Improved Reductions from Noisy to Bounded and Probing Leakages via Hockey-Stick Divergences

Maciej Obresmki[1], João Ribeiro[2⋆], Lawrence Roy[3], François-Xavier Standaert[4], and Daniele Venturi[5]

[1] National University of Singapore, Singapore, Singapore
obremski.math@gmail.com
[2] Instituto Superior Técnico, Universidade de Lisboa, Lisboa, Portugal
jribeiro@tecnico.ulisboa.pt
[3] Aarhus University, Aarhus, Denmark
ldr709@gmail.com
[4] Université catholique de Louvain, Louvain-la-Neuve, Belgium
fstandae@uclouvain.be
[5] Sapienza University of Rome, Rome, Italy
venturi@di.uniroma1.it

Abstract. There exists a mismatch between the theory and practice of cryptography in the presence of leakage. On the theoretical front, the *bounded leakage* model, where the adversary learns bounded-length but noiseless information about secret components, and the *random probing* model, where the adversary learns some internal values of a leaking implementation with some probability, are convenient abstractions to analyze the security of numerous designs. On the practical front, side-channel attacks produce long transcripts which are inherently noisy but provide information about all internal computations, and this noisiness is usually evaluated with closely related metrics like the mutual information or statistical distance. Ideally, we would like to claim that resilience to bounded leakage or random probing implies resilience to noisy leakage evaluated according to these metrics. However, prior work (Duc, Dziembowski and Faust, Eurocrypt 2014; Brian *et al.*, Eurocrypt 2021) has shown that proving such reductions with useful parameters is challenging.

In this work, we study noisy leakage models stemming from *hockey-stick divergences*, which generalize statistical distance and are also the basis of differential privacy. First, we show that resilience to bounded leakage and random probing implies resilience to our new noisy leakage model with improved parameters compared to models based on the statistical distance or mutual information. Second, we establish composition theorems for our model, showing that these connections extend to a setting where multiple leakages are obtained from a leaking implementation. We complement our theoretical results with a discussion of practical relevance, highlighting that (i) the reduction to bounded leakage applies to realistic leakage functions with noise levels that are decreased by several orders of magnitude compared to Brian *et al.*, and (ii) the reduction to random

⋆ Work done while at NOVA LINCS and Universidade Nova de Lisboa.

probing usefully generalizes the seminal work of Duc, Dziembowski, and Faust, although it remains limited when the field size in which masking operates grows (i.e., hockey-stick divergences can better hide the field size dependency of the noise requirements, but do not annihilate it).

1 Introduction

Side-channel attacks leverage properties of cryptographic implementations to obtain partial information about supposedly secret components, such as the long-term keys of authentication or encryption schemes. Several textbook versions of well-known algorithms are easily broken in practice via side-channel attacks. For example, textbook RSA is vulnerable to timing attacks, whereby an adversary measures the time elapsed during encryption and/or decryption [26]. Over the past two decades, various types of (usually simple) side-channel attacks have been employed with devastating effects on most (symmetric and asymmetric) cryptographic algorithms, including also tracking power consumption [25], the emission of electromagnetic radiation [1], and cache-based attacks [32]. Small embedded devices are natural targets, but side-channel attacks have been extended to hardware implementations [30] and high-frequency devices [2]. They can also be applied remotely [29], and new attacks keep on being discovered [27]. In general, more complex and high-frequency targets and more remote and less invasive adversarial conditions make the side-channel measurements less informative.

The devastating effect of these attacks have led to the study of generic solutions to prevent them, which we can roughly divide in two directions:

- *Primitive-level* countermeasures aim to design cryptographic algorithms of which (parts of) the implementation, that are usually denoted as leakage-resilient [19], remain secure even in the presence of bounded leakage. Such countermeasures typically leverage the frequent refreshing of the algorithms' secret state, which limits the side-channel attack surface and makes it more realistic to expect that a state's leakage is (intrinsically) bounded.
- *Implementation-level* countermeasures rather aim to limit the leakage for the parts of the cryptographic algorithms that are not leakage-resilient, such as the initialization of a secret state with a long-term secret key. In this case, where the adversary can continuously accumulate information on the same secret, masking (a.k.a. secret sharing) [12] is usually considered as the most viable option.[6] It allows amplifying the implementation noise exponentially in the number of shares at the cost of (roughly) quadratic overheads.

These solutions can then be combined so that leakage-resistant modes of operation can efficiently mix parts of the implementation where bounded leakage is obtained via cheap countermeasures (or no countermeasures at all) and a limited number of calls to parts of the implementation that require masking [6].

[6] There are, however, primitive-level alternatives to this initialization problem, such as using a leakage-resilient PRF for this part of the computation [20, 5, 13].

Most works on the formal study of leakage-resilience conveniently assume that the adversary is allowed to learn arbitrary bounded-length information about secret components. In particular, the adversary is allowed to choose a function $f : \{0,1\}^* \to \{0,1\}^\ell$, for a predetermined leakage bound ℓ, and to learn the bounded leakage $f(sk) \in \{0,1\}^\ell$, where sk is a secret key. We will refer to this model as the *bounded leakage model*. The survey of Kalai and Reyzin [24] is an excellent source on prior work on bounded leakage-resilience.

One of the main reasons behind the widespread usage of the bounded leakage model is that formally proving the security of a cryptographic algorithm in this model is more approachable than for most other leakage models. However, bounded leakage does not directly capture real-world side-channel attacks [39]. For example, transcripts produced via power analysis are typically much longer than the secret key under attack but, unlike bounded leakage, are inherently *noisy*. Motivated by this limitation, several models for noisy leakage have been studied in the literature. On the practical front, the most popular measure of a given leakage's "noisiness" is mutual information [38, 34]. More precisely, if X denotes the secret and Z is leakage from X, then $\mathsf{MI}(X;Z)$, the mutual information between X and Z, captures the mutual dependence between X and Z. Ideally, we would like to design cryptographic schemes that are secure against all noisy leakages Z satisfying $\mathsf{MI}(X;Z) \leq \delta$ for δ as large as possible.

Another closely related noise measure is the statistical distance [14] (a.k.a. the total variation distance) between P_{XZ} (the joint distribution of X and Z) and $P_X \otimes P_Z$ (the product distribution of X and Z, i.e., $(P_X \otimes P_Z)(x,z) = P_X(x) \cdot P_Z(z)$), denoted $\mathsf{SD}(P_{XZ}; P_X \otimes P_Z)$. The two measures are related via Pinsker's inequality, which implies that

$$\mathsf{SD}(P_{XZ}; P_X \otimes P_Z) \leq \sqrt{\frac{1}{2}\mathsf{MI}(X;Z)}.$$

This means that a scheme which is leakage-resilient against all leakages Z such that $\mathsf{SD}(P_{XZ}; P_X \otimes P_Z) \leq \delta$ is resilient against all leakages Z such that $\mathsf{MI}(X;Z) \leq 2\delta^2$. Other noise measures have been considered, including the average conditional min-entropy [31] and the average ℓ_2-norm between the marginal distribution X and the conditional distributions $X|Z=z$ [35].[7]

A similar situation can be observed in the context of implementation-level countermeasures and masking. There, one typically considers a stateful cryptographic circuit $\Gamma(k)$ (where k is the secret key) in the presence of adversaries that interact with the circuit via the input-output interface over several rounds, and continuously get leakage from the circuit wires in each round. Abstract leakage models have been introduced, such as the threshold probing model [23] (in which the adversary can probe a bounded number of wires in the circuit) and the random probing model [14] (in which the adversary can recover intermediate values in the circuit only with some probability). But despite the security of

[7] The statistical distance term $\mathsf{SD}(P_{XZ}; P_X \otimes P_Z)$ corresponds (up to a multiplicative 1/2 factor) to the ℓ_1-norm between P_{XZ} and $P_X \otimes P_Z$.

masked implementations is conveniently analyzed in these models, actual implementations are again better reflected by the noisy leakage model [35], which instead bounds the noisiness of the information retrieved from intermediate values based on the statistical distance and the mutual information metrics.

1.1 Reductions as a Bridge from Theory to Practice

As a result of the above discussion, on the one hand, there are many (primitive-level or implementation-level) cryptographic schemes that can be proven secure in the presence of bounded leakage or threshold/random probing. On the other hand, real-world side-channel attacks yield leakage whose noisiness can be measured by means of mutual information and statistical distance, but that is not bounded in length and leaks about all intermediate values. In this light, it is a fundamental question to study the connection between different leakage models, towards understanding whether cryptographic schemes formally proven secure under less realistic leakage assumptions remain secure against more realistic ones.

In the context of primitive-level countermeasures, progress towards answering the above question comes from a recent work of Brian, Faonio, Obremski, Ribeiro, Simkin, Skórski, and Venturi [9], which studied the relationship between the bounded leakage model and various notions of noisy leakage in a very general setting. More precisely, they consider a general *simulation paradigm*. Given a secret distribution X on $\mathcal{X}$ and a leakage Z from X, they ask if there exists a simulator Sim which is allowed to choose any bounded leakage function $g : \mathcal{X} \to \{0,1\}^\ell$, learns $g(X)$, and, after post-processing of $g(X)$, outputs a simulated leakage Z' such that

$$(X, Z) \approx_\varepsilon (X, Z'),$$

where $\approx_\varepsilon$ means that the two joint distributions are within statistical distance at most ε of each other, for a small error term ε. In other words, no adversary can distinguish (with non-negligible advantage) between the real secret-leakage pair (X, Z) and the fake pair (X, Z') where Z' is produced with only the help of a single query of ℓ-bounded leakage. On the positive side, using this paradigm, they showed that many cryptographic schemes resilient to ℓ bits of bounded leakage are also resilient to ℓ'*-min-entropy noisy leakage* [31] (i.e., the class of all leakages Z on a secret X such that Z drops the min-entropy of X by at most ℓ' bits), with $\ell' \approx \ell$ and little ε (as a function of the security parameter).[8]

In the context of implementation-level countermeasures, Duc, Dziembowski, and Faust showed an interesting reduction between the more abstract threshold probing model and the more realistic noisy leakage model, using random probing as a useful intermediate abstraction [14], which has then been (in part heuristically) connected to practical side-channel attacks [15].

[8] More precisely, $\widetilde{\mathbf{H}}_\infty(X|Z) \geq \mathbf{H}_\infty(X) - \ell'$ where $\mathbf{H}_\infty(X) = -\log(\max_x \Pr[X = x])$ denotes the *min-entropy* of X and $\widetilde{\mathbf{H}}_\infty(X|Z) = \mathbb{E}_{z \sim Z}\left[2^{-\mathbf{H}_\infty(X|Z=z)}\right]$ denotes the *average conditional min-entropy* of X given Z.

1.2 Limitations of Statistical Distance and Mutual Information

Although [9] derived positive results for some types of noisy leakages, they also showed that it is *impossible* to obtain non-trivial simulation theorems for noisy leakages based on statistical distance and mutual information via bounded leakage. The reason behind this is simple and instructive. Define the class of δ-SD-noisy leakages of X to be the set of all random variables Z such that

$$\mathsf{SD}(P_{XZ}; P_X \otimes P_Z) \leq \delta. \tag{1}$$

First, note that it is trivial to simulate Z with error δ *even without access to bounded leakage from* X. In fact, by Equation (1), the simulator can simply output Z' sampled independently according to the marginal P_Z. To complement this, [9] also shows that increasing the amount of bounded leakage available does not help in decreasing the error much compared to the trivial simulator. Indeed, there exist secret-leakage distributions P_{XZ} such that Z is δ-SD-noisy leakage from X, but Z cannot be simulated with error $\varepsilon < \delta/2$ *even with* $n-1$ *bits of leakage from* X. More precisely, let X be uniform over $\{0,1\}^n$, and consider what we call the *catastrophic* leakage Z from X defined as follows: with probability δ, set $Z = X$; otherwise, set $Z = \perp$.[9] It holds that Z is δ-SD-noisy leakage from X. To see intuitively why we cannot simulate Z with error below $\delta/2$ from $n-1$ bits of bounded leakage from X, suppose that we query X to learn the $(n-1)$-bounded leakage $(X(1), X(2), \ldots, X(n-1))$, where $X(i)$ is the i-th bit of X. If we wish to simulate Z, then we need to output X with probability approximately δ. However, this means that in that case we will have to guess $X(n)$, and we will fail and be caught by the adversary with probability approximately $\delta \cdot 1/2 = \delta/2$. A similar argument yields an impossibility result for simulating the analogous notion of δ-MI-noisy leakage (i.e., all random variables Z such that $\mathsf{MI}(X; Z) \leq \delta$).

