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

Maciej Obremski, João Ribeiro, Lawrence Roy, François-Xavier Standaert, and Daniele Venturi

***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 & J. Cryptology 2019; Brian et al., Eurocrypt 2021 & IEEE Trans. Inf. Theory 2022) 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. Then, 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 also show that our results generalize and improve on the main results of Brian et al. 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 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).

A preliminary version of this work was presented at CRYPTO 2024 [1].

M. Obremski was supported by the MOE tier 2 grant T2EP20223-0024, Breaking the Box – On Security of Cryptographic Devices. J. Ribeiro was supported in part by NOVA LINCS (ref. UIDB/04516/2020) and by national funds through FCT – Fundação para a Ciência e a Tecnologia, I.P., and, when eligible, co-funded by EU funds under project/support UID/50008/2025 – Instituto de Telecomunicações, with DOI 10.54499/UID/50008/2025. L. Roy was partially supported by the Danish Independent Research Council (grant DFF-0165-00107B "C3PO") and the DARPA SIEVE program (contract HR001120C0085 "FROMAGER"). F.-X. Standaert is a senior research associate of the Belgian fund for Scientific Research (FNRS-F.R.S). D. Venturi 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. This work was funded in part by the ERC Advanced Grant 101096871 (BRIDGE).

Maciej Obremski is with the School of Computing, National University of Singapore, Singapore 117543, Singapore. (email: obremski.math@gmail.com)

João Ribeiro was with NOVA LINCS and the Computer Science Department, NOVA School of Science and Technology, 2829-516 Caparica, Portugal. He is now with the Instituto de Telecomunicações, 1049-001 Lisboa, Portugal, and the Departamento de Matemática, Instituto Superior Técnico, Universidade de Lisboa, 1049-001 Lisboa, Portugal (email: jribeiro@tecnico.ulisboa.pt).

Lawrence Roy was with the Computer Science Department, Aarhus University, 8200 Aarhus, Denmark. He is now with IBM Research Europe, 8803 Rüschlikon, Switzerland (email: ldr709@gmail.com).

François-Xavier Standaert is with the Université catholique de Louvain, B-1348 Louvain-la-Neuve, Belgium (email: fstandae@uclouvain.be).

Daniele Venturi is with the Department of Computer Science, Sapienza University of Rome, 00185 Rome, Italy (email: venturi@di.uniroma1.it).

I. 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 [2]. 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 [3], the emission of electromagnetic radiation [4], and cache-based attacks [5]. Small embedded devices are natural targets, but side-channel attacks have been extended to hardware implementations [6] and high-frequency devices [7]. They can also be applied remotely [8], and new attacks keep on being discovered [9]. 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 [10], 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) [11] is usually considered as the most viable option.[1] 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 [15].

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 [16] 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 [17]. 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 [18], [19]. More precisely, throughout this paper we write $\log$ for the base-2 logarithm and $\ln$ for the natural logarithm, and define the Kullback-Leibler divergence between two distributions P and Q supported on a common finite set $\mathcal{X}$ as $D_{\mathsf{KL}}(P\|Q) = \sum_{x\in\mathcal{X}} P(x)\log\left(\frac{P(x)}{Q(x)}\right)$. If X denotes the secret and Z is leakage from X, then the mutual information between X and Z, defined as $I(X;Z) = D_{\mathsf{KL}}(P_{XZ}\|P_X \otimes P_Z)$ where P_{XZ} is the joint distribution of X and Z and $P_X \otimes P_Z$ is their product distribution (i.e., $(P_X \otimes P_Z)(x,z) = P_X(x)\cdot P_Z(z)$), 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 $I(X;Z) \le \delta$ for δ as large as possible.

Another closely related noise measure is the statistical distance [20] (a.k.a. the total variation distance) between the joint distribution P_{XZ} and the product distribution $P_X \otimes P_Z$, denoted as $\mathsf{SD}(P_{XZ}; P_X \otimes P_Z)$. For general distributions P and Q supported on a finite set $\mathcal{X}$, we define $\mathsf{SD}(P;Q) = \frac{1}{2}\sum_{x\in\mathcal{X}} |P(x) - Q(x)|$. The two measures are related via Pinsker's inequality, which implies that

$$\mathsf{SD}(P_{XZ}; P_X \otimes P_Z) \le \sqrt{\frac{\ln 2}{2}\cdot I(X;Z)}. \tag{1}$$

This means that a scheme which is leakage-resilient against all leakages Z such that $\mathsf{SD}(P_{XZ}; P_X \otimes P_Z) \le \delta$ is resilient against all leakages Z such that $I(X;Z) \le \frac{2\delta^2}{\ln 2}$. Other noise measures have been considered, including the average conditional min-entropy [21] and the average ℓ_2-norm between the marginal distribution P_X and the conditional distributions $P_{X|Z=z}$ [22].[2]

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 [20] (in which the adversary can recover intermediate values in the circuit only with some probability). But despite the security of masked implementations is conveniently analyzed in these models, actual implementations are again better reflected by the noisy leakage model [22], which instead bounds the noisiness of the information retrieved from intermediate values based on the statistical distance and the mutual information metrics.

A. 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 [24], 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

$$\mathsf{SD}(P_{XZ}; P_{XZ'}) \le \varepsilon,$$

for a small "simulation 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

[1] There are, however, primitive-level alternatives to this initialization problem, such as using a leakage-resilient PRF for this part of the computation [12], [13], [14].

[2] 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 $P_X \otimes P_Z$.

bits of bounded leakage are also resilient to ℓ'-*min-entropy noisy leakage* [21] (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 small ε (as a function of the security parameter).[3]

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 [20], which has then been (in part heuristically) connected to practical side-channel attacks [25].

B. Limitations of Statistical Distance and Mutual Information

Although [24] 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

$$\mathsf{SD}(P_{XZ}; P_X \otimes P_Z) \leq \delta. \tag{2}$$

First, note that it is trivial to simulate Z with error δ *even without access to bounded leakage from* X. In fact, by Equation (2), the simulator can simply output Z' sampled independently according to the marginal P_Z. To complement this, [24] 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 $Z = \perp$.[4] 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 $I(X;Z) \leq \delta$), see [24, Theorem 15].

From a practical perspective, the above is unsatisfactory because without countermeasures δ does not decrease fast enough with noise (e.g., see [25, 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 [20], 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 [19], [26].

C. 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* [27, Section VII], 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*[5]

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

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

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

$$P(\mathcal{S}) \leq 2^t \cdot Q(\mathcal{S}) + \delta \tag{3}$$

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[6] [29], something which we exploit in our results.

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

[4] This corresponds to the random probing model of [23], [20] in a large (n-bit) field.

[5] 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.

[6] A randomized mechanism $\mathcal{M}$ is (ε, δ)-approximately differentially private if and only if $\mathsf{SD}_{\varepsilon \log e}(P_{\mathcal{M}(x)}; P_{\mathcal{M}(x')}) \leq \delta$ for all pairs of databases (x, x') that differ in only one coordinate [28].

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 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. Crucially, proving (ii) relies on showing that (i) holds even for a more general leakage model than (t,δ)-SD-noisy leakage, which we discuss later.

We also argue how the (t,δ)-SD-noisy leakage model can be interpreted as an "average-case on X" (and, we believe, conceptually more natural) version of the dense leakage model of [24]. In particular, our improved analysis behind (i) leads to a more general simulation theorem compared to the main theorem of [24]. In turn, this implies improved simulation theorems for noisy leakage models that can be captured with better parameters as special cases of (t,δ)-SD-noisy leakage compared to the dense leakage model.

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 [20] (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 [24], 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 [20], 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. [19] 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). It nevertheless illustrates the unifying nature of hockey-stick divergences for cryptography in the presence of leakage.

II. 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 IV. 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.

A. 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 [24], 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 $\alpha+\delta$.*

For formal statements and proofs, see Section V. In that section, we show that this theorem holds for an even more general leakage model.

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 IX 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 [30]. 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 $t = \log(n) - \log\ln(1/\alpha)$.

B. Composition Theorems

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 VI-A. As mentioned before, establishing this result requires us to show that Theorem 1 actually holds for the more general model of GSD-noisy leakage that we introduce in Section IV.

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 [31], 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.

a) Advanced composition: Given the relationship between hockey-stick divergences and differential privacy, it is natural to wonder whether a composition theorem akin to *advanced composition* in differential privacy [32], which features improved scaling with the number of leakages, holds in some parameter regime. We prove such an advanced composition theorem for a natural symmetric strengthening of the (t, δ)-SD-noisy leakage model. This result has limitations analogous to advanced composition in differential privacy (it is only relevant when t is small), and so is less practically relevant than Theorem 2. Therefore, we see it as a consolidation of our theoretical understanding of the (t, δ)-SD-noisy leakage model. We discuss advanced composition in Section VI-B.

C. (t, δ)-SD-Noisy Leakage and Dense Leakage

It is interesting to compare Theorem 1 with the simulation result obtained alternatively by determining the parameters of (t, δ)-SD-noisy leakage with respect to the general *dense leakage* model of [24], and then applying their main simulation theorem for dense leakage. As we discuss in more detail in Section V, this "indirect" approach leads to a worse simulation theorem, which is due both to how dense leakage is defined in [24] (it is, in a sense, a "worst-case" leakage model) and to their sub-optimal analysis of rejection sampling simulators (which are also the basis of Theorem 1).

Motivated both by this and by the improved analysis behind Theorem 1, we explore the relationship between (t, δ)-SD-noisy leakage and dense leakage further in Section VII-A. We show that (t, δ)-SD-noisy leakage captures an "average-case on X" version of the dense leakage model of [24] (their main unifying leakage model) as a special case. To complement this, we also show that (t, δ)-SD-noisy leakage is captured by this average-case version of dense leakage up to a small constant loss in parameters. Therefore, the (t, δ)-SD-noisy leakage model is essentially equivalent to an average-case version of the dense leakage model.

As we discuss more carefully in Section VII-A, this relationship shows, in a precise sense, that our Theorem 1 is a more general simulation theorem than the main simulation theorem for dense leakage of [24]. Since the (t, δ)-SD-noisy leakage model captures other noisy leakage models with better parameters than dense leakage (and based on slightly cleaner proofs), our simulation theorem (t, δ)-SD-noisy leakage leads to improved simulation theorems for these models. We exemplify this through a detailed discussion of the "Uniform-Noisy" leakage model [33], [24]. Furthermore, through the relationship above, we see that Theorem 2 also

implies nice composition guarantees for the dense and average dense leakage models.

a) (t,δ)-SD-Noisy Leakage and Mutual Information: As further consolidation, in Section VII-B we discuss the relationship between (t,δ)-SD-noisy leakage and noisy leakages based on mutual information, which is a popular metric in practice. Pinsker's inequality (Equation (1)) implies that all "δ-MI-noisy" leakages Z from X (i.e., leakages satisfying $I(X;Z) \leq \delta$) are $(t=0,\delta'=\sqrt{\delta/2})$-SD-noisy leakages from X. Interestingly, the extension of Pinsker's inequality to hockey-stick divergences SD_t with $t>0$ [27, Theorem 33] provides a better relationship with mutual information and implies that, for any $t>0$, all δ-MI-noisy leakages from X are also $(t,\delta'=\delta/t)$-SD-noisy leakages from X. Coupled with Theorem 1, we get a simulation theorem for δ-MI-noisy leakages from bounded leakage where the simulation error decays linearly with the amount of bounded leakage. The fact that the simulation error decays only linearly is not surprising given that the δ-MI-noisy leakage model, like the $(t=0,\delta)$-SD-noisy leakage model, includes "catastrophic" leakage functions, as showcased in the negative result of Brian et al. [24, Theorem 15] (whose simulation error lower bound is consistent with the simulation error achieved by our simulation theorem).

D. 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 [22] and Duc, Dziembowski, Faust, and Standaert [20], [25] on compilers for leakage-resilient arithmetic circuits. Most relevant to our setting, Duc, Dziembowski, and Faust [20] 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.[7] The key lemma behind the main result of [20] (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 $p=\delta|\mathcal{X}|$.[8] 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 [25], [19], [26]. We extend the key lemma of [20] 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.

[7] 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.

[8] 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 $\Pr[Z=\perp]=1-p$.

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 [20, 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 VIII.

This result generally improves on [20, 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 IX nevertheless confirm that Theorem 3 can lead to sweet spots for practically-relevant leakage functions and noise levels.

We additionally present a reduction that trades the aforementioned field size penalty for positive statistical simulation error in Section VIII-B. We outline how to apply these reductions in order to obtain leakage-resilient circuit compilers tolerating RevSD-noisy leakage from circuit wires in Section VIII-C.

E. 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 [30] and its many follow-ups. Say, for example, that we want to ensure 128-bit security using the reduction of [24]. Ensuring 2^{-128} simulation error would require a noise variance in the $2^{128}\approx 10^{39}$ range, which no device offers intrinsically.[9] Even tolerating lower (e.g., 64-bit) security

[9] The noise requirements of a masked implementation are more accurately expressed in terms of a side-channel Signal-to-Noise Ratio (SNR) [34], which we defer to Section IX to keep this overview of our contributions concise.

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 IX 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 [24] 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 IX 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.

III. Preliminaries

A. Notation

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. 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 expected value by $\mathbb{E}[X]$, and its variance by $\mathbb{V}(X)$. We write $x \sim P$ to mean that x is sampled according to the distribution P. 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 $P_{XZ} = P_X \otimes P_Z$. For a finite set $\mathcal{S}$ and a distribution P, we write $P(\mathcal{S}) = \sum_{x\in\mathcal{S}} P(x)$.

B. 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 [24].

Definition 4 (Leakage simulation [24])**.** *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*

$$\mathsf{SD}(P_{(X,Z)}; P_{(X,\mathsf{Sim}_f^{\mathsf{Leak}(X,\cdot)})}) \le \varepsilon, \tag{4}$$

where $Z = f(X)$ and the oracle $\mathsf{Leak}(X,\cdot)$ accepts a single query $g \in \mathcal{G}(X)$ and outputs $g(X)$, and $\mathsf{Sim}_f^{\mathsf{Leak}(X,\cdot)}$ denotes the output of the simulator with access to this oracle.

Furthermore, when $\mathcal{G}(X)$ is the family of all ℓ-bounded leakage functions $g : \mathcal{X} \to \{0,1\}^\ell$ and Equation (4) holds, we say that $\mathcal{F}(X)$ is ε-simulatable from ℓ bits of bounded leakage.

