability mass function $\mathcal{P}$. For a distribution P over $\mathcal{X}$, we denote by $x \leftarrow P$ the action of sampling x from the distribution P.

For any distribution P and any function f over $\mathcal{X}$, we denote by $f(P)$ the distribution of $f(x)$ induced by sampling $x \leftarrow P$. We denote by $\mathrm{Supp}(X) := \{x \in \mathcal{X} \mid \mathcal{P}_X(x) > 0\}$ the support of a random variable X over $\mathcal{X}$ (and we define similarly the support of a distribution).

For any random variable X over $\mathcal{X}$ and a function $f : \mathcal{X} \to \mathcal{Y}$, we use the following notation:

$$\mathbb{E}_X[f(X)] = \sum_x f(x) \cdot \mathbb{P}[X = x] \ .$$

For two random variables X, Y over $\mathcal{X}$, the *statistical distance* between X and Y is defined as:

$$\Delta_{\mathrm{SD}}(X;Y) := \frac{1}{2} \sum_{x \in \mathcal{X}} |\mathbb{P}[X = x] - \mathbb{P}[Y = x]| \ .$$

Similarly, the Euclidean norm between X and Y is defined as:

$$\Delta_{\mathrm{EN}}(X;Y) = \sqrt{\sum_{x \in \mathcal{X}} \left(\mathbb{P}[X = x] - \mathbb{P}[Y = x]\right)^2} \ .$$

Finally, if X, Y have the same support, their relative error is:

$$\Delta_{\mathrm{RE}}(X;Y) := \max_{x \in \mathrm{Supp}(X)} \left| \frac{\mathbb{P}[X = x]}{\mathbb{P}[Y = x]} - 1 \right| \ .$$

We now recall these two definitions from [DDF14] and [PR13]:[5]

$$\mathrm{SD}(X|Y;X) = \sum_y \mathbb{P}[Y = y] \cdot \Delta_{\mathrm{SD}}(X|Y = y;X)$$
$$\mathrm{EN}(X|Y;X) = \sum_y \mathbb{P}[Y = y] \cdot \Delta_{\mathrm{EN}}(X|Y = y;X)$$

2.2 The Rényi Divergence

The Rényi divergence [Ré61] is a measure of divergence between distributions. In the recent years, it has found several applications in lattice-based cryptography [BLL+15,Pre17]. When used in security proofs, its peculiar properties allow designers of cryptographic schemes to set some parameters according to the number of queries allowed to an attacker, rather than to the security level, and this has often resulted in improved parameters. We first recall its definition as well as some standard properties.

[5] Instead of SD and EN, [DDF14] and [PR13] used the notations Δ and β; we prefer our notation as it avoids any confusion with greek letters denoting scalars.

Definition 1 (Rényi divergence). *Let P, Q be two distributions over X such that $\mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q)$. For $a \in (1, +\infty)$, their Rényi divergence of order a is:*

$$R_a(P\|Q) = \left(\sum_{x \in \mathrm{Supp}(P)} \frac{P(x)^a}{Q(x)^{a-1}} \right)^{\frac{1}{a-1}} .$$

In addition, the Rényi divergence of order $+\infty$ is

$$R_\infty(P\|Q) = \max_{x \in \mathrm{Supp}(P)} \frac{P(x)}{Q(x)} .$$

This definition is common in the lattice-based cryptography literature, whereas the information theory literature favors its logarithm as the definition. Classical properties of the Rényi divergence may be found in [FHT03], and cryptographic properties may be found in [BLL+15,Pre17]. In this paper, we use the following composition properties from [BLL+15].

Lemma 1. *For two distributions P, Q and two families of distributions $(P_i)_i, (Q_i)_i$, the Rényi divergence verifies the following properties:*

- ***Data processing inequality:*** *For any function f, $R_a(f(P)\|f(Q)) \leq R_a(P\|Q)$.*
- ***Multiplicativity:*** *$R_a(\prod_i P_i \| \prod_i Q_i) = \prod_i R_a(P_i\|Q_i)$.*
- ***Probability preservation:*** *For any event $E \subseteq \mathrm{Supp}(Q)$ and $a \in (1, +\infty)$,*

$$Q(E) \geq P(E)^{\frac{a}{a-1}} / R_a(P\|Q) ,$$
$$Q(E) \geq P(E) / R_\infty(P\|Q) .$$

2.3 Pointwise Mutual Information

The pointwise mutual information is a common tool in computational linguistics [CH89], where it serves as a measure of co-occurence between words. For example, the pmi of "Sean" and "Penn" is high because Sean Penn is a well-known person, whereas the pmi of "bankruptcy" and "success" is low because the two words are rarely used in the same sentence.

Formally, the pointwise mutual information is defined as follows.

Definition 2 (Pointwise mutual information). *Let X, Y be random variables over $\mathcal{X}$. Then, for any $(x, y) \in \mathrm{Supp}(X) \times \mathrm{Supp}(Y)$, we have:*

$$\mathrm{pmi}_{X,Y}(x, y) = \log \left(\frac{\mathbb{P}[X = x, Y = y]}{\mathbb{P}[X = x]\mathbb{P}[Y = y]} \right) .$$

We also define its exponential form as:

$$\mathrm{PMI}_{X,Y}(x, y) = e^{\mathrm{pmi}_{X,Y}(x,y)} - 1 = \frac{\mathbb{P}[X = x, Y = y]}{\mathbb{P}[X = x]\mathbb{P}[Y = y]} - 1 .$$

We note that when they are close to 0, $\mathrm{pmi}_{X,Y}(x,y) \sim \mathrm{PMI}_{X,Y}(x,y)$. *The mutual information between* X *and* Y *can be simply expressed from the pointwise mutual information, since we have:*

$$\mathrm{MI}(X;Y) = \mathbb{E}_{(X,Y)}\left[\mathrm{pmi}_{X,Y}\right] \ ,$$

where (X,Y) *denotes the joint distribution of* X *and* Y. *When* X *and* Y *are clear from context, we may omit the subscripts and simply note* pmi *and* PMI.

Interestingly, as we show in the next section, several metrics in leakage-resilient cryptography can be defined simply using the pointwise mutual information.

3 Unifying Leakage Models via the Pointwise Mutual Information

As already explained, in the noisy leakage model (defined below), an adversary learns noisy information $f(x)$ about every intermediate result x of a computation. The hope is that this leakage does not reveal much information about the actual value x, which is translated by the fact that the distribution X is *close* to the distribution $X|f(X)$. Two main notions of *closeness* (corresponding to two noisiness metrics) have been proposed, namely EN and SD.

3.1 Noisiness Metrics from Pointwise Mutual Information

It appears that the above noisiness metrics can easily be related to the pointwise mutual information, as we state in the following immediate proposition. Other natural metrics can also be derived from the pointwise mutual information, and we define two additional metrics in the subsequent definition.

Let us define the following four metrics with respect to the PMI.

Definition 3 (Noisiness metrics). *Let* X, Y *be random variables over sets* $\mathcal{X}, \mathcal{Y}$ *respectively. We define the following metrics based on the pointwise mutual information:*

- $\mathrm{SD}(X|Y) := \frac{1}{2} \cdot \mathbb{E}_X \mathbb{E}_Y \left[|\mathrm{PMI}|\right]$;
- $\mathrm{EN}(X|Y) := \mathbb{E}_Y \sqrt{\mathbb{E}_X \left[\mathbb{P}[X]\, \mathrm{PMI}^2\right]}$;
- $\mathrm{RE}(X|Y) := \max_{x,y} |\mathrm{PMI}|$;
- $\mathrm{ARE}(X|Y) := \mathbb{E}_Y \left[\max_x |\mathrm{PMI}|\right]$.

The four notions of noisiness defined here compute different norms of the $(\mathrm{PMI})_{x,y}$: SD compute the average value of $|\,\mathrm{PMI}\,|$, RE computes its max, and ARE computes something in between.