From a practical perspective, the above is unsatisfactory because without countermeasures δ decreases poorly with noise (e.g., see [15, Equation (7)]). Since good simulation can only be obtained by making δ exponentially small, it implies that formal security guarantees require extremely high noise levels that are not intrinsically present in actual implementations. As a result, the only way to exploit the reduction to bounded leakage is to rely on masking even for the leakage-resilient parts of an implementation. This goes against the aforementioned expectation that bounded leakage can be ensured without expensive countermeasures in this case, thanks to frequent state refreshing.

A similar limitation can be found in the reduction from noisy leakage to random probing of Duc, Dziembowsi and Faust [14], where δ-SD-noisy leakage from a secret supported on a set $\mathcal{X}$ can only be simulated with random probes having parameter $\delta \cdot |\mathcal{X}|$, although this "field size loss" does not seem to be observed for practically-relevant leakage functions [34, 4].

[9] This corresponds to the random probing model of [23, 14] in a large (n-bit) field.

1.3 A High-Level Overview of Our Contributions

In this paper, we show that the above limitations are not an insurmountable barrier towards general simulation theorems for practical noisy leakage models, but rather an invitation for further refining the statistical distance and mutual information metrics as empirical measures of quality for side-channel attacks.

Starting with the limitations of the simulation via bounded leakage, the issue with statistical distance and mutual information is that they cannot distinguish between innocent leakages such as "$Z = X(1)$ with probability 1" and catastrophic leakages such as "$Z = X$ with probability $1/n$ and $Z = \perp$ otherwise". Positing that such edge cases are the main impediment standing in front of practically useful simulation theorems, we explore ways to circumvent them in order to better match practical side-channel attacks. Towards this goal, we study noisy leakage models based on *hockey-stick divergences* [36], a well-known family of divergences that generalizes statistical distance (and is a special case of f-divergences).

Definition 1 (t-hockey-stick divergence). *For a real number $t \geq 0$, the* t-hockey-stick divergence *between two distributions P and Q supported on a discrete set $\mathcal{X}$, denoted by $\mathsf{SD}_t(P;Q)$, is defined as*

$$\mathsf{SD}_t(P;Q) = \sup_{\mathcal{S}}[P(\mathcal{S}) - 2^t \cdot Q(\mathcal{S})],$$

where the supremum is taken over all sets $\mathcal{S} \subseteq \mathcal{X}$.[10]

Equivalently, we have $\mathsf{SD}_t(P;Q) \leq \delta$ if and only if

$$P(\mathcal{S}) \leq 2^t \cdot Q(\mathcal{S}) + \delta \tag{2}$$

for all sets $\mathcal{S} \subseteq \mathcal{X}$. It is easy to see that $\mathsf{SD}_0 = \mathsf{SD}$, i.e., the 0-hockey-stick divergence is the statistical distance. These divergences form the basis of differential privacy [17] (approximate differential privacy is equivalent to an upper bound on a hockey-stick divergence [3]), something which we exploit in our results.

Following the previous approach for SD-noisy leakage, considering hockey-stick divergences leads to a noisy leakage model which is a two-parameter generalization of the SD-noisy leakage model: we say that Z is (t, δ)-SD-noisy leakage from X if $\mathsf{SD}_t(P_{XZ}; P_X \otimes P_Z) \leq \delta$. In a nutshell, the additional parameter t in our model allows us to avoid the catastrophic examples that sever the connection between bounded leakages and SD-noisy leakages. We use it to establish several properties of (t, δ)-SD-noisy leakage which we expect will be useful in practical applications. This includes: (i) a simulation theorem for (t, δ)-SD-noisy leakage from bounded leakage, and (ii) a composition theorem for (t, δ)-SD-noisy leakages, which allows one to argue about the combination of multiple (t, δ)-SD-noisy leakages.

[10] Hockey-stick divergences are usually defined with an e^t factor as opposed to the 2^t factor we use here. We opt for the latter because it leads to cleaner theorem statements; this change has no other consequences.

As a complement, we also study a natural *reverse* variant of (t, δ)-SD-noisy leakage, which we call (t, δ)-RevSD-noisy leakage, in which the roles of the distributions P_{XZ} and $P_X \otimes P_Z$ are swapped (i.e., we require that $\mathsf{SD}_t(P_X \otimes P_Z; P_{XZ}) \leq \delta$, and note that SD_t is not symmetric). We then show a simulation theorem for RevSD-noisy leakage from the random probing leakage model. This simulation theorem is a strict generalization of the main result of [14] (which we obtain as a special case by setting $t = 0$), and it allows us to mitigate the field size loss incurred in their simulation by random probing.

We conclude the paper by investigating the t and δ parameters that can be obtained for realistic leakage functions and noise levels. Compared to prior work [9], our concrete evaluations allow us to put forward considerable improvements of the simulation error for modest amounts of bounded leakage, both for the Hamming weight function and variants of which the deterministic part is bijective (ruling out trivial simulation). Combined with our composition theorems, these results can even be used to state formal guarantees for leakage-resilient modes of operation based on physical assumptions that can be matched by parallel hardware implementations (e.g., of the AES), confirming the intuition that bounded leakage can be ensured without (expensive) masking techniques.

We also discuss the practical impact of our improved reduction from (t, δ)-RevSD-noisy leakage to random probing. Although it remains conceptually contrasted since the δ parameter can only be used to hide the field size dependency in the reduction of [14], we show that the good scaling of the δ parameter in the noise level of realistic leakage functions makes this mitigation relevant, especially if masking is implemented in small fields (e.g., $\mathbb{F}_{2^8}$ for the AES). This contribution is a more consolidating one, since Prest *et al.* already proposed a noisy leakage model allowing to get rid of the field size penalty (at the cost of using a metric that scales worse with the noise than the mutual information or statistical distance) [34]. It nevertheless illustrates the unifying nature of hockey-stick divergences for cryptography in the presence of leakage.

2 More Detailed Overview of our Contributions

We now proceed with a more technical overview of our results, followed by a discussion about their practical implications. Our main new noisy leakage model is defined analogously to the notion of SD-noisy leakage as follows.

Definition 2 ((t, δ)-SD-noisy leakage). *Let X be a random variable over $\mathcal{X}$. Then, we say that a randomized function $f : \mathcal{X} \to \mathcal{Z}$ is a* (t, δ)-SD-noisy leakage function *from X if, denoting $Z = f(X)$, it holds that*

$$\mathsf{SD}_t(P_{XZ}; P_X \otimes P_Z) \leq \delta.$$

We denote the set of (t, δ)-SD-noisy leakage functions from X by $\mathsf{SD}_{t,\delta}(X)$, and we also say that $Z = f(X)$ is (t, δ)-SD-noisy leakage *from X.*

Since $\mathsf{SD}_0 = \mathsf{SD}$, we recover δ-SD-noisy leakage as $(t = 0, \delta)$-SD-noisy leakage. The useful properties (simulation via bounded leakage, composition) that we establish for (t, δ)-SD-noisy leakage actually hold as is for an even broader class of noisy leakages also inspired by hockey-stick divergences, which we call GSD-noisy leakage (the "G" standing for "Generalized"). We refrain from defining it formally here, and instead present the relevant definition later in Section 4. All of our results are established directly for (t, δ)-GSD-noisy leakage, as this leads to a much cleaner technical discussion, and they carry over automatically to (t, δ)-SD-noisy leakage which we use for our practical applications.

2.1 Simulation via Bounded Leakage

As discussed above, it is trivial to simulate δ-SD-noisy leakage from even 0 bits of bounded leakage with statistical error δ. Moreover, by [9], this cannot be improved much, even if we allow $n - 1$ bits of bounded leakage (assuming that $X \in \{0,1\}^n$). As our first technical result, we establish the following simulation theorem for (t, δ)-SD-noisy leakage from bounded leakage.

Theorem 1 (Informal). *For any X and $\alpha > 0$, it is possible to simulate the class of (t, δ)-SD-noisy leakage functions from X using $\lceil t + \log\ln(1/\alpha) \rceil$ bits of bounded leakage from X, with statistical error $\delta + \alpha$.*

For formal statements and proofs, see Section 5.

Given Theorem 1, we may see the parameter t as controlling the number of bits of bounded leakage required for simulation, and the parameter δ as controlling the statistical simulation error. At first sight, it may seem that we are not improving over the trivial simulator for δ-SD-noisy leakage, which also has error δ and uses 0 bits of bounded leakage. However, this is not the case as the additional parameter t now affords us significant freedom. In particular, we expect that when fitting concrete, widely used models for real-world side-channel attacks (e.g., Hamming weight leakages with additive Gaussian noise) into the (t, δ)-SD-noisy leakage model, we can significantly decrease δ by slightly increasing t, therefore trading some extra bits of bounded leakage for a much smaller statistical simulation error. Our empirical evaluation in Section 8, confirms this behavior.

Theorem 1 can be used to automatically establish that a broad class of cryptographic primitives resilient to bounded leakage are also resilient to (t, δ)-SD-noisy leakage for good choices of t and δ. As a concrete example, suppose that we have a symmetric-key PRNG that is γ-resilient to ℓ-bounded leakage with $\ell = \log(n)$ for some security parameter n [33]. This guarantees that no adversary with access to arbitrary $\log(n)$-bounded leakage from the secret key can predict the next pseudorandom block with advantage more than γ. Then, combining this with Theorem 1 (where X plays the role of the secret key) immediately implies that, given any parameters $\alpha, \delta > 0$, the same scheme is γ'-resilient to (t, δ)-SD-noisy leakage with $\gamma' = \gamma + \delta + \alpha$ and $t = \log(n) - \log\ln(1/\alpha)$.

2.2 Composition Theorem

There exist situations where the physical implementation of a cryptographic scheme may provide the adversary with several samples of noisy leakage. For example, a (round-based) hardware implementation of the AES will provide a few leakage samples per round, typically correlated with the Hamming weight of the intermediate value manipulated by the device. In such a case, it can be useful to have access to formal composition theorems for the noisy leakage model being used, so that we can formally argue about the combination of these multiple leakage samples. At an abstract level, consider the scenario where m noisy leakage samples $Z_1, \ldots, Z_m$ are computed from a secret random variable X. If we know that each Z_i is (t_i, δ_i)-SD-noisy leakage from X, and that for each $i \neq j$ it holds that Z_i and Z_j are conditionally independent given X, then what can we say about the noisiness of the *global leakage* $Z = (Z_1, \ldots, Z_m)$?

We prove the following composition theorem for (t, δ)-SD-noisy leakages that shows that such noisy leakages compose nicely, yielding a global leakage that is also simulatable via bounded leakage with good parameters.

Theorem 2 (Informal). *Suppose that $Z_1, \ldots, Z_m$ are conditionally independent given a secret random variable X and the samples Z_i are (t_i, δ_i)-SD-noisy leakage from X for $i \in [m]$. Then, for any $\alpha > 0$, the global leakage $Z = (Z_1, \ldots, Z_m)$ can be simulated using $\lceil \log \ln(1/\alpha) + \sum_{i=1}^{m} t_i \rceil$ bits of bounded leakage from X with statistical error $\alpha + \sum_{i=1}^{m} \delta_i$.*

For formal statements and proofs, see Section 6.