C. A Basic Property of Hockey-Stick Divergences

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

Lemma 1. *Let P and Q be any two distributions supported on $\mathcal{X}$. Define the set*

$$\mathcal{B} = \{x : P(x) > 2^t Q(x)\}.$$

Then,

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

Proof. First, note that for any fixed $t \ge 0$ we may write

$$\delta = \sup_{\mathcal{S}\subseteq\mathcal{X}} [P(\mathcal{S}) - 2^t Q(\mathcal{S})]. \tag{5}$$

Now, for any such set $\mathcal{S}$ we have that

$$\begin{aligned} &P(\mathcal{S}) - 2^t Q(\mathcal{S}) \\ &= \big(P(\mathcal{S}\setminus\mathcal{B}) - 2^t Q(\mathcal{S}\setminus\mathcal{B})\big) + \big(P(\mathcal{S}\cap\mathcal{B}) - 2^t Q(\mathcal{S}\cap\mathcal{B})\big) \\ &\le 0 + (P(\mathcal{B}) - 2^t Q(\mathcal{B})) \\ &= P(\mathcal{B}) - 2^t Q(\mathcal{B}). \end{aligned}$$

To see the inequality, first note that for any $x \in \mathcal{S}\setminus\mathcal{B}$ we have that $P(x) - 2^t Q(x) \le 0$, and so $P(\mathcal{S}\setminus\mathcal{B}) - 2^t Q(\mathcal{S}\setminus\mathcal{B}) \le 0$. Furthermore,

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

since each term $P(x) - 2^t Q(x)$ for $x \in \mathcal{B}$ 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

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

Finally, the rightmost equality in the lemma statement follows immediately by noting that $\max(0, P(x) - 2^t Q(x))$ is always non-negative, and is positive if and only if $x \in \mathcal{B}$. □

IV. The Generalized SD-Noisy Leakage Model

In this section we first recall the definitions of (t,δ)-SD-Noisy and (t,δ)-RevSD-Noisy leakage that we already discussed in Section II, and then introduce the more general (t,δ)-GSD-Noisy leakage model.

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.*

Intuitively, in the generalized definition below 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 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.*

We remark that other natural information measures also have such variational characterizations, i.e., can be characterized as the minimization over Q of an appropriate divergence. For example, the mutual information $I(X;Z)$ can be written as $I(X;Z) = \min_Q D_{\mathsf{KL}}(P_{XZ} \| P_X \otimes Q)$ (and in that case the minimizer is $Q = P_Z$) [35, Corollary 4.2].

In the next sections we establish useful properties of these leakage models. In Section V, 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. We explore connections to an average notion of the dense leakage model from [24] in Section VII-A. Then, in Section VI we prove composition theorems for these models, yielding Theorem 2 in particular. We explore advanced composition for a symmetric version of GSD-noisy leakage in Section VI-B. The relationship between RevSD-noisy leakage and the random probing model is studied in Section VIII. Empirical evaluations of these different leakage models are discussed in Section IX.

V. 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 $\alpha > 0$. More precisely, we prove the following, which 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

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

The simulator we use to establish Theorem 4 is based on rejection sampling. It is a close variant of the simulator used in [24] 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** $2^\ell - 2$ **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 $2^\ell - 1$
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 [24])**.** We outline the main differences with respect to the simulator from [24]. 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 [24] 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 $P_{Z|X=x}(\cdot)$ for any x 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 IX.

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

Lemma 2. *Let* $R(x) = 1 - \mathbb{E}_{z\sim Q}\left[\min\left(2^{-t}\cdot\frac{P_{XZ}(x,z)}{(P_X\otimes Q)(x,z)},1\right)\right]$ *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,*

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

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 $Q(z) \geq \min\left(2^{-t}P_{Z|X=x}(z), Q(z)\right)$, and 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

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

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

$$\begin{aligned}&1-\mathbb{E}_X[R(X)] \qquad (6)\\ &= \sum_{x\in\mathcal{X}}P_X(x)\sum_{z\in\mathcal{Z}}\min\left(2^{-t}P_{Z|X=x}(z),Q(z)\right)\\ &= 2^{-t}\sum_{x\in\mathcal{X},z\in\mathcal{Z}}\min\left(P_{XZ}(x,z), 2^t(P_X\otimes Q)(x,z)\right)\\ &= 2^{-t}\sum_{x\in\mathcal{X},z\in\mathcal{Z}}(P_{XZ}(x,z)\\ &\qquad - \max(0, P_{XZ}(x,z) - 2^t(P_X\otimes Q)(x,z)))\\ &= 2^{-t}\\ &\quad - 2^{-t}\sum_{x\in\mathcal{X},z\in\mathcal{Z}}\max\left(0, P_{XZ}(x,z) - 2^t(P_X\otimes Q)(x,z)\right)\\ &\geq 2^{-t}(1-\delta), \qquad (7)\end{aligned}$$

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. □

We are now ready to prove Theorem 4.

Proof of Theorem 4. 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 $P_{\mathsf{Sim}|X=x}$ and $P_{Z|X=x}$ using Lemma 2. Then, we use Lemma 3 to obtain the desired bound on the original statistical distance. By Lemma 2, we have

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

Notice that $P_{Z|X=x}(z) \ge 2^t(2^{-t}P_{Z|X=x}(z))$, so the second argument inside the max is always non-negative. Therefore,

$$\begin{aligned}
&D(x) && (8)\\
&\le \sum_{z\in\mathcal{Z}} P_{Z|X=x}(z) && (9)\\
&\quad - \sum_{z\in\mathcal{Z}} \min\left(\frac{1-R(x)^{2^\ell}}{1-R(x)}, 2^t\right) \min\big(2^{-t}P_{Z|X=x}(z), Q(z)\big) \\
&= 1 && (10)\\
&\quad - \min\left(\frac{1-R(x)^{2^\ell}}{1-R(x)}, 2^t\right) \sum_{z\in\mathcal{Z}} \min\big(2^{-t}P_{Z|X=x}(z), Q(z)\big) \\
&= 1 - \min\left(\frac{1-R(x)^{2^\ell}}{1-R(x)}, 2^t\right)\big(1-R(x)\big) \\
&= 1 - \min\Big(1-R(x)^{2^\ell}, 2^t\big(1-R(x)\big)\Big) \\
&= \max\Big(R(x)^{2^\ell}, 1-2^t\big(1-R(x)\big)\Big), && (11)
\end{aligned}$$

where the third line follows from Lemma 3.

Next, we upper bound Equation (11) by an affine function of $R(x)$ in order to more easily upper bound $\mathbb{E}_X[D(X)]$, which is the statistical distance between P_{XZ} and $P_{XZ'}$ with Z' the simulator's output, using $\mathbb{E}_X[R(X)]$. First, $R(x)^{2^\ell}$ is a convex function of $R(x)$, and so is upper bounded by a line drawn through the lower and upper bounds for $R(x)$. I.e.,

$$R(x)^{2^\ell} \le (1-2^{-t})^{2^\ell} + \frac{1-(1-2^{-t})^{2^\ell}}{2^{-t}}(R(x)-1+2^{-t}).$$

It turns out that this is also an upper bound on $1-2^t(1-R(x))$. Indeed, to see that

$$1-2^t(1-y) \le (1-2^{-t})^{2^\ell} + \frac{1-(1-2^{-t})^{2^\ell}}{2^{-t}}(y-1+2^{-t})$$

for every $y\in[0,1]$, it suffices to note that the two sides are equal when $y=1$, and that the derivative of the left-hand side expression is 2^t while the derivative of the right-hand side expression is $\frac{1-(1-2^{-t})^{2^\ell}}{2^{-t}} \le 2^t$ (so RHS$-$LHS is non-increasing in $[0,1]$ and equals 0 at $y=1$, hence is non-negative in $[0,1]$). Therefore,

$$D(x) \le (1-2^{-t})^{2^\ell} + \frac{1-(1-2^{-t})^{2^\ell}}{2^{-t}}(R(x)-1+2^{-t}). \quad (12)$$

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

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

The first inequality uses Equation (12). The second one follows from Lemma 3. The final inequality holds because $1+y \le e^y$ for any real y. This yields the first part of the theorem statement. To see the second part, set ℓ so that $\alpha \ge e^{-2^{\ell-t}}$. □

a) Direct vs. indirect approach: 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 simulation theorem for dense leakage from [24]. The main difference is that we would get worse simulation error through this 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 constant $c \ge 1$. 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 [24] 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 $P_{Z|X=x}(z) \le T\cdot P_Z(z)$ 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 $\mathbb{E}_{x\sim P_X}[\mathsf{SD}_t(P_{Z|X=x};Q)] \le \delta$ for some distribution Q. 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.

Another limitation of the indirect approach is that, while we can show that (t,δ)-SD-noisy leakage is captured by dense leakage (with sub-optimal parameters), it is not clear to us

whether this can also be done for GSD-noisy leakage in general because we may have $Q \neq P_Z$.

Motivated by these shortcomings and by our improved analysis of the rejection sampling simulator, we explore the relationship between GSD-noisy leakage and dense leakage further in Section VII-A. We discuss how we may interpret the GSD-noisy leakage model as imposing an "average on X" density constraint between the conditional leakage distributions $P_{Z|X=x}$ and the distribution Q. In particular, this means that GSD-noisy leakage is a more general model than the dense leakage model of Brian et al. [24], and we feel that it is the more conceptually appropriate definition of "dense leakage". We also discuss how our Theorem 4 generalizes the main simulation theorem for dense leakage from bounded leakage of [24], and Theorem 5 guarantees composition for (average) dense leakages. Moreover, existing noisy leakage models that were shown to be captured by dense leakage in [24] are captured by GSD-noisy leakage with better parameters. This leads to simulation theorems with practically significant improvements in simulation error. We discuss this for the Uniform-Noisy leakage model.

VI. Composition of GSD-Noisy Leakages

A. Main Composition Theorem for GSD-Noisy Leakages

In this section we 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 [36].

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 $Q_1 \otimes Q_2$.[10] 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

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

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,

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

Finally, take the expectation over X to get

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

The theorem statement follows. □

B. Advanced Composition for Two-Sided GSD-Noisy Leakage

In this section, we show that if leakages fall into the natural symmetric version of GSD-noisy leakage, which we call *two-sided GSD-noisy leakage*, then they satisfy a stronger composition theorem than that given by Theorem 5, akin to advanced composition in differential privacy. As in differential privacy, advanced composition of two-sided GSD-noisy leakages yields an improvement over standard composition only for a limited range of parameters, making it less practically relevant than Theorem 5. We begin by defining the two-sided GSD-noisy leakage model.

Definition 6 ((t, δ)-2GSD-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)-*2GSD-noisy leakage function *from X if, denoting $Z = f(X)$, there exists a distribution Q on $\mathcal{Z}$ such that $\mathsf{SD}_t(P_X \otimes Q; P_{XZ}) \leq \delta$ and $\mathsf{SD}_t(P_{XZ}; P_X \otimes Q) \leq \delta$.*

We may see 2GSD-noisy leakage as the natural symmetric variant of the GSD-noisy noisy leakage model. In fact, Z is 2GSD-noisy leakage from X exactly when it is both GSD-noisy leakage and RevGSD-noisy leakage from X.

It will be useful to rewrite the condition that Z is (t, δ)-2GSD-noisy leakage from X as follows: there exists a distribution Q on $\mathcal{Z}$ such that for every set $\mathcal{T} \subseteq \mathcal{X} \times \mathcal{Z}$ we have that

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

and

$$P_{XZ}(\mathcal{T}) \leq 2^t (P_X \otimes P_Q)(\mathcal{T}) + \delta.$$

[10] We remark that even when $Q_1 = P_{Z_1}$ and $Q_2 = P_{Z_2}$ (in which case each Z_i is (t_i, δ_i)-SD-noisy leakage from X) the resulting $Q = P_{Z_1} \otimes P_{Z_2}$ does not correspond to $P_{Z_1Z_2}$. Therefore, even in this case we cannot show that $Z = (Z_1, Z_2)$ is (t, δ)-SD-noisy leakage from X.

We will exploit this equivalent rephrasing in the proof of the following advanced composition theorem for 2GSD-noisy leakage.

Theorem 6. *Suppose that $Z_1, \ldots, Z_m$ are (t, δ)-2GSD-noisy leakages from X and conditionally independent given X. Then, for any $\gamma > 0$ it holds that the leakage tuple $Z = (Z_1, \ldots, Z_m)$ is (t', δ')-2GSD-noisy leakage from X with*

$$t' = 2t(1 - 2^{-t})m + t \cdot \sqrt{2m \ln(1/\gamma)} \quad \text{and} \quad \delta' = \gamma + 2m\delta.$$

We briefly discuss some relevant parameter regimes for this result. It is useful to keep in mind that $t(1 - 2^{-t}) \leq t^2$ for all $t \geq 0$. These correspond to when t is (i) a constant smaller than $1/2$, (ii) $t = m^{-c}$ with $0 < c < 1/2$, and (iii) $t = m^{-c}$ for $c \geq 1/2$. Set $\gamma = \delta$. If $t = 1/4$, then the global leakage Z is $(m/8 + \sqrt{2m \ln(1/\delta)}, (2m+1)\delta)$-2GSD-noisy leakage from X. In contrast, Theorem 5 only gives that Z is $(m/4, m\delta)$-GSD-noisy leakage from X for this choice of t. If $t = m^{-1/4}$ and $\gamma = \delta$, then the global leakage Z is $(2\sqrt{m} + m^{1/4}\sqrt{2\ln(1/\delta)}, (2m+1)\delta)$-2GSD-noisy leakage from X. In this setting, Theorem 5 would only guarantee that Z is $(m^{3/4}, m\delta)$-GSD-noisy leakage from X. If $t = m^{-1/2}$ and $\gamma = \delta$ then Theorem 6 ensures that the global leakage Z is $(O(t \cdot \sqrt{m \ln(1/\delta)}), (2m+1)\delta)$-2GSD-noisy leakage from X.

The behavior above features the same improved scaling with m that we observe in advanced composition for approximate differential privacy, and it has the same shortcoming in that it kicks in only when t is small. This makes Theorem 6 less practically relevant than Theorem 5 in our setting. As we discuss in more detail below, an inspection of the proof reveals that avoiding this shortcoming seems unlikely. Of course, if t is not small enough we can still enjoy the composition guarantees from Theorem 5. More precisely, it also holds that Z is $(mt, m\delta)$-GSD-noisy leakage.

Before we prove Theorem 6, we introduce a definition and a lemma.

Definition 7 ((t, δ)-indistinguishability)**.** *We say that two distributions P and Q are* (t, δ)-indistinguishable *if $\mathsf{SD}_t(P; Q) \leq \delta$ and $\mathsf{SD}_t(Q; P) \leq \delta$.*

The following decomposition lemma for this notion of indistinguishability will be useful. Various similar decompositions of this type are well known in the literature. We state a particular version with adapted notation.

Lemma 4 ([37, Lemma 23])**.** *If two distributions P and Q are (t, δ)-indistinguishable, then we can write $P = (1 - \delta)P' + \delta P''$ and $Q = (1 - \delta)Q' + \delta Q''$ for some distributions P', P'', Q', Q'' such that P' and Q' are $(t, 0)$-indistinguishable.*

We proceed to the proof of Theorem 6. We follow the proof sketch for advanced composition of approximate differential privacy in [38, Lemma 2.4] closely, adapting it to our scenario and notation and making the argument for $\delta > 0$ explicit for completeness.

Proof of Theorem 6. Since the Z_i's are (t, δ)-2GSD-noisy leakages from X, we know that there exist distributions $Q_1, \ldots, Q_m$ such that for any $x \in \mathsf{supp}(X)$ it holds that $\mathsf{SD}_t(P_{Z_i|X=x}; Q_i) \leq \delta_{i,x,L}$ and $\mathsf{SD}_t(Q_i; P_{Z_i|X=x}) \leq \delta_{i,x,R}$ for some $\delta_{i,x,L}, \delta_{i,x,R} \geq 0$ such that $\mathbb{E}_X[\delta_{i,X,L}], \mathbb{E}_X[\delta_{i,X,R}] \leq \delta$. In particular, this means that $P_{Z_i|X=x}$ and Q_i are $(t, \delta_{i,x} = \max(\delta_{i,x,L}, \delta_{i,x,R}))$-indistinguishable for all i and x. Note that $\mathbb{E}_X[\delta_{i,X}] \leq \delta + \delta = 2\delta$.