Note that this difference in their definition (average-case vs worst-case) is mirrored in the random probing models (average random probing vs random probing), so it is perhaps unsurprising that reductions between worst-case models (ARE-noisy leakage to random probing in Section 5) incur no loss, as well

as those between average-case models (SD-noisy leakage and average random probing in [DFS15b]), but that the worst-case-average-case reduction of [DDF14] incurs a loss by a factor $|\mathcal{X}|$.

We note that these definitions of SD and EN match the ones given in section 2.1: $\mathrm{SD}(X|Y) = \mathrm{SD}(X|Y;X)$ and $\mathrm{EN}(X|Y) = \mathrm{EN}(X|Y;X)$. This is done on purpose as we aim at introducing new noisiness metrics without discarding previously defined ones. We do so by expressing them all with a single common notion: the pointwise mutual information. The acronyms RE and ARE stand for Relative Error and Average Relative Error. We note that $\mathrm{RE}(X|Y) = \max_y \Delta_{\mathrm{RE}}(X|Y=y;X)$ and $\mathrm{ARE}(X|Y) = \mathbb{E}_Y \Delta_{\mathrm{RE}}(X|Y;X)$.

We now define a generic notion of noisy functions, parameterized by any of the above metrics.

Definition 4 (Noisy functions). *Let $D \in \{\mathrm{SD}, \mathrm{EN}, \mathrm{RE}, \mathrm{ARE}\}$ be one of the metrics defined in definition 3, X be a random variable over a set $\mathcal{X}$ and $\delta \geq 0$. We say that a function $f : \mathcal{X} \to \mathcal{Y}$ is δ-noisy for the metric D and the random variable X (or for short, δ-D-noisy for X) if:*

$$D(X|f(X)) \leq \delta \ .$$

If X follows the uniform distribution, we simply say that f is δ-D-noisy.

This definition highlights an important caveat of the noisy leakage model: the notion of noisy function is implicitly parameterized by an underlying distribution X. However, we will later show in Lemma 2 than for RE- and ARE-noisy functions, we can abstract ourselves from the underlying distribution at the cost of essentially a factor 2 in the noise parameter δ.

3.2 Basic properties

Before moving to the core results of the paper, we detail a few properties relating the above noisiness metrics to each other.

Proposition 1. *Let X, Y denote random variables over finite sets. Then we have:*

1. $\mathrm{SD}(X|Y) = \mathrm{SD}(Y|X)$;
2. $\mathrm{RE}(X|Y) = \mathrm{RE}(Y|X)$;
3. $2 \cdot \mathrm{SD}(X|Y) \leq \mathrm{ARE}(X|Y) \leq \mathrm{RE}(X|Y)$.

Moreover, if X follows the uniform distribution over a set $\mathcal{X}$ of size N, then:

$$\mathrm{ARE}(X|Y) \leq 2N \cdot \mathrm{SD}(X|Y) \ . \tag{1}$$

The above properties are immediate from Definition 3. We however provide a proof for the last one. Note that, as mentioned in the introduction, our reduction from the ARE-noisy leakage model to the random probing model described in Section 5 is tighter by a factor N compared to reduction from the SD-noisy leakage model to the random probing model from [DDF14]. Hence (1) implies that even in the worst case, our results give at least as good bounds (up to a factor 2) as prior reductions.

Proof. Since X is uniform, $\mathbb{P}[X = x] = \frac{1}{N}$ for any $x \in \mathcal{X}$. Hence for any fixed y:

$$\max_x |\operatorname{PMI}| \leq \sum_{x \in \mathcal{X}} |\operatorname{PMI}| = N \cdot \mathbb{E}_X |\operatorname{PMI}| \ .$$

Then (1) follows from the definitions of SD and ARE.

Remark 1. Note that the item 3 is tight. Indeed, considering the "checkerboard distribution" $Z = (X, Y)$ defined over $[\![0, m-1]\!] \times [\![0, n-1]\!]$ via:

$$\mathbb{P}[X = x, Y = y] = \frac{1}{mn}\left(1 + (-1)^{x+y}\delta\right) \ ,$$

One can easily check that $2\mathrm{SD}(X|Y) = \mathrm{ARE}(X|Y) = \mathrm{RE}(X|Y) = \delta$.

We can also relate the SD-noisiness and the RE-noisiness to the mutual information via the following inequalities, whose proofs are detailed in the full version [GMPP19]:

Proposition 2. *Let X, Y denote random variables over finite sets. Then, we have:*

$$2\mathrm{SD}(X|Y)^2 \leq \mathrm{MI}(X;Y) \leq 2\mathrm{RE}(X|Y)\mathrm{SD}(X|Y) \ .$$

The left inequality was already proven in [DFS15a, Theorem 1]. However, our proof relies on a completely different interpretation of the mutual information, and is arguably much simpler.[6] On the other hand, the right inequality improves a previous bound given in [DDF14] by a factor $\frac{N}{\ln(2)\mathrm{RE}(X|Y)}$. Overall, it allows to bound $\mathrm{MI}(X;Y)$ up to a factor $\frac{\mathrm{SD}(X|Y)}{\mathrm{RE}(X|Y)}$.

Finally, we provide a self-reducibility lemma for RE-noisy and ARE-noisy functions. We show that the underlying distribution is not too important, as a function f which is δ-noisy for a distribution X is also $\Theta(\delta)$-noisy for any other distribution X'.

Lemma 2 (Self-reducibility). *Let X, X' be two arbitrary distributions of support $\mathcal{X}$ and $f : \mathcal{X} \to \mathcal{Y}$ be a randomized function. Suppose that f is δ_{RE}-RE-noisy (resp. δ_{ARE}-ARE-noisy) for X. Then:*

1. *f is $\left(\frac{2\cdot\delta_{\mathrm{RE}}}{1-\delta_{\mathrm{RE}}}\right)$-RE-noisy for X';*
2. *f is $\left(\frac{2\cdot\delta_{\mathrm{ARE}}}{(1-\delta_{\mathrm{ARE}})(1-\delta_{\mathrm{RE}})}\right)$-ARE-noisy for X'.*

Lemma 2 is similar to [DFS16, Lemma 2], which shows that if f is δ-SD-noisy for X the uniform distribution, then it is $(3N\delta)$-SD-noisy for any distribution X'. Our proposition is more powerful than [DFS16, Lemma 2]: X can be any distribution, and the tightness loss is $O(1)$ as long as $\delta_{\mathrm{RE}} \leq 1 - c$ for a constant c. The proof of Lemma 2 is given in the full version [GMPP19].

[6] In addition, this interpretation of MI in terms of the Kullback-Leibler divergence gives us for free several bounds which are tighter for non-negligible values of SD: for example $\mathrm{MI} \geq \log\left(\frac{1+\mathrm{SD}}{1-\mathrm{SD}}\right) - \frac{2\mathrm{SD}}{1+\mathrm{SD}}$ [Vaj70] or $\mathrm{MI} \geq 2\mathrm{SD}^2 + \frac{4}{9}\mathrm{SD}^4 + O(\mathrm{SD}^6)$ [FHT03].

3.3 Noisy Leakage Adversary

We finally define the noisy leakage model. We consider an arbitrary sequence $(x_1, \ldots, x_\ell) \in \mathcal{X}^\ell$, with $\mathcal{X}$ being some finite set and ℓ being some parameter (typically the number of intermediate results in a computation). We denote by $\mathcal{A}$ a (possibly unbounded) adversary.

Definition 5 (Noisy Leakage Adversary). *Let $D \in \{\mathrm{SD}, \mathrm{EN}, \mathrm{RE}, \mathrm{ARE}\}$. For $0 \leq \delta \leq 1$, a δ-D-noisy adversary on $\mathcal{X}^\ell$ is a machine $\mathcal{A}$ that plays the following game against an oracle that knows $(x_1, \ldots, x_\ell) \in \mathcal{X}^\ell$:*

1. *$\mathcal{A}$ picks δ-D-noisy functions $(f_i)_{i\in[\ell]}$ with range $\mathcal{Y}$;*
2. *$\mathcal{A}$ receives $(f_i(x_i))_{i\in[\ell]} \in \mathcal{Y}^\ell$ and outputs $\mathsf{out}_{\mathcal{A}}(x_1, \ldots, x_\ell)$.*

4 From Concrete Leakage to Noisy Leakage Models

In order to have a full-fledged security proof of a circuit compiler with a leakage model, the first step consists in linking the concrete representation of the leakage to a noisy leakage model. This allows to ground firmly our metrics and models in the reality, and guarantee that the gains observed in subsequent sections are not artifacts of definitions.