For concrete leakages, the parameter t should be small, of the order $\log(n)$ for a security parameter n. On the other hand, δ will be negligible in the noise level. Therefore, the blow-up in the simulation error compared to the original δ_i's will also be small. Note that since practical leakage functions are often close to a deterministic function of X corrupted by additive noise [37], the conditional independence condition boils down to an independent noise one, which is a standard approximation. Note also that the t_i's in Theorem 2 do not need to be integer-valued. Not having to round each t_i to its ceiling can provide significant gains with respect to simulation when composing many noisy leakages.

In the full version of the paper we also study the notion of *strong composition* for a natural strengthening of the (t, δ)-SD-noisy leakage model.

2.3 Simulation via Random Probing

As already briefly mentioned above, a previous success story in linking practical noisy leakage models and theoretically-minded leakage models stems from work of Prouff and Rivain [35] and Duc, Dziembowski, Faust, and Standaert [14, 15] on compilers for leakage-resilient arithmetic circuits. Most relevant to our setting, Duc, Dziembowski, and Faust [14] showed that the leakage-resilient circuit compiler of Ishai, Sahai, and Wagner [23], which efficiently transforms any given arithmetic circuit into an equivalent circuit resilient to threshold probing leakage from the wires during computation, also yields a circuit resilient to SD-noisy

leakage on the wires.[11] The key lemma behind the main result of [14] (from which their applications to circuit computation easily follow) states that δ-SD-noisy leakage from a uniform secret X over $\mathcal{X}$ can be perfectly simulated by p-random probing leakage from X with $p = \delta|\mathcal{X}|$.[12] The linear dependence of p on the support size $|\mathcal{X}|$ in this simulation has been noted to be unsatisfactory and avoidable for concrete applications of this result [15, 34, 4]. We extend the key lemma of [14] for δ-SD-noisy leakage to a more general notion of *reverse* (t, δ)*-SD-noisy leakage*. In particular, this extension allows us to alleviate the "support size penalty" in the noisy-to-probing leakage simulation. The notion of reverse (t, δ)-SD-noisy leakage we use is similar to (t, δ)-SD-noisy leakage, and can also be seen as a natural generalization of δ-SD-noisy leakage.

Definition 3 ((t, δ)-RevSD-noisy leakage). *Let X be a random variable over $\mathcal{X}$. Then, we say that a randomized function $f : \mathcal{X} \to \mathcal{Z}$ is a* (t, δ)-RevSD-noisy leakage function *from X if, denoting $Z = f(X)$, it holds that*

$$\mathsf{SD}_t(P_X \otimes P_Z; P_{XZ}) \leq \delta.$$

We denote the set of (t, δ)-RevSD-noisy leakage functions from X by $\mathsf{RevSD}_{t,\delta}(X)$, and we also say that $Z = f(X)$ is (t, δ)-RevSD-noisy leakage *from X.*

We next highlight the connection we prove between RevSD-noisy leakage and random probing leakage, which generalizes the key lemma of [14, Lemma 2] mentioned above (which corresponds to the $t = 0$ case).

Theorem 3 (Informal). *Let X be uniform over $\mathcal{X}$ and suppose that Z is (t, δ)-RevSD-noisy leakage from X. Then, Z is perfectly simulatable by p-random probing leakage from X with $p = (1 - 2^{-t}) + \delta \cdot 2^{-t} \cdot |\mathcal{X}|$.*

For formal statements and proofs, see Section 7.

This result generally improves on [14, Lemma 2]. However, there still exists a tradeoff between the need to keep the t parameter small so that $(1 - 2^{-t})$ is small and the fact that the scaling of the δ parameter with respect to the noise level of the implementation gets worse for small t values (recall that for $t = 0$ we have that (t, δ)-RevSD-noisy leakage is equivalent to δ-SD noisy leakage). The empirical results of Section 8 nevertheless confirm that Theorem 3 can lead to sweet spots for practically-relevant leakage functions and noise levels.

In the full version of the paper we additionally present a reduction that trades the aforementioned field size penalty for positive statistical simulation error, and show how to apply the above reductions in order to obtain leakage-resilient circuit compilers tolerating RevSD-noisy leakage from the wires.

[11] A tuple $(Z_1, \ldots, Z_\ell)$ is τ-threshold probing leakage from $(X_1, \ldots, X_\ell)$ if $Z_i = X_i$ for at most τ indices $i \in [\ell]$, and $Z_i = \perp$ otherwise.

[12] Suppose that X is supported on $\mathcal{X}$. Then, $Z \in \mathcal{X} \cup \{\perp\}$ is p-random probing leakage from X if $\Pr[Z = X] = p$ and $\Pr[Z = \perp] = 1 - p$.

2.4 Practical Interpretation

Informally, the positive observations we obtain in the paper essentially stem from the fact that (t, δ)-SD-noisy and RevSD-noisy leakage scale much better with the implementation noise than δ-SD-noisy leakage (or the mutual information). This is because these former metrics are computed by integrating the (joint and product) leakage distributions over the whole leakage support. By contrast (t, δ)-SD-noisy (resp., RevSD-noisy) leakage are computed by integrating these distributions in regions where the joint (resp., product) distribution is 2^t times larger than the product (resp., joint) one. With modest t and realistic noise levels, these regions have small probability, explaining a faster decrease of δ.

This better scaling directly has strong impact for PRNGs like the one of [33] and its many follow-ups. Say, for example, that we want to ensure 128-bit security using the reduction of [9]. Ensuring 2^{-128} simulation error would require a noise variance in the $2^{128} \approx 10^{39}$ range, which no device offers intrinsically.[13] Even tolerating lower (e.g., 64-bit) security keeps the required parameters completely impractical. The only solution is then to use masking to "amplify" the noise to this level, which is expensive and contradicts the goal of leakage-resilience, where re-keying aims to maintain high physical security without masking.

In contrast, we highlight in Section 8 that for (t, δ)-SD-noisy and RevSD-noisy leakage it is possible to simulate with 2^{-128} simulation error by combining a modest amount of bounded leakage (typically, $\log(n)/c$ with c a small constant) with noise levels that are concretely reachable (e.g., in the 10^3 range) and may even be intrinsically present in hardware/parallel implementations.

To give a concrete illustration, assume for simplicity that masking with d shares raises the noise variance to a power d at the cost of quadratic implementation overheads. This means that for a leaking device with noise variance $\approx 10^3$ (which provides $\approx 2^{-128}$ simulation error with our reduction), the reduction of [9] would require 13-share masking to ensure the same simulation error (since $(10^3)^{13} = 10^{39}$), leading to a factor $13^2 = 169$ of implementation overheads.

Finally, despite our reduction to random probing being limited to smaller t values whenever one wants to ensure a low probing probability, we also show in Section 8 that Theorem 3 can lead to useful results in the case of small- to medium-sized fields (e.g., $\mathbb{F}_{2^8}$ for the AES), since reasonable noise levels can then be used to hide the field size dependency of the noise requirements with δ.

3 Preliminaries

3.1 Notation

Random variables are denoted by uppercase roman letters such as X, Y, and Z. Given a random variable X, we denote its probability distribution by P_X, its

[13] The noise requirements of a masked implementation are more accurately expressed in terms of a side-channel Signal-to-Noise Ratio (SNR) [28], which we defer to Section 8 to keep this overview of our contributions concise.

expected value by $\mathbb{E}[X]$, and its variance by $\mathbb{V}(X)$. We write $x \sim X$ to mean that x is sampled according to the distribution of X.

Given two random variables X and Z, we denote their joint probability distribution by P_{XZ} and their *product distribution* by $P_X \otimes P_Z$, i.e., $(P_X \otimes P_Z)(x, z) = P_X(x) \cdot P_Z(z)$, where P_X and P_Z are the marginal distributions of X and Z, respectively. Note that if X and Z are independent, then $P_{XZ} = P_X \otimes P_Z$. We use uppercase calligraphic letters, such as $\mathcal{S}$ and $\mathcal{T}$, to denote sets. We write log for the base-2 logarithm and ln for the natural logarithm.

3.2 The Leakage Simulation Paradigm

In this section, we formally define our notion of simulation of one family of leakages by another family. We follow the definition from [9].

Definition 4 (Leakage simulation [9]). *Given a random variable X supported on $\mathcal{X}$ and two families $\mathcal{F}(X)$ and $\mathcal{G}(X)$ of leakage functions from X, we say that $\mathcal{F}(X)$ is* ε-simulatable from $\mathcal{G}(X)$ *if for all $f \in \mathcal{F}(X)$ there is a (possibly inefficient) randomized algorithm Sim_f such that*

$$(X, Z) \approx_\varepsilon \left(X, \mathsf{Sim}_f^{\mathsf{Leak}(X,\cdot)}\right), \tag{3}$$

where $Z = f(X)$ and the oracle $\mathsf{Leak}(X, \cdot)$ accepts a single query $g \in \mathcal{G}(X)$ and outputs $g(X)$. Furthermore, when $\mathcal{G}(X)$ is the family of all ℓ-bounded leakage functions $g : \mathcal{X} \to \{0,1\}^\ell$ and Equation (3) holds, we say that $\mathcal{F}(X)$ is ε-simulatable from ℓ bits of bounded leakage.

3.3 A Basic Property of Hockey-Stick Divergences

We state here a basic but useful property of hockey-stick divergences, generalizing the analogous property for the statistical distance.

Lemma 1. *Let P and Q be two distributions supported on $\mathcal{X}$. Then,*

$$\mathsf{SD}_t(P; Q) = \sum_{x \in \mathcal{X}} \max(0, P(x) - 2^t Q(x)).$$

Proof. Looking ahead, this simple argument is implicit in our proof of Theorem 6. We isolate and reproduce it here for the sake of exposition.

Let $\mathcal{B} = \{x \in \mathcal{X} \mid P(x) - 2^t Q(x) > 0\}$. For any set $\mathcal{S} \subseteq \mathcal{X}$, it holds that

$$\begin{aligned} P(\mathcal{S}) - 2^t Q(\mathcal{S}) &= \big(P(\mathcal{S} \setminus \mathcal{B}) - 2^t Q(\mathcal{S} \setminus \mathcal{B})\big) + \big(P(\mathcal{S} \cap \mathcal{B}) - 2^t Q(\mathcal{S} \cap \mathcal{B})\big) \\ &\le 0 + \big(P(\mathcal{S} \cap \mathcal{B}) - 2^t Q(\mathcal{S} \cap \mathcal{B})\big) \\ &\le P(\mathcal{B}) - 2^t Q(\mathcal{B}) \\ &= \sum_{x \in \mathcal{X}} \max(0, P(x) - 2^t Q(x)), \end{aligned}$$

where the two inequalities and the last equality use the definition of $\mathcal{B}$. The desired statement now follows because $\mathsf{SD}_t(P; Q) = \sup_{\mathcal{S}}[P(\mathcal{S}) - 2^t Q(\mathcal{S})]$. □

4 Our Leakage Models

In this section we recall the definitions of (t, δ)-SD-Noisy and (t, δ)-RevSD-Noisy leakage, and introduce the more general (t, δ)-GSD-Noisy leakage model. Intuitively, in the generalized definition, we measure the leakage quality by bounding the hockey-stick divergence between the distributions P_{XZ} and $P_X \otimes Q$ for any suitable distribution Q over $\mathcal{Z}$ (not necessarily the marginal P_Z).