Invoking Lemma 4, for each i and x we can write $P_{Z_i|X=x} = (1-\delta_{i,x})P'_{i,x} + \delta_{i,x}P''_{i,x}$ and $Q_i = (1-\delta_{i,x})Q'_{i,x} + \delta_{i,x}Q''_{i,x}$ for some distributions $P'_{i,x}, P''_{i,x}, Q'_{i,x}, Q''_{i,x}$ such that $P'_{i,x}$ and $Q'_{i,x}$ are $(t, 0)$-indistinguishable. Now, consider the log-ratios

$$\mathcal{L}_{i,x}(z) = \log\left(\frac{P'_{i,x}(z)}{Q'_{i,x}(z)}\right) \tag{13}$$

defined for $z \in \mathsf{supp}(Q'_{i,x}) = \mathsf{supp}(P'_{i,x})$.

Since $P'_{i,x}$ and $Q'_{i,x}$ are $(t, 0)$-indistinguishable, we know that $|\mathcal{L}_{i,x}(z)| \leq t$ for all z. We will now show that

$$\mathbb{E}_{z \sim P'_{i,x}}[\mathcal{L}_{i,x}(z)] \leq 2t(1 - 2^{-t}). \tag{14}$$

First, note that $\mathbb{E}_{z \sim P'_{i,x}}[\mathcal{L}_{i,x}(z)] = D_{\mathsf{KL}}(P'_{i,x} \| Q'_{i,x})$ and $\mathbb{E}_{z \sim Q'_{i,x}}[-\mathcal{L}_{i,x}(z)] = D_{\mathsf{KL}}(Q'_{i,x} \| P'_{i,x})$, where we recall that $D_{\mathsf{KL}}(\cdot\|\cdot)$ denotes the Kullback-Leibler divergence (defined with respect to the base-2 logarithm $\log$) between two probability distributions, and so these quantities are non-negative. Therefore, to establish Equation (14) it actually suffices to show that

$$\mathbb{E}_{z \sim P'_{i,x}}[\mathcal{L}_{i,x}(z)] + \mathbb{E}_{z \sim Q'_i}[-\mathcal{L}_{i,x}(z)] \leq 2t(1 - 2^{-t}).$$

We can rewrite the left-hand expression as

$$\begin{aligned}
&\mathbb{E}_{z \sim P'_{i,x}}[\mathcal{L}_{i,x}(z)] + \mathbb{E}_{z \sim Q'_{i,x}}[-\mathcal{L}_{i,x}(z)] \\
&= \sum_z \left(P'_{i,x}(z) - Q'_{i,x}(z)\right) \cdot \mathcal{L}_{i,x}(z) \\
&\leq 2 \cdot \mathsf{SD}(P'_{i,x}; Q'_{i,x}) \cdot \max_z |\mathcal{L}_{i,x}(z)| \\
&\leq 2t(1 - 2^{-t}),
\end{aligned}$$

as desired. The second inequality follows from the fact that $(t, 0)$-indistinguishable random variables can be at most $(1 - 2^{-t})$ far apart in statistical distance[11] and that $\mathcal{L}_i(z) \leq t$ for all z.

For a vector $\vec{z} = (z_1, \ldots, z_m) \in \mathcal{Z}^m$, define

$$\mathcal{L}_x(\vec{z}) = \log\left(\frac{\prod_{i=1}^m P'_{i,x}(z_i)}{\prod_{i=1}^m Q'_{i,x}(z_i)}\right) = \mathcal{L}_{1,x}(z_1) + \cdots + \mathcal{L}_{m,x}(z_m).$$

Recall that $|\mathcal{L}_{i,x}(z_i)| \leq t$ for all i and z_i and that $\mathbb{E}_{z_i \sim P'_{i,x}}[\mathcal{L}_{i,x}(z_i)] \leq 2t(1 - 2^{-t})$ for all i by Equation (14). In particular, this implies that $\mathbb{E}_{\vec{z} \sim P'_{1,x} \otimes \cdots \otimes P'_{m,x}}[\mathcal{L}_x(\vec{z})] \leq 2t(1 - 2^{-t})m$. We can then apply Hoeffding's inequality[12]

[11] Let V and W be $(t, 0)$-indistinguishable distributions. Then, we know that $\mathsf{SD}(W; V) = \sum_x \max\{0; W(x) - V(x)\}$. But $W(x) \leq 2^t V(x)$, and so $2^{-t}W(x) \leq V(x)$. Plugging this in we get $\mathsf{SD}(W; V) \leq \sum_x \max(0; W(x)(1 - 2^{-t})) \leq 1 - 2^{-t}$.

[12] The version of Hoeffding's inequality that we use states that for $X_1, \ldots, X_n$ independent bounded random variables such that $\mathsf{supp}(X_i) \subseteq [a_i, b_i]$ for each i and any $w \geq 0$ we have $\Pr[\sum_{i=1}^m X_i - \sum_{i=1}^m \mathbb{E}[X_i] \geq w] \leq \exp\left(-2w^2 / \sum_{i=1}^m (b_i - a_i)^2\right)$ [39, Theorem 2.2.6]. We are applying Hoeffding's inequality to the random variables $\mathcal{L}_{i,x}(z_i)$ with $z_i \sim P'_{i,x}$, and so $b_i - a_i \leq 2t$ for each i and $\sum_{i=1}^m \mathbb{E}_{z_i \sim P'_{i,x}}[\mathcal{L}_{i,x}(z_i)] \leq \sum_{i=1}^m 2t(1 - 2^{-t}) = 2t(1 - 2^{-t})m$.

to conclude that for any $\gamma > 0$ it holds that

$$\Pr_{\vec{z}\sim P'_{1,x}\otimes\cdots\otimes P'_{m,x}}\left[\mathcal{L}_x(\vec{z}) \le 2t(1-2^{-t})m + t\cdot\sqrt{2m\ln(1/\gamma)}\right] \ge 1-\gamma. \tag{15}$$

Set $t' = 2t(1-2^{-t})m + t\cdot\sqrt{2m\ln(1/\gamma)}$. By definition of $\mathcal{L}$, for any set $\mathcal{T}\subseteq\mathcal{Z}^m$ we have that

$$\begin{aligned}
&(P'_{1,x}\otimes\cdots\otimes P'_{m,x})(\mathcal{T})\\
&\le \sum_{\vec{y}\in\mathcal{T}:\mathcal{L}(\vec{y})\le t'} (P'_{1,x}\otimes\cdots\otimes P'_{m,x})(\vec{y})\\
&+ \Pr_{\vec{z}\sim P'_{1,x}\otimes\cdots\otimes P'_{m,x}}[\mathcal{L}(\vec{z}) > t']\\
&\le \sum_{\vec{y}\in\mathcal{T}:\mathcal{L}(\vec{y})\le t'} 2^{t'}(Q'_{1,x}\otimes\cdots\otimes Q'_{m,x})(\vec{y}) + \gamma\\
&\le 2^{t'}(Q'_{1,x}\otimes\cdots\otimes Q'_{m,x})(\mathcal{T}) + \gamma.
\end{aligned}$$

This means that

$$\mathsf{SD}_{t'}(P'_{1,x}\otimes\cdots\otimes P'_{m,x}; Q'_{1,x}\otimes\cdots\otimes Q'_{m,x}) \le \gamma.$$

By symmetry, repeating the argument above with $Q'_{i,x}$ in place of $P'_{i,x}$ and vice-versa allows us to conclude that

$$\mathsf{SD}_{t'}(Q'_{1,x}\otimes\cdots\otimes Q'_{m,x}; P'_{1,x}\otimes\cdots\otimes P'_{m,x}) \le \gamma.$$

Therefore, $P'_{1,x}\otimes\cdots\otimes P'_{m,x}$ and $Q'_{1,x}\otimes\cdots\otimes Q'_{m,x}$ are (t',γ)-indistinguishable.

For any z, by the conditional independence of the Z_i's given X we can write

$$\begin{aligned}
P_{Z|X=x}(z) &= \prod_{i=1}^m P_{Z_i|X=x}(z_i)\\
&= \prod_{i=1}^m [(1-\delta_{i,x})P'_{i,x}(z_i) + \delta_{i,x}P''_{i,x}(z_i)]\\
&= \left(\prod_{i=1}^m(1-\delta_{i,x})\right)(P'_{1,x}\otimes\cdots\otimes P'_{m,x})(z)\\
&+ \left(1-\prod_{i=1}^m(1-\delta_{i,x})\right)V_x(z)
\end{aligned} \tag{16}$$

for some distribution V_x. Likewise, for $Q = Q_1\otimes\cdots\otimes Q_m$ we can write

$$\begin{aligned}
Q(z) = \left(\prod_{i=1}^m(1-\delta_{i,x})\right)&(Q'_{1,x}\otimes\cdots\otimes Q'_{m,x})(z)\\
&+ \left(1-\prod_{i=1}^m(1-\delta_{i,x})\right)W_x(z)
\end{aligned} \tag{17}$$

for some distribution W_x.

Define $\alpha_x = 1-\prod_{i=1}^m(1-\delta_{i,x})$. By the generalized Bernoulli inequality, we have that $0\le\alpha_x\le\sum_{i=1}^m\delta_{i,x}$. Then, for any set $\mathcal{T}\subseteq\mathcal{Z}^m$ we have

$$\begin{aligned}
&P_{Z|X=x}(\mathcal{T})\\
&= (1-\alpha_x)(P'_{1,x}\otimes\cdots\otimes P'_{m,x})(\mathcal{T}) + \alpha_x V_x(\mathcal{T})\\
&\le (1-\alpha_x)(2^{t'}(Q'_{1,x}\otimes\cdots\otimes Q'_{m,x})(\mathcal{T}) + \gamma) + \sum_{i=1}^m\delta_{i,x}\\
&= (1-\alpha_x)\left(2^{t'}\left(\frac{Q(\mathcal{T}) - \alpha_x W_x(\mathcal{T})}{1-\alpha_x}\right) + \gamma\right) + \sum_{i=1}^m\delta_{i,x}\\
&\le 2^{t'}Q(\mathcal{T}) + \gamma + \sum_{i=1}^m\delta_{i,x},
\end{aligned}$$

and so $\mathsf{SD}_t(P_{Z|X=x}; Q)\le\gamma+\sum_{i=1}^m\delta_{i,x}$. The first equality uses Equation (16). The first inequality follows from the fact that $P'_{1,x}\otimes\cdots\otimes P'_{m,x}$ and $Q'_{1,x}\otimes\cdots\otimes Q'_{m,x}$ are (t',γ)-indistinguishable and $\alpha_x\le\sum_{i=1}^m\delta_{i,x}$. The second equality uses Equation (17). Analogously, it holds that

$$Q(\mathcal{T}) \le 2^{t'}P_{Z|X=x}(\mathcal{T}) + \gamma + \sum_{i=1}^m\delta_{i,x}$$

for all sets $\mathcal{T}$, and so $\mathsf{SD}_t(Q; P_{Z|X=x})\le\gamma+\sum_{i=1}^m\delta_{i,x}$.

Finally, taking the expectation over X yields

$$\begin{aligned}
\mathsf{SD}_t(P_{XZ}; P_X\otimes Q) &= \mathbb{E}_{x\sim P_X}[\mathsf{SD}_t(P_{Z|X=x}; Q)]\\
&\le \gamma + \sum_{i=1}^m\mathbb{E}_X[\delta_{i,X}]\\
&\le \gamma + 2m\delta.
\end{aligned}$$

and

$$\begin{aligned}
\mathsf{SD}_t(P_X\otimes Q; P_{XZ}) &= \mathbb{E}_{x\sim P_X}[\mathsf{SD}_t(Q; P_{Z|X=x})]\\
&\le \gamma + \sum_{i=1}^m\mathbb{E}_X[\delta_{i,X}]\\
&\le \gamma + 2m\delta.
\end{aligned}$$

We conclude that Z is $(t',\gamma+2m\delta)$-2GSD-noisy leakage from X, as desired. □

a) Can we get better scaling for larger t?: We discuss whether one can prove an advanced composition theorem akin to Theorem 6 but guaranteeing $t' = o(m)$ even when t is not small. In short, this seems unlikely beyond improving constant factors. We sketch why below. Notice that the key steps of the proof consist in deriving the upper bound in Equation (14), and in using this bound in an application of Hoeffding's inequality in Equation (15). Concentration inequalities like Hoeffding's inequality or Chernoff bounds are known to be optimal up to a constant (e.g., see [40, Lemma 5.2]), which leaves little room for improvement. Our only hope would be to improve our upper bound $\mathbb{E}_{z\sim P'_{i,x}}[\mathcal{L}_{i,x}(z)] = D_{\mathsf{KL}}(P'_{i,x}\|Q'_{i,x}) \le 2t(1-2^{-t})$. However, this is not possible in general, as there exist families of $(t,0)$-indistinguishable distributions P and Q for growing t such that $D_{\mathsf{KL}}(P\|Q) = \mathbb{E}_{x\sim P}\left[\log\left(\frac{P(x)}{Q(x)}\right)\right] \ge t(1-2^{-t})/2$.

For example, consider t's of the form $t = \log(2^r-1)$ for an integer $r\ge 1$, define P as the uniform distribution over $\{0,1\}^n$, and let $\mathcal{S}\subseteq\{0,1\}^n$ be some set of size $|\mathcal{S}| = \frac{2^n}{1+2^t}$. Then, we define the distribution Q as

$$Q(x) = \begin{cases} 2^{t-n}, & \text{if } x\in\mathcal{S},\\ 2^{-t-n}, & \text{if } x\notin\mathcal{S}.\end{cases}$$

Note that $\sum_{x\in\{0,1\}^n} Q(x) = |\mathcal{S}|\cdot 2^{t-n} + (2^n - |\mathcal{S}|)\cdot 2^{-t-n} = 1$, and so Q is a valid probability distribution. Furthermore, P and Q are $(t,0)$-indistinguishable, since

$$2^{-t}P(x) = 2^{-t-n} \le Q(x) \le 2^{t-n} = 2^t P(x)$$

for all $x \in \{0,1\}^n$. Finally, we have

$$\begin{aligned} D_{\mathsf{KL}}(P\|Q) &= \frac{1}{1+2^t}\cdot(-t) + \left(1 - \frac{1}{1+2^t}\right)\cdot t \\ &= \left(\frac{1-2^{-t}}{1+2^{-t}}\right)\cdot t \\ &\ge \frac{t(1-2^{-t})}{2}. \end{aligned}$$

VII. GSD-Noisy Leakage and Other Leakage Models

A. GSD-Noisy Leakage and Average Dense Leakage

In this section we explore the relationship between GSD-noisy leakage and the main dense leakage model of Brian et al. [24].

We begin by introducing the "average dense leakage" model in Section VII-A1. Then, in Section VII-A2 we recall the dense leakage model of [24] and show how it is captured by the average dense leakage model. In Section VII-A3 we show that average dense leakage is a special case of GSD-noisy leakage, and discuss how Theorems 4 and 5 generalize the main results of [24]. We show that GSD-noisy leakage is captured by average dense leakage with only a small loss in parameters in Section VII-A4, meaning that GSD-noisy leakage and average dense leakage are almost equivalent. Finally, in Section VII-A5 we discuss an example of how Theorem 4 affords practically significant improvements over the main simulation theorem of [24].