4.1 A concrete modelization of the leakage

A common representation of the leakage $f(X)$ corresponding to the manipulation of an intermediate variable X is a function $l(X)$ tempered by the addition of a Gaussian noise $\mathcal{N}(0, \sigma)$. The function l is then defined by the consumption model. The most widely used consumption model is the Hamming weight model initially used by Brier, Clavier, and Olivier in [BCO04], namely:

$$f(x) = \mathsf{HW}(x) + \mathcal{N}(0, \sigma) \ ,$$

Our goal is now to determine how (RE/ARE/SD/EN)-noisy the function f is. We consider that x is distributed according to a uniformly random variable X over the set $[\![0, N-1]\!]$, where $N = 2^n$ is a power of two. This assumption is realistic since in a cryptographic algorithm, the diffusion of the random private key throughout the computation makes any intermediate variable looks like a random variable. As an illustration, we give in Figure 2 a toy example for the distributions of $f(X)$ and $f(X)|(\mathsf{HW}(X) = k)$.

4.2 A visual interpretation of the noisiness metrics

We give an intuition on how the different noisiness metrics are connected to the Hamming weight consumption model with the help of Figure 2. Let $Y = f(X)$ and $Y_k = f(X)|(\mathsf{HW}(X) = k)$. By Definition 3, we can link the four metrics to the pointwise mutual information. The pointwise mutual information can be depicted as the ratio between one of the Y_k curves and the Y curve, minus 1: $\mathrm{PMI}(x, y) = \frac{Y_{\mathsf{HW}(x)}(y)}{Y(y)} - 1$. With this in mind, we can provide a visual interpretation of the four metrics as follows:

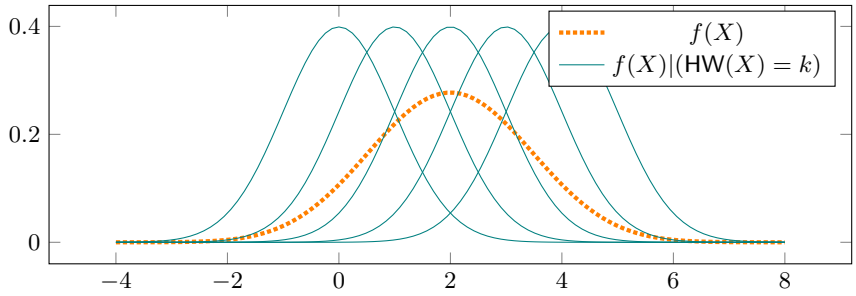

Figure 2: Distribution of the noisy function $f(X) = \mathsf{HW}(X) + \mathcal{N}(0,1)$ when X is uniformly distributed over $[\![0, 2^4 - 1]\!]$. The conditional distributions $f(X)|(\mathsf{HW}(X) = k)$ (for $k = 0, \ldots, 4$) are also represented.

- **SD**. The metric SD simply computes a ponderated mean of $|\,\mathrm{PMI}\,|$.
- **RE**. Since RE is the max of $|\,\mathrm{PMI}\,|$, it is essentially the maximum, minus 1, of the ratio Y_0/Y: this maximum is reached at the far right of Figure 2, and imposes a tailcut for the reasons detailed in Remark 2;
- **ARE**. Since ARE computes the mean (over Y) of the max (over X) of $|\,\mathrm{PMI}\,|$, it can be visually interpreted as the mean on the right side of the Figure 2 of "the ratio Y_0/Y, minus 1".
- **EN**. The visual interpretation of EN is a little more complex. Since here X is uniform, we have $\mathrm{EN}(X|Y) = \frac{1}{\sqrt{N}}\mathbb{E}_Y\sqrt{\mathbb{E}_X\left[\mathrm{PMI}^2\right]}$, so EN is essentially the scaled expected value (over Y) of Euclidean norm (over X) of the PMI.

Remark 2. We note that $\mathrm{RE}(X|f(X))$ is not formally defined as the value $|\,\mathrm{PMI}(x,y)|$ can be arbitrarily large. We overcome this issue by observing, see the full version [GMPP19], that with overwhelming probability $f(X)$ lies in the interval $[-\tau\sigma, \tau\sigma + \log N]$, where $\tau = \sqrt{-2\log(2^{-\lambda}\sqrt{2\pi})} = \Theta(\sqrt{\lambda})$. We can then define $\mathrm{RE}(X|f(X))$ with a tailcut argument.

4.3 Estimating the noisiness metrics in practice

In order to estimate the noisiness of f (with respect to RE, ARE, SD, and EN), we derive asymptotic bounds as shown in Proposition 3. To back up our theoretical results, we used a Sage implementation (which source code is given in the full version [GMPP19]) and obtained numerical values which match exactly our results.

Proposition 3. *Let X be a uniformly random variable over the set $\mathcal{X} = [\![0, N-1]\!]$, where $N = 2^n$ is a power-of-two. Let $\mathcal{Y} = \mathbb{R}$, and $f : \mathcal{X} \to \mathcal{Y}$ be defined with the Hamming weight model, namely:*

$$f(x) = \mathsf{HW}(x) + \mathcal{N}(0, \sigma) \ .$$

Let $\tau \in [1;\sigma]$ be a tailcut rate such that $|\mathcal{N}(0,\sigma)| \leq \tau \cdot \sigma$ with overwhelming probability. Then, for sufficiently large values of σ and N it holds that:

$$\begin{aligned}
\mathrm{RE}(X|f(X)) &\sim C_1 \cdot \tfrac{1}{\sigma} \cdot \tau \cdot \log N \ , \text{ with } C_1 = \tfrac{1}{2} \\
\mathrm{ARE}(X|f(X)) &\sim C_2 \cdot \tfrac{1}{\sigma} \cdot \log N \ , \text{ with } C_2 = \tfrac{1}{\sqrt{2\pi}} \\
\mathrm{SD}(X|f(X)) &\sim C_3 \cdot \tfrac{1}{\sigma} \cdot \sqrt{\log N} \ , \text{ with } C_3 = \tfrac{1}{2\pi} \\
\mathrm{EN}(X|f(X)) &\sim C_2 \cdot \tfrac{1}{\sigma} \cdot \sqrt{\tfrac{\log N}{N}} \ .
\end{aligned}$$

The proof of Proposition 3 can be found in the full version [GMPP19]. We note that a different model of the concrete leakage (say, x added to binomial noise) could lead to completely different equations.

RE vs ARE. The noisiness metric RE incurs an overhead of $O(\tau)$ compared to ARE. All other parts being equal, it is therefore more desirable to use the latter than the former. This observation is the ground motivation behind the use of ARE to show the reduction between the noisy leakage model and the probing model in Section 5.

ARE vs SD. Since ARE incurs an overhead of $O(\sqrt{\log N})$ compared to SD, one could be tempted to say that the latter leads to tighter bounds. However, we show in section 5 that when reducing to the random probing model, SD incurs an overhead of $O(N)$ compared to ARE. When linking the random probing model to a concrete model of leakage, ARE therefore allows a total gain of $O(N/\sqrt{\log N})$ compared to SD.

EN vs others. Unlike the other noisiness metrics, EN is $\tilde{O}(1/\sqrt{N})$. This suggest that this metric should lead to the most efficient discrimination of the four, but we see in Section 6 that in amplification-based proofs, the EN currently incurs a total overhead which is polynomial in N (compared to RE).

On the definition of EN. The presence in practice of a factor $\tilde{O}(1/\sqrt{N})$ in EN (as highlighted in item 4.3) suggests that the definition of EN is perhaps not the right one, along with other circumstantial evidence:

- In proposition 3, the definition of EN in terms of the pointwise mutual information is not as clean as for the other metrics;
- Several noise amplification theorems in [PR13] have an overhead $O(N^{O(d)})$. One could think this overhead is an artifact of the proof, but in some cases (such as [PR13, Theorem 1]), it is in fact an artifact of the definition.