Definition 2 ((t, δ)-SD-noisy leakage). *Let X be a random variable over $\mathcal{X}$. Then, we say that a randomized function $f : \mathcal{X} \to \mathcal{Z}$ is a* (t, δ)-SD-noisy leakage function *from X if, denoting $Z = f(X)$, it holds that*

$$\mathsf{SD}_t(P_{XZ}; P_X \otimes P_Z) \leq \delta.$$

We denote the set of (t, δ)-SD-noisy leakage functions from X by $\mathsf{SD}_{t,\delta}(X)$, and we also say that $Z = f(X)$ is (t, δ)-SD-noisy leakage *from X.*

Definition 3 ((t, δ)-RevSD-noisy leakage). *Let X be a random variable over $\mathcal{X}$. Then, we say that a randomized function $f : \mathcal{X} \to \mathcal{Z}$ is a* (t, δ)-RevSD-noisy leakage function *from X if, denoting $Z = f(X)$, it holds that*

$$\mathsf{SD}_t(P_X \otimes P_Z; P_{XZ}) \leq \delta.$$

We denote the set of (t, δ)-RevSD-noisy leakage functions from X by $\mathsf{RevSD}_{t,\delta}(X)$, and we also say that $Z = f(X)$ is (t, δ)-RevSD-noisy leakage *from X.*

Definition 5 ((t, δ)-GSD-noisy leakage). *Let X be a random variable over $\mathcal{X}$. Then, we say that a randomized function $f : \mathcal{X} \to \mathcal{Z}$ is a* (t, δ)-GSD-noisy leakage function *from X if, denoting $Z = f(X)$, there exists a distribution Q on $\mathcal{Z}$ such that*

$$\mathsf{SD}_t(P_{XZ}; P_X \otimes Q) \leq \delta.$$

We denote the set of (t, δ)-GSD-noisy leakage functions from X by $\mathsf{GSD}_{t,\delta}(X)$, and we also say that $Z = f(X)$ is (t, δ)-GSD-noisy leakage *from X.*

In the next sections we establish useful properties of these leakage models. In Section 5, we establish simulation theorems for (t, δ)-GSD-noisy leakage (and thus for (t, δ)-SD-noisy leakage too) from bounded leakage. In particular, this yields Theorem 1. Then, in Section 6, we prove composition theorems for these models, yielding Theorem 2. The relationship between RevSD-noisy leakage and the random probing model is studied in Section 7. Empirical evaluations of these different leakage models are finally discussed in Section 8.

In the full version, we study the relationship between the (t, δ)-SD-noisy leakage model and the *dense leakage* model studied by [9].

5 Simulating GSD-Noisy Leakage via Bounded Leakage

In this section we prove our main simulation theorem, which states (using the language from Definition 4) that the class of (t, δ)-GSD-noisy leakages is $(\alpha + \delta)$-simulatable from $\ell = t + \log\ln(1/\alpha)$ bits of bounded leakage for any $\alpha > 0$. This immediately implies Theorem 1. The simulator we use to establish this result is based on rejection sampling. It is a close variant of the simulator used in [9] with a (key) new, more streamlined and tighter, analysis. The rejection sampling simulator is described in Algorithm 1 for some (t, δ)-GSD-noisy leakage Z from X witnessed by a distribution Q in the sense that for all sets $\mathcal{S}$ it holds that

$$P_{XZ}(\mathcal{S}) \leq 2^t \cdot (P_X \otimes Q)(\mathcal{S}) + \delta.$$

Function $\mathsf{Leak}(x, r)$
  for $i := 0$ **to** $2^\ell - 1$ **do**
    Sample z according to Q using the random tape r
    with probability $\min\left(2^{-t} \cdot \frac{P_{XZ}(x,z)}{(P_X \otimes Q)(x,z)}, 1\right)$ **do**
      return i
    end
  end
  return 2^ℓ
end
Function $\mathsf{Sim}^{\mathsf{Leak}(x,\cdot)}$
  $r \leftarrow$ a random tape
  $i := \mathsf{Leak}(x, r)$
  $z' \leftarrow$ the i-th sample according to Q using random tape r
  return z'
end

Algorithm 1: The (t, ℓ)-*rejection sampling simulator* for the (t, δ)-GSD-noisy leakage $Z = f(X)$, where Q is a distribution on $\mathcal{Z}$ such that $P_{XZ}(\mathcal{S}) \leq 2^t \cdot (P_X \otimes Q)(\mathcal{S}) + \delta$ for all sets $\mathcal{S}$.

Remark 1 (Differences with respect to the simulator from [9]). We next outline the main differences with respect to the simulator from [9]. First, in our simulator the z_i's are sampled according to Q, and not necessarily P_Z. Moreover, we always output the last sample if we have rejected all previous samples. Finally, and of particular importance to our improved analysis, we accept a given sample z and stop with probability $\min\left(2^{-t} \cdot \frac{P_{XZ}(x,z)}{(P_X \otimes Q)(x,z)}, 1\right)$. This means that if $2^{-t} \cdot \frac{P_{XZ}(x,z)}{(P_X \otimes Q)(x,z)} \geq 1$ then we accept z and stop with probability 1. In contrast, the simulator from [9] rejected z automatically in this case.

Remark 2 (Complexity of our simulator). We discuss the computational complexity of our simulator, as it may be relevant for some (non-information-theoretic) reductions from noisy leakage-resilience to bounded leakage-resilience. Computing the ℓ leakage bits in Algorithm 1 requires sampling and rejecting 2^ℓ samples in the worst case. Assuming that we have efficient procedures for sampling according to Q and for computing the functions $P_{XZ}(\cdot,\cdot)$, $P_X(\cdot)$, and $Q(\cdot)$, which is a reasonable assumption when $Q = P_Z$ (i.e., when focusing on (t,δ)-SD-noisy leakage) for the noise distributions commonly used to model real-world side-channel attacks, we conclude that our simulator is efficient whenever ℓ is logarithmic in our parameter of interest. According to our simulation theorem, this holds when t is logarithmic, which is also the setting we study empirically in Section 8.

We begin by proving the following two lemmas which are stating useful properties of our rejection sampling simulator in Algorithm 1.

Lemma 2. *Let* $R(x) = 1 - \mathbb{E}_Q\left[\min\left(2^{-t}\cdot\frac{P_{XZ}(x,Z)}{(P_X\otimes Q)(x,Z)},1\right)\right]$ *be the sample rejection probability for the* (t,ℓ)*-rejection sampling simulator on input* $X = x$*, and let* $P_{\mathsf{Sim}|X=x}$ *be the conditional distribution for the simulator's output on input* $X = x$*. Then,*

$$\begin{aligned} P_{\mathsf{Sim}|X=x}(z) &= \sum_{i=0}^{2^\ell-2} R(x)^i \min\left(2^{-t}P_{Z|X=x}(z), Q(z)\right) + R(x)^{2^\ell-1}Q(z) \\ &\geq \frac{1-R(x)^{2^\ell}}{1-R(x)} \min\left(2^{-t}P_{Z|X=x}(z), Q(z)\right). \end{aligned}$$

Proof. In the first iteration, the simulator samples a given z and accepts it with probability

$$\begin{aligned} p_x(z) &= \min\left(2^{-t}\frac{P_{XZ}(x,z)}{(P_X\otimes Q)(x,z)},1\right)\cdot Q(z) \\ &= \min\left(2^{-t}\frac{P_{XZ}(x,z)}{P_X(x)}, Q(z)\right) \\ &= \min\left(2^{-t}P_{Z|X=x}(z), Q(z)\right), \end{aligned}$$

and rejects otherwise. The probability that the first round does not result in an "accept" is $1 - \mathbb{E}_{z\sim Q}[p_x(z)] = R(x)$. Extending this, the probability of accepting and outputting z in the first round is $p_x(z)$, the probability of rejecting in the first round and accepting and outputting z in the second round is $R(x)\cdot p_x(z)$, and, in general, the probability of rejecting in the first $r-1$ rounds and accepting and outputting z in the r-th round is $R(x)^{r-1}\cdot p_x(z)$. However, in the last iteration the sample is always output, whether it would be rejected or accepted – the probability of reaching this stage and observing output z is $R(x)^{2^\ell-1}\cdot Q(z)$. Summing over the 2^ℓ stages of the algorithm gives the first equation for $P_{\mathsf{Sim}|X=x}(z)$.

For the inequality, notice that $Q(z) \geq \min\left(2^{-t}P_{Z|X=x}(z), Q(z)\right)$, so

$$P_{\mathsf{Sim}|X=x}(z) \geq \sum_{i=0}^{2^\ell-1} R(x)^i \min\left(2^{-t}P_{Z|X=x}(z), Q(z)\right).$$

We obtain the desired inequality by summing this partial geometric series. □

Lemma 3. *Let f be a (t,δ)-GSD-noisy leakage function from X and $Z = f(X)$. Let Q be the associated distribution. Then, the (t,ℓ)-rejection sampling simulator's rejection probability equals*

$$R(x) = 1 - \sum_{z\in\mathcal{Z}} \min\left(2^{-t}P_{Z|X=x}(z), Q(z)\right),$$

and satisfies $1-2^{-t} \leq R(x) \leq 1$ and $\mathbb{E}_X[R(X)] \leq 1-2^{-t}(1-\delta)$.

Proof. The acceptance probability $1-R(x)$ is

$$\begin{aligned}
1-R(x) &= \mathbb{E}_Q\left[\min\left(2^{-t}\frac{P_{XZ}(x,Z)}{(P_X\otimes Q)(x,Z)}, 1\right)\right] \\
&= \sum_{z\in\mathcal{Z}} \min\left(2^{-t}\frac{P_{XZ}(x,z)}{(P_X\otimes P_Q)(x,z)}\cdot Q(z), Q(z)\right) \\
&= \sum_{z\in\mathcal{Z}} \min\left(2^{-t}P_{Z|X=x}(z), Q(z)\right) \\
&\leq \sum_{z\in\mathcal{Z}} 2^{-t}P_{Z|X=x}(z) \\
&= 2^{-t},
\end{aligned}$$

which gives the first equation and the lower bound on $R(x)$. On the other hand, we have $R(x) \leq 1$ because it is a probability. Taking expectation over X gives

$$\begin{aligned}
1-\mathbb{E}_X[R(X)] &= \sum_{x\in\mathcal{X}} P_X(x) \sum_{z\in\mathcal{Z}} \min\left(2^{-t}P_{Z|X=x}(z), Q(z)\right) \\
&= 2^{-t} \sum_{x\in\mathcal{X},z\in\mathcal{Z}} \min\left(P_{XZ}(x,z), 2^t(P_X\otimes Q)(x,z)\right) \\
&= 2^{-t} \sum_{x\in\mathcal{X},z\in\mathcal{Z}} \left(P_{XZ}(x,z) - \max\left(0, P_{XZ}(x,z) - 2^t(P_X\otimes Q)(x,z)\right)\right) \\
&= 2^{-t}\left(1 - \sum_{x\in\mathcal{X},z\in\mathcal{Z}} \max\left(0, P_{XZ}(x,z) - 2^t(P_X\otimes Q)(x,z)\right)\right) \\
&\geq 2^{-t}(1-\delta),
\end{aligned}$$

where the final inequality holds by Lemma 1, since $\mathsf{SD}_t(P_{XZ}; P_X\otimes Q) \leq \delta$ as f is a (t,δ)-GSD-noisy leakage function from X. □

The following result immediately implies Theorem 1.