1) The Average Dense Leakage Model: We provide some motivation before defining the average dense leakage model. Using Lemma 1, we can equivalently rewrite the inequality $\mathsf{SD}_t(P_{XZ}; P_X \otimes Q) \le \delta$ as

$$\begin{aligned} &\sum_{x,z} \min(P_{XZ}(x,z), 2^t(P_X\otimes Q)(x,z)) \\ &= \mathbb{E}_{x\sim P_X}\left[\sum_z \min(P_{Z|X=x}(z), 2^t Q(z))\right] \ge 1-\delta. \end{aligned} \quad (18)$$

Therefore, the (t,δ)-GSD-noisy leakage model is "average-case" over P_X in the sense that it only requires that for each x,

$$\sum_z \min(P_{Z|X=x}(z), 2^t Q(z)) \ge 1-\delta_x$$

for some non-negative real numbers $(\delta_x)_{x\in\mathcal{X}}$ such that $\mathbb{E}_X[\delta_X] \le \delta$.

Following [24], for two distributions P and Q over $\mathcal{X}$, we say that P is *t-dense* in Q if $P(x) \le 2^t Q(x)$ for all $x \in \mathcal{X}$. Motivated by the "average-case on X" property of the GSD-noisy leakage model, our average dense leakage model will impose (approximate) density constraints between the conditional distributions $P_{Z|X=x}$ and Q in expectation over X. First, we introduce a notion of approximate density that is weaker than the one of Brian et al. [24, Definition 3].

Definition 8 (Approximate density)**.** *We say that P is (t,δ)-*dense *in Q if*

$$P(\{z : P(z) \le 2^t Q(z)\}) \ge 1-\delta.$$

We define average dense leakage based on this notion of approximate density.

Definition 9 (Average Dense leakage)**.** *We say that a random variable Z supported on $\mathcal{Z}$ is (t,δ)-*average dense leakage *from X if there exists a distribution Q on $\mathcal{Z}$ such that P_{XZ} is (t,δ)-dense in $P_X \otimes Q$.*

We call this model "average dense" because it is equivalent to requiring that the conditional distributions $P_{Z|X=x}$ are approximately dense in P_Z with good parameters in expectation over $x \sim P_X$. More precisely, Z is (t,δ)-average dense leakage from X witnessed by Q if and only if $P_{Z|X=x}$ is (t,δ_x)-dense in Q for all x and some non-negative real numbers $(\delta_x)_{x\in\mathcal{X}}$ satisfying $\mathbb{E}_X[\delta_X] \le \delta$. To see why, first recall that P_{XZ} is (t,δ)-dense in $P_X \otimes Q$ if and only if $P_{XZ}(\mathcal{G}) \ge 1-\delta$ for $\mathcal{G} = \{(x,z) : P_{XZ}(x,z) \le 2^t(P_X\otimes Q)(x,z)\}$. Now, for each x let $\mathcal{G}_x = \{z : P_{Z|X=x}(z) \le 2^t Q(z)\}$ and $\delta_x = 1 - P_{Z|X=x}(\mathcal{G}_x)$. By definition, $P_{Z|X=x}$ is (t,δ_x)-dense in Q. Furthermore, $\mathcal{G} = \bigcup_x \{x\}\times\mathcal{G}_x$. As a result,

$$\begin{aligned} P_{XZ}(\mathcal{G}) &= \mathbb{E}_{x\sim P_X}[P_{Z|X=x}(\{z : P_{Z|X=x}(z) \le 2^t Q(z)\})] \\ &= 1 - \mathbb{E}_X[\delta_X], \end{aligned} \quad (19)$$

and so $P_{XZ}(\mathcal{G}) \ge 1-\delta$ if and only if $\mathbb{E}_X[\delta_X] \le \delta$.

2) Dense Leakage as Average Dense Leakage: For completeness, we begin by recalling the dense leakage model of [24] and discuss how it can be captured as a special case of average dense leakage. We present a definition of dense leakage with a slight change in notation compared to [24] for consistency and to avoid notational conflicts.

Definition 10 (Dense leakage [24])**.** *Let X and Z be random variables with supports $\mathcal{X}$ and $\mathcal{Z}$, respectively. We say that Z is (t,p,γ)-*dense leakage *from X if there exists a set $\mathcal{T} \subseteq \mathcal{X}$ and sets $\mathcal{S}_x \subseteq \mathcal{Z}$ satisfying*

- *$P_X(\mathcal{T}) \ge 1-p$ and $P_Z(\mathcal{S}_x), P_{Z|X=x}(\mathcal{S}_x) \ge 1-\gamma$ for each $x \in \mathcal{X}$;*
- *$P_{Z|X=x}(z) \le 2^t P_Z(z)$ whenever $x \in \mathcal{T}$ and $z \in \mathcal{S}_x$.*

The following simple result places dense leakage inside our average dense leakage model.

Theorem 7. *If Z is (t,p,γ)-dense leakage from X, then Z is also $(t,\delta = p+\gamma-p\gamma)$-average dense leakage from X.*

Proof. Let $\mathcal{T} \subseteq \mathcal{X}$ and $\mathcal{S}_x \subseteq \mathcal{Z}$ for $x \in \mathcal{X}$ be the sets guaranteed by the definition of Z as (t,p,γ)-dense leakage from X. Then, it suffices to note that

$$\begin{aligned} &P_{XZ}(\{(x,z) : P_{Z|X=x}(z) \le 2^t P_Z(z)\}) \\ &= \sum_{x\in\mathcal{X}} P_X(x)\cdot P_{Z|X=x}(\{z : P_{Z|X=x}(z) \le 2^t P_Z(z)\}) \\ &\ge \sum_{x\in\mathcal{T}} P_X(x)\cdot P_{Z|X=x}(\{z : P_{Z|X=x}(z) \le 2^t P_Z(z)\}) \\ &\ge \sum_{x\in\mathcal{T}} P_X(x)(1-\gamma) \end{aligned}$$

$$\geq (1-p)(1-\gamma)$$
$$= 1-(p+\gamma-p\gamma).$$

The second inequality uses the fact that $P_{Z|X=x}(\{z : P_{Z|X=x}(z) \leq 2^t P_Z(z)\}) \geq P_{Z|X=x}(\mathcal{S}_x) \geq 1-\gamma$. The third inequality holds since $P_X(\mathcal{T}) \geq 1-p$. □

3) Average Dense Leakage as GSD-Noisy Leakage, and Consequences: In this section we begin by showing that average dense leakage is a special case of GSD-noisy leakage. Then, we discuss consequences of this fact with respect to simulation and composition.

Theorem 8. *If Z is (t,δ)-average dense leakage from X, then Z is also (t,δ)-GSD-noisy leakage from X.*

Proof. Let Q be the distribution witnessing that Z is (t,δ)-average dense leakage from X, and define $\mathcal{G} = \{(x,z) : P_{XZ}(x,z) \leq 2^t(P_X \otimes Q)(x,z)\}$. Note that $P_{XZ}(\mathcal{G}) \geq 1-\delta$ since Z is (t,δ)-average dense leakage from X. Then, we have that

$$\sum_{x,z} \min(P_{XZ}(x,z), 2^t(P_X \otimes Q)(x,z))$$
$$= \sum_{(x,z)\in\mathcal{G}} P_{XZ}(x,z) + \sum_{(x,z)\notin\mathcal{G}} 2^t(P_X \otimes Q)(x,z)$$
$$\geq \sum_{(x,z)\in\mathcal{G}} P_{XZ}(x,z)$$
$$\geq 1-\delta.$$ □

The following corollary is an immediate consequence of Theorems 4 and 8.

Corollary 1. *For any $\alpha > 0$, the family of (t,δ)-average dense leakages from X is $(\delta+\alpha)$-simulatable using $\lceil t+\log\ln(1/\alpha)\rceil$ bits of bounded leakage from X.*

By the connection in Theorem 7, we have that Corollary 1 generalizes the main simulation theorem for dense leakage of Brian et al. [24, Theorem 3]. Indeed, as a special case, by combining Theorem 7 and Corollary 1 we immediately obtain the following simulation theorem for dense leakage itself, which improves slightly on the main simulation theorem for dense leakage of Brian et al.

Corollary 2. *For any $\alpha > 0$, the family of (t,p,γ)-dense leakages from X is $(\alpha+p+\gamma-p\gamma)$-simulatable using $\lceil t+\log\ln(1/\alpha)\rceil$ bits of bounded leakage from X.*

Because other noisy leakage models can be cast as special cases of average dense leakage and GSD-noisy leakage with better parameters than with respect to dense leakage, the more general simulation theorem in Corollary 1 yields practically significant improvements in the simulation of those noisy leakage models from bounded leakage. We discuss an example of this in detail in Section VII-A5.

On another note, combining Theorem 5 and Theorem 8 yields the following analogous composition theorem for average dense leakage (which also implies a composition theorem for dense leakage via Theorem 7).

Corollary 3. *Suppose that $Z_1,\ldots,Z_m$ are (t_i,δ_i)-average dense leakages from X, and are conditionally independent given X. Then, for any $\alpha > 0$, the global leakage $Z = (Z_1,\ldots,Z_m)$ is $(\alpha+\sum_{i=1}^m \delta_i)$-simulatable using $\lceil \log\ln(1/\alpha) + \sum_{i=1}^m t_i\rceil$ bits of bounded leakage from X.*

4) GSD-Noisy Leakage as Average Dense Leakage: To further consolidate our viewpoint that GSD-noisy leakage is similar to average dense leakage, we show that GSD-noisy leakage is captured by average dense leakage with only a small constant loss in parameters. Therefore, GSD-noisy leakage and average dense leakage are essentially equivalent models.

Theorem 9. *Let Z supported on $\mathcal{Z}$ be (t,δ)-GSD-noisy leakage from X supported on $\mathcal{X}$. Then, Z is also $\left(t+\Delta, \frac{2^\Delta}{2^\Delta-1}\cdot\delta\right)$-average dense leakage from X for any $\Delta \geq 1$. In particular, Z is $(t+1,2\delta)$-average dense leakage from X.*

Proof. Since Z is (t,δ)-GSD-noisy leakage from X, we know that there exists a distribution Q such that for every set $\mathcal{S} \subseteq \mathcal{X}\times\mathcal{Z}$ it holds that

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

For each x in the support of X, define δ_x as the smallest δ such that for every $\mathcal{A} \subseteq \mathcal{Z}$ we have that

$$P_{Z|X=x}(\mathcal{A}) \leq 2^t Q(\mathcal{A}) + \delta.$$

By Lemma 1, we have

$$P_{Z|X=x}(\mathcal{B}_x) = 2^t Q(\mathcal{B}_x) + \delta_x$$

for $\mathcal{B}_x = \{z : P_{Z|X=x}(z) > 2^t Q(z)\}$, and

$$P_{XZ}(\mathcal{B}) = 2^t(P_X \otimes Q)(\mathcal{B}) + \delta$$

for $\mathcal{B} = \{(x,z) : P_{XZ}(x,z) > 2^t(P_X \otimes Q)(z)\} = \{(x,z) : z \in \mathcal{B}_x\}$. In particular, this implies that $\mathbb{E}[\delta_X] = \delta$, since

$$\begin{aligned} 2^t(P_X \otimes Q)(\mathcal{B}) + \delta &= P_{XZ}(\mathcal{B}) \\ &= \mathbb{E}_{x\sim P_X}[P_{Z|X=x}(\mathcal{B}_x)] \\ &= \mathbb{E}_{x\sim P_X}[2^t Q(\mathcal{B}_x) + \delta_x] \\ &= 2^t(P_X \otimes Q)(\mathcal{B}) + \mathbb{E}_X[\delta_X]. \end{aligned}$$

We now wish to show that $P_{Z|X=x}$ is $\left(t+\Delta, \frac{2^\Delta}{2^\Delta-1}\cdot\delta_x\right)$-dense in Q for each x, which would conclude the argument. For a given $C > 0$, let $\mathcal{B}_{x,C} = \{z : P_{Z|X=x}(z) > C\cdot Q(z)\}$. By definition of $\mathcal{B}_{x,C}$, we have that

$$P_{Z|X=x}(\mathcal{B}_{x,C}) > C\cdot Q(\mathcal{B}_{x,C}).$$

Moreover, it holds that

$$P_{Z|X=x}(\mathcal{B}_{x,C}) \leq 2^t Q(\mathcal{B}_{x,C}) + \delta_x.$$

Combining these two inequalities implies that

$$P_{Z|X=x}(\mathcal{B}_{x,C}) < \frac{\delta_x}{1-2^t/C}.$$

Choosing $C = 2^{t+\Delta}$ for $\Delta \geq 1$ yields $P_{Z|X=x}(\mathcal{B}_{x,C}) < \frac{\delta_x}{1-2^{-\Delta}} = \frac{2^\Delta}{2^\Delta-1}\cdot\delta_x$, which implies that $P_{Z|X=x}$ is $\left(t+\Delta, \frac{2^\Delta}{2^\Delta-1}\cdot\delta\right)$-dense in Q. □

5) Uniform-Noisy Leakage as GSD-Noisy Leakage: We give an example of how our results for the GSD-noisy leakage model lead to improved simulation theorems compared to [24], and with slightly cleaner arguments. We begin by recalling the "Uniform-Noisy" leakage model of Dodis, Haralambiev, López-Alt, and Wichs [33]. We also recall the notions of min-entropy and average conditional min-entropy that were briefly discussed in Section I-A. The *min-entropy* of a random variable X is defined as $\mathbf{H}_\infty(X) = -\log(\max_x P_X(x))$, and $\widetilde{\mathbf{H}}_\infty(X|Z) = \mathbb{E}_{z\sim P_Z}\left[2^{-\mathbf{H}_\infty(X|Z=z)}\right]$ is the *average conditional min-entropy* of X given Z.

Definition 11 (ℓ-U-Noisy-leakage)**.** *A function $f\colon \mathcal{X} \to \mathcal{Z}$ is an ℓ-*U*-Noisy leakage function if $\widetilde{\mathbf{H}}_\infty(U|f(U)) \geq \mathbf{H}_\infty(U) - \ell$, where U is the uniform distribution on $\mathcal{X}$. We say that Z is ℓ-U-Noisy leakage from X if $Z = f(X)$ for an ℓ-U-Noisy function f.*

We now analyze the parameters of U-Noisy leakage as a special case of average dense leakage.