5 From ARE-Noisy Leakage Model to Threshold-Probing Model

While noisy-leakage models defined in Section 3.3 capture well what leaks from an actual computation on physical devices, it is fairly hard to build cryptosystems that achieve security in these complex models. Therefore, simpler and more idealistic models are often considered for constructing leakage-resilient cryptography. The most common model is the threshold-probing model, introduced by Ishai, Sahai, and Wagner in [ISW03]. In this model, an adversary can learn a bounded number of *exact* intermediate results of the computation (instead of noisy information about every intermediate results). This probing model being much simpler, it is easy to immunize any computation against such adversaries, and the hope is that secure constructions in this model offer some guarantees against more realistic forms of leakage.

Fortunately, it was recently proven in [DDF14] that this intuition is correct: Duc et al. proved that a construction secure in the threshold probing model is also secure in the SD-noisy leakage model. However, the reduction comes with an overhead in the size of the field. In [DFS15a], the authors showed with empirical methods that this overhead can be significantly reduced. In this section, we aim to demonstrate an improvement of [DFS15a] by using the ARE-noisy leakage model[7] instead of the SD-noisy-leakage model. Our proof follows a similar strategy as the original proof in [DDF14]. As an outcome, the reduction between the two leakage models produces a tighter bound compared to the previous results in the state-of-the-art, thus providing stronger security guarantees for probing-secure constructions in the real world.

5.1 Probing Models

We first recall standard models of adversaries relevant in our context, as defined in [DDF14].

Random-Probing Model. For $0 \leq \epsilon \leq 1$, we denote by $\mathsf{id}_\epsilon : \mathcal{X} \to \mathcal{X} \cup \{\perp\}$ the function that on input $x \in \mathcal{X}$ outputs x with probability ϵ and $\perp$ otherwise. For $0 \leq \epsilon \leq 1$, an ϵ-random-probing adversary on $\mathcal{X}^\ell$ is a machine $\mathcal{A}$ that plays the following game against an oracle that knows $(x_1, \ldots, x_\ell) \in \mathcal{X}^\ell$:

1. $\mathcal{A}$ picks a $(\epsilon_1, \ldots, \epsilon_\ell) \in [0; \epsilon]^\ell$;
2. $\mathcal{A}$ receives $(\mathsf{id}_{\epsilon_i}(x_i))_{i \in [\ell]} \in (\mathcal{X} \cup \{\perp\})^\ell$ and outputs $\mathsf{out}_{\mathcal{A}}(x_1, \ldots, x_\ell)$.

Threshold-Probing Model. For $0 \leq d \leq \ell$, a d-threshold-probing adversary on $\mathcal{X}^\ell$ is a machine $\mathcal{A}$ that plays the following game against an oracle that knows $(x_1, \ldots, x_\ell) \in \mathcal{X}^\ell$:

[7] Note that the reduction can also work with the RE-noisy leakage model. However, as shown in previous section using the ARE metric always induces tighter reduction than the RE metric.

1. $\mathcal{A}$ picks a set $\mathcal{I} \subseteq [\ell]$ with $|\mathcal{I}| \leq d$;
2. $\mathcal{A}$ receives $(x_i)_{i\in\mathcal{I}} \in \mathcal{X}^{|\mathcal{I}|}$ and outputs $\mathsf{out}_{\mathcal{A}}(x_1, \ldots, x_\ell)$.

Following the methodology of [DDF14], our proof proceeds in two steps:

1. reduction from the ARE-noisy leakage model to the random-probing model (Section 5.2);
2. reduction from the random-probing model to the threshold probing model (Section 5.3).

5.2 From ARE-Noisy Leakage Model to Random-Probing Model

The first step consists in reducing the ARE-noisy leakage model to the random-probing model. The main technicality consists in proving the following lemma, which is the ARE-noisy version of [DDF14, Lemma 2] that was given in the SD-noisy setting. The proof of this lemma is analogous to its SD-noisy counterpart and is detailed in the full version of the paper. We denote the equality between two distributions P and Q by $P \overset{d}{=} Q$.

Lemma 3. *Let $f\colon \mathcal{X} \to \mathcal{Y}$ denote a δ-ARE-noisy function for some distribution X. Then, there exists a (randomized) function $f^{\perp}\colon \mathcal{X} \cup \{\perp\} \to \mathcal{Y}$ such that for all $x \in \mathcal{X}$:*

$$f(x) \overset{d}{=} f^{\perp}(\mathsf{id}_\delta(x)) \ .$$

Moreover, if f is poly-time-noisy[8], then $f^{\perp}$ is efficiently computable.

We then obtain the following corollary:

Corollary 1. *Let $\mathcal{A}$ be a δ-ARE-noisy adversary on $\mathcal{X}^\ell$. Then there exists a δ-random-probing adversary $\mathcal{S}$ on $\mathcal{X}^\ell$ such that for all $(x_1, \ldots, x_\ell) \in \mathcal{X}^\ell$:*

$$\mathsf{out}_{\mathcal{S}}(x_1, \ldots, x_\ell) \overset{d}{=} \mathsf{out}_{\mathcal{A}}(x_1, \ldots, x_\ell) \ .$$

Moreover, if $\mathcal{A}$ is poly-time-noisy[9], then $\mathcal{S}$ runs in polynomial time.

Proof. It immediately follows from Lemma 3. $\mathcal{S}$ simply runs $\mathcal{A}$ which it provides with $(f_i^{\perp}(\mathsf{id}_\delta(x)))_{i\in[\ell]} \overset{d}{=} (f_i(x))_{i\in[\ell]}$ as inputs. When $\mathcal{A}$ halts, so does $\mathcal{S}$ with the same output. □

Interestingly we have an opposite reduction from random probing model to ARE-noisy leakage model. However this reduction comes with a loss in tightness by a factor $N - 1$.

Lemma 4. *If $\mathcal{A}$ is a δ-random probing adversary on $\mathcal{X}^\ell$, then it is also a $(|\mathcal{X}| - 1) \cdot \delta$-ARE-noisy leakage adversary on $\mathcal{X}^\ell$.*

Proof. From the definitions, it is immediate that the δ-identity id_δ is also a $(|\mathcal{X}| - 1) \cdot \delta$-ARE-noisy function for any distribution.

[8] By poly-time-noisy, we mean that f is poly-time computable, produces outputs in a finite set $\mathcal{Y}$, and $\mathbb{P}[f_i(x) = y]$ is poly-time computable for all x, y, i.

[9] By poly-time-noisy, we mean that $\mathcal{A}$ queries only poly-time-noisy functions $(f_i)_i$.

5.3 From Random-Probing Model to Threshold-Probing Model

The second step consists in reducing the random-probing model to the threshold-probing model. This step follows immediately from the results in [DDF14] and is independent of the metric.

Lemma 5 (Lemma 4 of [DDF14]). *Let $\mathcal{A}$ be a δ-random-probing adversary on $\mathcal{X}^\ell$. Then, there exists a $(2\delta\ell - 1)$-threshold-probing adversary $\mathcal{S}$ on $\mathcal{X}^\ell$ with similar running-time such that $\forall (x_1, \ldots, x_\ell) \in \mathcal{X}^\ell$:*

$$\mathsf{out}_{\mathcal{A}}(x_1, \ldots, x_\ell) \stackrel{d}{=} \mathsf{out}_{\mathcal{S}}(x_1, \ldots, x_\ell) \ ,$$

as long as $out_{\mathcal{S}}(x_1, \ldots, x_\ell) \neq \perp$. Moreover, the latter happens with probability:

$$\mathbb{P}[\mathsf{out}_{\mathcal{S}}(x_1, \ldots, x_\ell) \neq \perp] \geq 1 - \exp\left(-\frac{\delta\ell}{3}\right) \ .$$

The proof immediately follows from the fact that with probability at least $1 - \exp\left(-\frac{\delta\ell}{3}\right)$ (thanks to the Chernoff bound), a δ-random-probing adversary on $\mathcal{X}^\ell$ obtain at most $2\delta\ell - 1$ of the x_i's.