Theorem 4. *Let f be a (t,δ)-GSD-noisy leakage function from X. Let $Z = f(X)$ and Z' denote the output of the* (t,ℓ)-rejection sampling simulator *on input X. Then, we have that*

$$(X, Z) \approx_{\varepsilon} (X, Z')$$

for $\varepsilon = e^{-2^{\ell-t}} + \delta$. In particular, for any $\alpha > 0$ the class of (t,δ)-GSD-noisy leakage functions from X is $(\alpha+\delta)$-simulatable from ℓ bits of leakage when

$$\ell \geq t + \log \ln(1/\alpha).$$

Proof. We must bound the statistical distance between the true secret-leakage joint distribution P_{XZ} and the fake joint distribution $P_{XZ'}$, where Z' denotes the simulator's output. This will be achieved by first bounding, for any given x, the statistical distance $D(x)$ between the conditional distributions $(\mathsf{Sim}|X=x)$ and $(Z|X=x)$ using Lemma 2. Then, we use Lemma 3 to obtain the desired bound on the original statistical distance. We have that

$$\begin{aligned}
D(x) &= \sum_{z\in\mathcal{Z}} \max\left(0, P_{Z|X=x}(z) - P_{\mathsf{Sim}|X=x}(z)\right) \\
&\leq \sum_{z\in\mathcal{Z}} \max\left(0, P_{Z|X=x}(z) - \frac{1-R(x)^{2^\ell}}{1-R(x)} \min\left(2^{-t}P_{Z|X=x}(z), Q(z)\right)\right) \\
&\leq \left(1 - \frac{1-R(x)^{2^\ell}}{1-R(x)} \cdot 2^{-t}\right) \sum_{z\in\mathcal{Z}} \max\left(0, P_{Z|X=x}(z)\right) \\
&+ \frac{1-R(x)^{2^\ell}}{1-R(x)} \sum_{z\in\mathcal{Z}} \max\left(0, 2^{-t}P_{Z|X=x}(z) - \min\left(2^{-t}P_{Z|X=x}(z), Q(z)\right)\right) \\
&= 1 - \frac{1-R(x)^{2^\ell}}{1-R(x)} \cdot 2^{-t} \\
&+ \frac{1-R(x)^{2^\ell}}{1-R(x)} \left(\sum_{z\in\mathcal{Z}} 2^{-t}P_{Z|X=x}(z) - \sum_{z\in\mathcal{Z}} \min\left(2^{-t}P_{Z|X=x}(z), Q(z)\right)\right) \\
&= 1 - \frac{1-R(x)^{2^\ell}}{1-R(x)} 2^{-t} + \frac{1-R(x)^{2^\ell}}{1-R(x)} \left(2^{-t} - 1 + R(x)\right) \\
&= R(x)^{2^\ell},
\end{aligned}$$

where the first inequality follows from Lemma 2, and the second to last equality from Lemma 3. Next, notice that $R(x)^{2^\ell}$ is a convex function of $R(x)$, and so we can upper bound this by a line drawn through the lower and upper bounds for $R(x)$. Therefore,

$$D(x) \leq (1-2^{-t})^{2^\ell} + \frac{1-(1-2^{-t})^{2^\ell}}{2^{-t}}(R(x) - 1 + 2^{-t}). \tag{4}$$

Finally, we can use Lemma 3 to get a bound on the statistical distance between P_{XZ} and $P_{XZ'}$, where Z' is the simulator's output, which equals $\mathbb{E}_X[D(X)]$. We have that

$$\begin{aligned}
\mathbb{E}_X[D(X)] &\leq (1-2^{-t})^{2^\ell} + \frac{1-(1-2^{-t})^{2^\ell}}{2^{-t}}(\mathbb{E}_X[R(X)] - 1 + 2^{-t}) \\
&\leq (1-2^{-t})^{2^\ell} + \frac{1-(1-2^{-t})^{2^\ell}}{2^{-t}}(1 - 2^{-t}(1-\delta) - 1 + 2^{-t}) \\
&= (1-2^{-t})^{2^\ell} + \left(1-(1-2^{-t})^{2^\ell}\right)\delta \\
&= (1-2^{-t})^{2^\ell}(1-\delta) + \delta \\
&\leq e^{-2^{\ell-t}} + \delta.
\end{aligned}$$

The first inequality uses Equation (4). The second one follows from Lemma 3. The final inequality holds because $1 + y \leq e^y$ for any real y. This yields the first part of the theorem statement. To see the second part, set ℓ so that $\alpha \geq e^{-2^{\ell-t}}$. □

It is natural to wonder how this analysis compares to the indirect one in which we first establish the parameters of (t, δ)-SD-noisy leakage as *dense leakage*, and then apply the known simulator for dense leakage in [9]. The main difference is that we would get worse simulation error through the indirect approach. More precisely, while Theorem 4 guarantees simulation of (t, δ)-GSD-noisy leakage with error $\alpha + \delta$ using $t + \log\ln(1/\alpha)$ bits of bounded leakage, the indirect approach above would only yield simulation error $\alpha + c \cdot \sqrt{\delta}$ using the same amount of bounded leakage, for a small constant $c \geq 1$. Reducing the $\sqrt{\delta}$ term in the simulation error to δ is a significant improvement for practical applications.

Intuitively, the reason why the indirect approach via dense leakage can only yield a $\sqrt{\delta}$ term in the simulation error is that the definition of dense leakage in [9] imposes a "with high probability" constraint on X and Z. Namely, if Z is dense leakage from X, then with high probability over the choices $X = x$ and $Z = z$ we must have $P_{Z|X=x}(z) \leq T \cdot P_Z(z)$ for an appropriate "density parameter" T. On the other hand, GSD-noisy leakage imposes an "in expectation" constraint on X and Z. Namely, if Z is (t, δ)-GSD-noisy leakage from X, then we only require that $\mathbb{E}_{x \sim P_X}[\mathsf{SD}_t(P_{Z|X=x}; Q)] \leq \delta$. One can move from the "in expectation" constraint to the "with high probability" constraint via Markov's inequality. However, this incurs a loss, which causes exactly the δ *vs.* $\sqrt{\delta}$ difference between the two approaches. Our direct analysis of the simulator relies only on the "in expectation" constraint of GSD-noisy leakage, avoiding this loss.

6 Composition of GSD-Noisy Leakages

We now prove our main composition theorem. The theorem below is for two conditionally independent leakages, and applying it $m - 1$ times combined with Theorem 4 directly implies Theorem 2. The approach we take is an adaptation of Dwork and Lei's proof of basic composition for differential privacy [16].

Theorem 5. *Suppose that f_1 and f_2 are (t_1, δ_1)-GSD-noisy and (t_2, δ_2)-GSD-noisy leakage functions from X, respectively, and that the random variables $Z_1 = f_1(X)$ and $Z_2 = f_2(X)$ are independent when conditioned on X. Then $f(X) = (f_1(X), f_2(X))$ is a $(t_1 + t_2, \delta_1 + \delta_2)$-GSD-noisy leakage function from X.*

Proof. Let Q_1 and Q_2 be the distribution on $\mathcal{Z}_1$ and $\mathcal{Z}_2$ (the supports of $Z_1 = f_1(X)$ and $Z_2 = f_2(X)$, respectively) that establish f_1 and f_2 as GSD-noisy leakages, respectively. Then, set Q to be the distribution $Q_1 \otimes Q_2$. To prove our result, we must show that for any set $\mathcal{S} \subseteq \mathcal{X} \times \mathcal{Z}_1 \times \mathcal{Z}_2$,

$$P_{XZ_1Z_2}(\mathcal{S}) \leq 2^{t_1+t_2}(P_X \otimes Q)(\mathcal{S}) + \delta_1 + \delta_2.$$

Using Lemma 1, for $i \in \{1, 2\}$ let

$$\delta_i(x) = \mathsf{SD}_t(P_{Z_i|X=x}; Q_i) = \sum_{z_i \in \mathcal{Z}_i} \max(0, P_{Z_i|X=x}(z_i) - 2^{t_i} Q_i(z_i)).$$

In particular, $\mathbb{E}[\delta_i(X)] = \mathsf{SD}_t(P_{XZ_i}; P_X \otimes Q_i) \leq \delta_i$ because f_i is a (t_i, δ_i)-GSD-noisy leakage from X. Let $\mathcal{S}_x = \{(z_1, z_2) \mid (x, z_1, z_2) \in \mathcal{S}\}$ and $\mathcal{S}_{x,z_1} = \{z_2 \mid (x, z_1, z_2) \in \mathcal{S}\}$. Then,

$$\begin{aligned}
P_{Z_1Z_2|X=x}(\mathcal{S}_x) &= \mathbb{E}_{Z_1|X=x}[P_{Z_2|X=x}(\mathcal{S}_{x,Z_1})] \\
&= \mathbb{E}_{Z_1|X=x}\left[\min\left(1, P_{Z_2|X=x}(\mathcal{S}_{x,Z_1})\right)\right] \\
&\leq \mathbb{E}_{Z_1|X=x}\left[\min\left(1, 2^{t_2} Q_2(\mathcal{S}_{x,z_1}) + \delta_2(x)\right)\right] \\
&\leq \delta_2(x) + \sum_{z_1 \in \mathcal{Z}_1} P_{Z_1|X=x}(z_1) \min\left(1, 2^{t_2} Q_2(\mathcal{S}_{x,z_1})\right) \\
&\leq \delta_2(x) + \sum_{z_1 \in \mathcal{Z}_1} 2^{t_1} Q_1(z_1) \min\left(1, 2^{t_2} Q_2(\mathcal{S}_{x,z_1})\right) \\
&\quad + \sum_{z_1 \in \mathcal{Z}_1} \max\left(0, P_{Z_1|X=x}(z_1) - 2^{t_1} Q_1(z_1)\right) \min\left(1, 2^{t_2} Q_2(\mathcal{S}_{x,z_1})\right) \\
&\leq \delta_2(x) + 2^{t_1+t_2} \sum_{z_1 \in \mathcal{Z}_1} Q_1(z_1) Q_2(\mathcal{S}_{x,z_1}) \\
&\quad + \sum_{z_1 \in \mathcal{Z}_1} \max\left(0, P_{Z_1|X=x}(z_1) - 2^{t_1} Q_1(z_1)\right) \\
&= 2^{t_1+t_2} Q(\mathcal{S}_x) + \delta_1(x) + \delta_2(x).
\end{aligned}$$

Finally, take the expectation over X to get

$$\begin{aligned}
P_{XZ_1Z_2}(\mathcal{S}) &= \mathbb{E}_X[P_{Z_1Z_2|X}(\mathcal{S}_x)] \\
&\leq \mathbb{E}_X\left[2^{t_1+t_2} Q(\mathcal{S}_x) + \delta_1(x) + \delta_2(x)\right] \\
&\leq 2^{t_1+t_2}(P_X \otimes Q)(\mathcal{S}) + \delta_1 + \delta_2.
\end{aligned}$$

The theorem statement follows. □

Remark 3. It is well known that differential privacy enjoys even stronger composition theorems in which parameters do not scale linearly with number of queries, but instead scale with its square root. Given that our leakage model is closely connected to the metric used in differential privacy, it is natural to wonder whether we can derive a similar improvement in the context of GSD-noisy leakages. We show that the answer is positive for a natural restriction of the GSD-noisy leakage model in the full version of the paper.

7 Simulating RevSD-Noisy Leakage via Random Probing

In their seminal work, Duc, Dziembowski, and Faust [14] showed that δ-SD-noisy leakage can be perfectly simulated in the probing leakage model of Ishai, Sahai, and Wagner [23]. An unsatisfactory and unavoidable feature of this connection is that the probing noise required to simulate δ-SD-noisy leakage grows linearly with the field size of the secret [15]. In this section, we generalize this connection to (t, δ)-RevSD-noisy leakage, and show that in this alternative model we can alleviate the field size penalty for simulation by random probing leakage. Before stating our main result in this direction, we define p-random probing leakage.

Definition 6 (p-random probing leakage [14]). *Let X be some random variable supported on $\mathcal{X}$. We say that a random variable $Z \in \mathcal{X} \cup \{\perp\}$ is* p-random probing leakage from X *if* $\Pr[Z = X] = p$ *and* $\Pr[Z = \perp] = 1 - p$.

We have the following result.

Lemma 4. *Let X be uniformly distributed over $\mathcal{X}$ and suppose that Z is (t, δ)-RevSD-noisy leakage from X. Then, Z is 0-simulatable by p-random probing leakage from X with $p = (1 - 2^{-t}) + \delta 2^{-t}|\mathcal{X}|$.*

Duc, Dziembowski, and Faust [14, Lemma 2] proved this result only for the special case $t = 0$, which corresponds to δ-SD-noisy leakage.