Theorem 10. *If Z is ℓ-U-Noisy leakage from X, then Z is also $(t = \ell + \eta, \delta = 2^{-\eta})$-SD-noisy leakage from X for any $\eta > 0$.*

Proof. This argument is analogous (but slightly cleaner than) the proof of [24, Theorem 6]. By Theorem 8, it suffices to show that P_{XZ} is $(t = \ell + \eta, \delta = 2^{-\eta})$-dense in $P_X \otimes P_Z$. We can rewrite $\widetilde{\mathbf{H}}_\infty(U|f(U)) \geq \mathbf{H}_\infty(U) - \ell$ as

$$2^\ell \geq \mathbb{E}_{y\sim P_{f(U)}}\left[\max_x \frac{P_{U|f(U)=y}(x)}{P_U(x)}\right].$$

We also have

$$\begin{aligned}
&\mathbb{E}_{y\sim P_{f(U)}}\left[\max_x \frac{P_{U|f(U)=y}(x)}{P_U(x)}\right] \\
&= \sum_y P_{f(U)}(y) \max_x \frac{P_{U|f(U)=y}(x)}{P_U(x)} \\
&= \sum_y \max_x P_{f(U)|U=x}(y) \\
&\geq \sum_y \max_{x\in\mathsf{supp}(X)} P_{f(U)|U=x}(y) \\
&= \sum_y \max_{x\in\mathsf{supp}(X)} P_{Z|X=x}(y) \\
&= \mathbb{E}_{z\sim P_Z}\left[\max_{x\in\mathsf{supp}(X)} \frac{P_{Z|X=x}(z)}{P_Z(z)}\right].
\end{aligned}$$

Therefore, we get that

$$2^\ell \geq \mathbb{E}_{z\sim P_Z}\left[\max_{x\in\mathsf{supp}(X)} \frac{P_{Z|X=x}(z)}{P_Z(z)}\right].$$

Fix any $\eta > 0$. By an averaging argument, the inequality above implies that there exists a set $\mathcal{S}$ such that $P_Z(\mathcal{S}) \geq 1 - 2^{-\eta}$ and $\frac{P_{Z|X=x}(z)}{P_Z(z)} \leq 2^{\ell+\eta}$ for all x whenever $z \in \mathcal{S}$. Letting $\mathcal{G} = \mathcal{X} \times \mathcal{S}$, we get that $P_{XZ}(\mathcal{G}) = P_Z(\mathcal{S}) \geq 1 - 2^{-\eta}$, and so we conclude that P_{XZ} is $(t = \ell + \eta, \delta = 2^{-\eta})$-dense in $P_X \otimes P_Z$. □

The following immediate corollary of Theorems 4 and 10 presents a better simulation error vs. bounded leakage tradeoff than the corresponding simulation theorem of Brian et al. [24, Corollary 3] (more precisely, the simulation error is now $2^{-\eta} + \alpha$ instead of $2^{-\eta/2} + \alpha$).

Corollary 4. *For any $\alpha, \eta > 0$, the family of ℓ-U-noisy leakages from X is $(2^{-\eta} + \alpha)$-simulatable using $\lceil \ell + \eta + \log\ln(1/\alpha)\rceil$ bits of bounded leakage from X.*

B. GSD-Noisy Leakage and Mutual Information

In this section, we discuss the relationship between the GSD-noisy leakage model and noisy leakages based on mutual information, which is a popular metric in the literature on side-channel attacks [18]. We begin by recalling the δ-MI-noisy leakage model.

Definition 12 (MI-noisy leakage)**.** *We say that Z is δ-*MI*-noisy leakage from X if $I(X;Z) \leq \delta$.*

Pinsker's inequality (Equation (1)) shows that δ-MI-noisy leakages are $(t = 0, \delta' = \sqrt{\frac{\delta \ln 2}{2}})$-SD-noisy leakages. There exists a more general version of Pinsker's inequality that allows us to extend this connection to hockey-stick divergences with $t > 0$.

Lemma 5 ([27, Theorem 30, adapted and specialized], see also discussion after [27, Remark 40])**.** *For every $t > 0$ there exists a constant $c_t < 1/t$ such that for every random variables X and Z it holds that*[13]

$$\mathsf{SD}_t(P_{XZ}; P_X \otimes P_Z) \leq c_t \cdot I(X;Z) < \frac{1}{t} \cdot I(X;Z).$$

Remarkably, when $t > 0$ the square root in the upper bound in Pinsker's inequality disappears and is replaced by linear scaling with $I(X;Z)$ (and with $1/t$). We obtain the following statement as a direct consequence of Lemma 5.

Corollary 5. *Every δ-MI-noisy leakage from X is also $(t, \delta' = \delta/t)$-SD-noisy leakage from X.*

Note that the δ' term in Corollary 5 decays only linearly with t. This is not surprising, as the MI-noisy leakage model is a very general model that, similarly to the $(t = 0, \delta)$-SD-noisy leakage model, encompasses pessimistic leakage functions. Combining Corollary 5 and Theorem 4 yields the following simulation theorem for δ-MI-noisy leakages.

Corollary 6. *For any $\alpha > 0$, the class of δ-MI-noisy leakages from X is $(\alpha + \delta/t)$-simulatable from ℓ bits of bounded leakage when $\ell \geq t + \log\ln(1/\alpha)$.*

Brian et al. [24] focused only on the "trivial" simulator for MI-noisy leakage, which ignores the bounded leakage (i.e., sets $\ell = 0$) and outputs z' independently distributed according to the marginal P_Z. The linear tradeoff between t and the simulation error in Corollary 6 is consistent with the lower bound in [24, Theorem 15], which states that the class of δ-MI-noisy leakages from X uniformly distributed on $\{0,1\}^n$ is not ε-simulatable from $n-1$ bits of bounded leakage with any error $\varepsilon < \frac{\delta}{2n}$.

[13] Sason and Verdú [27] use the notation E_γ to denote the hockey-stick divergence SD_t with $t = \log\gamma$.

VIII. Simulating RevSD-Noisy Leakage via Random Probing

In their seminal work, Duc, Dziembowski, and Faust [20] showed that δ-SD-noisy leakage from a uniform distribution 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 [25]. 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. For completeness, following [20], we discuss applications of these simulation theorems to leakage-resilient circuit compilers in Section VIII-C.

Before stating our results in this direction, we define p-random probing leakage.

Definition 13 (p-random probing leakage [20])**.** *Let X be some random variable supported on $\mathcal{X}$. We say that a random variable Z supported on $\mathcal{X}\cup\{\perp\}$ is* p-random probing leakage from X *if* $\Pr[Z=x|X=x]=p$ *and* $\Pr[Z=\perp|X=x]=1-p$ *for all* $x\in\mathcal{X}$.

A. Zero-Error Simulation of Reverse SD-Noisy Leakage via Random Probing

We have the following result.

Lemma 6. *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 [20, Lemma 2] proved this result for the special case $t=0$, which corresponds to δ-SD-noisy leakage.

Proof of Lemma 6. Our argument follows the proof of [20, 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. To get some intuition for why the RevSD-noisy leakage model is useful towards establishing this, first note that this is clearly true in the special case where Z is independent of X. Therefore, if we have some control over how often $\min_x P_{Z|X=x}(z)$ is significantly smaller than $P_Z(z)$, then we can lower bound $\sum_z\pi(z)$ appropriately. The RevSD-noisy leakage model provides a guarantee of this type.

More precisely, we have that

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

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 $p=1-\sum_{z\in\mathcal{Z}}\pi(z)$. 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,

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

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

B. Low-Error Simulation of Reverse SD-Noisy Leakage via Random Probing

We provide another extension of the key lemma from [20] by allowing simulation of RevSD-noisy leakage from random probing leakage with positive simulation error. In contrast, [20] exclusively considered $t=0$ and the zero simulation error setting for a uniform secret X.

Lemma 7. *Fix $t>0$ and suppose that Z is (t,δ)-RevSD-noisy leakage from X. Then, Z is ε-simulatable by p-random probing leakage from X for $\varepsilon=2^{-t}\delta$ and $p=1-2^{-t}$.*

Proof. For all $x\in\mathcal{X}$ define the unnormalized distribution $\pi(z|x)$ given by

$$\pi(z|x)=\frac{1}{p}\cdot\max(P_{Z|X=x}(z)-2^{-t}P_Z(z),0)$$

and normalize it to get a probability distribution $\pi'(z|x)$ given by

$$\pi'(z|x) = \frac{\pi(z|x)}{\sum_{z' \in \mathcal{Z}} \pi(z'|x)}.$$

Note that

$$\begin{aligned}\sum_{z' \in \mathcal{Z}} \pi(z'|x) &\geq \frac{1}{p} \cdot \sum_{z' \in \mathcal{Z}} [P_{Z|X=x}(z') - 2^{-t} P_Z(z')] \\ &= \frac{1 - 2^{-t}}{p} \\ &= 1,\end{aligned}$$

and so $\pi'(z|x) \leq \pi(z|x)$ for all $z \in \mathcal{Z}$.

Now, define the simulation of Z to check for the p-random probing leakage, and if it equals $x \neq \perp$ to sample Z according to $\pi'(z|x)$, and otherwise sample Z according to P_Z. Let $P_{\mathsf{Sim}|X=x}$ be the conditional distribution of the simulator's output given $X = x$. Then,

$$\begin{aligned}P_{\mathsf{Sim}|X=x}(z) &= p \cdot \pi'(z|x) + (1-p) P_Z(z) \\ &\leq p \cdot \pi(z|x) + 2^{-t} P_Z(z) \\ &= \max(P_{Z|X=x}(z), 2^{-t} P_Z(z)). \qquad (20)\end{aligned}$$

The inequality uses the fact that $\pi'(z|x) \leq \pi(z|x)$ always, as discussed above. Then, we bound the simulation error $\mathsf{SD}(P_{X\mathsf{Sim}}; P_{XZ})$ as

$$\begin{aligned}&\mathsf{SD}(P_{X\mathsf{Sim}}; P_{XZ}) \\ &= \mathbb{E}_{x \sim P_X} \left[\mathsf{SD}(P_{\mathsf{Sim}|X=x}; P_{Z|X=x})\right] \\ &= \mathbb{E}_{x \sim P_X} \left[\sum_{z \in \mathcal{Z}} \max(P_{\mathsf{Sim}|X=x}(z) - P_{Z|X=x}(z), 0)\right] \\ &\leq \mathbb{E}_{x \sim P_X} \left[\sum_{z \in \mathcal{Z}} \max(\max(P_{Z|X=x}(z), 2^{-t} P_Z(z)) \right. \\ &\qquad \left. - P_{Z|X=x}(z), 0)\right] \\ &= \mathbb{E}_{x \sim P_X} \left[\sum_{z \in \mathcal{Z}} \max(2^{-t} P_Z(z) - P_{Z|X=x}(z), 0)\right] \\ &= 2^{-t} \cdot \mathsf{SD}_t(P_X \otimes P_Z; P_{XZ}) \\ &\leq 2^{-t} \delta,\end{aligned}$$

where the first inequality uses Equation (20) and the second and fourth equalities use Lemma 1. □

C. Application to Private Circuits

For completeness, we discuss how to use our zero-error reduction from RevSD-noisy leakage to random probing (Lemma 6) in order to obtain *private circuits*, i.e., stateful cryptographic circuits that maintain privacy even in the presence of an adversary observing RevSD-noisy leakage on the intermediate values produced during the computation. We start by recalling the definition of threshold probing leakage from [23].

Definition 14 (Vector τ-threshold probing leakage [23])**.** *Consider a random variable $X = (X_1, X_2, \ldots, X_\ell)$ where each $X_i \in \mathcal{X}$. We say that $Z = (Z_1, \ldots, Z_\ell)$ is τ-*threshold probing leakage from X *if each Z_i is either X_i or $\perp$, and $Z_i \neq \perp$ for at most τ indices $i \in [\ell]$.*

We can also easily extend the definition of random probing leakage to length-ℓ vectors.

Definition 15 (Vector p-random probing leakage [20])**.** *Consider a random variable $X = (X_1, \ldots, X_\ell)$ where each $X_i \in \mathcal{X}$. We say that a random variable $Z = (Z_1, \ldots, Z_\ell)$, where each $Z_i \in \mathcal{X} \cup \{\perp\}$, is* p-random probing leakage from X *if the Z_i's are conditionally independent given X and for each $i \in [\ell]$ we have $\Pr[Z_i = X_i] = p$ and $\Pr[Z_i = \perp] = 1 - p$.*

The following result from [20] links the random and threshold probing leakage models, and follows from a standard application of concentration inequalities.

Lemma 8 ([20, Lemma 6], adapted)**.** *Let $X = (X_1, \ldots, X_\ell)$ be an arbitrary random vector supported on $\mathbb{F}^\ell$. Suppose that $Z = (Z_1, \ldots, Z_\ell)$ is p-random probing leakage from $X = (X_1, \ldots, X_\ell)$ and let $\tau = 2p\ell - 1$. Then, there exists a simulator Sim running in time polynomial in $|\mathbb{F}|$ and ℓ with access to one query of τ-threshold probing leakage from X that either aborts or outputs Z' such that $\Pr[Z' = z | X = x] = \Pr[Z = z | X = x]$ for all x and z and $\Pr[\mathsf{Sim}$ aborts$] \leq e^{-\frac{p\ell}{3}}$.*

1) Leakage-Resilient Stateful Arithmetic Circuits: We proceed to define the circuit model we consider, which corresponds to the one from [23], [20]. Our presentation follows [20, Section 5.1] essentially verbatim.

a) Stateful arithmetic circuits: A *stateful arithmetic circuit* Γ is a directed graph whose nodes represent gates over a finite field $\mathbb{F}$. These gates can be input and output gates (with fan-in 0 and fan-out 0, respectively), addition and subtraction gates (with fan-in 2), multiplication gates (with fan-in 2), constant gates, random gates (with fan-in 0, producing a uniformly random element of $\mathbb{F}$ in each round), and memory gates (with fan-in 1, which on a given invocation output the previously stored value and update it to the current input). As in [23], [20], the fan-out of any of these gates is assumed to be at most 3. There may be cycles in Γ, but they must contain exactly one memory gate. We denote the number of gates in Γ by $|\Gamma|$.

The computation of Γ proceeds by rounds. Let k be the string containing the symbols stored in the memory gates before the first round in some predefined order. In the first round, the input gates of Γ are loaded with an input string a_1. Then, Γ produces a (possibly randomized) output string b_1, and the values of the memory gates are updated to some string k_1. The computation in the second round will use some new input string a_2 and memory gates with values from the updated string k_1. In general, in the i-th round Γ receives an input string a_i, outputs a random string b_i, and updates its memory from k_{i-1} to k_i. We denote the behavior of Γ with initial memory state k by $\Gamma(k)$, and its output given inputs $(a_1, \ldots, a_r)$ and initial memory state k by $\Gamma(k, a_1, \ldots, a_r)$.

b) Adversarial models: We will consider adversaries that interact with a circuit $\Gamma(k)$ via the input-output interface over

several rounds, and possibly get additional leakage from the circuit wires in each round. A *black-box circuit adversary* $\mathcal{A}$ interacts with $\Gamma(k)$ only through its input-output interface. We denote the output of $\mathcal{A}$ after such an interaction by $\mathsf{out}(\mathcal{A} \overset{\mathsf{bb}}{\rightleftarrows} \Gamma(k))$.