5.4 Putting Everything Together

Combining Corollary 1 and Lemma 5, we then obtain the following theorem:

Theorem 1. *Let $\mathcal{A}$ be a δ-ARE-noisy adversary on $\mathcal{X}^\ell$. Then, there exists a $(2\delta\ell - 1)$-threshold-probing adversary $\mathcal{S}$ on $\mathcal{X}^\ell$ such that:*

$$\mathsf{out}_{\mathcal{A}}(x_1, \ldots, x_\ell) \stackrel{d}{=} \mathsf{out}_{\mathcal{S}}(x_1, \ldots, x_\ell) \ ,$$

as long as $\mathsf{out}_{\mathcal{S}}(x_1, \ldots, x_\ell) \neq \perp$, which happens with probability:

$$\mathbb{P}[\mathsf{out}_{\mathcal{S}}(x_1, \ldots, x_\ell) \neq \perp] \geq 1 - \exp\left(-\frac{\delta\ell}{3}\right) \ .$$

Moreover, if $\mathcal{A}$ is poly-time-noisy, then $\mathcal{S}$ runs in polynomial time.

For comparison, the main theorem from [DDF14] states that a δ-SD-noisy adversary can be simulated by a $(2\delta\ell \cdot |\mathcal{X}| - 1)$-threshold probing adversary, with success probability at least $1 - \exp\left(-\delta\ell/(3|\mathcal{X}|)\right)$. Hence, we gain a multiplicative factor $\mathcal{X}$ in the number of probes and reduce the failure probability by an exponential factor in $\mathcal{X}$.

5.5 Circuit leakage resilience

Let us define a circuit compiler as in [DDF14]. Let us consider an adversary able to probe at most $\lfloor (d-1)/2 \rfloor$ wires from each gadget (i.e. masked operations) of the implementation. We define a (δ, ζ)-noise resilient implementation as follows:

Definition 6 (D-noise resilient implementation). *Let Γ be an stateful arithmetic circuit over $\mathcal{X}$ and Γ' denote the resulting masked circuit obtained via applying the compiler. Let Enc denote a randomized encoding function (i.e. that transform an input into a masked input). Let $D \in \{\mathrm{SD}, \mathrm{EN}, \mathrm{RE}, \mathrm{ARE}\}$ be a noisiness metric. We say that Γ' is a (δ, ζ)-D-noise resilient implementation of Γ with respect to Enc if the following properties hold for every input k:*

1. *the input-output behavior of $\Gamma(k)$ and $\Gamma'(Enc(k))$ is identical, i.e. for every sequence of inputs $a_1, \ldots, a_m$ and outputs $b_1, \ldots, b_m$ we have*

$$\mathbb{P}[\Gamma(k, a_1, \ldots, a_m) = (b_1, \ldots, b_m)] = \mathbb{P}[\Gamma'(Enc(k), a_1, \ldots, a_m) = (b_1, \ldots, b_m)]$$

2. *for every δ-D-noisy adversary $\mathcal{A}$ there exists a black-box circuit adversary $\mathcal{S}$ such that*

$$\Delta_{\mathrm{SD}}\left(\mathsf{out}\left(\mathcal{A} \overset{noisy}{\leftrightarrows} \Gamma'(Enc(k))\right); \mathsf{out}\left(\mathcal{S} \overset{bb}{\leftrightarrows} \Gamma(k)\right)\right) \leq \zeta$$

Then we have the following theorem.

Theorem 2. *Let Γ be an arbitrary stateful arithmetic circuit over $\mathcal{X}$. Let Γ' be the masked circuit. Then Γ' is a $(\delta, |\Gamma| exp(-d/12))$-RE-noise-resilient implementation of Γ with efficient simulation where*

$$\delta = \frac{1}{28d + 16} = O(1/d)$$

The proof is the exact same as the one given in [DDF14] with a numerical gain of a factor $|\mathcal{X}|$ in δ due to the use of ARE in Theorem 1.

6 A New Amplification-Based Proof for Block Ciphers

In this section, we revisit the approach initiated in [PR13] by Prouff and Rivain. Recall that in the latter work, the authors propose a solution to immunize block-ciphers in the noisy-leakage model (with the Euclidian norm EN measuring noisiness). They propose a secret-sharing based immunization for block-ciphers, basically by replacing every operations (linear functions and s-box evaluations) by one that operates on additive shares of the inputs and produce additive shares of the output. They analyze the security by decomposing the resulting protocol into 4 types of basic subsequences of operations: two types corresponding to simple subsequences and two types corresponding to more complex subsequences. The overall protocol is then proven secure by composition, assuming leak-free refresh gates can be used between each subsequence to refresh the additive shares. We refer the reader to Section 4 of [PR13] for the details about how to construct the secure subprotocols. The 4 types of subsequences needed for the analysis are recalled below. We propose a different security analysis in the noisy-leakage model using the RE metric instead of the Euclidian norm. Doing so, we are able to prove much tighter bounds for the security.

The 4 types of subsequences to consider are:

T1. $(z_i \leftarrow g(x_i))_{0\leq i\leq d}$, with g being a linear function (of the block-cipher);
T2. $(z_i \leftarrow g(x_i))_{0\leq i\leq d}$, with g being an affine function (of an s-box evaluation);
T3. $(v_{i,j} \leftarrow a_i \times b_j)_{0\leq i,j\leq d}$ (first step of secure non-linear multiplication);
T4. $(t_{i,j} \leftarrow t_{i,j-1} \oplus v_{i,j})_{0\leq i,j\leq d}$ (fourth step of secure non-linear multiplication).

While type 1 is obviously a particular case of type 2, we treat them separately as we are able to prove a better bound for linear functions than for affine functions.

In the rest of this section, we first provide several basic properties on the RE metric (Section 6.1). Then, in Section 6.2 we analyze the leakage of each type of subsequences. Next, we argue about the security of a complete evaluation of the block-cipher in Section 6.3. Finally, in Section 6.4 we apply the Rényi divergence to get a tight amplification-based proof and overcome the limitations in [PR13].

6.1 Basic Properties and Amplification for the Relative Error

First, we give several basic properties of RE-noisy functions and of the RE-noisiness metric that are used in our proofs. We essentially show that the relative error is preserved under function, application, projection and lifting on X. We also prove an amplification result (Lemma 6) that is central throughout our security analysis.

Proposition 4. *Let X, Y, W denote random variables over finite sets $\mathcal{X}, \mathcal{Y}$ and $\mathcal{W}$ respectively. Then we have the following:*

1. ***Data processing.*** *Let $f : \mathcal{X} \to \mathcal{Y}$ be a δ-RE-noisy function for X, and $g : \mathcal{X} \to \mathcal{X}$ be a (non necessarily deterministic) function. It holds that:*

$$\mathrm{RE}(X|f \circ g(X)) \leq \frac{2\delta}{1-\delta} \underset{\delta\to 0}{\sim} 2\delta \ .$$

   In addition, if g is deterministic and bijective, then $\mathrm{RE}(X|f \circ g(X)) = \delta$.
2. ***Conservation under projection and lifting.***

$$\mathrm{RE}(X|Y) \leq \mathrm{RE}((X,W)|Y) \ . \tag{2}$$

   In addition, if X and W are independent and $f : \mathcal{W} \to \mathcal{Y}$ is a RE-noisy function for W, then:

$$\mathrm{RE}((X,W)|f(W)) = \mathrm{RE}(W|f(W)) \ . \tag{3}$$

The proof of Proposition 4 is detailed in the full version of the paper.

Remark 3. Note that Inequality 1 is tight: Indeed, if we consider the checkerboard distribution of remark 1, take $f(X) = Y$, $g(0) = 0$ and for any $x > 0$, $g(x) = 1$, then:

$$\mathrm{RE}(X|f(X)) = \delta \text{ and } \mathrm{RE}(X|f \circ g(X) = \frac{2(1-1/m)\delta}{1-(1-2/m)\delta} \underset{m\to\infty}{\sim} \frac{2\delta}{1-\delta} \ .$$

Note that Inequality 2 is also tight via (3).