Proof (Lemma 4). Our argument follows the proof of [14, Lemma 2] closely. For any given leakage z, we define

$$\pi(z) = \min_{x \in \mathcal{X}} P_{Z|X=x}(z).$$

Note that $\pi(z) \geq 0$ for all z and $\sum_z \pi(z) \leq \sum_z P_Z(z) = 1$. We will also assume that Z is not independent of X, in which case there is a z such that $\pi(z) < P_Z(z)$, and so $\sum_z \pi(z) < 1$. When Z is independent of X it is clear that we can perfectly simulate it using 0-random probing leakage.

The main component of this argument consists in showing that π is "almost" a probability distribution, in the sense that $\sum_z \pi(z)$ is approximately equal to 1. More precisely, we have that

$$1 - \sum_z \pi(z) = \sum_z P_Z(z) - \sum_z \min_{x \in \mathcal{X}} P_{Z|X=x}(z)$$

$$
\begin{aligned}
&= \sum_z (1-2^{-t}) P_Z(z) + \sum_z [2^{-t} P_Z(z) - \min_{x \in \mathcal{X}} P_{Z|X=x}(z)] \\
&= (1-2^{-t}) + \sum_z \max_x [2^{-t} P_Z(z) - P_{Z|X=x}(z)] \\
&\leq (1-2^{-t}) + \sum_z \max_x \max(0, 2^{-t} P_Z(z) - P_{Z|X=x}(z)) \\
&\leq (1-2^{-t}) + \sum_z \sum_x \max(0, 2^{-t} P_Z(z) - P_{Z|X=x}(z)) \\
&= (1-2^{-t}) \\
&+ 2^{-t} \cdot |\mathcal{X}| \cdot \sum_z \sum_x \max(0, (P_X \otimes P_Z)(x,z) - 2^t P_{XZ}(x,z)) \\
&\leq (1-2^{-t}) + 2^{-t} \cdot |\mathcal{X}| \cdot \delta.
\end{aligned}
$$

The last equality uses the fact that X is uniform, and so $P_X(x) = 1/|\mathcal{X}|$ for all $x \in \mathcal{X}$. The last inequality uses the fact that Z is (t, δ)-RevSD-noisy leakage from X and Lemma 1. Let $p = 1 - \sum_z \pi(z)$. By the computation above, we know that $0 < p \leq (1-2^{-t}) + 2^{-t} \cdot |\mathcal{X}| \cdot \delta$. We proceed to show that Z can be perfectly simulated by p-random probing leakage from X. Denote the p-random probing leakage from X by W. For each x, we have that $P_{W|X=x}(x) = p$ and $P_{W|X=x}(\perp) = 1 - p$. Consider the randomized function g which receives $w \in \mathcal{X} \cup \{\perp\}$ as input and acts as follows:

- If $w = x$ for some $x \in \mathcal{X}$, then $g(w) = z$ with probability $\frac{P_{Z|X=x}(z) - \pi(z)}{p}$;
- If $w = \perp$, then $g(\perp) = z$ with probability $\frac{\pi(z)}{1-p}$.

Note that g is well-defined, since $\sum_z P_{g(\perp)}(z) = \sum_z \frac{\pi(z)}{1-p} = \frac{1-p}{1-p} = 1$ and $\sum_z P_{g(x)}(z) = \sum_z \frac{P_{Z|X=x}(z) - \pi(z)}{p} = \frac{1-(1-p)}{p} = 1$. We claim that $g(W)$ and Z have the same distribution conditioned on $X = x$. In fact,

$$
P_{g(W)|X=x}(z) = p \cdot \frac{P_{Z|X=x}(z) - \pi(z)}{p} + (1-p) \cdot \frac{\pi(z)}{1-p} = P_{Z|X=x}(z).
$$

This implies that $(X, Z) \equiv (X, g(W))$, and so Z is 0-simulatable by p-random probing leakage. □

We refer the reader to the full version of the paper for an alternative reduction that avoids the field size penalty at the cost of a positive statistical simulation error and for the application of our reductions to leakage-resilient circuits.

8 Empirical Evaluations

We complete the paper by investigating and discussing the practical implications of our findings. For this purpose, we start by describing how to compute the parameters t and δ of our new leakage model in Section 8.1. We then describe our evaluation settings in Section 8.2 and use them to discuss reductions to bounded leakage and random probing in Section 8.3 and Section 8.4, respectively.

8.1 Parameter Computation for Noisy Leakages

Given P_{XZ} for two random variables X and Z, we want to determine for which parameters t and δ we have that Z is (t, δ)-SD-noisy leakage from X. We prove the following result, which may be seen as a generalization of the fact that for statistical distance

$$\mathsf{SD}(P; Q) = \sup_{\mathcal{S}} |P(\mathcal{S}) - Q(\mathcal{S})|$$

the supremum is attained by the set $\mathcal{B} = \{x \mid P(x) > Q(x)\}$.

Theorem 6. *Let X and Z be any two random variables. Define the set*

$$\mathcal{B} = \{(x, z) \mid P_{XZ}(x, z) > 2^t (P_X \otimes P_Z)(x, z)\}.$$

Then, we have that Z is (t, δ)-SD-noisy leakage from X with

$$\delta = P_{XZ}(\mathcal{B}) - 2^t (P_X \otimes P_Z)(\mathcal{B}).$$

Proof. First, note that for any fixed t we may write

$$\delta = \sup_{\mathcal{S}} [P_{XZ}(\mathcal{S}) - 2^t (P_X \otimes P_Z)(\mathcal{S})], \tag{5}$$

where the supremum is taken over all subsets $\mathcal{S}$ of $\mathcal{X} \times \mathcal{Z}$. Now, for any such set $\mathcal{S}$ we have that

$$\begin{aligned}
&P_{XZ}(\mathcal{S}) - 2^t (P_X \otimes P_Z)(\mathcal{S}) \\
&= \big(P_{XZ}(\mathcal{S} \setminus \mathcal{B}) - 2^t (P_X \otimes P_Z)(\mathcal{S} \setminus \mathcal{B})\big) + \big(P_{XZ}(\mathcal{S} \cap \mathcal{B}) - 2^t (P_X \otimes P_Z)(\mathcal{S} \cap \mathcal{B})\big) \\
&\leq 0 + (P_{XZ}(\mathcal{B}) - 2^t (P_X \otimes P_Z)(\mathcal{B})).
\end{aligned}$$

To see the inequality, first note that for any $(x', z') \in \mathcal{S} \setminus \mathcal{B}$ we have that $P_{XZ}(x, z) - 2^t (P_X \otimes P_Z)(x, z) \leq 0$. Then, note also that

$$P_{XZ}(\mathcal{S} \cap \mathcal{B}) - 2^t (P_X \otimes P_Z)(\mathcal{S} \cap \mathcal{B}) = \sum_{(x,z) \in \mathcal{S} \cap \mathcal{B}} (P_{XZ}(x, z) - 2^t (P_X \otimes P_Z)(x, z))$$

and that each term in this sum is positive by construction of $\mathcal{B}$. This shows that the set $\mathcal{B}$ is the worst case scenario, and so, by Equation (5), we conclude that

$$\delta = P_{XZ}(\mathcal{B}) - 2^t (P_X \otimes P_Z)(\mathcal{B}),$$

as desired. □

The same result can be used to compute the parameters of RevSD-noisy leakages, by just swapping the roles of the product and the joint distributions.

In many scenarios the process laid out in Theorem 6 (i.e., computing the δ parameter in practice) can be further optimized. For example, if the deterministic part of Z takes on only a small amount of values we can go over all fixings of $Z = z$, compute $\delta_z = P_{XZ|Z=z}(\mathcal{B}) - 2^t (P_X \otimes P_{Z|Z=z})(\mathcal{B})$, and recombine as

$\delta = \sum_{z \in \mathcal{Z}} P_Z(z) \cdot \delta_z$. Moreover, note that Theorem 6 also provides an upper bound for the δ parameter for Z as (t, δ)-GSD-noisy leakage from X.

In certain cases we may obtain an even smaller δ value by choosing the distribution Q carefully. In the following, we nevertheless focus on the (t, δ)-SD-noisy model, which leads to simple and intuitive results for our leakage application, and we leave the study of improved parameter estimation algorithms for GSD-noisy leakage as an interesting problem for future work.

8.2 Evaluation settings

As a usual starting point, we considered the setting where leakages are written as the sum of a deterministic function d and a Gaussian noise R [37]:

$$Z = \mathsf{d}(X) + R. \tag{6}$$

In this setting, the amount of noise in the leakages is conveniently captured by the Signal-to-Noise Ratio (SNR) [28], defined as the ratio between the variance of the leakage function's deterministic part and the variance of the noise:

$$\mathrm{SNR} = \frac{\mathbb{V}(\mathsf{d}(X))}{\mathbb{V}(R)}. \tag{7}$$

As a complement to the textbook Hamming weight leakages, we considered noisy linear leakages where the deterministic function can be written as

$$\mathsf{d}(X) = \sum_{i=1}^{n} \beta_i \; X(i),$$

with $X(i)$ the i-th bit of X and the β_i's are real-valued coefficients. It generalizes the Hamming weight function where $\beta_i = 1$ for all i's. In order to evaluate the impact of leakage models that significantly deviate from the Hamming weight model, we considered two linear functions with coefficients that gradually deviate from one, and measured the distance between these models and the Hamming weight one with Pearson's correlation coefficient. The least variable model (with correlation 0.9) is illustrated and compared to the Hamming weight one in Figure 1, for $n = 8$. The more variable model (with correlation 0.5) goes significantly beyond the deviations experimentally observed in [22].

8.3 Simulating SD-Noisy Leakage via Bounded Leakage

We first computed the δ parameter (i.e., the simulation error) as a function of the SNR, for target values X of different bit sizes n and different amounts of bounded leakage t in the simulation for Hamming weight leakages.

This enables straightforward optimizations since $\mathsf{d}(X)$ can only take $n + 1$ values and has variance $n/4$ in this case. The δ parameter can therefore be easily evaluated for large (e.g., up to 128-bit) values, which we report in Figure 2.

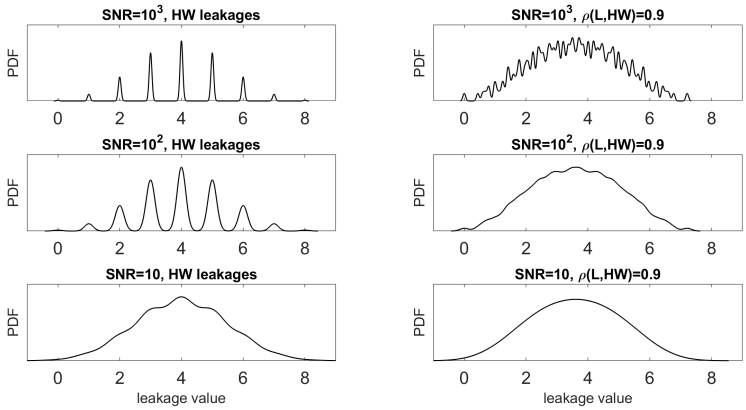

Fig. 1. Joint distribution of the noisy Hamming weight leakage function and exemplary noisy linear leakage function for different SNR values (with bit size $n = 8$).

Comparing the three first plots with the lower right one allows us to put forward the massive advantage of the (t, δ)-SD-noisy leakage model over δ-SD-noisy leakages (i.e., the $t = 0$ case). As outlined in introduction, reducing the simulation error using the techniques from [9] can only be done by reducing the SNR. But this scales badly because the MI and SD metrics of unprotected implementations decrease linearly with the noise variance and standard deviation, respectively [15]. The introduction of the t parameter circumvents this issue since as the noise increases, it allows limiting the area where the joint distribution is 2^t times larger than the product one to the extreme Hamming weights (i.e., the set $\mathcal{B}$ in Section 8.1), which only occur with exponentially small probability.