Turning to leakage models, we consider the following adversaries:

- An (r,τ)-*threshold probing circuit adversary* $\mathcal{A}$ interacts with $\Gamma(k)$ through its input-output interface over r rounds, and also learns threshold probing leakage from the wires of $\Gamma(k)$ in each round. More precisely, letting $X = (X_1,\ldots,X_w)$ denote the values of the wires of $\Gamma(k)$ in the i-th round, $\mathcal{A}$ learns any τ-threshold probing leakage $Z = (Z_1,\ldots,Z_w)$ from X of its choice. We denote the output of $\mathcal{A}$ after such an interaction by $\mathsf{out}(\mathcal{A} \overset{\mathsf{thres}}{\rightleftarrows} \Gamma(k))$.
- An (r,p)-*random probing circuit adversary* $\mathcal{A}$ behaves similarly to a threshold probing circuit adversary, except that in each round it learns p-random probing leakage $Z = (Z_1,\ldots,Z_w)$ from the wire values $(X_1,\ldots,X_w)$ in each of the r rounds of interaction. We denote the output of $\mathcal{A}$ after such an interaction by $\mathsf{out}(\mathcal{A} \overset{\mathsf{rand}}{\rightleftarrows} \Gamma(k))$.
- An (r,t,δ)-*Uniform-RevSD-noisy circuit adversary* $\mathcal{A}$ behaves similarly to a threshold or random probing circuit adversary, except that in each of the r rounds it specifies randomized functions $f_1,\ldots,f_w$ such that each f_i is a (t,δ)-RevSD-noisy leakage function from $U_{\mathbb{F}}$ (the uniform distribution on $\mathbb{F}$), and learns the leakage $(Z_1,\ldots,Z_w)$ where $Z_i = f_i(X_i)$ with X_i denoting the i-th wire value in that round, and the f_i's use independent randomness. We denote the output of $\mathcal{A}$ after such an interaction by $\mathsf{out}(\mathcal{A} \overset{\mathsf{noisy-unif}}{\rightleftarrows} \Gamma(k))$. The special case with $t = 0$ corresponds to the δ-noisy adversary of [20].
- An (r,t,δ)-*RevSD-noisy circuit adversary* $\mathcal{A}$ behaves similarly to the adversaries above. Suppose that in the j-th round, $j \in [r]$, the adversary $\mathcal{A}$ selects input a_j (which may be adaptively chosen based on inputs to, outputs from, and leakages from the wires of the circuit in previous rounds) and the circuit's memory state is k_j. Let $X_{i,j}$ denote the value of the i-th wire of the circuit in this round conditioned on a fixing of the input a_j and memory state k_j. Then, $\mathcal{A}$ learns leakage $Z_{i,j}$ from $X_{i,j}$, where $Z_{i,j}$ is some (t,δ)-RevSD-noisy leakage from $X_{i,j}$ and the $Z_{i,j}$'s for $i \in [w]$ are conditionally independent given $(X_{1,j},\ldots,X_{w,j})$. We denote the output of $\mathcal{A}$ after such an interaction by $\mathsf{out}(\mathcal{A} \overset{\mathsf{noisy}}{\rightleftarrows} \Gamma(k))$.

c) Leakage-resilient implementations of circuits: Ishai, Sahai, and Wagner [23] studied compilers that turn an arbitrary stateful arithmetic circuit Γ into an equivalent circuit that is resilient to leakage.

Definition 16 (Leakage-resilient implementation of a circuit)**.** *We say that Γ' is an (r,τ,ε)-*threshold-probing-leakage-resilient implementation of Γ *if there exists an encryption function* Enc *such that the following two properties hold:*

- ***Equivalence:*** *For any inputs $(a_1,\ldots,a_r)$, outputs $(b_1,\ldots,b_r)$, and initial memory state k we have that*

$$\Pr[\Gamma(k,a_1,\ldots,a_r) = (b_1,\ldots,b_r)] = \Pr[\Gamma'(\mathsf{Enc}(k),a_1,\ldots,a_r) = (b_1,\ldots,b_r)];$$

- ***Leakage-resilience****: For any initial memory state k and (r,τ)-threshold probing circuit adversary $\mathcal{A}$ there exists a black-box circuit adversary $\mathcal{S}$ interacting with Γ such that*

$$\mathsf{SD}(\mathsf{out}(\mathcal{S} \overset{\mathsf{bb}}{\rightleftarrows} \Gamma(k)); \mathsf{out}(\mathcal{A} \overset{\mathsf{thres}}{\rightleftarrows} \Gamma'(\mathsf{Enc}(k)))) \le \varepsilon.$$

We can define an (r,p,ε)-random-probing-leakage-resilient implementation of Γ, an (r,t,δ,ε)-Uniform-RevSD-noisy-leakage-resilient implementation of Γ, and an (r,t,δ,ε)-RevSD-noisy-leakage-resilient implementation of Γ analogously by replacing $\mathsf{out}(\mathcal{A} \overset{\mathsf{thres}}{\rightleftarrows} \Gamma'(\mathsf{Enc}(k)))$ by $\mathsf{out}(\mathcal{A} \overset{\mathsf{rand}}{\rightleftarrows} \Gamma'(\mathsf{Enc}(k)))$, $\mathsf{out}(\mathcal{A} \overset{\mathsf{noisy-unif}}{\rightleftarrows} \Gamma'(\mathsf{Enc}(k)))$, or $\mathsf{out}(\mathcal{A} \overset{\mathsf{noisy}}{\rightleftarrows} \Gamma'(\mathsf{Enc}(k)))$, respectively, in the leakage-resilience condition above.

Ishai, Sahai, and Wagner [23] described and analyzed an efficient compiler that transforms an arbitrary stateful arithmetic circuit Γ into another stateful arithmetic circuit Γ' equivalent to Γ that is resilient to threshold probing leakage. A bit more precisely, this compiler replaces each gate of Γ by an appropriate *gadget* consisting of multiple gates. For an integer parameter $d > 0$ which controls the probing threshold, such a gadget contains at most $3.5d^2 + d$ gates, and this compiler yields the following theorem.

Theorem 11 (Circuits resilient to threshold probing leakage [23], see also [20, Section 5.3])**.** *Let Γ be an arbitrary stateful arithmetic circuit over $\mathbb{F}$ and fix an integer parameter $d > 0$. Then, there is a procedure running in time polynomial in $|\Gamma|$ and $|\mathbb{F}|$ that compiles Γ into an $(r, \tau = \lfloor \frac{d-1}{2} \rfloor \cdot |\Gamma|, \varepsilon = 0)$-threshold-probing-leakage-resilient implementation Γ' for any r, provided that the adversary does not probe each gadget (containing at most $3.5d^2 + d$ gates) more than $\lfloor \frac{d-1}{2} \rfloor$ times in each round.*

The next result, observed in [20], follows by combining Lemma 8 and Theorem 11. We include a proof for completeness, since [20] considers only the single-round setting and argues directly for noisy leakage.

Corollary 7 (Circuits resilient to random probing leakage, implicit in [20, Theorem 1])**.** *Let Γ be an arbitrary stateful arithmetic circuit over $\mathbb{F}$ and fix an integer parameter $d > 0$. Then, there is a procedure running in time polynomial in $|\Gamma|$ and $|\mathbb{F}|$ that compiles Γ into an (r,p,ε)-random probing leakage-resilient implementation Γ' for $p = \frac{1}{28d+8}$ and $\varepsilon = r|\Gamma| \cdot e^{-d/12}$.*

Proof. For a given $d > 0$, let Γ' be the threshold-probing-leakage-resilient implementation of Γ guaranteed by Theorem 11. Each of the $|\Gamma|$ gadgets of Γ' contains at most $2(3.5d^2 + d) = 7d^2 + 2d$ wires since each gate has fan-in at most 2.

For an arbitrary gadget g of Γ', define ℓ_g to be its number of wires. Define also $p_g = \frac{d}{4\ell_g}$ and $\tau_g = 2p_g\ell_g - 1$. Note

that $\tau_g = 2p_g\ell_g - 1 = \frac{d}{2} - 1 \leq \lfloor \frac{d-1}{2} \rfloor$. Furthermore, since $\ell_g \leq 7d^2 + 2d$, we have $p_g = \frac{d}{4\ell_g} \geq \frac{1}{28d+8} = p$. Therefore, we can perfectly simulate p-random probing leakage from X using p_g-random probing leakage from X since $p_g \geq p$.

Let $\mathcal{A}$ be an (r, p)-random probing circuit adversary for Γ'. Consider the $(r, \tau = \lfloor \frac{d-1}{2} \rfloor)$-threshold probing adversary $\mathcal{A}'$ for Γ' which applies the simulator of Lemma 8 instantiated with $\ell = \ell_g$ and $p = p_g$ to $(\tau_g \leq \tau)$-threshold probing leakage from each gadget g, which simulates p_g-random probing leakage from the wires of g, and uses this to perfectly simulate p-random probing leakage from the wires of g, as $p_g \geq p$. Then, $\mathcal{A}'$ emulates the behavior of $\mathcal{A}$ given the resulting simulated leakages.

By Lemma 8, $\mathcal{A}'$ perfectly simulates $\mathcal{A}$ unless some of the simulators abort. Since the simulator used in gadget g aborts with probability $e^{-p_g\ell_g/3} = e^{-d/12}$, a union bound over the $|\Gamma|$ gadgets of Γ' and the r rounds of interaction gives that the probability that at least one simulator aborts is at most $r|\Gamma| \cdot e^{-d/12}$, and so

$$\mathsf{SD}(\mathsf{out}(\mathcal{A}' \overset{\mathsf{thres}}{\rightleftarrows} \Gamma'(\mathsf{Enc}(k))); \mathsf{out}(\mathcal{A} \overset{\mathsf{rand}}{\rightleftarrows} \Gamma'(\mathsf{Enc}(k)))) \leq r|\Gamma| \cdot e^{-d/12}. \quad (21)$$

Furthermore, by Theorem 11 we also know that there exists a black-box adversary $\mathcal{S}$ such that

$$\mathsf{SD}(\mathsf{out}(\mathcal{S} \overset{\mathsf{bb}}{\rightleftarrows} \Gamma'(k)); \mathsf{out}(\mathcal{A}' \overset{\mathsf{thres}}{\rightleftarrows} \Gamma'(\mathsf{Enc}(k)))) = 0. \quad (22)$$

Combining Equations (21) and (22) with a triangle inequality yields the desired result. □

2) RevSD-Noisy Leakage-Resilient Circuit Compilers: We combine Corollary 7 with Lemmas 6 and 7 to obtain the following compilers for noisy-leakage-resilient circuits. These corollaries extend the main result of [20, Theorem 1], who considered only the case $t = 0$. They are only relevant when t is quite small.

Corollary 8. *Let Γ be an arbitrary stateful arithmetic circuit over a finite field $\mathbb{F}$ and fix an integer parameter $d > 0$. Then, there is a procedure running in time polynomial in $|\Gamma|$ and $|\mathbb{F}|$ that compiles Γ into an $(r, t, \delta, \varepsilon)$-Uniform-RevSD-noisy-leakage-resilient implementation Γ' for $\delta = \frac{2^t(28d+8)^{-1}-(2^t-1)}{|\mathbb{F}|}$ and $\varepsilon = r|\Gamma| \cdot e^{-d/12}$.*

Proof. By Corollary 7, we can obtain an $(r, p = \frac{1}{28d+8}, \varepsilon = r|\Gamma| \cdot e^{-d/12})$-random-probing-leakage-resilient implementation Γ' of Γ in time polynomial in $|\Gamma|$ and $|\mathbb{F}|$. Moreover, by the discussion in Section VIII-C1, Γ' is obtained by transforming each gate of Γ into a different gadget consisting of at most $\ell = 3.5d^2 + d$ gates. Since each gate has fan-in at most 2, there are at most $2\ell = 7d^2 + 2d$ wires in each gadget.

Let $\mathcal{A}$ be an (r, t, δ)-Uniform-RevSD-noisy circuit adversary for Γ'. This means that for the i-th wire value in the j-th round, $X_{i,j}$, the adversary $\mathcal{A}$ learns the leakage $Z_{i,j} = f_{i,j}(X_{i,j})$, where $f_{i,j}$ is a (t, δ)-RevSD-noisy leakage function from $U_{\mathbb{F}}$. By Lemma 6, we can perfectly simulate each $Z_{i,j}$ by $(1 - 2^{-t} + \delta 2^{-t}|\mathbb{F}| = p)$-random probing leakage from $X_{i,j}$, even given $X_{i,j}$.[14] Therefore, the (r, p)-random probing circuit adversary $\mathcal{A}'$ for Γ' which applies this simulator to the p-random probing leakage from each wire and then emulates the behavior of $\mathcal{A}$ given the outputs of the simulators satisfies

$$\mathsf{SD}(\mathsf{out}(\mathcal{A}' \overset{\mathsf{rand}}{\rightleftarrows} \Gamma'(\mathsf{Enc}(k))); \mathsf{out}(\mathcal{A} \overset{\mathsf{noisy-unif}}{\rightleftarrows} \Gamma'(\mathsf{Enc}(k)))) = 0. \quad (23)$$

Since Γ' is an (r, p, ε)-random-probing-leakage-resilient implementation of Γ, there exists a black-box adversary $\mathcal{S}$ such that

$$\mathsf{SD}(\mathsf{out}(\mathcal{S} \overset{\mathsf{bb}}{\rightleftarrows} \Gamma(k)); \mathsf{out}(\mathcal{A}' \overset{\mathsf{rand}}{\rightleftarrows} \Gamma'(\mathsf{Enc}(k)))) \leq \varepsilon. \quad (24)$$

Applying the triangle inequality to Equations (23) and (24) yields

$$\mathsf{SD}(\mathsf{out}(\mathcal{S} \overset{\mathsf{bb}}{\rightleftarrows} \Gamma(k)); \mathsf{out}(\mathcal{A} \overset{\mathsf{noisy-unif}}{\rightleftarrows} \Gamma'(\mathsf{Enc}(k)))) \leq \varepsilon,$$

as desired. □

Corollary 9. *Let Γ be an arbitrary stateful arithmetic circuit over $\mathbb{F}$ and fix an integer parameter $d > 0$. Then, there is a procedure running in time polynomial in $|\Gamma|$ and $|\mathbb{F}|$ that compiles Γ into an $(r, t, \delta, \varepsilon)$-RevSD-noisy-leakage-resilient implementation for any $t \in \left[0, \log\left(1 + \frac{1}{28d+7}\right)\right]$ and $\varepsilon = r|\Gamma|(e^{-d/12} + (7d^2 + 2d)2^{-t}\delta)$.*

Proof. By Corollary 7, we can obtain an $(r, p = \frac{1}{28d+8}, \varepsilon' = r|\Gamma| \cdot e^{-d/12})$-random-probing-leakage-resilient implementation Γ' of Γ in time polynomial in $|\Gamma|$ and $|\mathbb{F}|$, and the number of wires in Γ' is $w \leq (7d^2 + 2d)|\Gamma|$.

Let $\mathcal{A}$ be an (r, t, δ)-RevSD-noisy circuit probing adversary for Γ'. Fix a round $j \in [r]$ and suppose that $\mathcal{A}$ adaptively chose input a_j and that the circuit's memory state in that round is k_j. Let $X_{i,j}$ be the random variable corresponding to the value of the i-th wire of Γ' in the j-th round conditioned on circuit input a_j and memory state k_j. The corresponding leakage $Z_{i,j}$ learned by $\mathcal{A}$ from $X_{i,j}$ is (t, δ)-RevSD-noisy leakage from $X_{i,j}$, and the $Z_{i,j}$'s are conditionally independent given $X_j = (X_{i,j})_{i\in[w]}$. Consider the (r, p)-random probing circuit adversary $\mathcal{A}'$ for Γ' that in the j-th round perfectly simulates $(1 - 2^{-t})$-random probing leakage from $X_{i,j}$ given p-random probing leakage from $X_{i,j}$, which is possible since $p = \frac{1}{28d+8} \geq 1 - 2^{-t}$ when $t \leq \log\left(1 + \frac{1}{28d+7}\right)$, applies to the $(1-2^{-t})$-random probing leakage from $X_{i,j}$ the simulator $\mathsf{Sim}_{i,j}$ guaranteed by Lemma 7 instantiated with $Z_{i,j}$ and $X_{i,j}$, and then behaves like $\mathcal{A}$ based on the outputs $(Z'_{i,j})_{i\in[w]}$ of those simulators.