We also prove the following amplification lemma for the relative error. It is the relative error counterpart (though the proof is completely different) of an amplification lemma by Maurer, Pietrzak and Renner [MPR07, Lemma 1]. In the context of leakage-resilient cryptography, the latter result was used and improved by [DFS15b,DFS16].

Lemma 6. *Let $\mathbb{F}$ be a finite field. Let $Z = U(\mathbb{F})$ be the uniform distribution over $\mathbb{F}$, and $Z_1, \ldots, Z_d$ be d independent random variables over $\mathbb{F}$. It holds that:*

$$\Delta_{\mathrm{RE}}\left(\left(\sum_{i=1}^{d} Z_i\right); Z\right) \leq \prod_{i=1}^{d} \Delta_{\mathrm{RE}}\left(Z_i; Z\right) \quad .$$

6.2 Security Analysis of Subsequences

We now detail our security analysis for the 4 different types of subsequences to be considered.

Type 1 and type 2 subsequences. We first deal with the simple case of subsequences where all the shares of a secret value are processed separately, i.e. for linear and affine functions. From a security perspective, these are the simplest subsequences as each share only leaks partial information once.

Type 1 subsequences. We first prove the following theorem for type 1 subsequences, which follows almost immediately from Lemma 6. For the sake of completeness, we provide a proof in the full version [GMPP19].

Theorem 3. *Let X be a uniform random variable over a finite field $\mathcal{X}$ and $(X_i)_{i\in\{0,\ldots,d\}}$ be a $(d+1)$-additive sharing of X*[10]*. Let $\delta \in [0,1)$ and $f_0, f_1, \ldots, f_d$ be δ-RE-noisy-leakage functions over $\mathcal{X}$. Then, we have:*

$$\mathrm{RE}(X|f_0(X_0), \ldots, f_d(X_d)) \leq \delta^{d+1} \quad .$$

Unlike [PR13, Theorem 1], we do not get a overhead of $N^{d/2}$ in our amplification theorem. One could think that this overhead is an artifact of their proof, but circumstantial evidence such as the presence in practice of a factor $1/\sqrt{N}$ in $\mathrm{EN}(X_i|f_i(X_i))$ let us think that it is inherent to the use of the Euclidean norm.

Type 2 subsequences. We can now easily analyze the security of type 2 subsequences, i.e. affine functions of s-box evaluations. Such evaluations are handled via Lagrange interpolation in [PR13], so each elementary calculation processes a share G_i of an encoding of $g(X)$, where X is a uniform s-box input, and g is a polynomial function. This case is covered by Corollary 2, whose proof immediately follows from Theorem 3 and Proposition 4, and is detailed in the full version [GMPP19].

[10] Precisely, $\sum_{i=0}^{d} X_i = X$ and the distribution of any strict subset of the X_i's is uniform.

Corollary 2. *Let X be a uniform random variable over a finite field $\mathcal{X}$, $g : \mathcal{X} \to \mathcal{X}$ be a deterministic function, d be a positive integer and $(G_i)_{i\in\{0,\dots,d\}}$ be a $(d+1)$-additive sharing of $g(X)$. Let $\delta \in [0,1)$ and let $f_0, f_1, \dots, f_d$ be δ-RE-noisy leakage functions over $\mathcal{X}$. Then, we have:*

$$\mathrm{RE}(X|f_0(G_0), \dots, f_d(G_d)) \leq \frac{2\delta^{d+1}}{1-\delta^{d+1}} \underset{\delta\to 0}{\sim} 2\delta^{d+1} \ .$$

Type 3 and type 4 subsequences. We now consider more complex subsequences, where a share is processed several times, and therefore may leak several times in the same subsequence. We first give a generic theorem regarding the bias induced by multiple leakages. We then use this theorem (whose proof is given in the full version [GMPP19]) to bound the leakage of subsequences of type 3 and type 4.

Theorem 4. *Let X be a uniform random variable over a finite field $\mathcal{X}$ and t be a strictly positive integer. Let $\delta \in [0,1)$ and $L_1, \dots, L_t$ be t random variables such that $\mathrm{RE}(X|L_i) \leq \delta$ for every i. We further assume that the random variables $(L_i|X=x)$ are mutually independent for every $x \in \mathcal{X}$. Then, we have:*

$$\mathrm{RE}(X|L_1, \dots, L_t) \leq \left(\frac{1+\delta}{1-\delta}\right)^t - 1 \underset{t\delta\to 0}{=} 2\cdot t\delta + O((t\delta)^2) \ .$$

In addition, if $\delta \leq 1/t$, then:

$$\mathrm{RE}(X|L_1, \dots, L_t) \leq \frac{t\delta}{1-(t-1)\delta} \underset{t\delta\to 0}{=} t\delta + O((t\delta)^2) \ .$$

Depending on the situation, we use one bound or the other in what follows.

Type 4 subsequences. We start by analyzing subsequences of type 4. Each elementary computation of these subsequences computes $T_{i,j} \leftarrow T_{i,j-1} \oplus V_{i,j}$, with $0 \leq i, j \leq d$ and $T_{i,0} = V_{i,0}$. At the end, the shares $(Z_i)_i = (T_{i,d})_i$ form an additive sharing of $g(X)$, where X is a uniform s-box input and g is a polynomial function over $\mathcal{X}$. Our goal here is to bound the bias of X given the leakages of all these elementary computations. We give a first theorem (whose proof is given in the full version [GMPP19]) which bounds the bias of the shares $(Z_i)_i = (T_{i,d})_i$.[11]

Theorem 5. *Let $T_0, T_1, \dots, T_d$ be $d+1$ independent uniformly random variables over a finite set $\mathcal{X}$. Let $\delta \in \mathbb{R}$ such that $\delta \leq \frac{1}{2d+1}$ and $f_1, f_2, \dots, f_d$ be a family of δ-RE-noisy functions defined over $\mathcal{X}\times\mathcal{X}$. We have:*

$$\mathrm{RE}(T_d|f_1(T_0,T_1), \dots, f_1(T_{d-1},T_d)) \leq \frac{d\delta}{1-(d-1)\delta} \ .$$

This implies the following corollary for the security of a subsequence of type 4:

[11] For concision, Theorem 5 omits the subscript i and writes $(T_j)_j$ instead of $(T_j)_{i,j}$.

Corollary 3. *The leakage of type 4 subsequences is upper bounded by:*

$$\mathrm{RE}(X|(f_{i,j}(T_{i,j-1}, V_{i,j})_{0\le i,j\le d}) \le \frac{2\delta'^{d+1}}{1-\delta'^{d+1}}, \text{ with } \delta' = \frac{d\delta}{1-(d-1)\delta} .$$

Type 3 subsequences. Only the case of type 3 subsequences remains, which is the most delicate one. As a preliminary result, we provide an upper bound on the bias for a uniform pair (A, B) given the leakage $(f_{i,j}(A_i, B_j))_{i,j}$.

Theorem 6. *Let A, B be two uniform random variables over a finite field $\mathcal{X}$, d a positive integer, and $(A_i)_i, (B_i)_i$ be $d+1$-additive-sharings of A and B respectively. Let $\delta \in \mathbb{R}$ such that $\delta \le \frac{1}{2d+1}$, and $(f_{i,j})_{i,j}$ be a family of randomized and mutually independent functions such that each $f_{i,j} : \mathcal{X} \times \mathcal{X} \to \mathcal{Y}$ is δ-RE-noisy. We have:*

$$\mathrm{RE}((A, B)|(f_{i,j}(A_i, B_j))_{i,j}) \le 3\left(\frac{(d+1)\delta}{1-d\delta}\right)^{d+1} .$$

The proof of Theorem 6 essentially combines Theorems 3 and 4, and is detailed in the full version [GMPP19].