Quite naturally, a simulation using $t = \log(n)$ bits of bounded leakage is not specially impressive for (noiseless) Hamming weight leakages since a trivial simulator perfectly succeeds in this case. As a first step towards confirming the generality of our results, the figure also shows that simulation with negligible errors can also be obtained with $t = \log(n)/2$ or $t = \log(n)/3$ bits of bounded leakage, at the cost of increasing the noise (i.e., decreasing the SNR).

For example, for $n = 128$, SNR $= 10^{-3}$ and $t = \log(n)/2$, we have $\delta \approx 2^{-128}$ with $t = 3.5$ and Theorem 1 indicates that we can simulate with statistical error $2^{-128} + \alpha$ with $3.5 + \log \ln(1/\alpha)$ bits of bounded leakage from X. Comparing the right plots of Figure 2, we can see that for the same SNR, using the SD (i.e., $t = 0$) would lead to $\delta \approx 2^{-7}$, and SNRs in the 2^{-128} range would be required to reach a 2^{-128} simulation error. Plugging in these numbers in our PRNG example of Section 2.1 finally shows that our results have direct application to leakage-resilient constructions under reasonable noise requirements.

We similarly evaluated the aforementioned linear leakage models that deviate from the Hamming weight one. Those models are interesting abstractions since

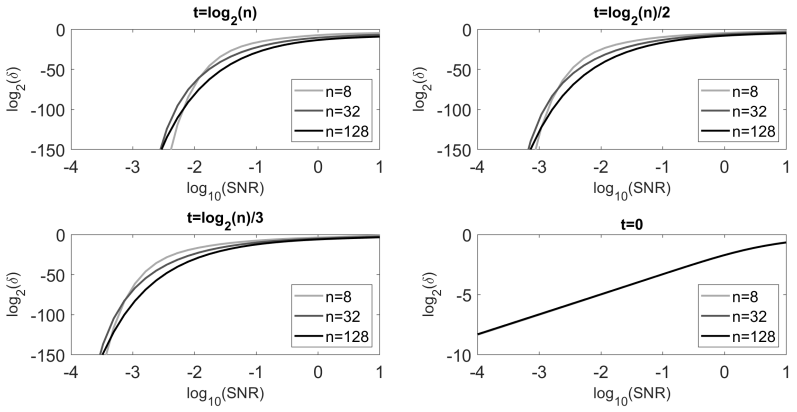

Fig. 2. Estimation of the δ parameter for SD-noisy leakages, in function of the SNR for Hamming weight leakages (with bit sizes n and an amount of bounded leakage t).

they are bijective without noise, meaning that the trivial simulation would require n bits of bounded leakage to succeed. Nevertheless, Figure 3 shows results that are very similar to Figure 2. This can be explained by looking at Figure 1 where it is clear that the amount of noise needed to "hide" the deviation of the linear model from the Hamming weight one is much lower than the amount of noise needed to simulate. For example, the lower plots of Figure 1 correspond to a SNR of 10 which is the rightmost point of the plots in Figure 3. This confirms that our simulation theorem applies to broad classes of leakage functions.[14]

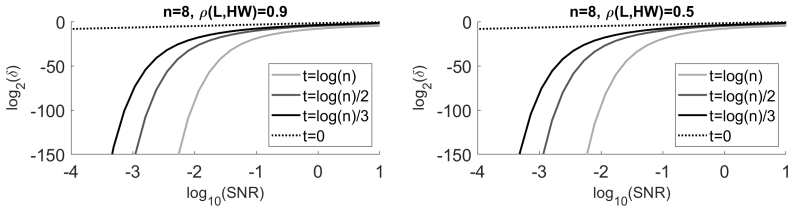

Fig. 3. Estimation of the δ parameter for SD-noisy leakages, in function of the SNR for linear leakages (with bit sizes n and an amount of bounded leakage t).

[14] This time we only computed the δ parameter for $n = 8$ because computing it (exactly) for larger n values is computationally intensive. By approximating the product distribution as a Gaussian, it is nevertheless possible to obtain efficient approximations of the δ parameter for larger n values, which should become accurate as the noise increases, and which we leave as a scope for further investigations.

Discussion. Based on the previous results, the last mile for implementers is to ensure SNRs in the 10^{-3} range. Under the (heuristic but usual) assumption that side-channel adversaries are computationally-bounded and can only exploit the signal of small (e.g., 8-bit to 32-bit) targets, a round-based hardware implementation of the AES, as can be found on off-the-shelf microcontrollers, should already be enough for this purpose [40]. Assuming (unrealistic) computationally unbounded adversaries able to characterize a full 128-bit state, one should consider more specialized architectures such as the unrolled ones in [7], where low SNRs are due to physical reasons (i.e., the weak leakage of the combinatorial logic) rather than algorithmic ones (i.e., computational limitations).

Similar observations can be made about composition. Taking the AES case study again, a round-based implementation will produce a ciphertext in 10 cycles, and each cycle will provide the adversary with a few leakage samples (typically correlated with the Hamming weight of the intermediate value). Denoting the intermediate AES results after i rounds as $X_i = \rho^i(P, K)$, with P the plaintext, K the master key and ρ the round function, we can assume for simplicity that the adversary will collect leakage samples of the form $Z_i = f(X_i)$ and that every Z_i is (t, δ)-SD-noisy. Since the X_i's are bijectively connected to K, the application of Theorem 2 implies that one would need 10 times more bounded leakage to simulate in this case (with simulation error multiplied by 10). Based on such a (worst-case) analysis, one should favor (low-latency) unrolled implementations to ensure high security levels. But this theorem again assumes that the leakage of all computations in an implementation are equally easy to exploit, which is not true for computationally-bounded adversaries [21]. So a reasonable rule-of-thumb to obtain less conservative results would be to apply composition results with only a fraction of the AES rounds, in which case round-based implementations should already lead to high security levels at lower implementation cost.

Note that the practical estimations in this section leverage two additional assumptions. First, the estimation of t and δ assume a uniformly distributed X. This is a natural assumption in side-channel analysis since the adversary has in general no efficient ways to force intermediate computations to values of her choice (e.g., extreme Hamming weights). This is even enforced in leakage-resilient constructions where the block cipher inputs are fixed by design [19, 6, 10]. But, of course, our theoretical results are applicable to non-uniform distributions as well. Besides, we recall that our composition theorem assumes the noise part of the leakage samples Z_i to be independent, which is a standard approximation.

So, overall, we can conclude that the requirements that our simulation and composition theorems impose are reachable for actual hardware implementations using known techniques and at non-negligible but affordable cost. Besides, and most importantly, they formally confirm that it is possible to simulate noisy leakages from bounded leakage with exponentially small error without masking (as witnessed by Figures 2 and 3), which in turn formally confirms the interest of the re-keying techniques used in leakage-resilient cryptography.

8.4 Simulating RevSD-Noisy Leakage via Random Probing

As a final investigation, Figure 4 reports the t and δ parameters corresponding to RevSD-noisy leakage, in a setting similar to Figure 2. The upper left plot is for $t = \log(n)/2$ and it is used to confirm that the trends for this model are similar to the ones of SD-noisy leakages (essentially for the same reason that increasing the t parameter leads to computing δ by integrating over low-probability areas, where the product distribution is 2^t times larger than the joint distribution).

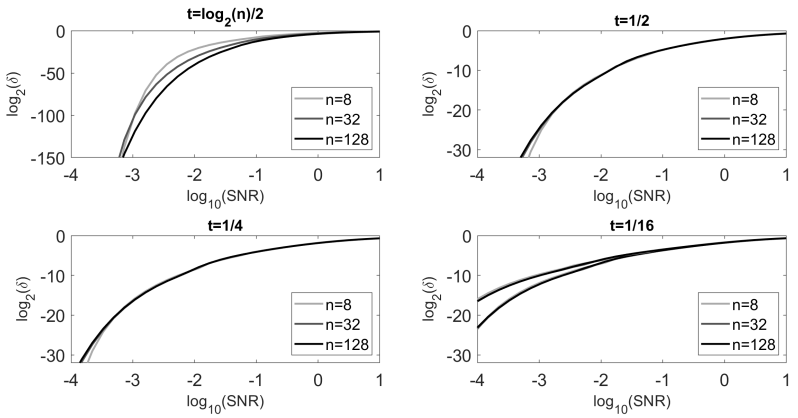

Fig. 4. Estimation of the δ parameter for RevSD-noisy leakages, in function of the SNR for Hamming weight leakages (with bit sizes n and an amount of bounded leakage t).

Concretely, though, the relevant t values are lower than in the simulation via bounded leakage. This is because the p parameter of the random probes in Theorem 3 is at least $(1 - 2^{-t})$. Hence, Figure 4 provides values for $t = 0.5$ (which corresponds to $p > 0.3$), $t = 0.25$ (which corresponds to $p > 0.15$) and $t = 0.125$ (which corresponds to $p > 0.08$). Assuming $n = |\mathcal{X}| = 256$ (as when masking the AES S-box) and a SNR of 10^{-3}, we see that even for $t = 0.125$ we have $\delta \approx 2^{-13}$, which is significantly below the field size and therefore amortizes the penalty term $\delta \cdot 2^{-t} \cdot |\mathcal{X}| \approx 0.02$, only impacting the security level mildly. Assuming $|\mathcal{X}| = 2$ as in a bitslice cipher, this penalty term falls down to $2 \cdot 10^{-4}$.

As mentioned in introduction, Prest *et al.* already proposed a noisy leakage model that is tightly connected to the random probing model, using the Average Relative Error (ARE) metric [34]. They provide an approximate closed-form formula for this metric in the context of Hamming weight leakages with Gaussian noise (that becomes accurate for large noise levels / low SNRs):

$$\mathrm{ARE}(X|Z) = \frac{n}{\sigma\sqrt{2\pi}},$$

where σ is the leakage noise's standard deviation. Since the SNR of the Hamming weight leakage function equals $\frac{n/4}{\sigma^2}$, we can directly compare the two approaches in this case. For this purpose, we plot in Figure 5 the random probing probability p in function of the SNR using the ARE and our reduction, for different values of the t parameter. It leads to the following main observations:

- By adapting the t parameter to the SNR, the $1 - 2^{-t}$ term (reflected by the plateau's on the left parts of the plots) is not dominating.
- The loss compared to the ARE increases with the field size, but is smaller than the field size (e.g., for $n = 8$, we lose a factor ≈ 2 rather than 2^8).

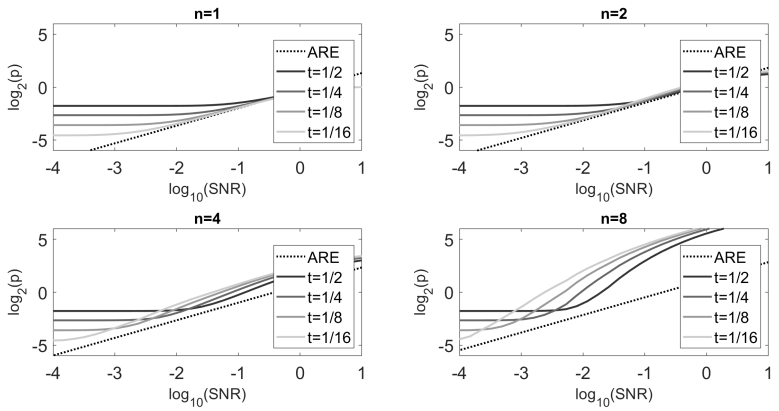

Fig. 5. Reductions to random probing using the ARE and RevSD metrics in function of the SNR for Hamming weight leakages and bit sizes $n = 1, 2, 4$ and 8.

So despite not improving the state of the art for such a realistic leakage function (as in the case of bounded leakage), our reduction gets reasonably close while improving the seminal one of Duc, Dziembowski and Faust with new techniques, confirming the unifying nature of hockey-stick divergences for cryptography in the presence of leakage. Besides, it is worth recalling that the ARE is a worst-case metric whereas the (G)SD and Rev(G)SD metrics are average-case metrics. So the results of Prest et al. and our results conceptually differ in the sense that the former deal with the field size loss in the metric whereas the latter deal with it in the reduction to the random probing model. Therefore, both types of models shed different light on the same issue.