We argue that

$$\mathsf{SD}(\mathsf{out}(\mathcal{A}' \overset{\mathsf{rand}}{\rightleftarrows} \Gamma'(\mathsf{Enc}(k))); \mathsf{out}(\mathcal{A} \overset{\mathsf{noisy}}{\rightleftarrows} \Gamma'(\mathsf{Enc}(k)))) \leq r|\Gamma|(7d^2 + 2d) \cdot 2^{-t}\delta \quad (25)$$

via a hybrid argument over the r rounds. For $j \in [r]$, consider the hybrid experiment Hyb_j where we interact with Γ' as follows. First, in rounds $1, \ldots, j-1$ we learn (t, δ)-RevSD-noisy leakages from the wires and behave like $\mathcal{A}$. Then, in

[14] Lemma 6 implies that for any fixing $X_{i,j} = x_{i,j}$ we can perfectly simulate the leakage $f_{i,j}(x_{i,j})$ given p-random probing leakage from $x_{i,j}$.

rounds $j, \ldots, r$ we learn p-random probing leakage from the wires and behave like $\mathcal{A}'$. The output of Hyb_j consists of the adaptively chosen circuit inputs $a_1, \ldots, a_r$, the corresponding circuit outputs $b_1, \ldots, b_r$, the state of the circuit's memory in each round $k_1, \ldots, k_r$, and for rounds $j' \in [j-1]$ the true noisy leakages $(Z_{i,j'})_{i\in[w]}$ and for rounds $j' \geq j$ the simulated noisy leakages $(Z'_{i,j'})_{i\in[w]}$ as defined in the previous paragraph.

The initial hybrid Hyb_1 corresponds to the interaction between $\mathcal{A}'$ and Γ'. On the other hand, Hyb_{r+1} corresponds to the interaction between $\mathcal{A}$ and Γ'. Furthermore, the only thing that changes between Hyb_{j-1} and Hyb_j is how the noisy leakages from the wires are generated in the j-th round. These noisy leakages are conditionally independent of everything else conditioned on the wire values in the j-th round. In turn, these wire values depend on Hyb_{j-1} only through the j-th input a_j and the j-th memory state k_j. Let $X_{i,j}$ be random variables corresponding to the wire values of Γ' conditioned on an arbitrary fixing of the input $a_j = a_j^\star$ and memory $k_j = k_j^\star$, and write $Z_{i,j,x} = (Z_{i,j}|X_{i,j} = x)$ and $Z'_{i,j,x} = (Z'_{i,j}|X_{i,j} = x)$. Write also $X_j = (X_{i,j})_{i\in[w]}$, $Z_{j,x_j} = (Z_{i,j,x_{i,j}})_{i\in[w]}$, and $Z'_{j,x_j} = (Z'_{i,j,x_{i,j}})_{i\in[w]}$. For brevity, let $\mathsf{SD}(\mathsf{Hyb}_{j-1}; \mathsf{Hyb}_j | a_j^\star, k_j^\star)$ denote the statistical distance between Hyb_{j-1} and Hyb_j conditioned on $a_j = a_j^\star$ and $k_j = k_j^\star$. Then, this discussion shows that

$$\mathsf{SD}(\mathsf{Hyb}_{j-1}; \mathsf{Hyb}_j | a_j^\star, k_j^\star) \leq \mathsf{SD}(P_{X_j Z_{j,X_j}}; P_{X_j Z'_{j,X_j}}) \quad (26)$$

for any $a_j^\star$ and $k_j^\star$.

We now prove that

$$\mathsf{SD}(P_{X_j Z_{j,X_j}}; P_{X_j Z'_{j,X_j}}) \leq |\Gamma|(7d^2 + 2d) \cdot 2^{-t}\delta. \quad (27)$$

Combined with Equation (26), this implies that Hyb_{j-1} and Hyb_j are $(|\Gamma|(7d^2+2d)\cdot 2^{-t}\delta)$-close in statistical distance for every j, which then yields Equation (25) via r applications of the triangle inequality across the hybrids $\mathsf{Hyb}_1, \ldots, \mathsf{Hyb}_{r+1}$, since $\mathsf{out}(\mathcal{A}' \overset{\mathsf{rand}}{\rightleftarrows} \Gamma'(\mathsf{Enc}(k)))$ is a function of Hyb_1 and $\mathsf{out}(\mathcal{A} \overset{\mathsf{noisy}}{\rightleftarrows} \Gamma'(\mathsf{Enc}(k)))$ is a function of Hyb_{r+1}.

If we denote $\varepsilon_{i,j,x} = \mathsf{SD}(P_{Z_{i,j,x}}; P_{Z'_{i,j,x}})$, then we know that

$$\mathbb{E}_{x\sim P_{X_{i,j}}}[\varepsilon_{i,j,x}] = \mathsf{SD}(P_{X_{i,j}Z_{i,j}}; P_{X_{i,j}Z'_{i,j}}) \leq 2^{-t}\delta, \quad (28)$$

where the last inequality follows from Lemma 7. Consider an arbitrary fixing $X_j = x_j = (x_{i,j})_{i\in[w]}$. Since the $Z_{i,j,x_{i,j}}$'s are independent and so are the $Z'_{i,j,x_{i,j}}$'s, from w applications of the triangle inequality for statistical distance across the w wires of Γ' we obtain

$$\mathsf{SD}(P_{(Z_{i,j,x_{i,j}})_{i\in[w]}}; P_{(Z'_{i,j,x_{i,j}})_{i\in[w]}}) \leq \sum_{i=1}^{w} \varepsilon_{i,j,x_{i,j}}. \quad (29)$$

Therefore, we have

$$\begin{aligned}
\mathsf{SD}(P_{X_j Z_{j,X_j}}; P_{X_j Z'_{j,X_j}}) &= \mathbb{E}_{x_j\sim P_{X_j}}[\mathsf{SD}(P_{Z_{j,x_j}}; P_{Z'_{j,x_j}})] \\
&\leq \mathbb{E}_{x_j\sim P_{X_j}}\left[\sum_{i=1}^{w} \varepsilon_{i,j,x_{i,j}}\right] \\
&= \sum_{i=1}^{w} \mathbb{E}_{x_j\sim P_{X_j}}[\varepsilon_{i,j,x_{i,j}}] \\
&= \sum_{i=1}^{w} \mathbb{E}_{x_{i,j}\sim P_{X_{i,j}}}[\varepsilon_{i,j,x_{i,j}}] \\
&\leq \sum_{i=1}^{w} 2^{-t}\delta \\
&\leq |\Gamma|(7d^2 + 2d) \cdot 2^{-t}\delta, \quad (30)
\end{aligned}$$

where the first inequality uses Equation (29) and the second to last inequality uses Equation (28). This establishes Equation (27), and hence Equation (25) as discussed above.

Finally, since Γ' is an (r, p, ε')-random-probing-leakage-resilient implementation of Γ, there exists a black-box adversary $\mathcal{S}$ such that

$$\mathsf{SD}(\mathsf{out}(\mathcal{S} \overset{\mathsf{bb}}{\rightleftarrows} \Gamma(k)); \mathsf{out}(\mathcal{A}' \overset{\mathsf{rand}}{\rightleftarrows} \Gamma'(\mathsf{Enc}(k)))) \leq \varepsilon' = r|\Gamma| \cdot e^{-d/12}.$$

Combining this inequality with Equation (25) via the triangle inequality yields the desired result. □

IX. 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 IX-A. We then describe our evaluation settings in Section IX-B and use them to discuss reductions to bounded leakage and random probing in Section IX-C and Section IX-D, respectively.

A. 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. To this end, we may use Lemma 1. More precisely, this lemma implies that for any given $t \geq 0$ the corresponding minimal $\delta \in [0, 1]$ is obtained by computing $\delta = P_{XZ}(\mathcal{B}) - 2^t(P_X \otimes P_Z)(\mathcal{B})$, where $\mathcal{B} = \{(x, z) \mid P_{XZ}(x, z) > 2^t(P_X \otimes P_Z)(x, z)\}$. The same lemma can be used to compute the parameters of RevSD-noisy leakages analogously by just swapping the roles of the product and the joint distributions.

In many scenarios the process described above (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 this procedure also provides an upper bound for the δ parameter for Z as (t, δ)-GSD-noisy leakage from X by the choice $Q = P_Z$.

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.

B. Evaluation Settings

As a usual starting point, we considered the setting where leakages are written as the sum of the Hamming weight function HW and a Gaussian noise R [41]:

$$Z = \mathsf{HW}(X) + R, \tag{31}$$

where X is uniformly distributed on $\{0,1\}^n$.[15] In this setting, the amount of noise in the leakages is captured by the Signal-to-Noise Ratio (SNR) [34], defined as the ratio between the variance of the leakage function's deterministic part and the variance of the noise:

$$\mathrm{SNR} = \frac{\mathbb{V}(\mathsf{HW}(X))}{\mathbb{V}(R)} \approx \frac{n/4}{\mathbb{V}(R)}. \tag{32}$$

As a complement to textbook Hamming weight leakages, we considered noisy linear leakages where the Hamming weight function is replaced by a linear function d [31], 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:

$$\rho(\mathsf{d}(X), \mathsf{HW}(X)) = \frac{\mathbb{E}\big(\mathsf{d}(X) \cdot \mathsf{HW}(X)\big) - \mathbb{E}\big(\mathsf{d}(X)\big) \cdot \mathbb{E}\big(\mathsf{HW}(X)\big)}{\sqrt{\mathbb{V}\big(\mathsf{d}(X)\big) \cdot \mathbb{V}\big(\mathsf{HW}(X)\big)}},$$

which is a typical metric for univariate side-channel models evaluation – see for example [43, Section 2.2.2].

Concretely, in the noisy linear leakage case we first sampled many real vectors $(\gamma_1, \ldots, \gamma_n)$ at random (each γ_i sampled independently according to a gaussian distribution). We obtained vectors of coefficients $(\beta_1, \ldots, \beta_n)$ as the normalization of $(\gamma_1, \ldots, \gamma_n)$ so that for $\mathsf{d}(X) = \sum_{i=1}^{n} \beta_i \; X(i)$ we have variance $\mathbb{V}\big(\mathsf{d}(X)\big) = n/4$, which equals that of the Hamming weight leakage model. We then kept two sets of coefficients such that $\rho(\mathsf{d}(X), \mathsf{HW}(X)) \approx 0.5$ and $\rho(\mathsf{d}(X), \mathsf{HW}(X)) \approx 0.9$, respectively, and ran our experiments with the noisy linear leakages defined by these sets of coefficients. The model that correlates more with the Hamming weight one (with $\rho = 0.9$) is illustrated and compared to the Hamming weight one in Figure 1, for $n = 8$. The less correlated model (with $\rho = 0.5$) goes significantly beyond the deviations experimentally observed in [44].

[15] 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 [10], [15], [42]. But, of course, our theoretical results are applicable to non-uniform distributions as well.

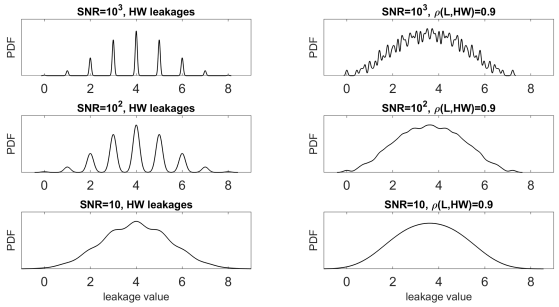


Fig. 1: Probability Density Function (PDF) of the leakage Z, for the noisy Hamming weight leakage function and a noisy linear leakage function, for different SNR values (with $n = 8$).

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

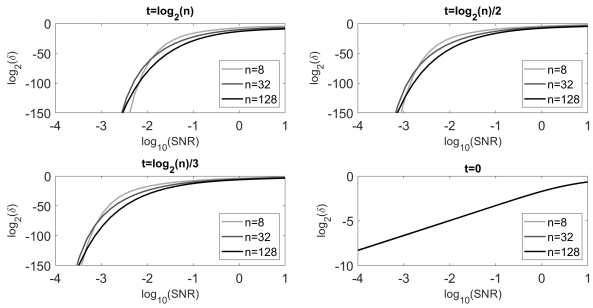


Fig. 2: Estimation of the δ parameter for SD-noisy leakages, in function of the SNR for Hamming weight leakages (with bit sizes n and an amount of bounded leakage t).

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 [24] 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 [25]. 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 IX-A), 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

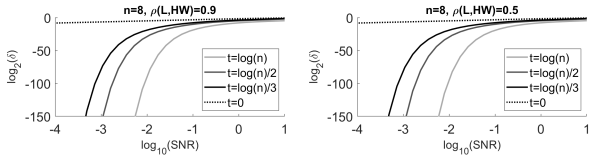


Fig. 3: Estimation of the δ parameter for SD-noisy leakages, in function of the SNR for linear leakages (with bit sizes n and an amount of bounded leakage t).

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 $n = 128$, $\mathrm{SNR} = 10^{-3}$ 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 II-A 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 investigated in Figure 2. Those linear models are interesting abstractions since 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.[16]

a) 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 [45]. 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 [46], 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).

[16]This time we only computed the δ parameter for $n = 8$ only, 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.

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, and further assuming the noise part of the leakage samples Z_i to be independent, 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 [47]. 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.

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.

D. 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).

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

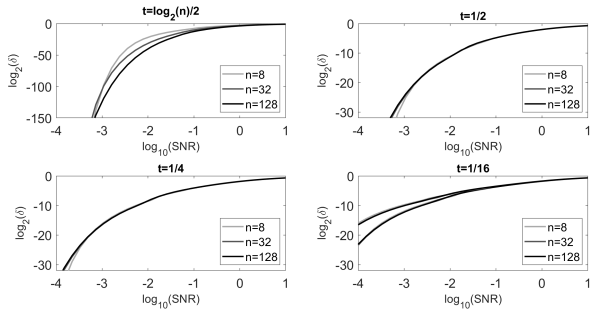


Fig. 4: Estimation of the δ parameter for RevSD-noisy leakages, in function of the SNR for Hamming weight leakages (with bit sizes n and an amount of bounded leakage t).

probing model, using the Average Relative Error (ARE) metric [19]. 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).

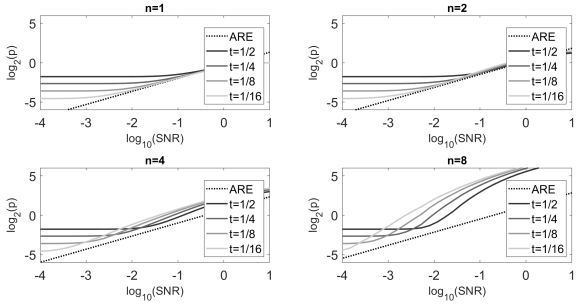


Fig. 5: Reductions to random probing using the ARE and RevSD metrics in function of the SNR for Hamming weight leakages and bit sizes $n = 1, 2, 4$ and 8.

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 [48], Brian, Dziembowski, and Faust 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 [49]. Second, in [50], Béguinot, Cheng, Guilley, and Rioul study a variant of the ARE metric (coined "Doeblin coefficients") that is better connected to the attacks' success.[17] 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

FXS is a research director of the Belgian Fund for Scientific Research (FNRS-F.R.S.). This work was funded in part by the European Union via the ERC project BRIDGE (101096871) and by DARPA. 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 those of the EU, the ERC or DARPA. Neither the EU nor the granting authorities can be held responsible for them. Distribution Statement "A" (Approved for Public Release, Distribution Unlimited).