We now give the leakage of type 3 subsequences. The difference with Theorem 6 is that A and B are not uniformly random, but rather $A = g(X)$ and $B = h(X)$ for some polynomial functions g, h. We then have the following corollary:

Corollary 4. *Let X be a uniform random variable over a finite field $\mathcal{X}$, let g, h be two deterministic functions from $\mathcal{X}$ to $\mathcal{X}$, d be a positive integer, and $(G_i)_i, (H_i)_i$ be $d+1$-additive-sharings of $g(X)$ and $h(X)$ respectively. Let $\delta \in \mathbb{R}$ such that $\delta \le \frac{1}{2d+1}$, and $(f_{i,j})_{i,j}$ be δ-RE-noisy functions over $\mathcal{X} \times \mathcal{X}$. We have:*

$$\mathrm{RE}(X|(f_{i,j}(G_i, H_j))_{i,j}) \le \frac{2\delta'}{1-\delta'}, \text{ with } \delta' = 3\left(\frac{(d+1)\delta}{1-d\delta}\right)^{d+1} .$$

Corollary 4 results from combining Theorem 6 with Proposition 4. It is detailed in the full version [GMPP19].

6.3 From Subsequences to a Complete Computation

Now that we have bounded the leakages of the individual subsequences, the next step is to bound the leakage of a single complete execution of a block cipher.

Modeling a block cipher. We use the same notations as in [PR13] and consider the resulting block cipher (after applying their compiler to the original block cipher), hereafter referred to as the *masked block cipher*. An evaluation of the masked block cipher gives $\mathcal{I} = (\mathsf{C}_i, f_i)_i$, where the C_i's denote elementary computations (or gates) of the masked block cipher, each C_i being associated to an RE-noisy function f_i. We assume that the (original) cipher involves t_{lin} linear transformations (corresponding to as many type 1 subsequences), t_{aff} affine functions (type 2 subsequences), and t_{nlm} nonlinear multiplications (types 3, 4).

Uniformity of the key and of the subsequence inputs. The block cipher is parameterized by a secret key k which is sampled from the uniform distribution K over $\mathcal{K} = \mathcal{X}^m$. Each subsequence subseq_j operates on an additive-sharing of a random variable X_j. We can write $X_j = g_j(K, \mathsf{msg})$, where msg denotes the message being processed by the block cipher, and g_j is a publicly-known function such that $g_j(\cdot, \mathsf{msg})$ maps the uniform distribution over $\mathcal{K}$ to the uniform distribution over $\mathcal{X}$, which is the case for block ciphers in practice. Therefore the inputs X_j's of each subsequence are uniformly random variables. Alternatively, one could rely on Lemma 2.

Leakage of a block cipher evaluation. For a given subseq_j, let L_j denote its leakage. Since each $g_j(\cdot, \mathsf{msg})$ maps the uniform distribution to the uniform distribution, we have $\mathrm{RE}(K|L_j) = \mathrm{RE}(X_j|L_j)$ from Proposition 4. Since the $t = t_{\mathsf{lin}} + t_{\mathsf{aff}} + 2t_{\mathsf{nlm}}$ subsequences composing the circuit are interleaved with leak-free refresh gates (by assumption), each of them operates on fresh random shares, therefore the leakages $(L_j|K = k)$ are mutually independent.

We suppose that there exists a $\delta_{\mathsf{subseq}} \geq 0$ such that $\forall j, \mathrm{RE}(X_j|L_j) \leq \delta_{\mathsf{subseq}}$. Theorems 3 and 5 as well as Corollaries 2 and 4 give us explicit conditions to fulfill this bound for each subsequence. Via Theorem 4, the leakage δ_{circ} of the whole secure evaluation is bounded by:

$$\delta_{\mathsf{circ}} \leq \frac{t\delta_{\mathsf{subseq}}}{1 - (t-1)\delta_{\mathsf{subseq}}} \approx t\delta_{\mathsf{subseq}}$$

which is non-vacuous as long as $\delta_{\mathsf{subseq}} \leq 1/(t-1)$.

6.4 Overall Security Proof with the Rényi Divergence

Now that we have bounded the overall leakage of one evaluation of the block cipher, we want to analyze the impact of this leakage on the concrete security of the block cipher. This last section corresponds somehow to the end of [PR13], where the leakage of an evaluation is translated into a bound on the mutual information provided by the leakage. Yet, this incurs the following limitations.

Limitations of the Prouff-Rivain approach. The use of the mutual information is somewhat problematic in the sense that it provokes paradoxical situations like the fact that a single pair of plaintext-ciphertext can information-theoretically reveal the key. The authors circumvent this by considering a random-plaintext attack where plaintexts and ciphertexts are both unknown. This does not cover many situations encountered in cryptography and is highly unusual compared to most works, which consider at least a chosen-plaintext attack. Finally, while not stated explicitly, for concrete security we need the mutual information to be upper bounded by $2^{-O(\lambda)}$), where λ is the targeted security parameter of the block cipher, hence the masking order depends only on λ and in particular does not depend on the amount of leakage the adversary can observe.

A proof based on the Rényi divergence. In this section, we provide an alternative security proof based on the Rényi divergence instead of the mutual information. This provides two main benefits compared to the previous work: (1) We can consider classical chosen-plaintext attacks, and (2) the requirement on the noise is much lower because it does not depend on the security level anymore but *on the number of leakages*, denoted q in what follows.

Description of the games. We consider two games. The first game models a black-box interaction of an attacker with an encryption oracle and corresponds to the standard security model such as the IND-CPA security game model. The second game models a grey-box interaction where the attacker has, in addition, access to leakage. This grey-box interaction captures the behavior of a block cipher in the real-world. We also introduce a third artificial (but easier) game which we use to connect the latter two games. These three games are summarized in Figures 3, 4, and 5 and are precisely described below.

Let $\mathcal{A}$ be an adversary interacting with an (encryption and decryption) oracle $\mathcal{O}$ in the following fashion:

1. $\mathcal{O}$ draws a secret key $k \leftarrow K$, where K denotes the uniform distribution over a finite set $\mathcal{K}$;
2. $\mathcal{A}$ makes a finite number q of queries to $\mathcal{O}$. This is the part where the three games differ:
   - Game 1 (*black-box*): $\mathcal{A}$ sends q plaintexts $\mathsf{msg}_1, \ldots, \mathsf{msg}_q$ to $\mathcal{O}$, who sends back the q corresponding ciphertexts $\mathsf{ctxt}_1 = E_k(\mathsf{msg}_1), \ldots, \mathsf{ctxt}_q = E_k(\mathsf{msg}_q)$;
   - Game 2 (*grey-box*): $\mathcal{A}$ sends q plaintexts $\mathsf{msg}_1, \ldots, \mathsf{msg}_q$ to $\mathcal{O}$. For each plaintext msg_i, $\mathcal{O}$ sends back the corresponding ciphertexts ctxt_i but also some value L_i which modelizes the physical leakage occurring during the computation $\mathsf{ctxt}_i \leftarrow E_k(\mathsf{msg}_i)$, and gets recorded by $\mathcal{A}$;
   - Game 3 (*hybrid*): This is a 2-stage game:
     (a) first, $\mathcal{A}$ sends q plaintexts $\mathsf{msg}'_1, \ldots, \mathsf{msg}'_q$, and $\mathcal{O}$ sends back the corresponding leakages L_i *but not* the ciphertexts $E_k(\mathsf{msg}'_i)$.
     (b) second, $\mathcal{A}$ sends q plaintexts $\mathsf{msg}_1, \ldots, \mathsf{msg}_q$, and $\mathcal{O}$ sends back the ciphertexts $E_k(\mathsf{msg}'_i)$ *but not* the corresponding leakages L_i.
3. After the query-reply phase, $\mathcal{A}$ outputs a value k'. $\mathcal{A}$ wins the game if $k = k'$.

Relationships between the games. It is clear that any attacker $\mathcal{A}$ that succeeds in Game 1 also does in Game 2, since $\mathcal{A}$ can choose to discard the additional leakage $L_1, \ldots, L_q$, in which case Game 2 becomes identical to Game 1. Similarly, any attacker that succeeds in Game 2 also does in Game 3 by simply querying $\mathsf{msg}'_i = \mathsf{msg}_i, \forall i$. Hence, it is sufficient to prove that the success probability of an adversary in Game 3 is close (in a precisely quantifiable way) to its success probability in the ideal Game 1 to argue that security holds in the real-world (Game 2). This is what we do in the rest of this section.