We finally mention two recent works that tackled the tightness of the reduction from the noisy leakage model to the random probing model. First, in [8], Brian et al. show how to get rid of the field size loss at the cost of a quadratic loss on the noise parameter, leveraging the average random probing model of [18].

Second, in [11], Béguinot et al. study a variant of the ARE metric (coined Doeblin coefficients) that is better connected to the attacks' success.[15] They additionally show that a loss when moving from the (average-case) noisy leakage model to the (worst-case) random probing model is in general unavoidable.

Acknowledgements. MO was supported by the MOE tier 2 grant T2EP20223-0024, Breaking the Box – On Security of Cryptographic Devices. JR was supported in part by NOVA LINCS (ref. UIDB/04516/2020) with the financial support of FCT - Fundação para a Ciência e a Tecnologia. FX is a senior research associate of the Belgian fund for Scientific Research (FNRS-F.R.S). DV was supported by project SERICS (PE00000014) and by project PARTHENON (B53D23013000006), under the MUR National Recovery and Resilience Plan funded by the European Union—NextGenerationEU. LR was partially supported by the Danish Independent Research Council (grant DFF-0165-00107B "C3PO") and the DARPA SIEVE program (contract HR001120C0085 "FROMAGER"). Work funded in part by the ERC Advanced Grant 101096871 (BRIDGE). Views and opinions are those of the authors only and do not necessarily reflect those of the European Union or the ERC. Neither the European Union nor the granting authority can be held responsible for them. Opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of DARPA. Distribution Statement "A" (Approved for Public Release, Distribution Unlimited). Part of this work was done while JR was visiting the Simons Institute for the Theory of Computing.

References

1. Agrawal, D., Archambeault, B., Rao, J.R., Rohatgi, P.: The EM side-channel(s). In: Kaliski, B.S., Koç, ç.K., Paar, C. (eds.) CHES 2002. pp. 29–45. Springer Berlin Heidelberg, Berlin, Heidelberg (2003)
2. Balasch, J., Gierlichs, B., Reparaz, O., Verbauwhede, I.: DPA, bitslicing and masking at 1 GHz. In: CHES. LNCS, vol. 9293, pp. 599–619. Springer (2015)
3. Barthe, G., Olmedo, F.: Beyond differential privacy: Composition theorems and relational logic for f-divergences between probabilistic programs. In: Fomin, F.V., Freivalds, R., Kwiatkowska, M., Peleg, D. (eds.) Automata, Languages, and Programming. pp. 49–60. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
4. Béguinot, J., Cheng, W., Guilley, S., Liu, Y., Masure, L., Rioul, O., Standaert, F.: Removing the field size loss from Duc et al.'s conjectured bound for masked encodings. In: COSADE. LNCS, vol. 13979, pp. 86–104. Springer (2023)
5. Belaïd, S., Santis, F.D., Heyszl, J., Mangard, S., Medwed, M., Schmidt, J., Standaert, F., Tillich, S.: Towards fresh re-keying with leakage-resilient PRFs: cipher design principles and analysis. J. Cryptogr. Eng. **4**(3), 157–171 (2014)
6. Bellizia, D., Bronchain, O., Cassiers, G., Grosso, V., Guo, C., Momin, C., Pereira, O., Peters, T., Standaert, F.: Mode-level vs. implementation-level physical security in symmetric cryptography - A practical guide through the leakage-resistance jungle. In: CRYPTO (1). LNCS, vol. 12170, pp. 369–400. Springer (2020)

[15] Doeblin coefficients actually appear in our proof of Lemma 4 as $\sum_z \pi(z)$.

7. Bhasin, S., Guilley, S., Sauvage, L., Danger, J.: Unrolling cryptographic circuits: A simple countermeasure against side-channel attacks. In: CT-RSA. Lecture Notes in Computer Science, vol. 5985, pp. 195–207. Springer (2010)
8. Brian, G., Dziembowski, S., Faust, S.: From random probing to noisy leakages without field-size dependence. In: Joye, M., Leander, G. (eds.) Advances in Cryptology – EUROCRYPT 2024. pp. 345–374. Springer Nature Switzerland, Cham (2024)
9. Brian, G., Faonio, A., Obremski, M., Ribeiro, J., Simkin, M., Skórski, M., Venturi, D.: The mother of all leakages: How to simulate noisy leakages via bounded leakage (almost) for free. In: Canteaut, A., Standaert, F.X. (eds.) EUROCRYPT 2021. pp. 408–437 (2021)
10. Bronchain, O., Momin, C., Peters, T., Standaert, F.: Improved leakage-resistant authenticated encryption based on hardware AES coprocessors. IACR Trans. Cryptogr. Hardw. Embed. Syst. **2021**(3), 641–676 (2021)
11. Béguinot, J., Cheng, W., Guilley, S., Rioul, O.: Formal security proofs via Doeblin coefficients: Optimal side-channel factorization from noisy leakage to random probing. Cryptology ePrint Archive, Paper 2024/199 (2024), https://eprint.iacr.org/2024/199
12. Chari, S., Jutla, C.S., Rao, J.R., Rohatgi, P.: Towards sound approaches to counteract power-analysis attacks. In: Wiener, M. (ed.) Advances in Cryptology — CRYPTO' 99. pp. 398–412. Springer Berlin Heidelberg, Berlin, Heidelberg (1999)
13. Dobraunig, C., Eichlseder, M., Mangard, S., Mendel, F., Mennink, B., Primas, R., Unterluggauer, T.: Isap v2.0. IACR Trans. Symmetric Cryptol. **2020**(S1), 390–416 (2020)
14. Duc, A., Dziembowski, S., Faust, S.: Unifying leakage models: From probing attacks to noisy leakage. J. Cryptol. **32**(1), 151–177 (2019)
15. Duc, A., Faust, S., Standaert, F.X.: Making masking security proofs concrete. In: Oswald, E., Fischlin, M. (eds.) Advances in Cryptology – EUROCRYPT 2015. pp. 401–429. Springer Berlin Heidelberg, Berlin, Heidelberg (2015)
16. Dwork, C., Lei, J.: Differential privacy and robust statistics. In: Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing. p. 371–380. STOC '09, Association for Computing Machinery, New York, NY, USA (2009)
17. Dwork, C., McSherry, F., Nissim, K., Smith, A.D.: Calibrating noise to sensitivity in private data analysis. In: Halevi, S., Rabin, T. (eds.) TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings. LNCS, vol. 3876, pp. 265–284. Springer (2006)
18. Dziembowski, S., Faust, S., Skorski, M.: Noisy leakage revisited. In: Oswald, E., Fischlin, M. (eds.) Advances in Cryptology - EUROCRYPT 2015. pp. 159–188. Springer Berlin Heidelberg, Berlin, Heidelberg (2015)
19. Dziembowski, S., Pietrzak, K.: Leakage-resilient cryptography. In: 2008 49th Annual IEEE Symposium on Foundations of Computer Science. pp. 293–302 (2008)
20. Faust, S., Pietrzak, K., Schipper, J.: Practical leakage-resilient symmetric cryptography. In: CHES. Lecture Notes in Computer Science, vol. 7428, pp. 213–232. Springer (2012)
21. Guo, Q., Grosso, V., Standaert, F., Bronchain, O.: Modeling soft analytical side-channel attacks from a coding theory viewpoint. IACR Trans. Cryptogr. Hardw. Embed. Syst. **2020**(4), 209–238 (2020)
22. Hoffmann, C., Méaux, P., Momin, C., Rotella, Y., Standaert, F., Udvarhelyi, B.: Learning with physical rounding for linear and quadratic leakage functions. In: CRYPTO (3). Lecture Notes in Computer Science, vol. 14083, pp. 410–439. Springer (2023)

23. Ishai, Y., Sahai, A., Wagner, D.: Private circuits: Securing hardware against probing attacks. In: Boneh, D. (ed.) Advances in Cryptology - CRYPTO 2003. pp. 463–481. Springer Berlin Heidelberg, Berlin, Heidelberg (2003)
24. Kalai, Y.T., Reyzin, L.: A survey of leakage-resilient cryptography. In: Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali, pp. 727–794. ACM (2019)
25. Kocher, P., Jaffe, J., Jun, B.: Differential power analysis. In: Wiener, M. (ed.) Advances in Cryptology — CRYPTO' 99. pp. 388–397. Springer Berlin Heidelberg, Berlin, Heidelberg (1999)
26. Kocher, P.C.: Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In: Koblitz, N. (ed.) CRYPTO '96. pp. 104–113. Springer Berlin Heidelberg, Berlin, Heidelberg (1996)
27. Liu, C., Chakraborty, A., Chawla, N., Roggel, N.: Frequency throttling side-channel attack. In: CCS. pp. 1977–1991. ACM (2022)
28. Mangard, S.: Hardware countermeasures against DPA ? A statistical analysis of their effectiveness. In: CT-RSA. Lecture Notes in Computer Science, vol. 2964, pp. 222–235. Springer (2004)
29. Martínez-Rodríguez, M.C., Delgado-Lozano, I.M., Brumley, B.B.: SoK: Remote power analysis. In: ARES. pp. 7:1–7:12. ACM (2021)
30. Moradi, A., Barenghi, A., Kasper, T., Paar, C.: On the vulnerability of FPGA bitstream encryption against power analysis attacks: extracting keys from xilinx Virtex-II FPGAs. In: CCS. pp. 111–124. ACM (2011)
31. Naor, M., Segev, G.: Public-key cryptosystems resilient to key leakage. SIAM Journal on Computing **41**(4), 772–814 (2012)
32. Osvik, D.A., Shamir, A., Tromer, E.: Cache attacks and countermeasures: The case of AES. In: Pointcheval, D. (ed.) CT-RSA 2006. pp. 1–20. Springer Berlin Heidelberg, Berlin, Heidelberg (2006)
33. Pietrzak, K.: A leakage-resilient mode of operation. In: EUROCRYPT. Lecture Notes in Computer Science, vol. 5479, pp. 462–482. Springer (2009)
34. Prest, T., Goudarzi, D., Martinelli, A., Passelègue, A.: Unifying leakage models on a Rényi day. In: Boldyreva, A., Micciancio, D. (eds.) Advances in Cryptology – CRYPTO 2019. pp. 683–712. Springer International Publishing, Cham (2019)
35. Prouff, E., Rivain, M.: Masking against side-channel attacks: A formal security proof. In: Johansson, T., Nguyen, P.Q. (eds.) Advances in Cryptology – EUROCRYPT 2013. pp. 142–159. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
36. Sason, I., Verdú, S.: f-divergence inequalities. IEEE Transactions on Information Theory **62**(11), 5973–6006 (2016)
37. Schindler, W., Lemke, K., Paar, C.: A stochastic model for differential side channel cryptanalysis. In: CHES. Lecture Notes in Computer Science, vol. 3659, pp. 30–46. Springer (2005)
38. Standaert, F.X., Malkin, T.G., Yung, M.: A unified framework for the analysis of side-channel key recovery attacks. In: Joux, A. (ed.) Advances in Cryptology - EUROCRYPT 2009. pp. 443–461. Springer Berlin Heidelberg, Berlin, Heidelberg (2009)
39. Standaert, F.X., Pereira, O., Yu, Y.: Leakage-resilient symmetric cryptography under empirically verifiable assumptions. In: Canetti, R., Garay, J.A. (eds.) Advances in Cryptology – CRYPTO 2013. pp. 335–352. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
40. Udvarhelyi, B., van Wassenhove, A., Bronchain, O., Standaert, F.: On the security of off-the-shelf microcontrollers: Hardware is not enough. In: CARDIS. Lecture Notes in Computer Science, vol. 12609, pp. 103–118. Springer (2020)