DV is a member of the Gruppo Nazionale Calcolo Scientifico Istituto Nazionale di Alta Matematica (GNCS-INdAM). His research was partially supported by project BEAT—funded by Sapienza University of Rome—and by the AI-Sec Lab initiative between Sapienza University of Rome and the CYBER 4.0 Competence Center.

Part of this work was done while JR was visiting the Simons Institute for the Theory of Computing. We thank the anonymous CRYPTO 2024 reviewers for insightful questions that motivated Section VII-A.

References

[1] M. Obresmki, J. Ribeiro, L. Roy, F.-X. Standaert, and D. Venturi, "Improved reductions from noisy to bounded and probing leakages via hockey-stick divergences," in *Advances in Cryptology – CRYPTO 2024*, L. Reyzin and D. Stebila, Eds. Cham: Springer Nature Switzerland, 2024, pp. 461–491.

[2] P. C. Kocher, "Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems," in *CRYPTO '96*, N. Koblitz, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 1996, pp. 104–113.

[3] P. Kocher, J. Jaffe, and B. Jun, "Differential power analysis," in *Advances in Cryptology — CRYPTO' 99*, M. Wiener, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 1999, pp. 388–397.

[17] Doeblin coefficients actually appear in our proof of Lemma 6 as $\sum_z \pi(z)$.

[4] D. Agrawal, B. Archambeault, J. R. Rao, and P. Rohatgi, "The EM side-channel(s)," in *CHES 2002*, B. S. Kaliski, ç. K. Koç, and C. Paar, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003, pp. 29–45.

[5] D. A. Osvik, A. Shamir, and E. Tromer, "Cache attacks and countermeasures: The case of AES," in *CT-RSA 2006*, D. Pointcheval, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 1–20.

[6] A. Moradi, A. Barenghi, T. Kasper, and C. Paar, "On the vulnerability of FPGA bitstream encryption against power analysis attacks: extracting keys from xilinx Virtex-II FPGAs," in *CCS*. ACM, 2011, pp. 111–124.

[7] J. Balasch, B. Gierlichs, O. Reparaz, and I. Verbauwhede, "DPA, bitslicing and masking at 1 GHz," in *CHES*, ser. LNCS, vol. 9293. Springer, 2015, pp. 599–619.

[8] M. C. Martínez-Rodríguez, I. M. Delgado-Lozano, and B. B. Brumley, "SoK: Remote power analysis," in *ARES*. ACM, 2021, pp. 7:1–7:12.

[9] C. Liu, A. Chakraborty, N. Chawla, and N. Roggel, "Frequency throttling side-channel attack," in *CCS*. ACM, 2022, pp. 1977–1991.

[10] S. Dziembowski and K. Pietrzak, "Leakage-resilient cryptography," in *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, 2008, pp. 293–302.

[11] S. Chari, C. S. Jutla, J. R. Rao, and P. Rohatgi, "Towards sound approaches to counteract power-analysis attacks," in *Advances in Cryptology — CRYPTO' 99*, M. Wiener, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 1999, pp. 398–412.

[12] S. Faust, K. Pietrzak, and J. Schipper, "Practical leakage-resilient symmetric cryptography," in *CHES*, ser. Lecture Notes in Computer Science, vol. 7428. Springer, 2012, pp. 213–232.

[13] S. Belaïd, F. D. Santis, J. Heyszl, S. Mangard, M. Medwed, J. Schmidt, F. Standaert, and S. Tillich, "Towards fresh re-keying with leakage-resilient PRFs: cipher design principles and analysis," *J. Cryptogr. Eng.*, vol. 4, no. 3, pp. 157–171, 2014.

[14] C. Dobraunig, M. Eichlseder, S. Mangard, F. Mendel, B. Mennink, R. Primas, and T. Unterluggauer, "Isap v2.0," *IACR Trans. Symmetric Cryptol.*, vol. 2020, no. S1, pp. 390–416, 2020.

[15] D. Bellizia, O. Bronchain, G. Cassiers, V. Grosso, C. Guo, C. Momin, O. Pereira, T. Peters, and F. Standaert, "Mode-level vs. implementation-level physical security in symmetric cryptography - A practical guide through the leakage-resistance jungle," in *CRYPTO (1)*, ser. LNCS, vol. 12170. Springer, 2020, pp. 369–400.

[16] Y. T. Kalai and L. Reyzin, "A survey of leakage-resilient cryptography," in *Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali*. ACM, 2019, pp. 727–794.

[17] F.-X. Standaert, O. Pereira, and Y. Yu, "Leakage-resilient symmetric cryptography under empirically verifiable assumptions," in *Advances in Cryptology – CRYPTO 2013*, R. Canetti and J. A. Garay, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 335–352.

[18] F.-X. Standaert, T. G. Malkin, and M. Yung, "A unified framework for the analysis of side-channel key recovery attacks," in *Advances in Cryptology - EUROCRYPT 2009*, A. Joux, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 443–461.

[19] T. Prest, D. Goudarzi, A. Martinelli, and A. Passelègue, "Unifying leakage models on a Rényi day," in *Advances in Cryptology – CRYPTO 2019*, A. Boldyreva and D. Micciancio, Eds. Cham: Springer International Publishing, 2019, pp. 683–712.

[20] A. Duc, S. Dziembowski, and S. Faust, "Unifying leakage models: From probing attacks to noisy leakage," *J. Cryptol.*, vol. 32, no. 1, pp. 151–177, 2019.

[21] M. Naor and G. Segev, "Public-key cryptosystems resilient to key leakage," *SIAM Journal on Computing*, vol. 41, no. 4, pp. 772–814, 2012.

[22] E. Prouff and M. Rivain, "Masking against side-channel attacks: A formal security proof," in *Advances in Cryptology – EUROCRYPT 2013*, T. Johansson and P. Q. Nguyen, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 142–159.

[23] Y. Ishai, A. Sahai, and D. Wagner, "Private circuits: Securing hardware against probing attacks," in *Advances in Cryptology - CRYPTO 2003*, D. Boneh, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003, pp. 463–481.

[24] G. Brian, A. Faonio, M. Obremski, J. Ribeiro, M. Simkin, M. Skórski, and D. Venturi, "The mother of all leakages: How to simulate noisy leakages via bounded leakage (almost) for free," *IEEE Transactions on Information Theory*, vol. 68, no. 12, pp. 8197–8227, 2022, preliminary version in Eurocrypt 2021.

[25] A. Duc, S. Faust, and F.-X. Standaert, "Making masking security proofs concrete," in *Advances in Cryptology – EUROCRYPT 2015*, E. Oswald and M. Fischlin, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2015, pp. 401–429.

[26] J. Béguinot, W. Cheng, S. Guilley, Y. Liu, L. Masure, O. Rioul, and F. Standaert, "Removing the field size loss from Duc et al.'s conjectured bound for masked encodings," in *COSADE*, ser. LNCS, vol. 13979. Springer, 2023, pp. 86–104.

[27] I. Sason and S. Verdú, "f-divergence inequalities," *IEEE Transactions on Information Theory*, vol. 62, no. 11, pp. 5973–6006, 2016.

[28] G. Barthe and F. Olmedo, "Beyond differential privacy: Composition theorems and relational logic for f-divergences between probabilistic programs," in *Automata, Languages, and Programming*, F. V. Fomin, R. Freivalds, M. Kwiatkowska, and D. Peleg, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 49–60.

[29] C. Dwork, F. McSherry, K. Nissim, and A. D. Smith, "Calibrating noise to sensitivity in private data analysis," in *TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, ser. LNCS, S. Halevi and T. Rabin, Eds., vol. 3876. Springer, 2006, pp. 265–284.

[30] K. Pietrzak, "A leakage-resilient mode of operation," in *EUROCRYPT*, ser. Lecture Notes in Computer Science, vol. 5479. Springer, 2009, pp. 462–482.

[31] W. Schindler, K. Lemke, and C. Paar, "A stochastic model for differential side channel cryptanalysis," in *CHES*, ser. Lecture Notes in Computer Science, vol. 3659. Springer, 2005, pp. 30–46.

[32] C. Dwork, G. N. Rothblum, and S. Vadhan, "Boosting and differential privacy," in *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, 2010, pp. 51–60.

[33] Y. Dodis, K. Haralambiev, A. López-Alt, and D. Wichs, "Cryptography against continuous memory attacks," in *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, 2010, pp. 511–520.

[34] S. Mangard, "Hardware countermeasures against DPA ? A statistical analysis of their effectiveness," in *CT-RSA*, ser. Lecture Notes in Computer Science, vol. 2964. Springer, 2004, pp. 222–235.

[35] Y. Polyanskiy and Y. Wu, *Information Theory: From Coding to Learning*. Cambridge University Press, 2024, available at https://people.lids.mit.edu/yp/homepage/data/itbook-export.pdf.

[36] C. Dwork and J. Lei, "Differential privacy and robust statistics," in *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, ser. STOC '09. New York, NY, USA: Association for Computing Machinery, 2009, p. 371–380.

[37] T. Steinke, "Composition of differential privacy & privacy amplification by subsampling," *arXiv e-print*, Oct. 2022, https://arxiv.org/abs/2210.00597.

[38] S. Vadhan, "The complexity of differential privacy," *Tutorials on the Foundations of Cryptography: Dedicated to Oded Goldreich*, pp. 347–450, 2017.

[39] R. Vershynin, *High-Dimensional Probability*. Cambridge University Press, 2018.

[40] P. Klein and N. E. Young, "On the number of iterations for dantzig–wolfe optimization and packing-covering approximation algorithms," *SIAM Journal on Computing*, vol. 44, no. 4, pp. 1154–1172, 2015. [Online]. Available: https://doi.org/10.1137/12087222X

[41] S. Mangard, E. Oswald, and T. Popp, *Power analysis attacks - revealing the secrets of smart cards*. Springer, 2007.

[42] O. Bronchain, C. Momin, T. Peters, and F. Standaert, "Improved leakage-resistant authenticated encryption based on hardware AES coprocessors," *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, vol. 2021, no. 3, pp. 641–676, 2021.

[43] F.-X. Standaert, *Side-Channel Analysis and Leakage-Resistance*. Version 1.2, September 2024, available at https://perso.uclouvain.be/fstandae/SCALEbook_v12.pdf.

[44] C. Hoffmann, P. Méaux, C. Momin, Y. Rotella, F. Standaert, and B. Udvarhelyi, "Learning with physical rounding for linear and quadratic leakage functions," in *CRYPTO (3)*, ser. Lecture Notes in Computer Science, vol. 14083. Springer, 2023, pp. 410–439.

[45] B. Udvarhelyi, A. van Wassenhove, O. Bronchain, and F. Standaert, "On the security of off-the-shelf microcontrollers: Hardware is not enough," in *CARDIS*, ser. Lecture Notes in Computer Science, vol. 12609. Springer, 2020, pp. 103–118.

[46] S. Bhasin, S. Guilley, L. Sauvage, and J. Danger, "Unrolling cryptographic circuits: A simple countermeasure against side-channel attacks," in *CT-RSA*, ser. Lecture Notes in Computer Science, vol. 5985. Springer, 2010, pp. 195–207.

[47] Q. Guo, V. Grosso, F. Standaert, and O. Bronchain, "Modeling soft analytical side-channel attacks from a coding theory viewpoint," *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, vol. 2020, no. 4, pp. 209–238, 2020.

[48] G. Brian, S. Dziembowski, and S. Faust, "From random probing to noisy leakages without field-size dependence," in *Advances in Cryptology –*

EUROCRYPT 2024, M. Joye and G. Leander, Eds. Cham: Springer Nature Switzerland, 2024, pp. 345–374.

[49] S. Dziembowski, S. Faust, and M. Skorski, "Noisy leakage revisited," in *Advances in Cryptology - EUROCRYPT 2015*, E. Oswald and M. Fischlin, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2015, pp. 159–188.

[50] J. Béguinot, W. Cheng, S. Guilley, and O. Rioul, "Formal security proofs via Doeblin coefficients: Optimal side-channel factorization from noisy leakage to random probing," in *Advances in Cryptology – CRYPTO 2024*, L. Reyzin and D. Stebila, Eds. Cham: Springer Nature Switzerland, 2024, pp. 389–426.

Maciej Obremski received the Ph.D. degree from the University of Warsaw, Poland, in 2013, under the supervision of Stefan Dziembowski. He is currently a Research Assistant Professor at the School of Computing, National University of Singapore. Before joining NUS, he was a Post-Doctoral Researcher at Aarhus University, Denmark, with Ivan Damgård and Jesper Buus Nielsen, and the University of War- saw. His research interests mostly lie in information-theoretic cryptography, randomness and randomness extractors, and non-malleable codes and secret sharing.

João Ribeiro is an Assistant Professor in the Department of Mathematics of Instituto Superior Técnico, Universidade de Lisboa, Portugal and a researcher at Instituto de Telecomunicações, Lisboa, Portugal. Previously, he was an Assistant Professor in the Department of Computer Science of the NOVA School of Science and Technology, Caparica, Portugal, a researcher at NOVA LINCS, Caparica, Portugal, and a Post Doctoral Fellow in the Computer Science Department of Carnegie Mellon University, Pittsburgh, PA, USA. João received the Ph.D. degree in computing from Imperial College London, UK, in 2021. His research interests lie in coding theory, information theory, cryptography, and pseudorandomness.

Lawrence Roy received his Ph.D. degree in Computer Science from Oregon State University in 2022. From 2022 to 2026, he was a postdoctoral researcher at the Department of Computer Science, Aarhus University, and he is currently a postdoctoral researcher at IBM Research Zürich. His research is focused on secure multi-party computation and its building blocks, such as homomorphic secret sharing, garbled circuits, and leakage-resilience, as well as applications such as post-quantum signatures based on VOLE-in-the-head. He was a co-recipient of an Honorable Mention for Best Paper at CRYPTO 2021.

François-Xavier Standaert received the Electrical Engineering degree and PhD degree from the Universite catholique de Louvain, respectively in 2001 and 2004. In 2004-2005, he was a Fulbright visiting researcher at Columbia University and at the MIT Medialab. From 2005 to 2008, he was a post-doctoral researcher of the Belgian Fund for Scientific Research. Since 2008 (resp., 2017, 2024), he is associate researcher (resp. senior associate researcher, research director) of the Belgian Fund for Scientific Research (FNRS-F.R.S). Since 2013 (resp., 2018, 2026), he is associate professor (resp., professor, full professor) at UCLouvain. His research interests include cryptographic hardware and embedded systems, physical security issues including side-channel fault attacks, and the design and analysis of cryptographic primitives that can cope with physical attack vectors.

Daniele Venturi (Senior Member, IEEE) received the B.Sc. degree in electrical engineering from Roma Tre University and the M.Sc. degree in telecommunication engineering and the Ph.D. degree in information and communication engineering from the Sapienza University of Rome. He is currently a Full Professor at the Department of Computer Science, Sapienza University of Rome. Prior to joining Sapienza as a Professor, he was a Post-Doctoral Researcher at the Department of Computer Science at Aarhus University and Sapienza University. He was an Assistant Professor at the University of Trento. His research is focused on theoretical and applied cryptography at large, in particular information-theoretic cryptography, public-key cryptography, non-malleability, zero-knowledge, and multi-party computation. He was a co-recipient of the Best Paper Award at EUROCRYPT 2011.