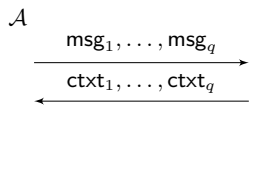

Figure 3: Game 1 Black-box

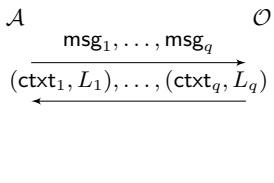

Figure 4: Game 2: Grey-box

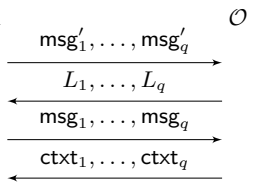

Figure 5: Game 3: Hybrid

Applying the Rényi divergence. At the end of the first step of Game 3, $\mathcal{A}$ has learnt leakages $L_1, \ldots, L_q$. These leakages imply a bias in the distribution of possible secret keys K (which was originally the uniformly random distribution). We denote by K' the distribution $(K|L_1, \ldots, L_q)$. Hence, after the first step of Game 3, the vision of $\mathcal{A}$ is the same as playing Game 1 with the secret key being taken from distribution K' (instead of uniformly at random).

Suppose that $\forall i, \mathrm{RE}(K|L_i) \leq \delta_{\mathsf{circ}}$ for some $\delta_{\mathsf{circ}} \in [0, 1)$. Assuming leak-free refresh gates, it follows from Theorem 4 that:

$$\Delta_{\mathrm{RE}}(K'; K) = \mathrm{RE}(K|L_1, \ldots, L_q) \leq \left(\frac{1 + \delta_{\mathsf{circ}}}{1 - \delta_{\mathsf{circ}}}\right)^n - 1 \ . \tag{4}$$

Let $E \subseteq \mathrm{Supp}(K)$ be an arbitrary event. We recall that $K(E)$ denotes the probability of E occurring under the distribution K. First, from the probability preservation property of the Rényi divergence (Lemma 1):

$$K'(E) \leq K(E) \cdot R_\infty(K' \| K). \tag{5}$$

On the other hand, from the definition of the Rényi divergence:

$$R_\infty(K' \| K) \leq 1 + \Delta_{\mathrm{RE}}(K'; K) \tag{6}$$

Combining (4), (5) and (6) yields

$$K'(E) \leq K(E) \cdot \left(\frac{1 + \delta_{\mathsf{circ}}}{1 - \delta_{\mathsf{circ}}}\right)^q$$

Practical implications. The consequence of this security proof is that as long as the number of leakage queries q is in $O(1/\delta_{\mathsf{circ}})$, an adversary does not have significantly larger chances to break a leaking block cipher implementation than it does for the black-box implementation.

For example, let E be the event that $\mathcal{A}$ solves a search problem (finding a secret key, forging a signature, decrypting a message, etc). If we take $q \leq 1/\delta_{\mathsf{circ}}$, then $K'(E) \leq e^2 K(E)$, which means that the leakages do no improve the probability of $\mathcal{A}$ solving the search problem by more than a factor e^2; this means that less than 3 bits of security have been lost between the black-box (Game 1)

and leaking (Game 2) implementations. In contrast, an analysis based on the statistical distance or the mutual information would require $\delta_{\mathsf{circ}} = 2^{-O(\lambda)}$.

We note that this Rényi-divergence based analysis is only valid for search problems: achieving the same efficiency for decision problems is still an open question [BLL+15,Pre17] .

References

ADF16. Marcin Andrychowicz, Stefan Dziembowski, and Sebastian Faust. Circuit compilers with o(1/\log (n)) leakage rate. In *EUROCRYPT (2)*, volume 9666 of *Lecture Notes in Computer Science*, pages 586–615. Springer, 2016.

AIS18. Prabhanjan Ananth, Yuval Ishai, and Amit Sahai. Private circuits: A modular approach. Cryptology ePrint Archive, Report 2018/566, 2018. https://eprint.iacr.org/2018/566.

BCO04. Eric Brier, Christophe Clavier, and Francis Olivier. Correlation power analysis with a leakage model. In Marc Joye and Jean-Jacques Quisquater, editors, *Cryptographic Hardware and Embedded Systems - CHES 2004*, pages 16–29, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.

BLL+15. Shi Bai, Adeline Langlois, Tancrède Lepoint, Damien Stehlé, and Ron Steinfeld. Improved security proofs in lattice-based cryptography: Using the rényi divergence rather than the statistical distance. In *ASIACRYPT (1)*, volume 9452 of *Lecture Notes in Computer Science*, pages 3–24. Springer, 2015.

CH89. Kenneth Ward Church and Patrick Hanks. Word association norms, mutual information and lexicography. In *ACL*, pages 76–83. ACL, 1989.

CJRR99. Suresh Chari, Charanjit S. Jutla, Josyula R. Rao, and Pankaj Rohatgi. Towards sound approaches to counteract power-analysis attacks. In 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 *Lecture Notes in Computer Science*, pages 398–412. Springer, 1999.

DDF14. Alexandre Duc, Stefan Dziembowski, and Sebastian Faust. Unifying leakage models: From probing attacks to noisy leakage. In *EUROCRYPT*, volume 8441 of *Lecture Notes in Computer Science*, pages 423–440. Springer, 2014.

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.

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.

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.

DP08. Stefan Dziembowski and Krzysztof Pietrzak. Leakage-resilient cryptography. In *FOCS*, pages 293–302. IEEE Computer Society, 2008.

FHT03. A. A. Fedotov, P. Harremoes, and F. Topsoe. Refinements of pinsker's inequality. *IEEE Transactions on Information Theory*, 49(6):1491–1498, June 2003.

GJR17. Dahmun Goudarzi, Antoine Joux, and Matthieu Rivain. How to securely compute with noisy leakage in quasilinear complexity. Cryptology ePrint Archive, Report 2017/929, 2017. https://eprint.iacr.org/2017/929.

GMPP19. Dahmun Goudarzi, Ange Martinelli, Alain Passelègue, and Thomas Prest. Unifying leakage models on a rényi day. *IACR Cryptology ePrint Archive*, 2019:138, 2019.

ISW03. Yuval Ishai, Amit Sahai, and David A. Wagner. Private circuits: Securing hardware against probing attacks. In *CRYPTO*, volume 2729 of *Lecture Notes in Computer Science*, pages 463–481. Springer, 2003.

KJJ99. Paul C. Kocher, Joshua Jaffe, and Benjamin Jun. Differential power analysis. In *CRYPTO*, volume 1666 of *Lecture Notes in Computer Science*, pages 388–397. Springer, 1999.

Koc96. Paul C. Kocher. Timing attacks on implementations of diffie-hellman, rsa, dss, and other systems. In *CRYPTO*, volume 1109 of *Lecture Notes in Computer Science*, pages 104–113. Springer, 1996.

MPR07. Ueli M. Maurer, Krzysztof Pietrzak, and Renato Renner. Indistinguishability amplification. In *CRYPTO*, volume 4622 of *Lecture Notes in Computer Science*, pages 130–149. Springer, 2007.

MR04. Silvio Micali and Leonid Reyzin. Physically observable cryptography (extended abstract). In *TCC*, volume 2951 of *Lecture Notes in Computer Science*, pages 278–296. Springer, 2004.

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.

Pre17. Thomas Prest. Sharper bounds in lattice-based cryptography using the rényi divergence. In *ASIACRYPT (1)*, volume 10624 of *Lecture Notes in Computer Science*, pages 347–374. Springer, 2017.

Ré61. Alfréd Rényi. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, pages 547–561, Berkeley, Calif., 1961. University of California Press.

Vaj70. I. Vajda. Note on discrimination information and variation (corresp.). *IEEE Transactions on Information Theory*, 16(6):771–773, November 1970.